Somatic evolution in human blood and colon

Henry Lee-Six
Trinity College
University of Cambridge

September, 2018

Dissertation submitted for the degree of Doctor of Philosophy

Summary

Somatic evolution in human blood and colon

Henry Lee-Six

All cancers were once normal cells. They became cancerous through the chance acquisition of particular somatic mutations that gave them a selective advantage over their neighbours. Thus, the mutations that initiate cancer occur in normal cells, and the normal clonal dynamics of the tissue determine a mutant cell's ability to establish a malignant clone; yet these remain poorly understood in humans. One tissue was selected for the exploration of each of these two facets of somatic evolution: blood for clonal dynamics; colon for mutational processes.

Blood presents an opportunity to study normal human clonal dynamics, as clones mix spatially and longitudinal samples can be taken. We isolated 140 single haematopoietic stem and progenitor cells from a healthy 59 year-old and grew them *in vitro* into colonies that were whole genome sequenced. Population genetics approaches were applied to this dataset, allowing us to elucidate for the first time the number of active haematopoietic stem cells, the rate at which clones grow and shrink, and the cellular output of stem cell clones.

Colonic epithelium is organised into crypts, at the base of which sit a small number of stem cells. All cells in a crypt ultimately share an ancestor in one stem cell that existed recently, and consequently share the mutations that were present in this ancestor. We exploited this natural clonal unit, isolating single colonic crypts through laser capture microdissection. 570 colonic crypts from 42 individuals were whole genome sequenced. We describe the burden and pattern of somatic mutations in these genomes and their variability across and within different people, identifying some mutational processes that are ubiquitous and others that are sporadic. Targeted sequencing of an additional 1,500 crypts allowed us to quantify the frequency of driver mutations in normal human colon.

Together, these two studies inform on the somatic evolution of normal tissues, describing new biology in human tissue homeostasis and providing a window into the processes that govern cancer incidence.

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except as specified in the text. It is not substantially the same as any that I have submitted, or is being concurrently submitted, for a degree or other qualification at the University of Cambridge or any other university. It does not exceed the prescribed limit of 60,000 words.

Description of contributions

This work was collaborative in nature. Here, I outline the contributions that others have made to work presented in this dissertation. I further draw attention to their contributions at relevant points throughout the text.

Results Chapter 1: Clonal dynamics of normal blood.

The general experimental design, as displayed in figure 1.1a, was formulated by Peter Campbell, David Kent, and Tony Green. The bone marrow aspirate was performed by Brian Huntly. Flow sorting of the bone marrow aspirate and culturing colonies was performed by the Kent lab. Subsequent peripheral blood draws and flow sorting of peripheral blood were performed by members of the Green and Kent labs.

All bioinformatic analyses were performed by me, with the following exceptions:

- Sebastian Grossman ran the Shearwater and HipSTR algorithms, and most iterations of the SCITE tree building algorithm.

- Robert Osborne wrote and ran an algorithm to tabulate reads from deep sequencing of peripheral blood.

- Peter Campbell wrote the mixed effects model to separate signal from noise in the targeted sequencing data from peripheral blood.

My contributions to the project were the following:

- Contributing to decisions about which peripheral blood draws should be sequenced.

- Devising the filters to clean the whole genome sequencing data and the method of calling mutation.

- Analyses of mutational signatures and embryology of blood, and the relationship between blood cells on the phylogeny.

- Conceiving the method of estimating stem cell numbers by performing a capture-recapture experiment of different peripheral blood samples.

- With substantial guidance from Peter Campbell and Kevin Dawson, writing the approximate Bayesian computation for estimating the number of active haematopoietic stem cells.

- Comparing the clonal contributions to granulocytes and lymphocytes.

Results Chapter 2: Mutational landscape of normal colon.

The idea of using laser capture microdissection to isolate normal crypts for sequencing was Mike Stratton's. Collaborators kindly provided frozen biopsies of colonic tissue. The organoid component of this study was set up by Sam Behjati and Sophie Roerink in collaboration with the Clevers lab, and the organoids were derived by the Clevers lab. A protocol to sequence small amounts of laser capture microdissected material was devised by Peter Ellis.

My contributions to the project were the following:

- Designing, through discussions with Mike Stratton, the experiment in terms of the number of crypts from each patient on which to perform targeted and whole genome sequencing.
- Setting up collaborations to obtain tissue for microdissection.
- Devising a method for fixing, staining, sectioning, and dissecting crypts.
- Microdissecting all the crypts.
- All analyses of the crypts and organoids, including, principally:
  - Calling mutations.
  - Driver analysis.
  - Mutational signature analysis.
  - Comparisons between normal and cancer.

For both projects, conclusions are either entirely my own or have arisen through frequent discussions over three years with many lab members, but principally Peter Campbell and Mike Stratton.

Acknowledgements

With warm and sincere gratitude, I thank my doctoral supervisors Peter Campbell and Mike Stratton. They gave me incredibly exciting projects to work on and put everything in place that I might carry them out. Together, we pored over novel findings, and spent, cumulatively, hundreds of hours discussing their implications and the next steps. Observing and imitating their distinctive styles of engaging with data was the best teaching I have ever had. I thank them for allowing me both to explore my own ideas and for guiding me back in the right direction when I had lost my way.

This thesis would not have been the same without the hundreds of conversations, in corridors and over coffee as much as in formal meetings, with many people. Inigo Martincorena was, in my first year, my 'go-to post doc' for understanding the deeper implications of what we were attempting to do; later, as a principal investigator, he was always willing to provide thoughtful and thought-provoking explanations for our discoveries. David Kent helped me navigate the complex literature of haematopoiesis when I was at sea, and so appreciate the broader implications of our findings. Rob Osborne was a friendly source of encouragement and level-headed advice who gave me a structure to follow when nothing was working in the lab; without him I may have given up on the approach that was ultimately fruitful. Kevin Dawson and I spent many hours on the blue sofa seemingly going around in circles over the mathematics of determining the number of haematopoietic stem cells; in retrospect, we were spiralling upwards to a conclusion. I learnt from Moritz Gerstung in our joint weekly statistics meetings (as well as from Peter, Inigo, and Kevin) how to look at statistics critically and constructively. I thank Grace Collord for her unfailing encouragement and advice (and cake), but especially for painstakingly reading this thesis; her comments improved it substantially. All were also good companions.

I am indebted to all those with whom I had the pleasure of working over the last three years. Their enthusiasm made me look forward to going into the lab. In no particular order, I thank: Seb, whose rational approach was invaluable to the construction of the stem cell phylogeny; Luiza, whose impetus drove the laser capture microdissection protocol; Simon and Heather, both good sounding boards and running companions; Sophie, who generously allowed

Contents

Discussion 197

Bibliography 205

List of abbreviations and acronyms 225

Appendix A Myeloid leukaemia and clonal haematopoiesis mutation hotspot loci included in bait-set for targeted sequencing of peripheral blood 226

Appendix B Colorectal cancer genes included in bait-set for targeted sequencing of colonic crypts 229

Appendix C Characteristics of informative crypts

232

INTRODUCTION

1. Somatic evolution

One general law, leading to the advancement of all organic beings, namely, multiply, vary, let the strongest live and the weakest die.

(Charles Darwin, *On the Origin of Species*).

Darwin's theory applies not only to whole organisms but also to the cells that constitute them; this is what is meant by 'somatic evolution'. Variation and selection occur within the body and so evolution must follow. Cancer represents the archetype of this process, although the principles apply beyond the disease, and it is through the study of cancer genetics that our and other laboratories have approached the concept of somatic evolution more broadly. Much of this dissertation is framed in terms of cancer, therefore, and, for these reasons, I begin with a discussion of cancer genomics.

2. Cancer genomics

2.a. Cancer is a disease of the genome

Cancer has been viewed as a disorder of cells since the nineteenth century. In 1840, Langenbeck stated 'every single carcinoma cell must now appear as an organism endowed with life-force and developmental ability' (cited in Bignold 2006). A series of discoveries over the following 150 years demonstrated that mutations were the molecular basis of cancer. Chromosomes were discovered in 1870, and the observation that the sum of the numbers of chromosomes in sperm and egg equalled that in the zygote suggested that they contributed the individuality of the tissue. In 1890, von Hansemann described abnormal mitoses in cancer cells and hypothesised that an imbalance in chromosome numbers might be important (von Hansemann 1890). A decade later, Boveri observed that sea urchins with multipolar mitoses (due to two spermatozoa fertilising an egg) had developmental defects and went on to propose that cancer

might be due to an 'abnormal chromosome complex' (Boveri 1914). He further postulated the existence of tumour suppressor genes and oncogenes:

> [...] in every normal cell there is a specific arrangement for inhibiting, which allows the process of division to begin only when the inhibition has been overcome by a special stimulus. To assume the presence of definite chromosomes which inhibit division, would harmonize best with my fundamental idea [...] Cells of tumours with unlimited growth would arise if those 'inhibiting chromosomes' were eliminated [...] On the other hand, the assumption of the existence of chromosomes which promote division, might satisfy this postulate [...] If three or four such chromosomes meet, the whole number of chromosomes being otherwise normal, then the tendency to rapid proliferation would arise.
>
> (Boveri 1914, cited in Knudson 2001)

It was observed that mutagenic exposures (first ionising radiation (Muller 1927) and then chemicals such as those in coal tar (reviewed in Loeb and Harris 2008)) were frequently carcinogenic. The discovery of the Philadelphia chromosome (Nowell and Hungerford 1960) showed that a specific chromosomal alteration could be associated with cancer, providing evidence for the concept of driver mutations. In 1971, Knudson fitted Poisson models to the number of tumours found in hereditary and sporadic cases of retinoblastoma and noted that the data were consistent with the need for two mutations occurring at the same rate to develop the disease (of which one was already inherited in familial cases). Identification of the *RB1* gene suggested that these were two inactivating hits in the same gene rather than mutations in two separate genes (Friend et al. 1986). This was the first tumour suppressor gene. In the same period, revival of work from the beginning of the century on tumour viruses indicated that specific genes carried by viruses, homologous to – but different from – host genes, were capable of causing cancer (Martin 1970; Stehelin et al. 1976). Similarly, it was noted that moving DNA from cancer cells into a phenotypically normal cell line could transform it, indicating the existence of dominantly-acting oncogenes (Krontiris and Cooper 1981; Shih et al. 1981). Isolation of the causal gene (*HRAS*) from a bladder carcinoma cell line and the identification of the particular base change with oncogenic activity resulted in the description of the first oncogenic point mutation (Reddy et al. 1982). In the decades since, through first targeted and later systematic analyses of thousands of genes across many different cancer histologies the number of cancer driver genes has risen to a few hundred

2

(and thus 1-3% of all genes), although this is still a matter of debate (Stratton et al. 2009; Vogelstein et al. 2013; Forbes et al. 2017). The number of driver mutations per tumour is also contentious. This is discussed in the context of colorectal adenocarcinoma in Results Chapter 2.

2.b. Mutational processes

Genome sequencing has revealed that only a small fraction of mutations in a tumour are driver mutations. The vast majority are passenger mutations, exerting no effect (or, at least, one that is too weak to be selected over the lifespan of a cancer) on the fitness of the cells that carry them. There is substantial variability between and within different cancer types in mutation burden, but most adult cancers have thousands to tens of thousands of single base substitutions, tens to hundreds of dinucleotide substitutions (where two adjacent bases both mutate), hundreds to thousands of small insertions and deletions (indels), and up to a few hundred structural variants (Alexandrov et al. 2018; Li et al 2017). In addition, there may be insertion of viral sequences or reactivation of endogenous mobile elements, and most cancers display changes in their epigenome relative to their cells of origin, but this is beyond the scope of the present dissertation and will not be discussed here.

Although most are innocuous, these scattered mutations nonetheless represent the substrate for evolution. Studying their properties tells us about the processes that cause the mutations that do exert a functional effect and provides a window into the basis of the 'variation' component of somatic evolution. The vast numbers of somatic mutations yielded by large scale sequencing efforts permits a systematic analysis of the mutational processes that cause cancer. To date, approximately 85 million somatic mutations of all classes have been described across thousands of cancers (Alexandrov et al. 2018; Li et al. 2017), from which signatures of mutational processes (or mutational signatures) can be extracted. Here, I explain mutational signatures in the context of single base substitutions, but the same methodology can be applied to all mutation types.

The analysis of mutational signatures rests upon the idea that different mutational processes mutate the genome in distinct ways, such that their activity leaves a characteristic 'signature' in the genome. To begin to describe signatures, we must first classify mutations. For example, substitutions can be grouped according the identity of the base change, presented with the

pyrimidine of the mutated Watson-Crick base pair as the original base, into six classes: C>A, C>G, C>T, T>A, T>C, and T>G substitutions. It has long been recognised that particular mutational processes have a predilection for some base changes over others. For example, *TP53* mutations in skin cancers are often C>T, while *TP53* in lung cancers from smokers – but not non-smokers – has a large number of C>A mutations (Alexandrov and Stratton 2014). These observations have a molecular basis. C>T mutations in skin cancers are thought to be due to a combination of, firstly, ultraviolet irradiation inducing cyclobutane-pyrimidine dimers in which the C is unstable and spontaneously deaminates to a uracil, opposite which an A is added in replication, and secondly, error-prone translesion polymerase bypass of ultraviolet-induced photolesions (Ikehata and Ono 2011). C>A mutations associated with smoking are thought to be due to bulky adducts such as benzo(a)pyrene metabolites binding guanines. The damaged guanine is frequently complemented by an adenine when replicated by the translesion polymerase POLH (Christmann et al. 2016).

Mutational processes with a reasonably well-characterised origin and molecular explanation, such as ultraviolet irradiation and smoking, were the exception rather than the rule before the advent of mutational signature analysis, and most somatic mutations in cancer genomes were of unknown cause. Even now, the aetiology of approximately half of mutational signatures is unknown (Petljak and Alexandrov 2016). It is not the case that the genome of one cancer represents one mutational process: most cancer genomes seem to have been sculpted by a number of co-occurring processes, which will have overlapping features (for example, there may be multiple processes that all cause C>T mutations). We do not know *a priori* the pattern of mutations associated with each process, nor how much each contributes to a cancer genome. The only information we have is the counts of mutations in every category that we have defined in each cancer. Remarkably, though, by comparing the mutational patterns across thousands of cancer samples one can deconvolute these mutations into their constituent signatures.

Humans can deconvolute simple cases intuitively. Imagine that three tumours have been sequenced, and the mutations in each tabulated. If the mutations in tumour 1 are 50% C>A, 30% C>G, and 20% C>T, tumour 2 only has C>A mutations, and tumour 3 has 60% C>G and 40% C>T mutations, we would naturally explain this with two mutational processes: process 1 is C>A and process 2 is C>G and C>T, and each process accounts for half of the mutations in tumour 1. For cancer genomes, which are far more complicated, computational methods are required. The most widely-applied method to date is non-negative matrix factorisation (NNMF), used to perform

the first comprehensive and field-defining analysis of mutational signatures (Alexandrov et al. 2013), but many others along similar lines have arisen since (EMu (Fischer et al. 2013), signeR (Rosales et al. 2016), SignatureAnalyzer (Kim et al. 2016), HDP (Roberts 2018)). In this thesis, NNMF and a hierarchical Dirichlet process (HDP) are both used, and a brief explanation of their inner workings is supplied in the methods section (Methods, section 12).

For greater resolution of the distinction between different mutational processes, mutations can be divided into a larger number of biologically meaningful classes. For the analysis of substitutions, the base change is further subcategorised by the bases 3' and 5' of the mutated base, considering, for example, C>T in an ACG (the mutated base is underlined) context differently from C>T in ACC. This allows, for example, discrimination between the C>T mutations in an NCG context associated with the hydrolytic deamination of 5-methylcytosine and the C>T mutations in an NCC context associated with ultraviolet irradiation and thus the separation of these two processes. Depending on the richness of the dataset and the behaviour of mutational processes (it may be, for example, that widening the context beyond a certain point adds no extra information), finer categories may be defined by integrating topological features of the genome such as transcription strand bias or positioning relative to histones.

So far, using a trinucleotide categorisation of substitutions, an analysis of 4,645 whole genomes and 19,184 exomes has resulted in the description of 49 single base substitution signatures, as well as 11 doublet base substitution signatures and 17 indel signatures (Alexandrov et al. 2018). Similarly, nine signatures of structural variation have been characterised (Li et al. 2017). Some signatures are of extrinsic origin (such as the one due to smoking) and some are of intrinsic origin (such as the spontaneous deamination of 5-methylcytosine). Particular signatures will be discussed at relevant points in this thesis.

It should be noted that mutational signatures represent sets of correlated features that need not be a result of a single molecular process. If two separate carcinogens that left different imprints on the genome were always present at exactly the same proportion in tobacco, their combined effect would always be extracted as one signature. They could only be separated if samples were found in which one process had predominated over the other. We may, therefore, expect some signatures to split into two as more samples are sequenced, which has indeed occurred when updating the TCGA analysis (Alexandrov et al. 2013) to the PCAWG analysis (Alexandrov et al.

2018). In due course, signatures should stabilise and their experimental validation (by, for example, exposing cell lines to putative carcinogens) will provide a final and robust set.

The exciting feature of mutational signatures is that they are a link between epidemiological observations and our molecular understanding of cancer. Their systematic exploration over time and across different tissues, people, and states of health, will be central to improving models of cancer and other diseases. Hopefully, mutational signatures will allow us to discover preventable exposures and ways to intervene in diseases of somatic evolution. For example, the mutational signature of aristolochic acid has been found in 47% of liver cancers in China, indicating that this exposure is likely to be a cause of significant morbidity (Ng et al. 2017).

The cancer genome has – in the decade since the first cancer genomes were sequenced (Ley et al. 2008, Pleasance et al. 2010a, Pleasance et al. 2010b) – largely been characterised (Campbell et al. 2017, Alexandrov et al. 2018, Li et al. 2017, Sabarinathan et al. 2017, Dentro et al. 2018, Gerstung et al. 2017), and bulk sequencing of additional common tumours is unlikely to change our understanding of it radically. This is not to say that all questions have been resolved – far from it – but perhaps more is to be gained by probing deeply into cancer biology through other means. This involves extending our understanding of the life history of cancer and the forces that have shaped its evolution.

2.c. Cancer as a multi-step branching evolutionary process

Foulds made the case in 1958 that cancer was the end result of a series of qualitatively different changes, based on the concepts of 'initiation' and 'promotion' that came out of rabbit skin carcinogenesis and the fact that mammary hyperplasia often preceded frank carcinomas in mouse models of breast cancer (Foulds 1958). By the 1970s, the clonal origin of tumours was widely recognised (Nowell 1976) based on three observations: all cells from tumours had the same karyotype (Sandberg and Hossfeld 1970); G6PD heterozygote women only expressed one allele per tumour (the gene is X-linked, and so in a given cell one copy is randomly inactivated) (Fialkow 1974); and plasma cell cancers only produced one immunoglobulin (discussed in Nowell 1976). Cairns related the cellular turnover of normal tissues to their acquisition of mutations, and discussed cancers as a product of natural selection, with sequential clonal expansions each driven

by a mutation (Cairns 1975). Nowell made the same argument, but explicitly added the component of selection, removing the least fit cells (Nowell 1976). Importantly, Nowell stated the capability of 'continued variation so long as the tumour persists', making the point that evolution has no endpoint. In this section, I attempt to cover succinctly the evidence for the multi-step nature and branching evolution of cancer, a brief discussion of mutation rates, and an acknowledgement of the importance of microenvironmental selection pressures. I will focus on the earliest stages of cancer and will not discuss metastasis or the response to therapy.

2.c.i. Multi-step tumorigenesis

Early evidence of cancer's multi-step nature came from the lag between carcinogenic exposures and the development of cancers, which suggested that other events must occur in the meantime, and modelling of age-incidence curves which indicated a number of rate-limiting events (Armitage and Doll 1954, Nordling 1953). Incidence modelling of colorectal cancer is discussed in Results Chapter 2. In the 1980s and 1990s, these steps were tied to a series of histopathological and molecular changes: the adenoma-carcinoma sequence. It was known that most adenomas arose from a single crypt (Ponder and Wilkinson 1986). Histopathology had shown that most carcinomas grew out of adenomas (Sugarbaker 1985), and their shared origin was confirmed by finding the same *RAS* mutation in both the adenoma and its linked cancer. The progression was further strengthened by epidemiological observations: for example, people whose adenomas were not removed were at greater risk of developing cancer. Once the histological progression was established it was possible to investigate its molecular basis. Fearon and Vogelstein's review of this work is the bedrock of our understanding of multi-step carcinogenesis (Fearon and Vogelstein 1990). Early adenomas, intermediate adenomas, late adenomas, and carcinomas were all sequenced, and the number of drivers in each tabulated. More drivers were found in more progressed lesions, with four to five (of the drivers known at the time and assayed) found in most frank carcinomas. Particular drivers tended to be associated with particular stages. For example, *RAS* mutations were found in the same proportion of intermediate adenomas as carcinomas, which suggested that *RAS* mutations' main effect was to reach the intermediate adenoma stage without driving further progression, while loss of 17p (containing *TP53)* was enriched in carcinomas. This

led Fearon and Vogelstein to propose the famous sequence of: 5q (containing *APC*) loss driving the formation of an early adenoma; *KRAS* mutations driving that of an intermediate adenoma; 18q (containing *SMAD4*) loss driving that of a late adenoma; and 17p (*TP53*) loss driving that of a carcinoma. The authors stress that it is the total sum of alterations that matters rather than their precise order, since this ordering of mutations is common but far from universal. Additionally, that certain mutations are found in more progressed lesions does not necessarily imply that they occurred late: *TP53* would also appear mostly in carcinomas if it were, in fact, the first mutation but encouraged extremely rapid progression to cancer such that very little time was spent as an intermediate lesion. This model has been extrapolated to most other cancer types and a similar multi-step progression has been observed. For example, in the progression to oesophageal adenocarcinoma, *TP53* mutations tend not to occur before high grade dysplasia and *SMAD4* mutations are not commonly found before full-blown cancer (Weaver et al. 2014).

Pre-malignancy in blood warrants some discussion, since the clonal dynamics of blood make up half of the present dissertation. In childhood acute lymphoblastic leukaemia, the observation that monozygotic twins can both develop the disease years after birth, but their leukaemias share an initiating clonotypic gene fusion that must have occurred prenatally, indicates that extra events are required to trigger the conversion to malignancy (Greaves et al. 2003, Ford et al. 1993, Greaves 2018). In addition, many children harbour the initiating fusion but never develop leukaemia (Greaves 2018). In adult blood, the myeloproliferative neoplasms typically only bear a small number of known driver mutations per neoplasm (two thirds only have one known driver), and acquisition of certain additional driver mutations is associated with conversion to an acute leukaemia (Nangalia and Green 2018). Intriguingly, the order in which these first few mutations occur influences both the clonal dynamics of the neoplasm and its clinical manifestation (Ortmann et al. 2015). By assaying single cell-derived colonies from patients with myeloproliferative neoplasms for common driver mutations it has been shown that if *JAK2* mutations occur before *TET2* mutations, most sampled haematopoietic stem and progenitor cells had both mutations, whereas if *TET2* occurred before *JAK2*, large numbers of cells with only the *TET2* mutation were found (Ortmann et al. 2015). This indicates that *TET2* mutation promotes a more rapid clonal expansion than *JAK2* mutation does.

More recently, much excitement has been generated by the discovery of known myeloid driver mutations in the blood of adults with normal blood counts and the fact that the bearers of

these clones are at greater risk of cancer. For such mutations to be detectable they need to be in a reasonable proportion of blood cells, indicating a clonal expansion, and so this phenomenon has been termed 'clonal haematopoiesis'. Skewing of the proportion of cells with each X chromosome inactivated in women without cancer had been observed for a long time and ascribed to various causes, but in 2012 it was found that some of these women had *TET2* mutations (Busque et al. 2012). Of seven women followed over five years, one developed essential thrombocythaemia. A series of articles in 2014 based on the targeted sequencing of thousands of normal blood samples demonstrated that clonal haematopoiesis was widespread, with mutations in multiple leukaemia and lymphoma-related genes reported (Xie et al. 2014, Jaiswal et al. 2014, Genovese et al. 2014). The proportion of people in whom clones were detected increased with age. For example, by ultra-deep sequencing of 15 hotspots, clonal haematopoiesis was detected in 1% of patients under 60 and 20% of patients aged over 90 (Mckerrell et al. 2015). Importantly, the presence of mutations in driver genes was associated with a hazard ratio of 11.1 (95% CI 3.9-32.6) of developing a leukaemia, indicating that the mutant clones might be pre-malignant (Jaiswal 2014). The definition of clonal haematopoiesis is somewhat in flux, however. Firstly, the number of clones detected depends on the sensitivity of the detection methodology: when mutations could be accurately called down to a frequency of one in 10,000 cells, clonal haematopoiesis could be detected in 19 out of 20 healthy people aged between 50 and 60 (Young et al. 2016). Secondly, clones have been found in which no known driver mutations were detected (Zink et al. 2017). It is therefore unclear whether clonal haematopoiesis represents a qualitatively distinct stage between normal haematopoiesis and cancer, or if the clones detected are merely a result of normal neutral drift that results occasionally in clones of a detectable size. We currently have a very limited understanding of the range of blood clone sizes in healthy humans. It is to be hoped that this will be resolved by studying further the clonal dynamics of normal human haematopoiesis.

2.c.ii. Branching evolution

Branching evolution implies that complete clonal sweeps through a tumour are infrequent. Experimentally, this means that if different samples are taken from a tumour, each will have a large number of private mutations. This can be depicted as a phylogeny with long private branches

leading from the samples to their coalescence into the trunk of shared, clonal mutations. This is frequently observed in cancers (Yates and Campbell 2012). The length of private branches is determined by a number of factors. Branches are lengthened by restricted movement within the founding clone preventing clonal sweeps (the most extreme example being that of monozygotic twins with concordant clonal acute lymphoblastic leukaemia), by an increased mutation rate, and by decreased competition for resources (such as space) between tumour cells.

In the last few years, much has been made of a "Big Bang" model of colorectal cancer (Sottoriva et al. 2015). The same sub-clonal copy number changes were found in different parts of colorectal carcinomas. In the absence of cell migration (which is perhaps a strong assumption given that the ability to migrate is a hallmark of a carcinoma), this would be consistent with the emergence of multiple subclones and a lack of complete selective sweeps since. Subsequent work showed that the distribution of clone sizes across many tumours, determined based on the allele fraction of mutations, was not significantly different from what one would expect in the absence of selection (Williams et al. 2016). Conversely, it has been shown that a similar distribution of allele fractions can also be observed in the presence of weak subclonal selection (Tarabichi et al. 2017). The authors of the latter study conclude that allele fractions alone are not sufficient to distinguish between neutrality and weak or occasional subclonal selection. Indeed, incomplete clonal sweeps within a tumour need not equate with neutrality: carcinomas have long been observed to grow out of adenomas and this has never been interpreted as neutral evolution. Orthogonal genetic approaches, such as dNdS (a method that compares the proportion of nonsynonymous and synonymous mutations in order to detect selection), have shown the presence of weak subclonal positive selection (Tarabichi et al. 2017). Furthermore, there is no reason that selection should stop operating during tumour evolution. Nonetheless, the similarity between observed allele fractions and simulations under a neutral model indicates that neutral drift is likely to play a large part in determining clone sizes in many tumours.

2.d. Selective pressures

The interactions between the tumour and its microenvironment are multifarious (Greaves and Maley 2012). I cannot hope to do them justice in this short section, but aim merely to

acknowledge their importance. The role of tissue architecture in regulating clonal competition is discussed in more detail below. Here, I provide a few brief vignettes on the effect of selection pressures in the earliest stages of cancer development.

Immune surveillance provides a well-established selection pressure. In aplastic anaemia, CD8 cytotoxic T lymphocytes attack one's own haematopoietic stem cells, reacting against antigens on their surface. In 50% of people with this condition, a cell that acquires a mutation in the X-linked *PIGA* gene that results in the loss of expression of glyosylphosphatidylinositol anchors (such that surface antigens can no longer be tethered to the cell surface) escapes immune attack and proliferates to form a clone. This is paroxysmal nocturnal haemoglobinuria, which is associated with a low risk of transformation to leukaemia (Hoffbrand et al. 2011). Indeed, even in the absence of a well-described pathology such as paroxysmal nocturnal haemoglobinuria, 15% of people with aplastic anaemia will develop myelodysplastic syndrome or acute myeloid leukaemia (Yoshizato et al. 2015), indicating that these microenvironmental changes create the conditions for a mutant cell to outcompete its wild-type neighbours. Even when we do not understand the selection pressure, there is evidence that a change in environment affects clonal expansions. For example, in clonal haematopoiesis, mutations in spliceosome components suddenly increase in frequency in the elderly (Mckerrell et al. 2015). This may be due to a change in selection pressures associated with old age (Mckerrell and Vassiliou, 2015).

Inflammation may provide a selective pressure that favours the outgrowth of certain clones. Dominant negative *P53* mutations in mouse colonic crypts conferred no advantage in a physiological setting but did in a mouse model of colitis (Vermeulen et al. 2013). This might be – at least in part – because of a loss of sensitivity to cellular damage which allows the mutant cells to continue to proliferate while their wild-type counterparts arrest for DNA repair (Breivik 2001). *TP53* mutations are enriched in colitis-associated neoplasia, including in non-dysplastic crypts that neighbour lesions (Leedham et al. 2009), and indeed, ulcerative colitis is associated with an increased risk of colorectal cancer. Similarly, in a model of paediatric acute lymphoblastic leukaemia, B cell progenitors with the *ETV6-RUNX1* fusion (a common initiating lesion) proliferate more slowly than their wild-type counterparts in homeostasis but gain a competitive advantage in the presence of the cytokine TGFβ (Ford et al. 2009).

One experiment produced such a peculiar result that it deserves mention. *Ctnnb1* mutations induced in small populations of cells in mouse skin were enveloped by wild-type skin, and after a

time, expelled from the tissue (Brown et al. 2017). When the proliferation of wild-type cells was inhibited, mutant cells were not expelled and were able to form large cysts. This indicates negative selection that is driven by wild-type cells. In a study of sequencing human skin, no negative selection was detected (Martincorena et al. 2015). If microenvironmental selection pressures arrested the growth of mutant cells, rather than eliminating them, however, no negative selection would be observed. The authors of the latter study did note that the range of mutant clone sizes detected in human skin was unexpectedly small, indicating that there may be some constraint on clonal expansion.

Clearly, there is a world of complexities of selection pressures in normal tissues that remains to be explored. Some have gone so far as to claim that the age incidence curves of cancer are due to changes in the microenvironment with age rather than the need to acquire multiple driver mutations (DeGregori 2017).

2.e. Mutation rates in cancer

There has been much debate over whether the mutation rates in normal tissues are sufficient to accumulate the number of driver mutations needed for cancer. A higher mutation rate allows for more rapid evolution, but it may be unnecessary and could even cause deleterious mutations. Loeb proposed the 'mutator phenotype' in 1974, based on the observation that polymerases isolated from acute lymphoblastic leukaemia cells replicated DNA less faithfully than those from normal lymphocytes. Various modelling approaches were used to determine whether an increase in mutation rate was necessary, some finding that it was (Loeb 1991) and others that it was not (Tomlinson et al. 1996). These models seem to be relatively sensitive to the number of hits that are necessary before the first clonal expansion and the extent of clonal expansions. Estimates of mutation rates were based on single gene reporter assays and so not representative of the whole genome.

The data, too, are mixed. By the 1970s, a number of studies had indicated that tumours made more mitotic errors than normal cells, and that the numbers increased as the tumour progressed (discussed in Nowell 1976). The discovery of Lynch syndrome and other genetic predisposition syndromes that result in an increased mutation rate certainly provides strong

evidence that in these circumstances an increased mutation rate is beneficial, but it is unclear whether this also applies to sporadic tumours where more hits would be necessary in order to increase the mutation rate. In colorectal cancers, it was observed that tumours that were mismatch repair deficient rarely exhibited chromosomal instability, and conversely tumours that were chromosomally unstable rarely lost mismatch repair, suggesting that there might be an advantage to increasing the mutation rate by whichever means (Fearon 2011).

With the advent of cancer genome sequencing, it has become clear that most tumours have tens of thousands of mutations due to a number of mutational signatures (Alexandrov et al. 2013, Alexandrov et al. 2018). Without sequencing normal tissues, however, one cannot know if this is abnormal. In one of the few tissues in which the comparison of tumour and normal has been performed, it was found that the mutation burden of acute myeloid leukaemia cells was not increased relative to healthy haematopoietic stem and progenitor cells (Welch et al. 2012). Nonetheless, blood cancers may represent a special case, since they are on the whole remarkably un-mutated and require few driver mutations. More recently, sequencing of organoids derived from normal colonic stem cells (discussed in Results Chapter 2 section I.4.) has shown that normal colonic cells have a lower mutation burden than colorectal cancers (Blokzijl et al. 2016, Alexandrov et al. 2018). Finally, a recent analysis shows that – with the exception of a few tens of genes in haploid regions which are very rarely mutated – negative selection is virtually absent in cancers, indicating that the mutation rates observed in cancers are not deleterious through the inactivation of essential genes (Martincorena et al. 2017).

Taken together, it seems that there is more evidence for an increased mutation rate in cancer than against it, in some tissues at least. While elevated mutation rates may not be necessary in order to acquire the sufficient number of driver mutations, an increased mutation rate may still increase the probability of this occurring. Definitive answers, however, can only come through the comparison of normal and cancerous cells.

In summary, cancer is a multi-step branching evolutionary process, with ongoing clonal competition and complex selection pressures. Despite these great advances in our conception of cancer, many questions remain unanswered. Two lacunae that are of relevance to this dissertation stand out. First, most descriptions of tumour progression are not quantitative (with a few exceptions, e.g. Bozic et al. (2010), Mitchell et al. (2018)). To understand cancer evolution fully,

one would need to measure the selective benefit conferred by each change in the context of its particular microenvironment. Second, we have a limited understanding of somatic evolution before the advent of histological changes.

3. The mutation rates and stem cell numbers of normal tissues affect cancer risk

Under a simple model of cancer, where a cell needs to acquire a defined number of driver mutations to become cancerous and driver mutations are independent events that occur at a constant rate over time, the incidence rates of cancer are closely related to the number of driver mutations required. If a single driver were required, the incidence rate of cancer at all ages would be the same, since the probability of acquiring the driver is the same in every decade. If two driver mutations were required, the incidence rate would increase linearly, since the probability of a cell having already acquired one driver increases linearly with time and the probability of acquiring a second driver is constant. If three drivers were required, the cancer incidence rate would increase proportionally to the square of age, and so on. For seven drivers:

$$\text{Incidence rate} = k * p1 * p2 * p3 * p4 * p5 * p6 * p7 * t^6$$

where k is a constant and t is age. p1 to p7 are the rate of acquisition of each of the driver mutations, which in this model will be entirely dependent on the mutation rate *in normal tissues*. Plotting the logarithm of cancer incidence *versus* age, the slope of the line should be one fewer than the number of driver mutations required to cause a cancer. This is the reasoning famously applied by Nordling (1953) and Armitage and Doll (1954) to infer the number of rate-limiting events necessary to cause cancer (commonly interpreted as driver mutations) from age incidence data (Nordling used mortality as a proxy for incidence). Driver mutations are unlikely to be independent, since they may cause clonal expansions and/or increase the mutation rate. Nonetheless, at the very least, cancer incidence will be related to the probability of acquiring the first driver mutation (p1), which in turn is likely to depend approximately linearly on both the mutation rate in normal tissues and the number of cells at risk of transformation.

14

A higher mutation rate in normal tissues, therefore, should increase the risk of cancer. In support of this, many risk factors for cancer increase the mutation rate, whether genetic (biallelic mismatch repair deficiency, xeroderma pigmentosum, Fanconi anaemia, or polymerase proof-reading polyposis syndrome, to name but a few examples) or environmental (such as smoking, ionising radiation, or aristolochic acid). Although many of these might act in other ways as well (such as by changing the clonal dynamics of the tissue through inducing inflammation and proliferation) in a study of cancers whose relative risk is elevated by smoking, the mutation burden of cancers from smokers was higher relative to that of cancers from non-smokers, which is consistent with smoking exerting its effect at least in part through elevated mutation rates (Alexandov et al 2016).

3.a. Stem cells are the cell-of-origin of many adult cancers

The more cells at risk of transformation, the greater the risk of cancer. Determining which are the cells at risk is therefore of primary importance. Given that the capability for self-renewal is shared by both cancers and adult tissue stem cells, the latter have been proposed as the cell-of-origin for a number of malignancies. Here, I will discuss the evidence for this in colon and blood.

Barker and colleagues inactivated *Apc* in cells that express *Lgr5* in colonic crypts (Barker et al. 2009). Lgr5 is a marker expressed at the base of the crypt, and a subset of Lgr5-expressing cells behave as functional stem cells (to be discussed in more detail in Results Chapter 2) (Barker et al. 2007, Kozar et al. 2013). Loss of Apc in the stem cell compartment resulted in the rapid formation of large adenomas, whereas loss of Apc in the more differentiated transit-amplifying cells produced microadenomas that were only very rarely observed to progress into macroscopic lesions, and these rare cases were explained by occasional unintended Apc inactivation in stem cells (Barker et al. 2009). Nonetheless, in inflammation or under the influence of certain driver mutations, more differentiated cells within the crypt have been observed to re-express stem cell markers and reacquire the ability to form adenomas (Schwitalla et al. 2013). Similarly, quiescent Paneth cell-precursors have been shown to produce multilineage output following injury (Buczacki et al. 2013), and if Lgr5+ stem cells are deleted by inducing the expression of diphtheria toxin only in those cells, more differentiated cells seem to be able to take their place (Tetteh et al. 2016). It is

possible, therefore, that in some circumstances initiating lesions in more differentiated cells could result in cancer.

The evidence that haematopoietic stem cells are the cell-of-origin of certain adult blood cancers is less direct. The leukaemic blasts from G6PD heterozygote women with chronic myeloid leukaemia only expressed one isoenzyme, and B lymphocyte populations also only expressed the same isoenzyme (Fialkow et al. 1978, Martin et al. 1980), indicating that the cell-of-origin of the leukaemia had multilineage potential. It is, of course, also possible that the cell-of-origin could have restricted lineage potential and the driver mutations themselves alter the cell's lineage output. In a model of the disease, self-renewal properties were not conferred by transduction of the pathognomonic *Bcr-Abl* fusion into progenitor cells, but they were by certain other gene fusions, indicating that differentiating cells could, in some cases, also act as the cell-of-origin (Huntly et al. 2004). In acute myeloid leukaemia, it was observed that the cells capable of causing leukaemia when transplanted into immunocompromised mice bore the same surface markers as haematopoietic stem cells (Bonnet and Dick 1997). The authors argued that for a stem cell to transform was more parsimonious than for a more differentiated cell to transform and reacquire stem-like properties and markers. Childhood acute lymphoblastic leukaemias, in contrast, may mostly arise from more differentiated cells, as myeloid cells are not usually found to share markers with leukaemic cells (Greaves 1993).

3.b. Stem cell divisions and cancer risk

It was recently and controversially asserted that two thirds of the variation in cancer risk is a result of the "bad luck" of the random replication errors of normal cells (Tomasetti and Vogelstein 2015). For 31 cancer types, the authors correlated the lifetime cancer risk in the US against estimates of the number of stem cell divisions in the tissue of origin. A correlation of 0.8 meant that 65% (95% CI 39% to 81%) of cancer risk was associated with stem cell divisions. Nothing proves that the risk associated with stem cell divisions is due to somatic mutations, but this is the most likely mechanism. Such a strong correlation is perhaps surprising given all the different factors that affect cancer development and vary across tissues: carcinogenic exposures,

microenvironmental and stem cell architecture differences (discussed below), the number of driver mutations needed by different cancers, and the proliferative advantage provided by each driver.

A number of criticisms can be levelled at this analysis. First, there is inaccuracy in the estimates of the number of stem cell divisions. The estimate of haematopoietic stem cell numbers is likely to be a significant overestimate (as will be discussed in Results Chapter 1), and other errors have been noted (Rozhok et al. 2015). Second, little attempt is made to explain some of the largest departures from this correlation. Most strikingly, the number of stem cell divisions in small intestine and colon is similar, and yet the incidence of colorectal cancer is two orders of magnitude higher. Third, it has been pointed out that this correlation cannot distinguish between extrinsic and intrinsic effects both because the effect of extrinsic mutational processes can be linked to the stem cell division rate if rapidly cycling cells are more likely to fix mutations, and because some extrinsic mutational processes might act upon multiple cell types to similar extents, thus preserving the correlation (Wu et al. 2016). An analysis of mutational signatures across normal tissues and exposures would clarify this. Even if valid, the conclusions of the Tomasetti and Vogelstein study were presented in an unhelpful way in terms of public health. The relative risk of cancer across *tissues* is quite different to that across *people* and the implications of this were not made clear enough in the mass media (e.g. Galagher 2015). Whatever my behaviours I will always be more likely to develop cancer of the colon than an osteosarcoma, but I can significantly reduce my risk of cancer relative to my identical twin who drinks like a fish and smokes like a chimney. The "bad luck" applies to my colon relative to my femur rather than to me relative to my debauched twin. Despite these criticisms, this study identifies those tissues that have more cancer than one would anticipate from the number of stem cell divisions. Similar approaches might represent an interesting way to explore cancer aetiology across different tissues and potentially identify preventable exposures, although perhaps better estimates of stem cell numbers and division rates are needed first.

A related point was raised by Tomasetti and Vogelstein's follow-up article (Tomasetti and Vogelstein 2017), in which a distinction is drawn between the amount of cancer that is due to an exposure and the proportion of driver mutations in a tumour that are caused by it. They calculate that even for an exposure such as smoking, which is responsible for 90% of lung adenocarcinomas, still 35% of total driver mutations are likely to be due to random errors of DNA replication. This is perhaps unsurprising from a perspective of species evolution. Natural selection has endowed us

with tumour suppressive mechanisms such that cancer is unlikely during a reproductive lifespan with the number of drivers that are likely to occur by replication errors alone across a tissue. Mutagens need only tip us over the threshold of the number of driver mutations that we can tolerate for them to have a pronounced carcinogenic effect.

3.c. Tissue architecture limits the selective advantage of mutant cells

Cairns, after noting the potential of natural selection to promote malignancy, wrote:

> We may therefore expect to find fast-multiplying tissues arranged in such a way that neighbouring stem cells (or sets of stem cells) are restricted to limited territories so that they cannot easily compete with each other.
>
> (Cairns 1975).

Going on to discuss intestinal crypts, he noted that both restricting the number of stem cells and restraining the competition between them would decrease the accumulation of stem cells with driver mutations (Cairns 1975). Here, I focus on the colon; the same mechanisms of restricting stem cell competition do not seem to be operative to the same extent in normal blood since, here, stem cells are able to recirculate (although many aspects of the haematopoietic stem cell niche are incompletely understood).

The organisation of the colonic epithelium into crypts means that even if a stem cell with a driver mutation manages to sweep through an entire crypt, its clone size is still limited. Clonal expansion cannot occur through proliferation alone: crypt fission is needed. Limiting clone sizes at this early stage would presumably significantly reduce the probability of another driver mutation occurring in the same clone. The role of driver mutations in promoting crypt fission is discussed in Results Chapter 2.

Less intuitively, the small number of competing stem cells in the intestinal stem cell niche strengthens the hand of chance in stem cell competition (Calabrese and Shibata 2010, Rozhok and DeGregori 2015). In population genetics terms, population structure decreases the effective population size, and the probability of fixation of a beneficial allele is related to the effective

18

population size divided by the census population size (Whitlock 2003). Stochastic factors are stronger in smaller niches, and so most driver mutations, even with a strong selective advantage, are likely to be lost from the crypt by chance. The other side of this coin, of course, is that there will be less negative selection as well, and colonic crypts are more likely to accumulate deleterious mutations than a system like blood. One can speculate that the need to balance positive and negative selection in tissues contributed to the evolution of different stem cell niche structures (in addition to obvious physiological reasons). Thus, the number of stem cells and their mutation rates are only one part of the story.

4. Somatic evolution in ageing and diseases beyond cancer.

Little is known about the role of somatic mutations beyond cancer, and so this section will necessarily be brief. The somatic mutation theory of ageing (Szilard 1959, Morley 1995) posits that the stochastic and progressive accumulation of mutations over life results in a loss of function of tissues that is responsible for ageing. It was based on observations that large scale irradiation of mammals seems to result in premature ageing (Henshaw et al. 1947). Several features of Szilard's model are now known to be wrong, such as the assumption that one mutation might inactivate a whole chromosome, but in the absence of a better explanation, the somatic mutation theory of ageing – with some alterations to Szilard's model – remains attractive.

Progeria syndromes, such as Cockayne syndrome, are one source of evidence that somatic mutations play a role in ageing (Garinis et al. 2008). In these, features of the ageing process are accelerated, and the underlying genetic defect is frequently in a component of a DNA damage repair pathway (for example, transcription-coupled nucleotide excision repair is impaired in Cockayne syndrome). Some syndromes, such as xeroderma pigmentosum, are associated with both progeria and an increased risk of cancer. Furthermore, long-term survivors of chemotherapy and radiotherapy show signs of premature ageing (Garinis et al. 2008). A recent study of single cell sequencing of individual neurones (the method is discussed below) from individuals with Cockayne syndrome, xeroderma pigmentosum, and controls, showed that the mutation rate was indeed increased in somatic tissues from individuals with both diseases (Lodato et al. 2018).

Nonetheless, it is questionable whether the mutation rate in normal tissues would be sufficient to bring about a functional defect. For example, only ~10 exonic single nucleotide variants are observed in a blood cell from a person in their 70s (Welch et al. 2012). Assuming a constant mutation rate, by age 100 a person is unlikely to have more than 20 protein-coding mutations per cell, which, scattered among the thousands of genes in the genome, most of which are diploid and non-essential, seems unlikely to have much functional effect in most cells. The mutation rates in other tissues are likely to be higher, but probably not high enough to cause a functional decline.

If the loss of function were associated with a selective advantage, however, an unlikely event affecting a small number of cells by chance could, after a period of clonal expansion, affect the whole tissue. Evidence of this has come from the recent discovery of clonal haematopoiesis, which increases the risk of all-cause mortality even after excluding malignancy (Jaiswal et al. 2014). Much of this risk seems to come from increased cardiovascular events, seemingly through the aberrant expression of atherogenic cytokines by mutant cells (Jaiswal et al. 2017). Other diseases may be clonal in origin too. There is evidence, for example, that somatic mutations may play a role in autoimmune diseases such as Sjögren's syndrome (Nocturne et al. 2013). While these are examples of specific diseases, the same may be true of ageing.

Thus, although the evidence is not yet fully convincing, the exploration of somatic mutations in normal tissues may reveal a somatic mutation basis for ageing and a number of diseases. Establishing a baseline of somatic mutation burden and processes for normal tissues from healthy people will provide a reference against which different disease states may be compared.

5. Massively parallel sequencing

All of the work in this thesis relies upon massively parallel sequencing in order to detect somatic mutations, and so a brief overview of the method is in order. Only Illumina/Solexa sequencing will be discussed, since it is the dominant technology and the only one used here.

After extraction of DNA from cells, the DNA is sheared (typically by sonication, but in some experiments by enzymatic fragmentation (see Methods)) into fragments that are less than 1,000 bases long. The ends of the fragments are repaired to remove overhangs, and adaptors are

20

ligated. The molecules are washed over a flow cell, where they anneal by a short sequence of the adaptor tail to oligonucleotides covalently bound to the surface of a glass slide. Each molecule of DNA is replicated by bridge amplification, resulting in the formation of DNA colonies that can be sequenced.

Illumina/Solexa sequencing uses the method of sequencing by synthesis. Reversible chain-terminating nucleotides, each coupled to a fluorophore, are washed over the flow cell. A free-floating nucleotide binds to extend a strand of DNA along the template if it is complementary to the next base of the template to be copied. This terminates the reaction. On excitation by a light source, the fluorophore is cleaved, emitting a light signal that is detected by a camera. The colour of the fluorophore gives the identity of the base that has just been added. Cleavage of the fluorophore also unblocks the end of the molecule such that the next base can be added. The process continues until a read of the desired length has been achieved (150 base pairs for the XTEN platform which was mostly used in this study). Both ends of the DNA molecule are sequenced, resulting in paired-end reads. The informatic representation of the DNA sequence is then aligned to the human reference genome and mutations can be called (Methods).

6. Methods for studying mutations in normal tissues

The polyclonal nature of normal tissues makes them difficult to study. Standard sequencing protocols require hundreds of nanograms of DNA and so tens of thousands of cells. Mutations can only be called accurately if they are present in a substantial proportion of these cells. Unlike a cancer in which large numbers of mutations are shared by most cells in a biopsy sample, a biopsy of tens of thousands of cells from most normal tissues will reveal very few shared (and so detectable) somatic mutations. A number of methods have been used in recent years to tackle this problem: single cell sequencing; heavily error-corrected deep sequencing; *in vitro* clone formation; and microbiopsy sequencing. The latter two are used in this thesis.

In single cell sequencing, the DNA from single nuclei is amplified using a method such as multiple displacement amplification (MDA) and sequenced. Clean single cell whole genome sequencing would be the perfect solution to the study of normal tissues, and all the methods discussed below would largely be obsolete. Unfortunately, the process of whole genome

21

amplification from such a small amount of input DNA both introduces large numbers of false positive errors and amplifies the genome unevenly, resulting in large portions of the genome that are not covered (Wang and Navin 2015). Despite these caveats, it has been used to determine the mutation rates in single neurones (Lodato et al. 2018). In order to reduce false positives, mutations were only called if they were correctly phased with a germline polymorphism; the mutation burden was then extrapolated to the whole genome. A correlation with age was observed, indicating that real mutations were being called. This method, cannot, however, be used to catalogue every mutation in the genome of a single cell, as mutations in a high proportion of the genome cannot be called.

Heavily error-corrected deep sequencing provides a way to study a polyclonal population. The most successful methods have been based on the principles of 'duplex sequencing' (Schmitt et al 2012, Kennedy et al 2014), where the sequences of either strand of a DNA duplex are identified separately and can be compared: artefacts introduced in library preparation or sequencing errors are not likely to affect both strands of a duplex in the same way. One difficulty of this sort of approach is that recapturing both strands of a DNA duplex is unlikely. The probability of this was improved by a strong dilution step before the polymerase chain reaction (Hoang et al. 2017), and the method could then be used to call mutations in polyclonal tissues. Identifying and reproducing the dilution sweet-spot is technically challenging, and, even if achieved, the full complement of mutations in a genome cannot be identified, since the information of which cell each DNA molecule comes from is lost. This makes analyses of clonal dynamics (such as phylogenetic reconstruction) more challenging.

In vitro clone formation involves the isolation of single cells and their expansion in culture, in the presence of growth factors, into colonies that are sufficiently large to be whole genome sequenced. The cell's own replicative machinery is less error prone than whole genome amplification methods, and the genome is naturally amplified evenly. Furthermore, the precise cell type of interest can be isolated prior to amplification (for example, by flow cytometry). This method has been used successfully for intestinal and liver cells (Blokzijl et al. 2016) and for blood stem and progenitor cells (Ortmann et al. 2015) amongst others. While the range of tissues for which this is possible is being broadened, not all cells can be expanded in this way and it seems unlikely that this approach will be possible for post-mitotic tissues without extensive manipulation. There are two other caveats of this approach: mutations can occur *in vitro* (although this can be

estimated and largely corrected: see Methods), and there is potential for positive or negative selection of mutations during growth (van de Wetering et al. 2015).

Microbiopsies exploit *in vivo* clonal expansions to identify mutations present in small numbers of adjacent cells. Microbiopsies may be blind, in which a random biopsy of a tissue is taken and one hopes to encounter a clone, or they may be targeted at histologically defined clonal units, such as a colonic crypt (discussed in greater detail in Results Chapter 2). Blind microbiopsies have been used successfully to identify mutations in human skin (Martincorena et al. 2015). Using the 'pigeonhole principle', which states that clones whose sum is greater than half of the biopsy size must be nested, it may, in some cases, be possible to infer clonal relationships, and the size of clonal patches can inform on selection. The technique only works for tissues in which clonal patches are of a detectable size, however, which may limit its use to squamous epithelia or tissues with a defined and large clonal unit. The mutations called are those that were present in the most recent common ancestor of the clone rather than those in extant cells.

7. Blood is well suited to studying the clonal dynamics, and colon the mutational processes, of normal tissues

Blood and colon are both well-studied tissues, with properties that lend themselves to the investigation of different aspects of somatic evolution. Blood can be sampled randomly and longitudinally with relatively little discomfort to the patient, making the investigation of clonal dynamics tractable. Although haematopoietic stem cells are perhaps the best-studied adult stem cell, we lack the answers to basic questions in humans, such as the number of stem cells and range of clone sizes in healthy humans. These are important parameters to model leukaemia and other diseases, such as those of bone marrow failure. Furthermore, the recent discovery of widespread clonal haematopoiesis in the elderly associated with an increased risk of malignancy and death from other causes adds urgency to the need to establish a baseline for the normal clonal dynamics of blood in humans.

The colon is well-suited to the investigation of mutational processes and driver events. Colonic crypts form visible clonal units, permitting a microdissection approach to identifying somatic mutations that have occurred in colonic crypts, without the caveats of selection for or

against certain mutations or the acquisition of *in vitro* mutations. Colonic mucosa is frequently sampled, allowing us to investigate the range of somatic mutational processes across a range of people. The archetypal model of the progression from normality to cancer was defined in the colon (Fearon and Vogelstein 1990), and it has historically been the battleground for debates on the mutator phenotype (Loeb 1991, Tomlinson et al. 1996). The investigation of normal colon complements these studies. Finally, colorectal cancer is a common and lethal disease, with 42,000 new cases diagnosed in the United Kingdom in 2015, and 16,000 deaths from it in 2016 (Cancer Research UK, accessed August 2018). Epidemiological observations (discussed in Results Chapter 2 section I.3.) indicate that some risk is preventable. The elucidation of the mutational processes that cause normal cells to become cancerous can hopefully contribute to the identification of the causes of excess risk and indicate opportunities for early intervention, as well as contribute to our understanding of the pathogenesis of the disease.

8. Thesis aims

This dissertation is divided into two chapters. The first investigates the clonal dynamics of normal blood, and the second the mutational landscape of normal colonic epithelium.

The principle aim of the study of normal haematopoiesis in Results Chapter One is to estimate the number of haematopoietic stem cells that actively contribute to blood in one healthy human, and the rate at which haematopoietic stem cells replace one another. Second, it aims to investigate clonal relationships between the myeloid and lymphoid lineages. Third, it seeks to establish the mutation burden per genome and mutational processes of normal haematopoietic stem and progenitor cells.

The main aim of the study of the colon in Results Chapter Two is to describe the repertoire of mutational processes and their mutational burden operative across different sites on the colon of a number of different people. Second, it seeks to establish the frequency of driver mutations in normal colon. Third, it investigates how the mutation rate and mutational processes change in the progression from normal to cancer.

Background specific to each tissue is presented at the beginning of the relevant results chapter.

METHODS

1. Samples

1.a. Primary bone marrow and peripheral blood samples

A bone marrow (BM) aspirate, peripheral blood (PB) sample, and a buccal swab were obtained in collaboration with the laboratories of David Kent and Tony Green from a 59 year-old man with normal blood counts and no history of blood disorders. Follow-up PB samples were taken every two months for the following six months, and thereafter once every six months by me, Carlos Gonzalez-Arias or Jacob Grinfeld.

1.b. Colon samples for laser capture microdissection

We obtained healthy colonic biopsies from four cohorts of patients. The first represents seven organ donors ranging in age from 36 to 67, from whom colonic biopsies were taken at the time of transplantation. These samples were provided by Kourosh Saeb-Parsy. The second represents patients aged 60 to 72 who were having a colonoscopy following a positive faecal occult blood test as part of the Bowel Cancer Screening Programme; we selected 16 patients who were not found to have either an adenoma or a carcinoma on colonoscopy, and 15 patients who were found to have a colorectal carcinoma (the normal biopsies that we use were distant from these lesions). These samples were provided by Nick Coleman. The third represents three paediatric patients who were having a colonoscopy for inflammatory bowel disease following suggestive symptoms, but in whom no diagnosis of inflammatory bowel disease was made. These samples were provided by Matthias Zilbauer. Finally, access to samples from one 78 year-old gentleman with oesophageal cancer who underwent a warm autopsy was provided by Rebecca Fitzgerald. A table of which samples come from which patients is provided in Appendix C.

1.c. Colon samples for organoid derivation

Tissue material was obtained in collaboration with the Clevers lab from The Diakonessen Hospital, Utrecht. From the resected colon segment, both normal and tumour tissues were isolated. The isolated tumour tissue was subdivided into 4-5 segments. Normal tissue was taken at least 5 cm away from the tumour.

2. Sample preparation

2.a. Isolation of haematopoietic stem and progenitor cells and peripheral blood fractions by the Kent lab

The following work was performed by the Kent lab. Mononuclear cells (MNCs) from BM and PB were isolated by density gradient centrifugation (Lymphoprep; Axis-Shield, Oslo, Norway), and enriched for CD34 positive cells (EasySep Human CD34+ enrichment kit, STEMCELL Technologies, Vancouver, Canada (STEMCELL)) as per the manufacturer's guidelines except that only one round of depletion in the magnet was performed. Cells were then stained with the following antibodies; CD38-FITC (Clone HIT2, BD Biosciences, San Jose, CA, USA (BD)), CD34-PerCPCy5.5 (Clone 581, Biolegend, San Diego, USA (Biolegend)), CD10-APC-Cy7 (Clone HI10a, Biolegend), CD90-APC (Clone 5E 10, Biolegend or BD), CD45RA-Violet450 (Clone HI100, BD) and CD135-PE (Clone BV10A4H2, Biolegend). Single $CD34^+CD38^-CD90^+CD45RA^-$ cells (hereafter "HSCs") from BM and PB were isolated for liquid culture using an Influx sorter (BD), equipped with the following lasers; 405nm, 488nm, 561nm, and 640nm, and filter sets; 530/40 (for FITC), 710/50 (for PerCPcy5.5), 750LP (for APC-Cy7), 670/30 (for APC), 460/50 (for Violet450), and 585/29 (for PE). Bulk HSCs and $CD34^+CD38^+CD90^-CD135^+CD45RA^-$ (common myeloid progenitors, CMPs), $CD34^+CD38^+CD90^-FLK2^+CD45RA^+$ (granulocyte-macrophage progenitors, GMPs) and $CD34^+CD38^+CD90^-FLK2^-CD45RA^-$ (megakaryocyte-erythroid progenitors, MEPs) from BM and PB were isolated for colony forming cell (CFC) assays using the same setup.

For liquid culture, single HSCs were sorted into 96-well plates containing StemSpan (STEMCELL), cc100 cytokine cocktail (STEMCELL), Penicillin/Streptomycin (Sigma-Aldrich, St. Louis, MO, USA (Sigma-Aldrich)), L-Glutamine (Sigma-Aldrich), 2-Mercaptoethanol (Life Technologies, Carlsbad, CA, USA (Life Technologies)). Ten days post-sort the medium was supplemented with additional recombinant growth factors (GM-CSF (Miltenyi Biotec, Bergisch Gladbach, Germany (Miltenyi)), G-CSF (Miltenyi), M-CSF (Miltenyi), Epo (Janssen Pharmaceutica, Beerse, Belgium), IL-3(Miltenyi), TPO (Miltenyi), and IL-6 (Miltenyi)). After 4-6 weeks in culture, single-cell derived clones were harvested into PBS for subsequent DNA extraction.

For colony assays, bulk HSCs, CMPs, GMPs and MEPs were sorted into 1.5ml Eppendorf tubes containing StemSpan before distribution and plating in MethoCult H4435 (STEMCELL). After 2-4 weeks in culture individual colonies were picked into PBS for subsequent DNA extraction.

For longitudinal time points, PB samples were collected (40-60 ml, Lithium-Heparin tubes) and MNCs and granulocytes were isolated using Lymphoprep. The granulocyte fraction was subject to two rounds of red blood cell lysis (Ammonium chloride, STEMCELL) and defined numbers of cells (1,000, 10,000 and 100,000 cells/well) were plated into 96-well plates for later DNA extraction (stored in PBS). Bulk CD4$^+$/CD8$^+$ T-cells and CD19$^+$ B-cells were sorted and plated in 96-well plated for later DNA extraction (1,000 cells/well).

DNA was extracted by the Cancer Genome Project Lab core facility from cell pellets using DNeasy kit (Qiagen 69504) for whole genome sequencing of colonies, and using the BioRobot according to the protocol for dried blood samples (Qiagen 965942) for targeted sequencing of bulk granulocytes and lymphocytes.

2.b. Laser capture microdissection of colonic crypts

Fresh frozen biopsies were embedded in optimal cutting temperature (OCT) compound. 30 micrometre sections were fixed in methanol for five minutes, washed three times with phosphate-buffered saline, and then stained with Gill's haematoxylin for 20 seconds. Crypts were isolated by laser capture microdissection, with every crypt falling into a separate well. They were lysed using

the Arcturus PicoPure Kit (Applied Biosystems) according to the manufacturer's instructions. DNA library prep then proceeded without clean-up or quantification.

In an attempt to avoid batch effects, most plates of crypts included samples from multiple patients. For example, rather than dissecting all the crypts from one patient into one plate, in most cases the samples from a given patient were split into three groups, and crypts were cut into separate plates, which they shared with samples from other patients. Almost all patients had some samples dissected on a different day and into a different well to other samples from the same patient. Batch effects that affect coverage should be accounted for by statistical analysis that explicitly takes coverage into account. Batch contamination with DNA of another species should be detectable at quality control, as a low proportion of reads mapping to the genome; in other work, for example, Mycoplasma contamination of sequenced cell lines was detected. Batch effects cannot be solely responsible for the presence of any novel signature, as all novel signatures have been detected in crypts from multiple plates and sequencing runs. Nonetheless, we cannot exclude batch effects that could affect crypts from the same section of tissue, as normally these are processed together.

2.c. Isolation and growth of colonic organoids by the Clevers lab

The following work was performed by the Clevers lab. Human normal and tumour colon organoids were established and maintained from isolated colonic epithelium. In brief, long term normal colonic organoid culture required Human Intestinal Stem Cell Medium (HISM) composed of Advanced DMEM/F12 (AdMEM) with penicillin/ streptomycin, 10 mM HEPES, 1xGlutaMAX,1xB27 (Invitrogen) and 1 µM N-acetylcysteine (SIGMA), supplemented with 50 ng mL^{-1} human recombinant EGF (Peprotech), 0.5 µM A83-01 (Tocris), 3 µM SB202190 (SIGMA), 1 µM Nicotinamide (SIGMA), 10 nM Prostaglandin E2,Wnt3A-condition medium (CM) (50% final concentration), Noggin-CM (10% final concentration), and R-Spondin1-CM (10% final concentration). Tumour organoids were cultured in medium containing only EGF, Noggin-CM, R-Spondin1-CM and A83-01.

For clonal organoids from normal crypts, single crypts were embedded in 10 µl Matrigel and cultured in HISM medium. For clonal tumour organoid cultures, tumour cell suspensions were

cultured for 7-14 days in HISM without Wnt3A-CM. 10-15 individual organoids were picked and separately dissociated into single cells by TryPLE express (Thermo Fisher), washed and suspended in AdMEM containing Propidium iodide (PI). 48 single cells were sorted into tumour organoid medium (HISM plus 10 µM ROCK inhibitor Y-27632 (Tocris BioScience); no Wnt-CM) from each tumour organoid. Sorting was based on FCS Area/FCS peak and PIneg/FCS Area using a Moflo machine (Beckman Coulter). Sorted cells were spun down at 1000 g and 4°C for 5 minutes, after which single cells were each embedded into 10 µl of Basement Membrane Extract (BME, Amsbio) and seeded into 96 well-plates at a ratio of 1 cell/well. The gel was left to solidify in a 37°C incubator after HISM (no Wnt3A-CM) was added. Y-27632 was added to the medium for the first a week after sorting. For each original tumor organoid, a single clonal organoid was selected and expanded for further study.

DNA was extracted from frozen tissue samples or organoid cultures using AllPrep DNA/RNA minikit (Qiagen 80204).

3. Library preparation and sequencing

3.a. Library preparation

Libraries for the blood study and for the colonic organoids were prepared by the Sanger Institute Core Pipelines. Library preparation of low-input microdissected crypts was performed by Peter Ellis and Chris Alder. Two library preparation methods were used for low-input material: initially, sonication was used to fragment DNA, and later, an enzymatic fragmentation method was implemented as it could make libraries from even lower input. Comparison of the two methods showed no difference in mutation calls once post-processing filters (described below) had been implemented. For sonication libraries, 20ul of LCM lysate was sheared with focused acoustics and DNA isolated with magnetic beads. End preparation and A tailing was performed and duplexed adaptors were ligated. After another bead-cleanup and 12 PCR cycles, a final bead-cleanup was performed and sequenced on the HiSeq X platform (Illumina). For enzymatic libraries a magnetic bead separation was performed up-front. DNA was not washed off beads, but a bead/DNA slurry

was incubated with Ultra II FS buffer and Ultra II FS enzyme (New England BioLabs). DNA that was fragmented and A-tailed was ligated with duplex enhancers and processed as for sonication.

3.b. Sequencing

For whole genome sequencing, paired end sequencing reads (150bp) were generated using Illumina XTEN® machines. For targeted sequencing, sequencing reads (75bp) were generated using Illumina HiSeq2000® machines. Sequences were aligned to the human reference genome (NCBI build37) using BWA-MEM for 150bp reads and BWA-align for 75bp reads by the Sanger Institute Core Informatics team.

4. Bait-set design for targeted sequencing of peripheral blood

Very deep sequencing of every single substitution detected in the clones would be cost-prohibitive. We therefore selected mutations to be included in the bait-set according to the following criteria:
- All mutations called by the Shearwater algorithm (see below) in two or more samples
- Mutations called in at least one clonal sample and also in at least one polyclonal sample
- We aimed to sample approximately 20 mutations private to each colony (for both clonal and polyclonal). These were selected in the following way:
  - Mutations were ranked according to their sequencing error rate from whole genome sequencing of a large panel (approximately 1000 tumours and normal samples). The Shearwater algorithm was used to calculate the error rate.
  - We then picked the 10 mutations on the X chromosome with the lowest error rate, since mutations on the X chromosome should provide greater sensitivity as our patient is male.
  - We then picked the 10 mutations from elsewhere in the genome with the lowest error rate.
  - When selecting mutations private to polyclonal colonies, we prioritised mutations with the highest VAF, that were likely to be in the dominant clone.

30

- We also included 100 mutations that are known hotspot mutations found in clonal haematopoiesis of indeterminate potential (CHIP) and leukaemias in order to detect CHIP should it arise.

We performed custom RNA bait design following manufacturer's guidelines (SureSelect®, Agilent®, UK). Of the 19336 mutations selected, only 9175 could have a high-quality bait designed for them (some mutations are not suitable for bait design because, for example, they fall in a repetitive region of the genome). Driver mutations that had been included to allow us to detect clonal haematopoiesis of indeterminate potential could not, of course, be assigned to the tree, and neither could mutations that were present in polyclonal samples and not in any clonal samples. After final data curation and re-assessment of mutation assignment to the tree, 7116 mutations could be assigned to the tree of 140 clonal samples unambiguously and used for analysis.

5. Mutation calling

5.a. Substitutions

For all studies, substitution calling was broken down into three steps: mutation discovery; post-hoc filtering to produce a list of clean sites; and, finally, genotyping, where the presence or absence of every mutation in every sample is evaluated. Mutations were initially discovered using the Cancer Variants through Expectation Maximisation (CaVEMan) algorithm (Jones et al. 2016). CaVEMan uses a naïve Bayesian classifier to derive the probability of all possible genotypes at each nucleotide. CaVEMan copy number options were set to major copy number 5 and minor copy number 2 for normal clones, as in our experience this maximises sensitivity, but for analysing tumour organoids the tumour copy number profile from ASCAT (see below) was used. For all studies, the algorithm was run using an unmatched normal in order to be able to derive phylogenies: had another sample from the same individual been treated as a matched normal, early embryonic mutations would have been treated as germline and discarded, resulting in incorrect trees. Some post-processing filters were used in all studies. These included filtering against a panel of 75 unmatched normal samples to remove common single nucleotide polymorphisms, and two

filters (only applied to whole genome sequencing data) designed to remove mapping artefacts associated with BWA-MEM: the median alignment score of reads supporting a mutation should be greater than or equal to 140, and fewer than half of these reads should be clipped. After these filters, a pileup of all the samples from a given patient was constructed, counting the number of mutant and wild type reads in every sample over every site that had been called in any sample from that patient. Only reads with a mapping quality of 30 or above and bases with a base quality of 30 or above were counted.

| Mutations called against unmatched normal with CaVEMan algorithm | → | Filtered against 75 unmatched normals | → | Post processing filters as described in Nik-Zainal et al 2012 | → | Median alignment score of reads (ASMD) >=140 | → | Fewer than half the reads should be clipped (CLPM=0) | → | Pileup: count mt and wt reads in all samples for subs called in any sample | → | Dataset-specific filtering |
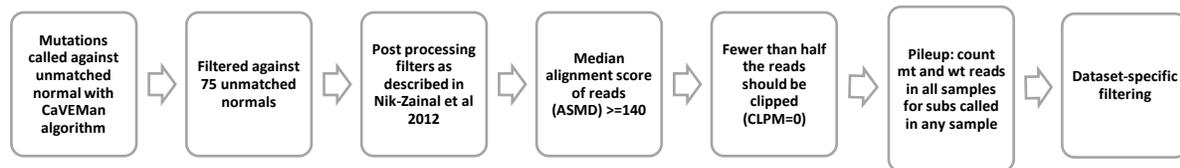
Different post-processing filters were applied depending on the dataset.

5.a.i. Substitution calling in blood colonies

To build a phylogeny from 140 colonies sequenced at low coverage requires a stringent set of mutation calls. To minimise the number of false positive mutations, the following post-processing filters were included in addition to those detailed above:

- To remove germline:
    - We knew from initial manual tree building that there was an early split in the tree into two sample groups with more than 20 clonal colonies in either group. Therefore, if a mutation had mutant reads in greater than 120 out of 140 clonal colonies we considered it as germline.
- To remove false positives, or mutations that would be uninformative for tree building:
    - Mutations that fell within 10 base pairs of each other or within 10 base pairs of indels were removed.

32

- Mutations that had a sequencing coverage in a given sample of less than 6 reads for autosomes and less than 3 reads for sex chromosomes were given a status of NA in that sample. Mutations with this NA status in more than 5 samples were removed.
- We observed that mapping artefacts could cause a mutation to be seen at low VAF in multiple samples. We therefore applied the following filters
  - the mean VAF for a mutation, across all samples in which there are any mutant reads at all, should be greater than 0.3.
  - exclude sites where more than 10% of the samples in which there are any mutant reads have a VAF of less than 0.1

This resulted in 134,411 sites that were carried forward for further evaluation using the Shearwater algorithm (Gerstung et al. 2014), which was run by Sebastian Grossmann. This algorithm compares allele frequencies of variants to a background error model derived from sequencing thousands of samples from unrelated studies on the same platform. Sequencing errors are known to occur at different frequencies across different sites of the genome. By obtaining a comprehensive view of the number of variant calls at each position in unrelated normal genomes, we built an error model for each nucleotide change for each position. This method has previously been employed to find variants supported by small numbers of reads (Gerstung et al. 2014; Martincorena et al. 2015). For this study, the position-specific error probabilities were inferred from 234 whole genome sequences from unrelated normal samples from various donors and tissues. The ShearwaterML model to compare observed variant frequencies with the background error model is described in the R-package deepSNV-1.21.4 (Gerstung et al. 2014). After correcting for multiple testing with the Benjamini-Hochberg procedure, only variants were kept that were significantly mutated over the error model, using a corrected p value cutoff of 0.05. Although most true variants largely exceed this threshold, this procedure maximizes the chance of retaining variants on a relatively low number of reads but that occur in parts of the genome that suffer from few false positive calls. As a further stringent filter to minimise false positive calls, variants had to be supported by at least 3 mutant reads to be considered by the algorithm. In this way, every mutation called in an individual was genotyped as being present or absent in each sample.

5.a.ii. Substitution calling in colon microbiopsies

In addition to the filters described in section 5.a., additional filters were designed by Mathijs Sanders to remove artefacts associated with low-input library construction. The library preparation protocol for microbiopsies produced shorter library insert sizes than standard methods. It was, therefore, common for paired-end reads from microbiopsy libraries to overlap, resulting in double counting of mutant reads. Mathijs Sanders therefore generated fragment-based statistics to prevent the calling of variants supported by a low number of fragments. Variants were annotated by ANNOVAR (Wang et al. 2010) and fragment-based statistics (fragment coverage, number of fragments supporting the variant, fragment-based VAF) were calculated for each variant after the exclusion of marked PCR duplicates. In the rare event of a disagreement in the called base at the variant position between overlapping paired-end reads, the base with the highest quality score was selected. Fragment-based statistics were calculated separately for all fragments, only counting those with alignment score ≥ 40 and base scores ≥ 30. Variants supported by at least 3 high quality fragments were retained and used for the next stage of variant filtering.

Examination of variants called in microbiopsies demonstrated that an excess was present within inverted repeats capable of forming hairpin structures, that they were supported by reads with very similar alignment start positions (i.e., not marked as PCR duplicates), and were primarily located close to the alignment start within the supporting reads. These variants were frequently within 1-30 base pairs of another variant. Filtering based on variant proximity alone would also remove actual kataegis events, and so could not be used. *In silico* modelling of the potential hairpin revealed that the variants called in the same read were aligning to each other in the stem of the structure, but could not form a base pair (i.e., mismatched), while all other bases could. Careful consideration indicated that the artefacts were the consequence of erroneous processing of cruciform DNA (existing either prior to DNA isolation or formed during library preparation) by the enzymatic digestion protocol applied. Mathijs Sanders considered modelling the hairpin structures to filter these variants, but given the fact that read clustering (i.e. similar alignment position) serves as a strong hallmark for these artefactual variants, he opted to use the proximity of the variant to the alignment start, and the standard deviation (SD) and median absolute deviation (MAD) of the variant position within the supporting reads, as features for filtering. These statistics were calculated separately for positive and negative strand aligned reads. In case the variant was

34

supported by a low number of reads (i.e., 0-1 reads) for one of the strands, the filtering was based only on the statistics generated for the other strand. Per variant, if one of the strands was supported by too few reads, it was required for the other strand that either: (I) there should be ≤ 90% of supporting reads to report the variant within the first 15% of the read starting from the alignment start, or (II) the statistics MAD > 0 and SD > 4. Per variant, if both strands were supported by a sufficient number of reads it was required for both strands separately that either: (I) there should be ≤ 90% of supporting reads to report the variant within the first 15% of the read, (II) the statistics MAD > 2 and SD > 2, or (III) that the other strand should have the statistics MAD > 1 and SD > 10 (i.e., the variant is retained if the other strand demonstrates strong measures of variance). The proposed strategy reduced the number of artefactual variants while retaining all other variants, as assessed by running the last filtering step on WGS data from non-LCM experiments.

After applying these filters, mutations were genotyped based on the number of mutant and wild type reads at each locus. Mutations were called based on a variant allele fraction (VAF) > 0.2, a depth > 7, and at least 4 mutant reads. If the depth over a locus was less than seven in a given sample, or if there was more than one mutant read but the other criteria were not met, the genotype was set to NA for tree construction purposes. Loci that were set to NA in more than one third of the samples were removed for construction of the phylogeny. Positions were called as germline if they were either called as present or NA in all of the samples from a given patient.

5.a.ii. Substitution calling in colonic organoids

As organoids were sequenced with normal library preparation methods at higher coverage than other samples, they were fully clonal, and there were fewer samples from each patient (which makes phylogeny construction easier), less stringent filtering was required. However, a calling method had to be developed that allowed the detection of a pre-malignant clone that could contain both the normal organoids (which were derived from >5cm away from the tumour) and the tumour (a field effect). For each patient, the only germline reference available was healthy colorectal tissue >5cm distant from the tumour, consisting of epithelial and connective tissue. We reasoned that if there were a field effect the somatic mutations in this tissue should not be fully clonal. We therefore deducted germline mutations on the basis that they were fully clonal in the bulk normal, while

mutations that were subclonal in the bulk were not removed from the analysis. To define a mutation as subclonal in the bulk, the probability of finding the observed number of mutant reads or fewer given the sequencing coverage had to be less than 0.005, based on the binomial distribution with a probability of 0.5 for autosomes. Mutations that failed to meet this criterion were considered to be germline and were removed. Mutations were then genotyped using the Shearwater algorithm, as for blood colonies. The Shearwater algorithm was run by Sebastian Grossmann.

5.b. Small insertions and deletions (indels)

As for substitutions, calling of indels was broken down into mutation discovery, filtering, and genotyping. Mutations were called with the Pindel algorithm (Raine et al. 2015) using an unmatched normal. Post processing filters were applied as in Nik-Zainal et al. (Nik-Zainal et al. 2012), and the number of mutant and wild-type reads was tabulated as above. The same dataset-specific filters were applied as for substitutions. Indels were then genotyped based on a VAF>0.2, a depth of at least 10, and support of at least 5 mutant reads.

5.c. Short tandem repeats (STRS)

Short tandem repeats were called only for the blood colonies in order to test the robustness of the tree. The HipSTR algorithm to call them was run by Sebastian Grossmann (Willems et al. 2017). HipSTR was run using *de novo* stutter estimation and STR calling with *de novo* allele generation as per the author's recommendation. The calls were filtered on a calling quality of at least 0.95 and at least six reads per sample spanning STR loci on autosomes and at least three reads on sex chromosomes. Calls were also filtered if more than 15% of the reads were affected by PCR stutter or featured indels in the STR flanking regions. If STR calls were removed in more than five samples, the locus was flagged for all samples.

5.d. Structural variants

Genomic rearrangements were called using the BRASS algorithm (Li et al. 2017, https://github.com/cancerit/BRASS). Abnormally paired read pairs from WGS were grouped and filtered by read remapping. Read pair clusters with ≥50% of the reads mapping to microbial sequences were removed, as were rearrangements where the breakpoint could not be reassembled. Candidate breakpoints were matched to copy number breakpoints defined by ASCAT within 10kb. Only structural variants where the two breakpoints were more than 1000 base pairs apart were considered. Blood colony structural variants were called against another colony, colonic microbiopsies against a matched normal when available and against another crypt when not, and colonic organoids against bulk normal epithelium.

5.e. Copy number changes

Copy number changes were called using the Allele-Specific Copy number Analysis of Tumours (ASCAT) algorithm (Van Loo et al. 2010, Nik-Zainal et al. 2012, https://www.crick.ac.uk/peter-van-loo/software/ASCAT). The same matched normal sample was used as for calling structural variants. For additional validation of copy number changes in normal colon, the QDNAseq algorithm (Sheinin et al. 2014) was run. ASCAT uses both the read depth and ratios of heterozygous single nucleotide polymorphisms to determine an allele-specific copy number, while the QDNAseq relies solely on variations in sequencing depth. To call amplifications and deletions in the colonic microbiopsy cohort, only those that were called by ASCAT and showed a clear departure from the background on QDNAseq were retained. To call copy neutral loss of heterozygosity in this cohort, all such events called by ASCAT were checked visually on Jbrowse (Buels et al. 2016) to verify an imbalance of parental snps.

6. Construction of phylogenies

Phylogenies were constructed using the matrix of genotypes (with samples as columns and mutations as rows) derived for every patient.

6.a. Derivation of the phylogeny of blood

The phylogeny of blood was built using the maximum likelihood implementation of SCITE (Jahn et al. 2016), a tree-inference algorithm that uses Markov chain Monte Carlo sampling with an error model that takes potential false positives and false negatives into account for tree scoring. The SCITE algorithm was run by Sebastian Grossmann. The false positive error rate was known as the probability of every mutation being real is provided by Shearwater. False negative error rates were estimated by visual inspection using jBrowse of over 1,000 sites that had been provided to Shearwater but not called as present. We ran five independent Markov chain Monte Carlo simulations with one million iterations to test for robustness of the obtained phylogeny that all resulted in the same tree. Furthermore, the results were consistent with manual tree building attempts on a subset of clonal colonies.

To test the robustness of phylogeny, we bootstrapped the substitution input matrix for SCITE 1,000 times and ran SCITE 1,000 times to generate bootstrapping p values for each node in the tree (figure 1.4a). Furthermore, we took the following steps to rebuild our tree using different data and/or different tree building methodology:

1. We repeated tree building from the same matrix of substitution calls by maximum parsimony, bootstrapping 100 times.
2. We repeated tree building from a combined matrix of substitution and indel calls by maximum parsimony, bootstrapping 100 times.
3. We used the HipSTR algorithm to call STRs and built a tree with neighbour joining, bootstrapping the input matrix of mutation calls 100 times.
4. We built the tree using all variants combined (substitutions + indels + STRs) using neighbour joining, bootstrapping 100 times.
5. We built the tree using substitutions and non-STR indels using maximum parsimony, bootstrapping 100 times.

These are detailed below, and the results of each tree-building approach are shown in figure 1.4. SCITE and HipSTR were run by Sebastian Grossmann, and all other work performed by me. It should be noted that the shape of our tree, with multiple short branches that are ancestral to many

cells, makes bootstrapping conservative. As some of the earliest splits in the tree are supported by a single mutation which may be omitted in a bootstrap replicate (a given mutation will be present on average in 630 out of 1000 replicates), some of the earliest splits may not be supported by a particularly high proportion of replicates, even though the mutation calls are highly confident (based on multiple colonies carrying 5-10 reads reporting the variant, while being completely absent from other colonies). Indeed, shorter branches have lower p values.

When comparing trees, the quantity of interest is whether high-confidence nodes from other datasets and tree-building methods agreed or disagreed with our tree. For the purpose of this analysis, we define a high confidence node as one that created the same split of samples into two groups in >=70% of bootstrap replicate trees built with a given method.

1. Substitutions using maximum parsimony and bootstrapping.

The input matrix of substitution calls was bootstrapped 100 times and trees were built with each bootstrap replicate using maximum parsimony. We used the mix programme from the phylip (Felsenstein 1989) suite of tools, using the Wagner method. Ignoring terminal nodes (which are uninformative), 58 out of 139 nodes had >70% support. The fact that such a low proportion of nodes was well-supported is most likely due to the fact that many of the branches at the top of the tree are supported by no mutations (hence the polytomies in our original tree) or a small number of mutations, which need not be present in every bootstrap replicate. Of the 58 well-supported nodes, 56 were present in our tree, and two were absent. The differences only involve a minor reordering of branches within a clade. The discrepancies are not supported by many mutations in the SCITE tree.

2. (Non-STR) Indels using maximum parsimony and bootstrapping

There are only 228 shared indels outside of short tandem repeats (STRs), and so it is not surprising that a tree built only from these should have a high degree of uncertainty and several polytomies. Only 22 nodes (out of 126 this time because of polytomies in the parsimony tree) are supported by >70% confidence. Two of these 22 disagree with our original SCITE tree. As with the substitutions, these two differences are minor departures from the original SCITE tree.

3. STRs with neighbour joining bootstrapping results

Short tandem repeat indels were called with HipSTR (see above). The phasing quality was low and so the sum of alleles at a given locus was used, ignoring the parental origin. Trees were built with neighbour joining using the absolute distance between cells. We bootstrapped the matrix of calls 100 times, recalculating the distance matrix each time. The resultant consensus tree was uncertain, with multiple polytomies and only 2 nodes supported by 70% of the trees. Of these two high confidence nodes, one agreed with our original SCITE tree and one disagreed. This discordant node sits at the top of one clade with 10 descendants in the STRs NJ tree. In our original SCITE tree, however, these cells are scattered around the tree. Some of these cells share large numbers of substitutions with cells not contained in the STRs NJ clade (such as BMH97 (inside the STR NJ clade), which shares ~500 mutations with BMH73 (outside the STR NJ clade)), which leads us to trust the substitution tree over the STR tree.

4. Combined substitutions, indels, and STRs using neighbour joining

We tried neighbour joining on the combined dataset of STRs, substitutions, and indels, bootstrapping the matrix of mutation calls 100 times. As there were 92,661 shared STRs, 9,982 substitutions, and 228 indels, the STRs dominate each bootstrap replicate. As with the tree built only from STRs using neighbour joining, few nodes are well-supported. Of 27 nodes that were present in 70% or more of the trees, 23 were present in our original SCITE tree and 4 were discordant. Three of these 4 nodes were nested inside each other, and so really represent only two discordant clades.

5. Combined substitutions and indels analysis using maximum parsimony

Because of the lack of consistency among trees built by neighbour joining from bootstrap replicates of STRs, we performed an analysis on combined substitutions and indels, excluding STRs. We used maximum parsimony, as different false positive and false negative rates for substitutions and indels meant that SCITE would not have been appropriate. Of 59 nodes that were present in 70% or more of the trees, 56 were present in our original SCITE tree and 3 were discordant. The arrangements of the nodes in these three discordant clades is similar to their arrangement in the original SCITE tree.

In summary, the vast majority of high confidence nodes found using other data types and tree building methods support the tree that we built originally with SCITE. Furthermore, with the exception of a couple of overlaps between maximum parsimony substitution and indel trees compared to maximum parsimony on combined subs and indels (i.e. the same tree building method on nested datasets), the disagreements between our tree and those built with different datasets or tree-building methods were not recurrent: *different disagreements were found with each alternative tree*. We are therefore confident that the tree that we have provided is the best tree that we could have built.

The accuracy of the reconstruction of the tree matters principally for four findings:
1. The uneven contribution to the embryo of the first division that we can reconstruct.
2. The relationship between different cell types (stem cell vs progenitor, and peripheral blood-derived stem cell vs bone marrow-derived stem cell).
3. The timing of branch points throughout life for phylodynamic inference (both for the stem cell population size trajectory and the estimation of the number of stem cells contributing actively to granulopoiesis).
4. Timing when in life mutations occur, such that we can assign our mutations to the tree. This is important both in terms of estimating stem cell pool size and for the analysis of stem cell clone contribution to different mature blood cell types.

For the embryology analysis, since mutations were checked manually we are confident that the uneven split into two groups at the top of the tree is correct. For the latter three points, the most important features of the tree are later branchpoints. The precise arrangement of short branches right at the top of the tree is less informative than the branches that we observe later in molecular time. These branches are necessarily supported by more mutations shared by a small number of colonies, and so we are very likely to reconstruct them correctly. Indeed, none of these later branch-points were called into question by our different tree-building analyses.

6.b. Derivation of the phylogeny of colonic microbiopsies

Derivation of a phylogeny of colonic crypts was less challenging than that of blood because there were fewer samples per patient, and accuracy in the precise arrangement of short early branches was less important because we perform no analyses of the relatedness of different crypts. All that phylogenies are used for in this analysis is for timing mutations. The most informative branches in this case are the long branches shared by a small number of crypts, which are very robust to all tree construction methods. For this reason, trees were built using maximum parsimony, as for the validation of the blood phylogeny. The tree was bootstrapped 100 times and the consensus taken.

6.c. Derivation of the phylogeny of colonic organoids

Phylogenies were constructed by maximum parsimony, as detailed in the validation of the phylogeny of blood. The SCITE (Jahn et al. 2016) and Sifit (Zafar et al. 2017) algorithms were run for validation and produced the same results (data not shown).

7. Assignment of mutations to the phylogeny

Phylogeny inference programmes provide the topology of the tree but not the assignment of mutations. Mutations from the input matrix of genotypes therefore have to be re-assigned to branches. In order to assign a set of mutation calls with no false negative and no false positives to a tree, each branch of the tree would be considered in turn. If a mutation was called in all the descendants of a given branch, and in no samples that were not descendants of the branch, mutations would be assigned to that branch. However, the complexity of the phylogeny and the number of false positive and false negative calls in the mutation matrix affect how many mutations fit the tree perfectly. Because of differences in the datasets, slightly different approaches were taken to assign mutations to each tree.

7.i. Assignment of mutations to the phylogeny of blood

Given the size of the phylogeny of blood, its shape, and non-zero false negative and false positive rates of mutation calls, we expected that a proportion of mutations would not fit our tree perfectly. Consider a mutation that is truly present in 50 colonies. With 15x coverage over the site in every sample, if we simulate resampling of mutant reads from the binomial distribution (with size of 15 and probability of 0.5), 17% of the time the mutation will have fewer than the 3 reads required to call it in at least one sample. Variation in sequencing depth and sequencing errors would further decrease the probability of calling the mutation perfectly in every sample. Mutations were therefore assigned in the following way:

- Private mutations, only called in a single sample, could be assigned perfectly to the branch ancestral just to that sample.
- 6,819 out of 9,982 mutations called in more than one clonal sample fit the tree perfectly. These were assigned to the branch of the tree that was ancestral to all the colonies that bore the mutation to none of the colonies that did not bear the mutation.
- The 3,163 mutations that did not fit the tree perfectly were assigned in a probabilistic manner.
  - True positive and true negative rates of 0.99 were chosen.
  - For every node, the probability that a mutation that was truly present in that node given the number of positive and negative calls in the descendants of the node and in all of the colonies that do not descend from that node, was calculated.
  - The mutation was assigned to the node which had the highest probability of the mutation being there. 1,306 mutations that had ambiguous placements (they fit two or more nodes in the tree equally well) were not assigned.

7.ii. Assignment of mutations to the phylogenies of colonic microbiopsies

Colonic microbiopsies also suffered from low coverage, and, what is more, some stromal contamination. For this reason we did not expect mutations to fit the tree perfectly. Unlike in the phylogeny of blood, where assignment of mutations to branches mattered for the recapture phase, here mutations were only assigned to the tree in order to determine the mutational processes active at a particular time. We reasoned, therefore, that it was preferable to assign only mutations that fit the tree perfectly, and rather adjust the branch lengths based on the power to call mutations at a given branch. Using the clonality and coverage of all descendants of a branch, the proportion of true substitutions or indels on the branch that would be first discovered (whether by CaVEMan or Pindel) and then genotyped as present according to the criteria described above was calculated. The observed branch length was then adjusted by dividing by this proportion. This was done for both substitutions and indels, but not for structural variants and for larger copy number changes due to a lack of data: most branches have no large variants and so could not be extended appropriately. Rearrangements and copy number changes were assigned to phylogenies manually.

7.iii. Assignment of mutations to the phylogenies of colonic organoids

The colonic organoid data was high coverage (~30X) and fully clonal, and with relatively small numbers of samples per patient. Ignoring private mutations, which necessarily fit any tree, 97.7% of shared mutations fitted the tree structure from patient 1 perfectly, 89.7% fitted the tree from patient 2 perfectly, and 88.1% fitted the tree from patient 3 perfectly. The lower concordance with the tree for patients 2 and 3 reflects the increased copy number changes that have occurred in these phylogenies. Examination of the copy number state at loci where there were discordant mutations showed that the majority could be explained by deletions of those mutations in a subclone. Substitutions that did not fit the tree perfectly were therefore assigned to the most recent common ancestor of the samples in which they were called.

The approach for assigning rearrangements was slightly different, as the same rearrangement may be called in related samples with slightly different breakpoints. To identify rearrangements that had been sequenced in related clones as the same, both the upstream and the

downstream breakpoints had to fall within 500bp of each other. The majority of rearrangements fitted the tree. Visualisation of discordant rearrangements using IGV (http://www.broadinstitute.org/igv) showed that often an overlapping rearrangement meant that the rearrangement was lost in a clone.

8. Analysis of relationships in phylogeny for blood

Testing for differences in the relatedness of cell types on the tree was carried out by analysis of molecular variance (AMOVA (Bird et al. 2011; Excoffier et al. 1992)), comparing stem cells derived from the bone marrow with those from peripheral blood, comparing stem cells relative to progenitors, and different progenitor types relative to one another. The phylogeny was first made ultrametric by extending private branches to the maximum branch length of 1210 mutations. Then the mutational distance between two samples (i.e. the number of mutations over which you would have to walk to go from one cell to another on the tree) was calculated for all sample pairs. Within-population and between population sum of squared distances were calculated, divided by their degrees of freedom, and used to estimate the fraction of total variation explained by differences between cell populations (the F statistic). To calculate p values, the population labels of cells were randomly re-assigned 30,000 times, and the same statistic was calculated. The p value is the proportion of random re-assignments that has an F statistic more extreme than the observed value.

9. Analysis of population size trajectory for tree of blood

Analysis of population size trajectory was performed using the Phylodyn package (Karcher et al. 2017, Lan et al. 2015), using a phylogeny made ultrametric by extending private branches to the maximum branch length of 1210 mutations, and with 70 grid points.

10. Targeted sequencing mutation analysis of peripheral blood

A stringent error-correction method of counting the number of mutant and wild type reads over a locus was devised and executed by Robert Osborne. BAM files were annotated with read coordinate (rc), mate coordinate (mc) and optical (od) auxiliary tags using biobambam2 (Tischler and Leonard 2014). Reads were included for analysis if they were marked as proper-pairs, had a minimum mapping quality $\geq$ 30, < 3 mismatches and were not marked as optical duplicates, supplementary, QC fail, unmapped or secondary alignments. We also restricted analyses to bases with base quality $\geq$ 30. Overlapping reads can result in double-counting of a base. If the base calls on read 1 and read 2 were not identical then the call was discarded. If the base calls on reads 1 and 2 matched then the call was assigned to the read with the highest base quality score. If the call on reads 1 and 2 matched and the base quality scores on both reads was identical then the call was randomly assigned to one or other reads. Instead of discarding PCR duplicate molecules, we generated consensus calls by grouping calls from reads with the same fragmentation breakpoints, read orientation and read number. A consensus base was called if >90% of the reads shared the same base. There were no minimum number of reads (if the group contained only one read then its call was retained).

For plotting purposes, signal was then separated from noise by using data from control cord blood from two individuals using a Bayesian generalised mixed effects Poisson model which was written by Peter Campbell, with minor amendments by me. The input data comprise the observed number of reads reporting each variant in the control sample and the corresponding total depth across that base; and the same for each test sample from our subject. A Poisson model was fitted in which the dependent variable was the number of reads reporting the variant, including an offset for log(total depth) (so that we are really estimating the fraction of reads). Random effects were included for the branch in the test sample (allowing for the VAF of mutations on the same branch to be correlated); the mutation (allowing for variable error rates among different mutations); and the interaction term of specific mutation by test sample (allowing for the test sample to have excess reads reporting the variant than the control sample). Parameter estimates and residuals are assumed to be drawn from a multivariate normal distribution with uninformative conjugate priors on the covariance. The model was fitted with the R package 'MCMCglmm', with 300,000 iterations, a burn-in of 10,000 iterations and thinning to 1/1000 iterations. Mutations were called as 'detected'

if more than 90% of iterations of the MCMC chain after burn-in estimated a positive term for the test sample.

11. Detection of driver mutations and positive selection

11.a. Detection of driver mutations and positive selection in blood

We searched both for particular mutations that could confer a selective advantage (driver mutations), and for global signals of positive selection in colonies that had undergone whole genome sequencing.

First, we first took all coding mutations called by CaVEMan or Pindel, but without applying post processing filters beyond those detailed in Nik-Zainal et al (2012) in order to maximise our sensitivity. We also searched for copy number changes and rearrangements that might affect known myeloid cancer genes. We then intersected this with a list of 348 genes that are found to be mutated in at least 10 haematopoietic and lymphoid neoplasms from the COSMIC database, found to be mutated in CHIP (Jaiswal et al. 2014; Genovese et al. 2014, Xie et al. 2014, Mckerrell et al. 2015), in AML (Cancer Genome Atlas Research Network et al. 2013), or in the COSMIC cancer gene census (Forbes et al. 2017). Any mutations in genes in our list were visually inspected on jBrowse as a form of verification, and we used prior knowledge of the gene's function and the ways in which it is mutated to decide whether the mutation that we observed was a likely driver mutation.

Second, we tested for positive and negative selection of mutations using dNdScv (Martincorena et al. 2015, Martincorena et al. 2017). In brief, this algorithm compares the number of synonymous and non-synonymous mutations in coding regions, adjusting for the pattern of mutational processes in the sample and local mutation rates, to detect if there is an excess or paucity of non-synonymous mutations relative to what is expected by chance under a neutral model. An excess of non-synonymous mutations indicates positive selection, while a paucity indicates negative selection. This method of driver detection is therefore independent of prior knowledge. This allowed us to search both for genes that might be under positive selection, and to see the

global dNdS of all the mutations, which tells us whether there is any excess of non-synonymous over synonymous mutations in the whole dataset.

Finally, 100 of the most frequent mutation hotspots in clonal haematopoiesis of indeterminate potential were included in the bait-set for targeted sequencing of peripheral blood, allowing us to detect driver mutations that may not have been present in the clones that were originally whole genome-sequenced (Appendix A).

11.a. Detection of driver mutations and positive selection in colonic microbiopsies

Again, driver mutations were detected both through an unbiased dNdS method and through manual annotation. For these analyses, the CaVEMan and Pindel calls were used without post-processing filters in order to maximise our sensitivity. All putative driver variants were visually inspected using Jbrowse (Buels et al. 2016), and so we could afford a higher false positive rate in the mutation discovery phase.

dNdScv (Martincorena et al. 2015, Martincorena et al. 2017) was used to conduct three tests: first, using only the whole genome sequencing data, an analysis of selection over all genes; second, using combined whole genome and targeted sequencing data, over all the genes covered by the bait-set; and finally, using again this combined dataset, over 90 selected cancer genes (Appendix B).

Manual annotation of driver variants based on prior knowledge complemented this. A list of 90 colorectal cancer genes (Appendix B) curated from the literature, that were also covered by the bait-set were intersected with the list of substitutions and indels from combined whole genome and targeted sequencing. Mutations were annotated as putative drivers if they were either missense mutations that fell in an oncogene hotspot (based on visualisation of the distribution of mutations in the gene on COSMIC (Forbest et al. 2017)), or if they were nonsense mutations that fell in a tumour suppressor gene.

Finally, structural variants that might act as drivers were assessed by intersection of genes involved in each structural variant with the twelve genes involved in gene fusions that have been reported in colorectal cancer in COSMIC (*VTI1A, TCF7L2, TPM3, NTRK1, PTPRK, RSPO3, ETV6, NTRK3, EIF3E, RSPO2, C2orf44*, and *ALK*). No fusion genes were found. The genes

48

involved in structural variants in our data did not overlap with the list of 90 cancer genes used for assessing substitutions and indels, and nor were there any genes that were affected by more than one structural variant.

11.a. Detection of driver mutations and positive selection in colonic organoids

Driver analysis in colonic organoids was performed by Sophie Roerink. To classify driver events in substitutions, indels and rearrangements the following criteria were used: 1) deleterious mutations in genes identified in CRC by TCGA (Cancer Genome Atlas Research Network et al. 2012) 2) all other known oncogenes carrying a canonical activating mutation 3) tumour suppressor genes with loss of function, and/or carrying two deleterious mutations.

12. Adjusting crypt mutation burden by the callable proportion of the genome

Whole genome sequences from crypts have variable coverage and clonality, which affect the proportion of the genome that can be called. In order to compare crypts, crypt mutation burdens throughout the text are adjusted for their coverage and clonality. The read depths for 1,000 sites were sampled from each crypt in order to capture the variability in its coverage. For each site, a mutation was simulated. Numbers of mutant reads over a given simulated position were drawn from a binomial distribution with probability equal to the clonality of the crypt, and the number of trials equal to the coverage that has been sampled for that site. The proportion of simulated sites for which mutations are callable serves as an estimate of the callable proportion of the genome. Raw mutation burdens per crypt are divided by this proportion to estimate the total mutation burden per crypt.

13. Extraction of mutational signatures

13.a. Categorisation of mutations

Mutations were categorised following the method used by the Mutational Signatures working group of the Pan Cancer Analysis of Whole Genomes (Alexandrov et al. 2018). Single base substitutions were categorised into 96 classes according the identity of the pyrimidine mutated base pair, and the base 5' and 3' to it. Doublet base substitutions were categorised into 78 classes according to the identity of the reference and alternative bases. Indels were classified according to whether they were an insertion or a deletion, the identity of the inserted/deleted base, the length of the mononucleotide tract in which they occurred, or the degree of homology with the surrounding sequence into 83 classes (figure 2.1).

13.b. Non negative matrix factorisation (NNMF)

Consider a matrix, with the name of every sample along one side, the name of every mutation category that we have defined (e.g. C>A in A$\underline{C}$A) on the other. The matrix is populated by the counts of mutations in each sample of each category. NNMF considers this observed matrix to be the product of a matrix of signatures and a matrix of exposures, plus some noise (Alexandrov et al. 2013b). The matrix of signatures is made up of the categories on one side, and the names of the signatures on the other, and populated by the proportion of a signature that is in a particular context. The matrix of exposures is made up of the samples on one side, the signatures on the other, and the amount that each signature contributes to each sample populates the matrix. The numbers that populate the matrices of signatures and of exposures are unknown. They are learnt by trying different combinations of numbers in these two matrices, and measuring the similarity of their product with the observed set of counts. In the classical version of the algorithm written by Ludmil Alexandrov, the distance metric used is the Frobenius norm (Alexandrov et al. 2013, Alexandrov et al. 2013b). The number of signatures that are used is not learnt from the data. Different numbers of signatures are tried, and a measure of the stability of the solution (calculated by bootstrapping the input data) and how well it approximates the observed data is used to select

50

the optimal number of signatures. With more signatures, the observed counts are approximated more accurately but the reconstruction is less stable.

13.c. Hierarchical Dirichlet Process

The use of a Hierarchical Dirichlet Process to extract mutational signatures was developed by Nicola Roberts (Roberts 2015, Roberts 2018). A Dirichlet Process is a non-parametric clustering method that takes as input a probability distribution and produces as output a more discretised probability distribution over the same domain. Signatures are treated as a multinomial probability distribution of the set of mutation classes (e.g. the trinucleotide context). The mutation counts per category in a sample are considered to be the result of draws from a sample-specific mixture of a shared set of signature probability distributions. The hierarchical nature comes from the fact the samples are put in groups (for example, in my analysis all samples from one patient were put in a group), and the amount that each signature contributes to a sample is drawn from parent node of the group; this has the effect that samples within a group are considered to be more similar to each other than samples from different groups.

Signatures are learnt by Gibbs sampling (Teh et al. 2006). A set of clusters is initialised with random counts in each category, all contributing randomly to every sample. In every iteration of the chain, mutations are shuffled between clusters. They are more likely to be drawn into clusters that resemble them (i.e. that have a high proportion of mutations in the same class). With some low probability, mutations can also form their own clusters, which means that the number of signatures in the dataset can be learnt rather than having to be specified. Averaging over posterior samples over the chain produces the number of signatures, their identity, and their contribution to each sample.

The algorithm can be conditioned on known signatures by including nodes of "fake data" that containing pseudocounts in the category distribution of a known signature. Mutations from the real datat that are similar to these known signatures are likely to be drawn into their clusters. Mutations that are not similar to any of the preconditioned clusters, however, can still form their own clusters, allowing simultaneous matching to known signatures and discovery of new ones.

14. Timing copy number changes in colonic microbiopsies

If substitutions are assumed to occur at a constant rate, the copy number of mutations over a particular copy number segment allows us to time when the copy number change occurred. For example, if one chromosome copy is amplified (i.e. a change from 1+1 to 2+1, where the number indicates the number of copies of each parental chromosome), then mutations that occurred on the amplified chromosome prior to its amplification will be on two copies, whereas those that occurred after the amplification will only be on one copy. Mutations were timed using the MutationR package (Gerstung et al. 2017, https://github.com/gerstung-lab/MutationTimeR), which accounts for lower detection sensitivity of mutations at lower copy number. To maximise reliability only copy number changes with that contained 100 substitutions at good coverage were timed. Trinucleotide profiles of substitutions estimated to occur before and after the copy number change were inspected and found to be very similar, suggesting that there has been no change in the mutational processes operative since the copy number change, and so supporting our assumption of a constant mutation rate.

15. Timing mutations relative to a whole genome duplication

A whole genome duplication (WGD) was observed in the trunk of the tumour for patient 2 in the organoid study. Timing as many mutations as possible relative to this allowed the investigation of the evolution of mutational processes.

15.a. Timing substitutions relative to the whole genome duplication

For every truncal substitution in every tumour clone from patient 2, the copy number segment (as called by ASCAT) in which that mutation fell was defined. Mutations could only be timed in samples in which there was a minor copy number of 0 and a major copy number greater than 1. Fortunately, because of the extensive copy number changes in this tumour, all mutations

fell in a region that met these criteria in at least one sample. For a given mutation that fell in such a copy number segment in a given sample, the VAF in that sample of known germline single nucleotide polymorphisms that fell in that segment (that necessarily occurred before the WGD) and the VAF of somatic mutations assigned to branches further down the tree (that necessarily occurred after the WGD) was examined. If, in a given sample, a mutation had a VAF greater than 90% of the VAFs of the mutations that were known to occur further down tree it was considered to have occurred before the WGD, whereas if it had a VAF less than 90% of the VAFs of the SNPs it was considered to have occurred after the WGD. If there was any overlap between the 90^{th} percentiles of the SNPs and the later mutations, or if the mutation fitted neither of these criteria, it was considered uninformative and was not used in the signature analysis. This accounted for 9,094 mutations (out of a total of 12,623 assigned to the trunk), that were not used in signature analysis. There is no reason to believe that mutations that were excluded for these reasons should be attributable to different mutational signatures than those that could be included, and indeed their trinucleotide mutation contexts are similar (data not shown). For each mutation, then, the number of samples in which it had been counted before and after the WGD was tallied. If a mutation was called as occurring before the WGD in some samples and after the WGD in others, the mutation was designated as conflicting and excluded from the analysis. 82 mutations fell into this category, and the remaining 3,447 could be timed unambiguously relative to the WGD and used in the signature analysis. In figure 1 we extrapolated the preWGD and postWGD fractions and their relative signature components to all mutations identified in the clonal trunk of P2.

15.b. Timing driver mutations relative to the whole genome duplication

Driver mutations in *TP53* and *APC* were timed relative to the WGD in patient 2. The *TP53* mutation was at VAF 1 in a region that was 2+0 in all samples, indicating that it occurred before the WGD. There were mutations in both alleles (which we will call mutation 1 and mutation 2) of *APC*. P2.T4.2 and P2.T5.1 both had the *APC* locus called as 2+2, and both mutations were at VAF 0.5. P2.T1.1, P2.T1.3, and P2.T6.2 were 2+1 in the *APC* region. Mutation 1 was at VAF 0.67 and mutation 2 at 0.33. In P2.T2.5 the region was also called as 2+1, but mutation 1 was at VAF 0.33 and mutation 2 at VAF 0.67. This shows biallelic inactivation of *APC* prior to the WGD.

15. Placing a lower bound on the age at which cells acquired aberrant mutational processes

The onset of signature 18 was timed in each patient relative to the rate of signature 1 (see Results Chapter 2, section R.7.b.). Calculations are found in table M.1.

	patient 1	patient 2	patient 3
branch to calculate ratio of signature 1 : signature 18	Ancestor of all clones except for PD21928c6	Segment after WGD but before the most recent common ancestor	Average of the two branches after the trunk
signature 1 mutations in branch	2088	975	2043
signature 18 mutations in branch	2926	410	2571
signature 1:signature 18 in branch	0.7	2.4	0.8
signature 18 mutations in trunk	1054	934	1617
signature 1 mutations in trunk	2508	3509	3050
estimate of signature 1 mutations after signature 18 onset	752	2221	1285
estimate of signature 1 mutations before signature 18 onset	1756	1288	1765
estimated years since signature 18 onset	24	20	22

Table M.1. Calculations to estimate the onset of signature 18.

RESULTS CHAPTER 1
CLONAL DYNAMICS OF NORMAL BLOOD

Introduction to this chapter

I.1. Haematopoietic stem cells

In adult homeostasis, hundreds of billions of blood cells must be produced every day, in the correct proportion of many specialised cell types, over the whole of life. Mature blood cell types include erythrocytes that transport oxygen, platelets that control haemostasis, and a battalion of cells with complementary immune functions: neutrophils, basophils, eosinophils, dendritic cells, macrophages, natural killer cells, B lymphocytes, T lymphocytes, and innate lymphocytes. Many of these have multiple subtypes. All mature blood cells ultimately derive from haematopoietic stem cells (HSCs). The conceptual definition of an HSC, therefore, is that it should have the potential to produce all blood cell types and to self-renew, producing daughters that have equal potential for the lifespan of the organism. In between stem cells and their mature progeny in the hierarchy of differentiation are progenitor cells, with a weaker capability for self-renewal and a more restricted set of possible differentiation fates. A progenitor will typically undergo multiple cell doubling and differentiating divisions to produce large numbers of functional mature blood cells.

The term *Stamzelle* was coined by Ernst Haeckel in 1868 as the first organism from which all life stemmed, in line with his Darwinian views (Haeckel 1868, Laurenti and Gottgens 2018). He later used it to describe the fertilised egg, from which all cells in the organism derive. Thus, as explained by Laurenti and Gottgens, the term *stem cell* has always had the connotation of being at the root of a phylogeny, whether germline or somatic (Laurenti and Gottgens 2018). Bone marrow cells that could generate mature blood cells were described independently by Bizzozero and Neumann (Bizzozero 1868, Neumann 1868) and a cell type that might give rise to all of blood was proposed in the late nineteenth century, when it was suggested that red and white blood cells might have a common ancestor (Pappenheim 1896). This idea remained controversial (Cooper 2011)

until 1945, when Owen noted that fraternal twin cattle that had shared a placenta also shared the blood cell types of both calves for life, thus providing evidence of a long-lived cell type that could be the source of blood (reviewed in Weissman and Shizuru 2008). Owen wrote:

> Since many of the twins in this study were adults when they were tested, and since the interchange of formed erythrocytes alone between embryos could be expected to result in only a transient modification of the variety of circulating cells, it is further indicated that the critical interchange is of embryonal cells ancestral to the erythrocytes of the adult animal. These cells are apparently capable of becoming established in the hemopoietic tissues of their co-twin hosts and continuing to provide a source of blood cells distinct from those of the host, presumably throughout his life.
>
> (Owen 1945)

In the 1950s, transplantation studies provided the means to begin to characterize these cells. Bone marrow was shown to be the site of residence of HSCs, as lethally irradiated mice could be saved by injecting them with bone marrow from other animals (Lorenz, 1951). Although regrowth of haematological tissues was observed in the salvaged mice, it was initially believed that this was due to a humoral factor in the marrow stimulating autochthonous regeneration. In 1956, however, a chromosomal (Ford et al. 1956) or enzymatic (Nowell et al. 1956) marker was shown to be shared by mature blood cells from both the donor and the recipient, indicating that cells themselves were transferred and that the restoration of haematopoiesis might be cellular. The approach described in this results chapter of using mutational markers to track stem cell clones traces its origin to these original experiments.

In a series of seminal experiments in the 1960s, Till and McCulloch showed the existence of multipotent stem cells. They induced chromosomal markers in one mouse by irradiation and transplanted them into a second mouse. Spleen colonies in the recipient were clonal for a given marker, indicating that a single cell could make all the cells in a colony (Becker et al. 1963). Re-transplanting cells from suspensions of spleen colonies into secondary recipients showed the capacity of some cells to self-renew through multiple rounds of transplantation and produce multilineage output (Siminovitch et al. 1963). Some chromosomal markers were found in both myeloid and lymphoid tissues (Wu et al. 1968), from which the authors concluded:

56

These findings […] indicate that hematopoietic colony-forming cells, erythroblasts, granulocytes, thymic cells, and the cells of lymph nodes may all belong to the same clone. Our studies, however, do not allow us to determine precise parent-progeny relationships within such a clone. For example, we do not know whether or not some or all hematopoietic colony-forming cells can differentiate to give rise to both myeloid and lymphoid descendants or, alternatively, if colony-forming stem cells and lymphoid cells have a common, as yet unidentified, precursor. The resolution of this problem will require further detailed analysis of patterns of differentiation within large hematopoietic clones.

(Wu et al. 1968)

These questions are still not fully answered.

The invention of multi-parameter flow cytometry in the late 1970s provided new tools to study HSC behaviour. In the 1980s, the field moved to the prospective isolation of HSCs by flow cytometry on cell surface markers, assaying the potential of different populations to make the various mature cell types by *in vitro* colony assays or transplantation experiments (Visser et al. 1981, Visser et al. 1984). This work led to the hierarchy of stem cell differentiation found in most textbooks today (reviewed in Weissman and Shizuru 2008). Crucially, in 1988 it was shown that functional stem cells could be isolated (Spangrude et al. 1988).

The current working experimental definition of a stem cell has not moved on very far: a cell is generally accepted to be a stem cell if it can give rise to cells of both the myeloid and lymphoid lineages for at least 16 weeks within an irradiated host (Bryder et al. 2006), although a more rigorous test is whether its progeny can do the same when re-transplanted into a secondary recipient (Eaves 2015). In the intervening decades, multiple properties of stem cells have been characterised: for most of the time they should be quiescent (Wilson 2008, Foudi 2009, Cabezas-Wallscheid 2017), glycolytic (Simsek et al. 2010, Ito and Suda 2014), rely on autophagy (Warr et al. 2013, Ho et al. 2017), and not synthesise as many proteins as their descendants (Signer et al. 2014). None of these is absolute, and most of the time functional assays are required to show stemness.

I.2. The embryological origins of blood

Our study permits us to examine retrospectively the ancestry of adult blood cells, and so a very brief discussion of the embryology of blood aims to put this in context. The first blood cells in the developing embryo appear in the yolk sac (Sabin 1920). Primitive erythrocytes derived from the yolk sac assure oxygen transport in the developing fetus. They derive from a progenitor that is shared with the endothelial cells of the yolk sac's blood islands, and consequently called a haemangioblast (Choi et al. 1998). Based on shared surface markers, haemangioblasts are thought to originate in the posterior region of the primitive streak (Huber et al. 2004). Yolk sac-derived cells, however, are not the ancestors of most adult haematopoietic tissue. This was elegantly demonstrated by grafting quail embryo bodies onto chicken yolk sacs before development of vascular connections between the chicken body and yolk sac (Dieterlen-Lievre 1975). Quail cells have a prominent nucleolus, which could be used as a lineage marker. Blood from these chimaeras was shown to derive exclusively from quail cells in 16 out of 17 animals. The site of development of the cells that go on to provide life-long haematopoiesis was established by transplanting various parts of the mouse conceptus into adult irradiated recipients (reviewed in Dzierzak and Speck 2008). The aorta-gonad-mesonephros (AGM) region and the vitelline and umbilical arteries contained cells with the capability of long-term multilineage reconstitution. Haematopoietic clusters were seen to emerge from the ventral aspect of the dorsal aorta only three days after the beginning of yolk sac haematopoiesis, which indicates that the AGM and yolk sac haematopoietic precursors are unlikely to be particularly closely related. Haematopoietic cells made in the AGM, the yolk sac, and the placenta, all migrate to and reside in the fetal liver for the duration of embryogenesis, where they proliferate (reviewed in Mikkola and Orkin 2006). They colonise the bone marrow around the time of birth.

I.3. Estimates of the numbers of active HSCs

One of the aims of this thesis is to estimate the number of active HSCs in a healthy human. Estimates have varied widely over the past few decades. Here, I review important contributions to the field.

I.3.a. Animal models

A feature of studies that aim to determine the identity and output of stem cells is the need to label cells such that their progeny can be traced. The induction of labels generally involves genetic manipulation, which is not feasible in humans. Most studies, therefore, have relied upon animal models.

I.3.a.i. Transplantation studies in animal models

Transplantation studies carry a triple benefit for tracking cells: first, the cell type of interest can be isolated by flow cytometry; second, it is easier to label these cells outside the body than inside it; and third, transplantation into a recipient allows an assay of stem cell function, based on the concept that stem cells are able to reconstitute long term multilineage haematopoiesis. These studies, however, bear two major and linked disadvantages: first, they provide an assay of stem cell *potential* rather than *fate* (cells that behave as stem cells physiologically may not do so in transplant and *vice versa*); second, the stress of cell labelling and transplantation affects stem cell dynamics.

Initial studies of transplanting retrovirally-barcoded cells into genetically anaemic or irradiated mice indicated that one to three haematopoietic stem and progenitor cells (HSPCs) were responsible for the majority of the blood produced in the salvaged animal (Dick et al. 1985, Keller et al. 1985). Barcodes could be found in all lineages and after secondary transplantation, proving that the cells capable of reconstitution were from high up the differentiation hierarchy (Lemischka et al. 1986). Remarkably, transplantation of even a single cell was sufficient to reconstitute long-term multilineage haematopoiesis (Osawa et al.1996).

More quantitative approaches have relied upon limiting dilution transplantation experiments in mice to determine the frequency of four to eight HSCs per 100,000 nucleated marrow cells (Abkowitz et al. 2000). As mentioned in the General Introduction in the context of determining the clonality of tumours, measuring the expression of an allele on the X chromosome

(X chromosome inactivation or XCI skewing) of female heterozygote mammals can be used to detect an imbalance in clone sizes: because of random X chromosome inactivation early in embryogenesis, under a stable polyclonal model of haematopoiesis half of cells should express the paternal allele and the other half the maternal one. This approach was used to estimate stem cell numbers in cats. Following autologous transplantation with different numbers of stem cells, the degree of XCI skewing was used to estimate the frequency of feline stem cells in the bone marrow (Abkowitz et al. 1996). Combining frequency estimates with counts of nucleated marrow cells per animal (determined from the distribution of radioactive transferrin) resulted in the estimate of 11,000-22,000 HSCs per mouse and 6,000-16,800 HSCs per cat (Abkowitz et al. 2002).

Transplantation studies of an animal model closer to humans physiologically and genetically have also suggested the existence of thousands of stem cells. Autologous transplantation of irradiated macaques with large numbers of retrovirally-marked HSPCs suggested that at least 1,000 HSPCs contributed to multilineage haematopoiesis in the first year following transplantation (Kim et al. 2000). Longer-term follow-up of macaques has suggested that the first year post-transplantation is one of clonal instability (Kim et al. 2014). Nevertheless, follow-up for years still supports highly polyclonal long term multilineage haematopoiesis, with hundreds to thousands of unique markers detected (Kim et al. 2014, Koelle et al. 2017). It should be noted that long-term clonal marking assays will all underestimate the total number of *cells* contributing to haematopoiesis due to sequencing detection thresholds (a 0.05% detection threshold is reported in Koelle et al. 2017): small clones are missed, and a detectable one is counted as a single clone rather than the hundreds of stem cells that may descend from it.

I.3.a.ii. Studies of native haematopoiesis in animal models

Due to the disadvantages of transplantation studies outlined above, recent work has exploited advances in inducible genetic labelling to mark cells *in vivo*, allowing assays of actual stem cell fate rather than stem cell potential.

In 2014, the Camargo group used an inducible sleeping beauty transposon system to label HSPCs indiscriminately at one time-point (Sun et al. 2014). Four months later (allowing the shortest-lived cells to exhaust), they searched for tags among different blood fractions. They found

that over 90% of granulocyte tags were only ever found at one timepoint for up to 12 months of chase. The authors ascribe this result to a model of clonal succession, where stem cell clones are recruited successively to drive haematopoiesis and then exhaust (Kay 1965). For a label to be detectable by their semi-quantitative method, it probably needs to be present in a number of stem cells, given that they estimate the sensitivity of their assay to be ~0.1% and there are likely to be thousands to tens of thousands of active stem cells. If this is indeed the case, then the model of clonal succession would require a number of closely related stem cells to enter synchronously into a phase of doubling and differentiating cell divisions to fuel blood production, until they all exhaust at the same time. The coordination of multiple stem cell members of a clone seems implausible, and such a conceptualisation of haematopoiesis is perhaps a consequence of the experimental approach of labelling cells at one particular time-point. This separates a large branching phylogeny into distinct clades (or clones) by drawing a horizontal line at the time of labelling and suggests that a clone is a meaningful biological entity. If labelling had occurred earlier, clones would have contained more stem cells, and if later, fewer; there is no reason for all members of a clone to share a pattern of behaviour when the level of grouping is so arbitrary. The original hypothesis of clonal succession (Kay 1965), proposed that coordination of clones might be spatial or involve a positive feedback loop. Little evidence of these feedback loops has emerged in the fifty years since the hypothesis was formulated. In the absence of this evidence it is hard to reconcile clonal succession with polyclonal haematopoiesis.

Sun and colleagues found that only 7% of bone marrow granulocyte tags were shared with nascent pro/pre-B cells, which was interpreted to mean that haematopoiesis was largely driven by unipotent progenitors. There are, however, other plausible interpretations of this finding. Firstly, analysis of their results is complicated by the fact that both stem and progenitor cells are labelled: tags from unipotent progenitors will necessarily be short-lived and lineage-restricted, and will far outnumber tags in stem cells and multipotent progenitors. Putting that aside, as a long chase period was allowed, these results can be explained by the polyclonality of blood without the need to invoke clonal succession or consign the bulk of steady-state haematopoiesis to lineage-restricted populations. Let us suppose that 15,000 stem cells are present in a mouse (Abkowitz et al. 2002). Allowing for some neutral drift of the population over the chase period, these might retrace their ancestry to perhaps 12,000 stem cells at the time of labelling. The labelling strategy tags ~30% of HSCs, so ~4,000 clones are tagged. Sun et al. generally recapture a few hundred tags, so let us say

that 400 tags from granulocytes and another 400 from B lymphocytes are captured. Assuming that these are derived from HSCs that make exactly equal proportions of B lymphocytes and granulocytes, these each represent independent samples of 10% of the labelled stem cell clones, and we would only expect a 1% overlap between the two of them. This 1% of the total pool of labelled clones represents 10% of the labelled granulocyte clones that were captured, and so the authors would present this finding as showing that only 10% of tags from granulocytes were recaptured in B lymphocytes – a number not dissimilar from the 7% that they found. The same argument can be used for why granulocyte tags are infrequently recaptured. The overlap could be further reduced if the number of stem cells were larger (which it could quite conceivably be), some stem cells were even slightly biased towards one lineage or another, there were an imbalance in the number of tags recaptured from different samples, or if there were a degree of transposon reactivation during the chase period. Indeed, the transposon was found to have remobilized in the absence of induction in one of 24 secondary colonies derived to test the leakiness of their assay. Nonetheless, this study remains an elegant demonstration of the vast polyclonality of blood in unperturbed haematopoiesis.

Busch et al. took the alternative lineage-tracing approach of permanently inducing the expression of yellow fluorescent protein (YFP) in 1% of immunophenotypic HSCs and all their descendants (Busch et al. 2015). The population is studied as a whole: individual cells are not marked and consequently any stem cell heterogeneity is missed. By limiting dilution, they estimated that ~30% of immunophenotypic HSCs contributed to haematopoiesis over most of a mouse's adult life. As most strategies for isolating HSCs also mark some progenitors, this proportion could well be higher: it is conceivable that all stem cells contribute. No attempt was made to estimate the number of HSCs, but rather this proportion was combined with estimates from limiting dilution transplantation experiments of ~17,000 stem cells to derive a lower bound of ~5,000 HSCs actively contributing to haematopoiesis in the investigated time period. Slow equilibration between the proportion of the label in the stem cell and progenitor compartments was taken to mean that the bulk of haematopoiesis must be sustained by cells downstream of the most undifferentiated HSCs. Finally, YFP+ cells were transplanted into recipients in order to investigate the differences between *in vivo* and transplant haematopoiesis. They found very variable engraftment, with a mean of one in 33 engrafting. This helps to explain some results of oligoclonal haematopoiesis from transplantation experiments and strengthens the case for *in vivo* studies.

A similar study was performed by Sawai et al. (2016), where a different system was used to lineage-trace a high proportion (80-90%) of HSCs with the least differentiated phenotype and slowest proliferation rate. Longitudinal analysis allowed a precise timing of when the label appeared in different compartments and the kinetics of differentiation. Broadly, their results agree with those of Busch et al. and Sun et al. on the large numbers of stem cell clones contributing to haematopoiesis, although they find that the rate of transition from stem cell to mature blood cell types is faster, estimating that 3-8% of the labelled HSCs enter differentiation every day. This supports a more traditional model in which progenitors are less long-lived.

Finally, a polylox recombination system which marks cells uniquely by randomly excising and inverting ten cassettes was used by Pei et al. (2017) to mark embryonic cells in mice at the emergence of HSCs. 10 months after birth, barcodes were detected, showing that the adult HSC compartment is a mosaic of hundreds of embryonic clones, ranging in size from 0.2-3.8%, mostly with multilineage output. Since adult clones are nested within embryonic clones, the number of embryonic clones forms an extreme lower bound for the number of active stem cells.

In summary, all studies of *in vivo* unperturbed haematopoiesis in mice support a polyclonal model of haematopoiesis where large numbers of adult stem cells, likely in the thousands, contribute to haematopoiesis. It could be argued that they have not changed our understanding of haematopoiesis dramatically relative to transplantation studies, but they do at least demonstrate similar findings in a more physiological model. It should be noted that even in animal models there are no direct estimates of the number of active stem cells at steady state.

I.3.b. Estimates of human HSC numbers

65 million years of evolutionary divergence, coupled with the long lifespan and large size of humans indicate that – even if we had good estimates from mice – it may not be sufficient to extrapolate from mouse studies to estimate human stem cell numbers and dynamics. Mice are typically studied under pathogen-free conditions, the proportion of peripheral blood cells that are myeloid in humans is larger, and there are many known differences between HSCs from the two species, including immunophenotypic definition, cytokine requirements, and differences in the HSC niche (Doulatov et al, 2012, Abkowitz et al.1996, Catlin et al. 2011, Larochelle 1996). Below,

I discuss experiments that aim to estimate stem cell numbers in humans, first by transplantation into mice or other humans, and then in unperturbed haematopoiesis by necessarily indirect means.

I.3.b.i. Xenotransplantation of human cells into animals

Xenotransplantation studies face all the caveats of autologous and allogeneic transplantation, with the added complication of differences between the human niche and that of the recipient immunosuppressed animal. Furthermore, they have produced varying results, with the estimate of the number of stem cells varying from 250,000 (Wang et al. 1997) to ~7,000 (discussed in Bystrykh et al. 2012). It is unclear how to reconcile these experiments, although differences in the degree of immunocompromise of the animal, the transplantation regimen, and the amount of *ex vivo* manipulation of stem cells may all be important.

I.3.b.ii. Transplantation of human cells into humans

The advent of gene therapy over the last few decades has allowed analysis of unique gene insertion sites as a side-product of clinical trials. Although initial attempts were leukaemogenic through insertional mutagenesis (Hacein-Bey-Abina et al. 2003), more recently it seems that insertion of the genes acts as a neutral tag. In addition to the caveat that transplantation measures potential rather than fate, these patients are by definition unwell. Despite the fact that analysis of viral insertion sites is only semi-quantitative, counting the number of unique insertion sites detected years after transplantation provides a lower bound on the number of active stem cells in this non-physiological setting. At least hundreds (Cartier et al. 2009) or thousands (Biasco et al. 2016) of stem cells contributed to long term multilineage haematopoiesis in these patients.

I.3.b.iii. Indirect methods of studying native human haematopoiesis

Methods for studying native human haematopoiesis have been based on the detection of markers that vary naturally between different somatic human cells. Firstly, in women heterozygous for an X-linked marker gene, detection of the proportion of cells that have inactivated either X chromosome (X chromosome inactivation or XCI skewing) provides insights into population dynamics. Early work inferred from the small degree of XCI skewing in most women that the number of active stem cells must be at least 400, based on binomial statistics (Buescher et al. 1985). A more recent study (Catlin et al. 2011), used the proportion of women with XCI skewing as a proxy for the rate of clonal drift. The rate of clonal drift is determined jointly by the population size and the rate of symmetric cell divisions (see below); knowing any two of these three terms allows the third to be inferred, but one is insufficient. Catlin et al. assumed that humans had the same number of stem cells as had been inferred in mice and cats by limiting dilution transplantation experiments (discussed above) and used that to infer the rate of symmetrical stem cell divisions. Furthermore, this study predates the description of clonal haematopoiesis (General Introduction), which may be responsible for a proportion of the skewing observed, particularly in the older of the two cohorts used. Finally, the model of how XCI skewing occurred is unusual: it was assumed to be due to hemizygous selection of one X allele over the other, with the parameter for the strength of the selective advantage drawn from a distribution based on the strength of hemizygous selection in Safari cats (cats that are the offspring of the distantly-separated South American and Eurasian breeds). There is little evidence of hemizygous selection in humans beyond the observation of skewed XCI. This Catlin et al. article is often cited in relation to the number of human stem cells despite the fact that no attempt is made to estimate it there (for example, Welch et al. 2012, Laurenti and Gottgens 2018, Young et al. 2016).

Secondly, analysis of telomere length distributions over the life of patients (Werner et al. 2015) provides a window into tracking unperturbed haematopoiesis. Telomerase is expressed in HSCs, but at insufficient levels to stop telomere attrition completely over the course of life (Lansdorp 2008). Werner et al. compared simulations to observed telomere data to infer the proportion of symmetric to asymmetric stem cell divisions over life. The underlying principle is that, under a model of purely asymmetric stem cell divisions, the telomeres shorten in all stem cells at the same rate, resulting in a Poisson distribution of telomere lengths over the population of cells, where the

lambda of the distribution decreases linearly with age. In contrast, in a model with mostly symmetric stem cell divisions, some stem cells have been through many more cell divisions than others. More stem cells are held in reserve, without cycling, and the decay in telomere length is logarithmic. Analysis of telomere data from a range of patients of different ages showed that the decrease in telomere length in the first decade of life is logarithmic and becomes linear later, indicating many symmetrical stem cell divisions and an increase in population size in childhood which then plateaus in adolescence. This, incidentally, argues against a model of clonal succession in adulthood.

Taken as a whole, this large body of work across humans and animal models suggests that the more physiological the model, the greater the evidence for a large number of multilineage clones driving blood production at steady state. Nonetheless, no studies have estimated *ab initio* the number of stem cells in humans, due to the ethical contraints on inducing clonal markers in healthy, unperturbed humans.

I.4. HSC clone lineage biases

Over the past fifteen years or so a number of findings have challenged the textbook model of haematopoiesis, in which a homogeneous pool of stem cells – through a rigid series of progenitors with defined lineage restrictions – produces the right balance of all cell types. In particular, the introduction of assays of single cells, whether colony-forming assays (Doulatov et al. 2010), transcriptomic analyses (Paul et al. 2015), or single cell barcoding experiments (Perie et al. 2015) have suggested alternative patterns of progressive lineage restriction and generally pointed towards lineage priming occurring higher up the hierarchy than anticipated by the original model, which was derived from observations of bulk populations that shared a surface marker (Laurenti and Gottgens 2018). However, this work is not the focus of the present chapter: the discussion here is restricted to assays of the cell types that descend from *clones* of stem cells (rather than the precise series of progenitors that lie between an individual stem cell and its mature progeny) as that is of greater relevance to the results presented below.

I.4.a. Transplantation studies in animals

Transplantation experiments in mice have shown that cells that meet the criteria for stemness can be further stratified based on their surface markers into those with greater or lesser self-renewal potential (Kent et al. 2009), indicating intrinsically-determined HSC heterogeneity. Furthermore, immunophenotypic HSCs can be biased towards producing one mature cell type or another (reviewed in Copley et al. 2012, and in Schroeder 2010). Some HSCs have been found to be biased towards the myeloid or lymphoid lineage (Muller Sieburg et al. 2004, Dykstra et al. 2007, Challen et al. 2010, Muller-Sieburg et al. 2012), and some towards the megakaryocyte-erythrocyte lineage (Sanjuan-Pla et al. 2013). Interestingly, HSC biases are recapitulated upon their transplantation, suggesting that this is a somatically heritable phenotype, although it may additionally be modulated by extrinsic signals such as growth factors (Muller-Sieburg et al. 2004, Dykstra et al. 2007, Challen et al. 2010).

I.4.b. *In vivo* studies in animals

As previously discussed, transplantation experiments and colony assays determine a cell's potential rather than its fate had it remained in its host under physiological conditions. HSCs *in vivo* could therefore be more or less biased than is suggested by *ex vivo* experiments, including the experiments that built the classical roadmap of haematopoiesis in the 1980s and 1990s. As described in the section on estimating stem cell number above, it has recently become possible to induce unique and heritable tags in haematopoietic stem and progenitor cells such that their clonal output can be determined in homeostasis. In a study that used the same sleeping beauty transposon system as Sun et al. (2014), after eight weeks of chase many tags were shared across all lineages except for megakaryocyte precursors, suggesting most stem cells are ancestral to multiple cell types, with the exception of megakaryocytes (Rodriguez-Fraticelli et al. 2018). Conversely, a large number of tags were found only in megakaryocytes, indicating that in unperturbed haematopoiesis a reasonable proportion of stem cells make megakaryocytes and negligible numbers of other cell

types. Importantly, some of these megakaryocyte-restricted HSCs produced multilineage output when transplanted into a recipient, highlighting the suggestibility of these biased HSCs.

4.c. Observations and transplantation studies in humans

In humans, it has been observed that with increasing age a myeloid bias develops (Pang et al. 2011). Two explanations have been put forward: stem cells represent one congruent population without bias in youth, but a bias towards the myeloid lineage develops with increasing age; or the stem cell pool represents at least two groups, one of which has a myeloid bias and one a lymphoid bias, and over the course of life the myeloid-biased group gradually predominates over the other (Schroeder 2010). Which of these is the case remains to be determined. Gene therapy studies have provided the opportunity to track transplanted cells in humans using the site of vector integration as a unique clonal mark. In a study of four patients with Wiskott-Aldrich syndrome who were transplanted with autologous genetically-modified HSPCs, integration sites were frequently shared between myeloid and lymphoid cells (unfortunately, no separate analysis of B and T lymphocytes was reported) showing the existence of multipotent stem cell clones (Biasco et al. 2016). In contrast, in one B-thalassaemia patient who received a genetically modified HSPC transplant, a long-lived (33 months) B and T lymphocyte-deficient clone was observed, suggestive of a myeloid biased stem cell (Cavazzana-Calvo et al. 2010). Intriguingly, the lentiviral vector in this clone had integrated into and increased the expression of *HMGA2*, a known regulator of HSC self-renewal durability (Copley et al. 2013). Its behaviour may therefore reflect normal biology or be due to insertional mutagenesis.

In summary, there is compelling evidence for some degree of HSC lineage bias, although most stem cell clones in humans and unperturbed mice still seem to have multipotent output. Very little is known about biases in unperturbed human haematopoiesis. What bias there is in animal models seems to be both cell-intrinsic and flexible in the case of stress. One can conjecture that the ability of HSCs to modulate output based on environmental stressors might be advantageous in evolutionary terms, given that the haematopoietic system is faced with the challenge of massively amplifying the production of one cell type – such as granulocytes in acute infection or platelets in haemorrhage – without necessarily scaling up the production of other cell types.

I.5. Spontaneous somatic mutations as natural barcodes

To track the clonal dynamics of stem cells in a healthy unperturbed human requires a label that occurs naturally in cells at a high rate, persists over life, and can be reliably detected. Somatic mutations provide such a marker. Induced mutations were famously used by Till and McCulloch to study normal haematopoiesis, and spontaneous mutations have been used to track the clonal dynamics of tumours (Ding et al. 2010, Ding et al. 2012, Tsao et al. 1999). We build on this work to use spontaneous somatic mutations to study normal blood.

In blood, all available evidence – limited though it is – indicates that the mutation rate over life is relatively constant. Welch and colleagues cultured and whole exome sequenced three single HSPCs from each of seven healthy individuals spanning in age from birth to the eighth decade of life (Welch et al. 2012). Although the total number of mutations per exome is low, these data show a linear accumulation of mutations. Furthermore, in acute myeloid leukaemia the mutation burden increases linearly with the age of diagnosis (Welch et al. 2012, Alexandrov et al. 2015). More noise is present in analyses of blood cancers, most likely due to a combination of occasional aberrant mutational processes in some cancers which inflates the mutation burden and a variable time to the most recent common ancestor (MRCA) of the tumour (a longer time to the MRCA decreases the number of mutations called since many subclonal mutations will not be detected). Not only does the mutation rate seem constant, but the mutation burden is sufficient for us to detect large numbers of potential clonal markers per cell. Welch et al.'s normal HSPC exomes acquired mutations at a mean of 0.13 exonic mutations per year of life. Assuming that approximately 1% of mutations occur in coding regions, this translates to 13 mutations per genome per year, or 780 in a 60 year-old. Most AMLs have approximately 400-500 mutations per genome (Cancer Genome Atlas Research Network et al. 2013, Welch et al. 2012). Thus, by sequencing 100-200 normal blood genomes in a 60 year-old, one could hope to detect over 100,000 nested clonal barcodes. Such an approach has become feasible in recent years because of the falling costs of sequencing and reliable methods of culturing single HSPCs into large colonies amenable to whole genome sequencing.

I.6. Population genetics models for studying stem cells

In this study, we use methods from the field of population genetics to study the clonal dynamics of blood. We conceive of blood stem cells as a population of asexually-reproducing individuals. It is therefore necessary to introduce briefly the population genetics models that we apply: the neutral Moran model and the neutral Wright-Fisher model, and the associated coalescent process.

The population of stem cells can be thought of as behaving like the neutral Moran model of population genetics (Moran 1958, Gladstein 1978). In this model, in every generation two different individuals (or stem cells, in our case) are chosen from the population at random; one dies (or differentiates), and the other reproduces (a symmetrical doubling stem cell division), leaving two daughter individuals which carry the set of mutations carried by the parent individual. The waiting time from the birth of an individual until it next gives birth or dies, in the Moran model, is equivalent to the waiting time from the origin of a stem cell (in a symmetrical cell division) until the first time its stem cell descendant undergoes a new symmetrical cell division. We are blind to asymmetric cell divisions, as they do not change the genotype composition of the stem cell population.

Associated with the Moran model, which evolves forwards in time, is a genealogical process, where a sample of individuals from the population retrace their ancestry backwards through time. Here, in every generation back in time, a coalescent event may or may not occur among the ancestors of the sample. The occurrence of coalescent events is independent in each generation, and therefore the number of Moran model generations (going backwards) until a coalescent event has a geometric distribution, just as the number of coin tosses until a head is seen follows a geometric distribution. Kingman's coalescent is a stochastic process that generates genealogies, for any specified sample size, exactly as if the forward time process is a neutral Moran model (Kingman 1982). The timings of the coalescent events and the random pairing of lines of descent together determine the topology of the phylogenetic tree and the branch lengths.

When the population size N is large, N/2 generations of the Moran model closely approximates to one generation of the Wright-Fisher model (Fisher 1922, Wright 1931), which is more convenient to simulate forwards in time. In every generation of the Wright-Fisher model, the genotypes of individuals are obtained by sampling with replacement from the parent generation.

70

As with the Moran model, a reverse time genealogical process is associated with the forward time Wright-Fisher model, where every cell in one generation is assigned to a parent in the previous generation, independently of other cells; thus, a coalescence is generated when two cells happen to be assigned to the same parent.

For a sample of cells from a given population, the probability of two cells picking the same parent in the preceding generation is dependent on the population size. If we only have two cells in our sample, and a population size of N, the probability that both will pick the same common ancestor in the preceding generation is 1/N. When the population size is large relative to the sample size, either zero or one coalescent event will occur in a given generation; that more than two cells should happen to pick the same common ancestor out of a large pool becomes very unlikely, making the model more similar to the Moran model. As with the Moran model, the number of generations until a coalescent event follows a geometric distribution. The probability of two cells not coalescing in a Wright-Fisher generation is $(1-(1/N))$. Therefore, the probability of coalescing G generations ago is the probability of not coalescing for (G-1) generations, and then of coalescing in the G^{th} generation: $P(G) = (1-(1/N))^G*(1/N)$.

It follows that if the population size is 10 times larger, the probability of two cells picking the same parent is 10 times smaller, and 10 times more generations are required to achieve the same probability of a coalescent event. This means that the pattern of branch points in a phylogeny derived from a certain population size having gone through a certain number of generations will be the same as that from a population 10 times as large having gone through 10 times as many generations.

It can be shown mathematically that the mean waiting time from the birth of an individual until it next gives birth or dies is the same as the generation time of the Wright-Fisher model (as shown by Kevin Dawson in the technical appendix of Lee-Six et al. (2018)). This in turn is equivalent to the mean time between symmetric cell divisions along a given line-of-descent. Because of this equivalence, we can simulate stem cell population drift using the Wright-Fisher model, which is computationally more convenient because fewer generations are required and each is of a fixed chronological duration.

Results

R.1. Experimental design

Here follows a brief overview of the experimental design. In order to exploit the somatic mutations of normal stem cells as clonal markers, our experiment was designed in two phases: a 'mutation discovery' phase, in which somatic mutations were discovered by whole genome sequencing of colonies derived from single HSPCs from one individual; and a 'mutation quantification' phase, in which the mutations that we had discovered by whole genome sequencing were detected and accurately quantified in mature blood fractions from the same individual, telling us the output of the clones that we had discovered. These data were then analysed with methods adapted from population genetics and ecology, effectively treating each stem cells as an individual within a population. An outline of the experimental workflow is provided in figure 1.1a.

R.2. Mutation discovery phase: HSPC isolation, clonal amplification, and sequencing

A bone marrow aspirate was performed (by Brian Huntly) on one healthy 59 year-old male volunteer with normal blood counts and no past history of blood disorders. Peripheral blood was taken on the same day. We also obtained a buccal swab as a germline control. An individual of this age was picked because he would be sufficiently old to have acquired a workable number of mutations (assuming the mutation rate reported by Welch and colleagues (2012)) but was young and healthy enough to allow, potentially, decades of longitudinal follow-up. Studying a male was advantageous because he would have only one copy of the X chromosome such that the allele fraction of X chromosome mutations could be evaluated more accurately. On a diploid chromosome, even a perfectly clonal mutation can still fluctuate in allele fraction because of binomial sampling of mutant and wild type reads, complicating the estimation of the true allele fraction, whereas a homozygous clonal mutation should be present on every single sequencing read.

Cells obtained from the bone marrow and peripheral blood harvest were sorted by flow cytometry. This served the dual purpose of isolating the immunophenotypic cell populations of

interest based on their cell surface markers and of obtaining single cells. Cells were cultured in methylcellulose or liquid culture for 2-6 weeks until colonies were derived that were sufficiently large to be whole genome sequenced. Overall, we harvested 285 stem cell-derived colonies and 390 progenitor-derived colonies, split across common myeloid progenitors (CMPs), granulocyte macrophage progenitors (GMPs), and megakaryocyte erythroid progenitors (MEPs). Progenitor cells that produced a colony with a different morphology to their immunophenotypic definition were not picked. All of the colony derivation work was performed by David Kent, Mairi Shepherd, Nina Friesgaard-Oebro, and Miriam Belmonte, with advice on culturing conditions from Elisa Laurenti.

When we began this study, whole genome sequencing required large amounts of DNA (~100ng), and not all colonies could be expanded to such a large size. This imposed a selection bias on the colonies that we sequenced towards colonies from less differentiated cell types that have greater expansion potential. It is also possible that colonies with certain mutations may not amplify well under normal culturing conditions relative to their wild-type counterparts, as has been observed for cancer-derived organoids from a number of tissues (van de Wetering et al. 2015). We chose to whole genome sequence at ~15X the 100 stem cell-derived colonies with the largest amount of DNA, and the 98 progenitor-derived colonies with the largest amount of DNA.

R.2.a. Quality control of colony whole genomes

The clonal origin of each HSPC colony was determined by visual examination of histograms of their mutant allele fractions. In a colony that is truly single cell-derived, somatic mutations should only be present on one of two copies of an autosome, and so their allele fraction should be binomially distributed around 0.5. For the X chromosome of our male individual, allele fractions should be exactly 1. For 140 colonies, this was indeed the case. However, for the remaining 58, the allele fractions were lower, indicating that these colonies were derived from one or more cells (figure 1.1b). This is probably a result of colonies growing into one another in methylcellulose. Because some of the stem cells were grown in liquid culture, while all progenitors

Figure 1.1. Experimental design and quality control. a, The experimental workflow. Following bone marrow aspiration, single haematopoietic stem cells (HSCs) and haematopoietic progenitor cells (HPCs) were isolated by fluorescence-activated cell sorting (FACS) and expanded in liquid culture or methylcellulose into colonies that could be whole genome sequenced. Somatic mutations discovered by sequencing were used to construct the phylogeny. This is the mutation discovery phase. In the mutation quantification phase, peripheral blood samples underwent deep targeted sequencing to quantify mutations that had been found in the discovery phase. **b**, histograms of the variant allele fraction (VAF) of mutations were used to identify whether colonies were derived from one cell or more. An example clonal sample (BMH66) and polyclonal sample (BMP12) are shown. In a clonal sample, the VAF of autosomal mutations should be distributed around 0.5, and that of sex chromosomes should be 1. For a subclonal sample it should be lower. Occasionally, even in a clonal sample lower VAFs are seen due to the failure to detect a mutation on a read, or a read from another locus being aberrantly mapped to the locus in question and lowering the apparent coverage, or a mutation acquired *in vitro*. **c**, histograms of the burden of substitutions and indels across all 140 clonal samples. **d**, the distribution of mutations from all 140 clonal samples around the genome, shown as a circos plot. The outermost ring of the circos plot depicts the karyotypic ideogram. Moving inwards, base substitutions are shown as rainfall plots where the height of the dot in the substitution ring is proportional to log10 of the distance to the next mutation and with the colour of the dot illustrating the base change, as shown in the key. (c) a comparison of the substitution burden between stem cells and progenitor cells. There were not significantly more mutations in progenitors than stem cells (p=0.14, Wilcoxon Rank Sum test).

a

Bone marrow aspirate → Isolation by FACS (HSCs, HPCs) → Clonal expansion → Whole genome sequencing of colonies → Phylogenetic tree construction → Bulk peripheral blood → Targeted sequencing

"Mutation discovery" phase

"Mutation quantification" phase

b

BMH66 autosomes

BMP12 autosomes

BMH66 sex chromosomes

BMP12 sex chromosomes

Substitutions

Variant Allele Fraction

c

Colonies

Substitutions

Indels

d

C>A   T>A
C>G   T>C
C>T   T>G

e

Substitutions

HSCs   HPCs

were grown in methylcellulose, the stem cells were less affected than the progenitors. Furthermore, MEPs grow as more compact balls than CMPs and GMPs, and so were relatively spared. For these reasons, after removal of 58 polyclonal colonies we were left with 63 bone marrow-derived HSCs, 16 peripheral blood-derived HSCs, 38 MEPs, five CMPs, and eight GMPs.

We could also use allele fractions to assess what proportion of mutations might have occurred *in vitro* during the process of colony derivation. Unless a selective sweep occurred *in vitro* (which seems unlikely), only mutations that were present in the sorted founder cell will be fully clonal in a colony. The proportion of mutations that are subclonal within a colony with a VAF peak at 0.5 informs on the mutation rate. Because of binomial sampling of mutant and wild type reads combined with the low sequencing depth, it can be difficult to tell if a mutation on an autosome is truly subclonal. This is further complicated by sequencing false negative artefacts, and more frequently mismapping of reads to this location in the genome and appearing as a wild type read. Mismapping events are relatively frequent with Illumina HiSeqX-sequenced and BWA-mem-mapped data. The bottom left-hand panel of figure 1.1d demonstrates two X chromosome mutations with a VAF below 1. Visual inspection of these mutations shows that one is present on 4 out of 8 reads. This mutation is in a relatively low complexity part of the genome and could represent either a sequencing artefact or a mutation acquired *in vitro*. The other mutation is present on 12 out of 13 reads. The read that does not support the mutation maps uniquely to that position in the genome, and so probably represents a sequencing error. The variant, however, is likely to be a true clonal mutation. Thus, even true clonal mutations on the X chromosome may not have a VAF of 1, but the mutation is likely still to be present on a large proportion of the reads. Variants acquired *in vitro*, however, should be present in half of the cells in a clone or fewer. Therefore, a more informative metric of the proportion of variants that might have occurred *in vitro* is the proportion of X chromosome variants at variant allele fraction of 0.5 or less (assuming equal amplification of both daughters of the first cell division). Considering only positions with depth 10 or above, at which the allele fraction can be estimated accurately, and considering only colonies with more than five high coverage X chromosome mutations called, the mean percentage of mutations per colony at a variant allele fraction of <0.5 is 5.6%.

An additional source of evidence that *in vitro* mutations do not contribute substantially to our mutation catalogue is that two colonies share 1,115 substitutions and have only 51 and 87

private mutations respectively. *In vitro* mutations will typically be private, so fewer than 7% (87/(87+1115)) of the total mutations called in these colonies could have occurred *in vitro*.

R.2.b. Genomics of normal blood cells

A mean of 1,023 (range 815-1,210) single base substitutions and 20 small insertions and deletions (indels) (range 2-37) were observed across the 140 colonies (figure 1.1c). No copy number changes or structural variants were detected. Progenitor cells had slightly more substitutions than stem cells (mean of 1,038 for progenitors and 1,024 for stem cells), but the difference was not statistically significant (p=0.14 by Wilcoxon rank sum test) (figure 1.1e). This result does not necessarily mean that progenitors do not have an increased mutation rate over stem cells, as the line-of-descent to any given progenitor will have spent nearly all its time as an HSC. Imagine that a progenitor that we sampled exited the stem cell compartment one month before sampling, and over that last month acquired mutations at a 10-fold greater rate than in stem cells. In this scenario, the progenitor would only have 10-20 more mutations than its stem cell cousins.

Comparison of the mutation burden that we detect in normal HSPCs and that reported in acute myeloid leukaemia (AML) supports at the level of the whole genome the findings by Welch et al. that AML exomes do not have more mutations than normal HSPC exomes. Indeed, many AML genomes have a lower mutation burden than that reported here: the mean number of mutations per AML genome in a cohort with a mean age of 55 was just over 400 (Cancer Genome Atlas Research Network et al. 2013). Some of this difference may be due to different sequencing platforms and mutation calling algorithms. In addition, in some patients, contamination of the matched normal tissue with leukaemic blasts (such that somatic mutations may appear to be germline) could result in a significant underestimation of the mutation burden. Another – and biologically interesting – component is that the most recent common ancestor of the leukaemia may have existed a long time ago; if a single leukaemic cell were sequenced, it might have the same mutation burden as a single normal cell, but bulk sequencing misses a large proportion of the subclonal mutations in the tumour. Depending on the strength of the effect of the first three factors, comparison with our data would indicate that the MRCA of many AMLs occurred years to decades before diagnosis.

Figure 1.2. Trinucleotide context of mutations in normal blood. a, the counts of substitutions in each of 96 trinucleotide contexts from all 140 clonal samples combined. **b**, **c,** the same figure for representative individual progenitor colonies (**b**) and stem cell colonies (**c**). **d,** Non-negative matrix factorisation of the blood trinucleotide context along with a panel of blood cancers deconvoluted the substitutions in the our samples into two signatures. Signature A accounted for 30% of mutations and bears a strong resemblance to signature 1 (see text). Signature B accounted for 70% of mutations and was novel. **e**, trinucleotide context of mutations from three blood cancers, chosen for their similarity to the trinucleotide context of normal blood colonies. **f**, results from mutational signature decomposition of a panel of cancers. The mutations burden attributable to each signature from each sample of a given cancer type is shown. Signatures that are not signature A or B are all coloured grey. CLLE, chronic lymphocytic leukaemia; COAD, colorectal adenocarcinoma; DLBC, diffuse large B cell lymphoma; KIRP, kidney renal papillary cell carcinoma; LAML, acute myeloid leukaemia; MALY, malignant lymphoma; MDS, myelodysplastic syndrome; MPN, myeloproliferative neoplasms; MM, multiple myeloma; PBCA, paediatric brain cancer.

a

b BMP121, BMP14, BMP184

c BMH97, BMH98, BMH99

d Signature A, Signature B

e MPN, CLLE, LAML

f Signature A, Signature B
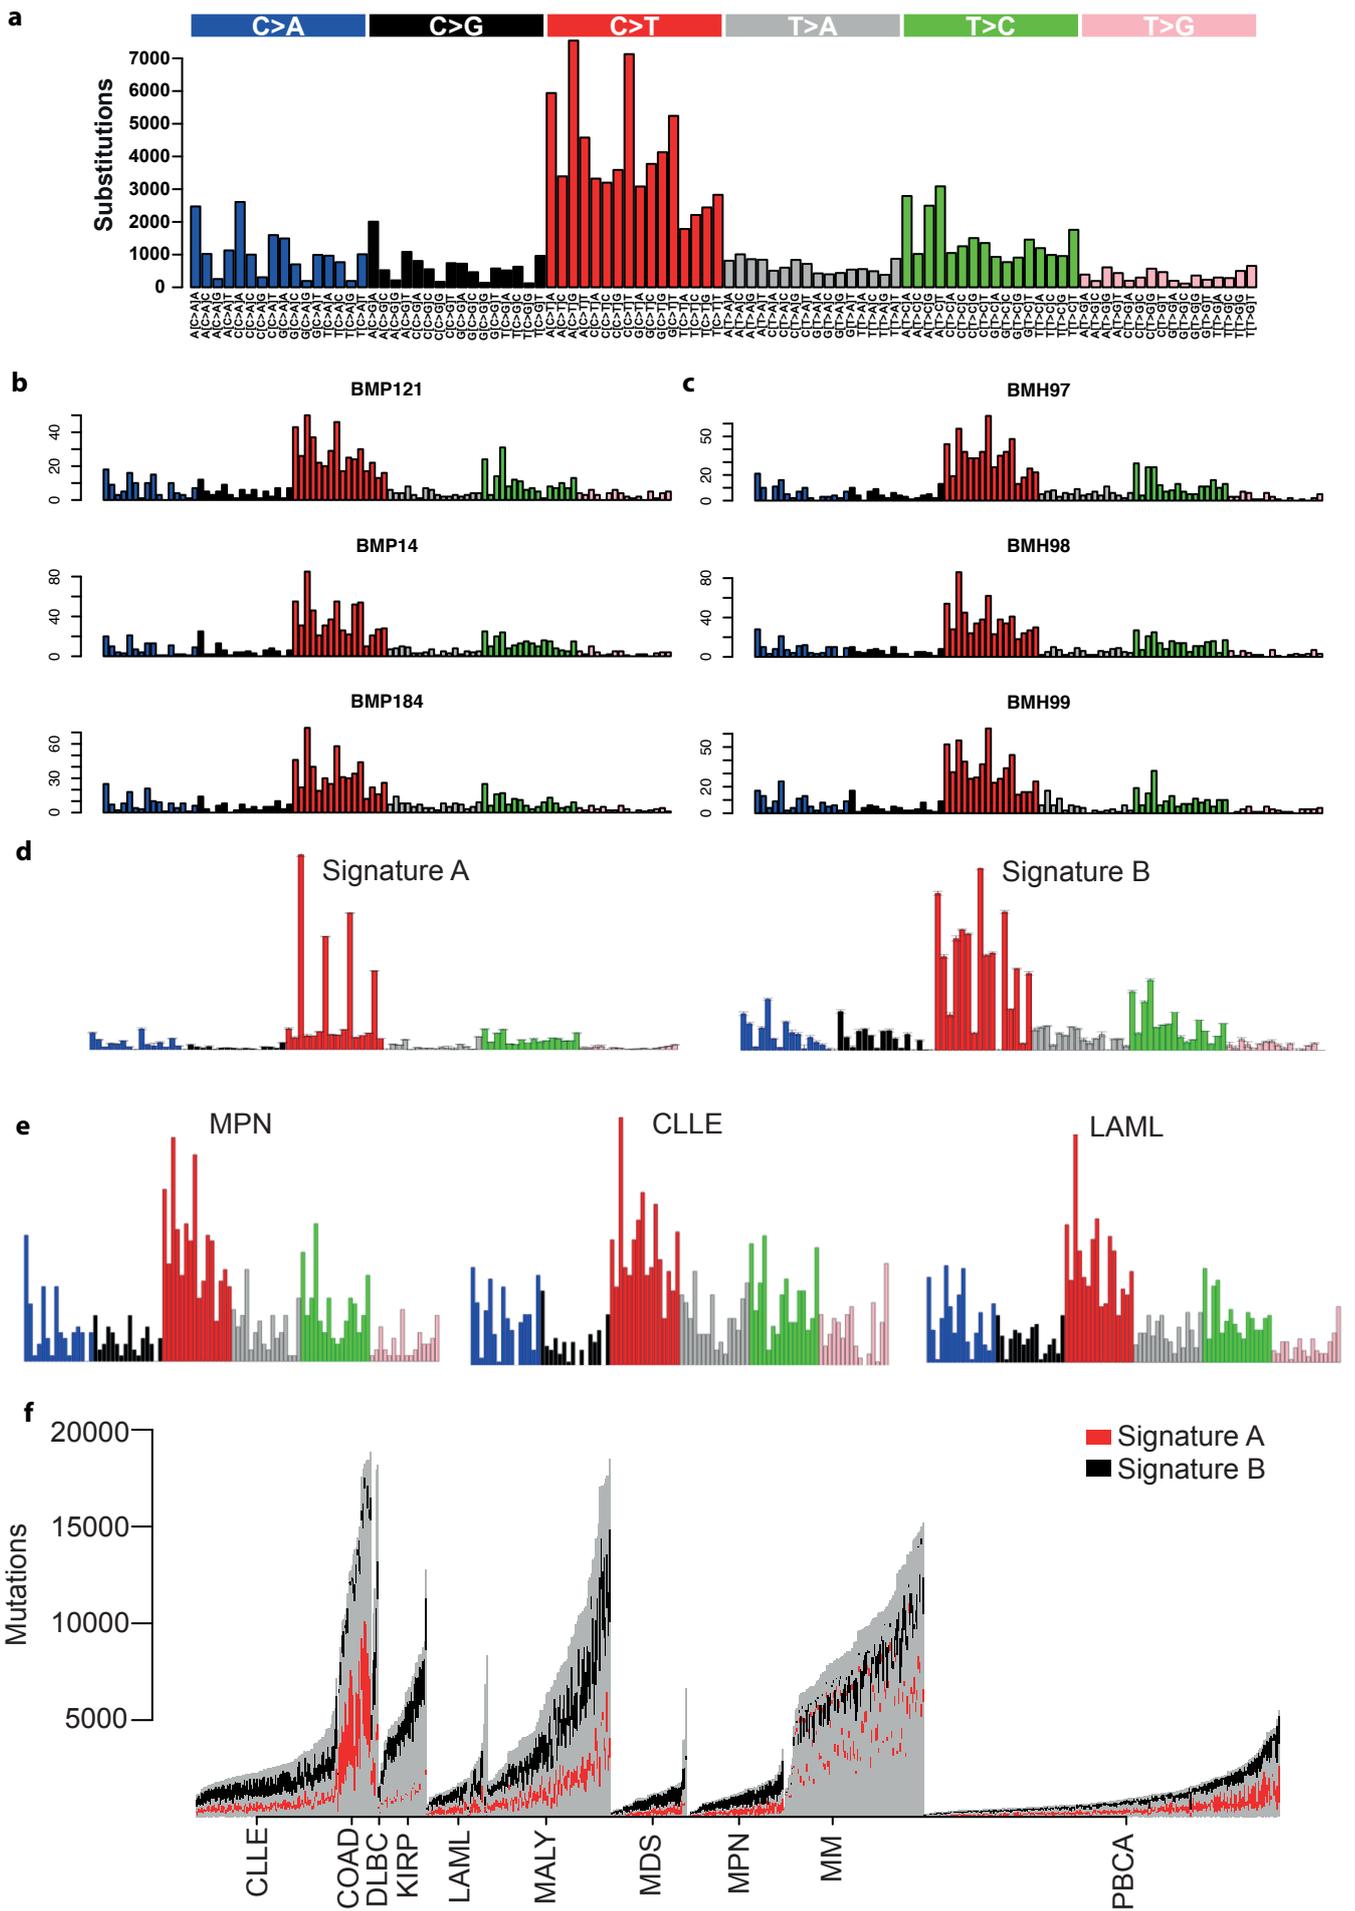
The somatic mutations in blood stem cell genomes that we observed were mostly C>T and T>C transitions, particularly in a context of NCT (the mutated base is underlined) (figure 1.2a). The trinucleotide spectrum of mutations in all of our samples was similar (figure 1.2b and 1.2c). A similar spectrum has been observed in a study of clonal haematopoiesis (Zink et al. 2017). Visual inspection of all blood cancer genomes included in the Pancancer Analysis of Whole Genomes study revealed that it was present in some cancers (figure 1.2e). A mutational signature (General Introduction) extraction was run with the combination of our samples and all of these cancers. The trinucleotide context of our sample was deconvoluted and found to be ~30% a signature that similar to signature 1, a known signature present in all tissues that is discussed in Results Chapter 2, and ~70% a novel mutational process (signature B in figure 1.2d). Quantification of the contribution of this mutational process to all blood cancer genomes analysed indicates that it is present at some level in most blood cancers. It is responsible for a large proportion of mutations in myeloproliferative neoplasms, myelodysplastic syndrome, and AML, which have relatively quiet genomes and, as discussed above, for the most part have not departed significantly from the somatic mutational landscape of normal blood (figure 1.2f). In more mutated cancers, such as certain lymphomas, the signature is obscured by a vast excess of additional process, which may explain why it has not been reported previously. The signature contributed proportionately less to solid tumour genomes included as a negative control, although it was detected in kidney cancers and paediatric brain cancers. The cause of this signature, as with many mutational signatures, remains unknown.

The mutations were scattered across the genome (figure 1.1d), with <1% causing non-synonymous changes in protein-coding genes. We searched our 140 genomes for driver mutations, (which would affect our interpretation of clonal dynamics in this experiment) in two ways: first, by manual annotation of known cancer genes, and second, using a dNdS approach which detects the enrichment of non-synonymous mutations in a gene above what would be expected by chance, indicating positive selection (Martincorena et al. 2017). The dNdS approach therefore does not rely on prior knowledge of which mutations are drivers. dNdS also allows us to estimate a global measure of selection across all the mutations in the dataset. The global dNdS across all 140 blood genomes is 1.0010 (95% CI 0.889-1.127), with values greater than one indicating positive selection and values less than one negative selection. No genes were found to be under positive selection by dNdS, and no known driver mutations were found by manual curation of known drivers of myeloid

neoplasms. Finally, in the second phase of the experiment in which we performed targeted sequencing, we included 100 common hotspot mutations frequently observed in clonal haematopoiesis and myeloid malignancies, and no drivers were detected despite deep sequencing. We cannot exclude the presence of as-yet undiscovered drivers that are rare in our dataset, but all the evidence that we have gathered indicates that this individual's blood is evolving neutrally.

R.2.c. Construction of a phylogeny of HSPCs

A phylogeny of HSPCs was constructed based on the sharing of somatic mutations across the 140 colonies. Our confidence in each of the branch-points was quantified by Felsenstein's non-parametric bootstrapping method (Felsenstein 1989), with 1,000 bootstrapping replicates. The tree was constructed independently based on indels and on short tandem repeats and using all types of data combined, using bootstrapping in all cases. We applied all three of a maximum likelihood, maximum parsimony, and neighbour joining methods (although not all on every dataset, as some are not appropriate for certain data types (Methods)). In general, the structure of the tree was supported with high confidence: few branch-points that were supported with high confidence using different methods disagree with the tree that we present (figure 1.4) (Methods). Tree-building was performed with substantial help from Sebastian Grossman (distribution of the work-load described in Methods).

The phylogeny that was reconstructed had a deep branching structure, with most mutations private to a given cell: of 129,582 substitutions on the tree, 8,676 were shared amongst different colonies (figure 1.3a). The construction of a phylogeny allowed us to perform multiple analyses: we could time mutations relative to one another (mutations on a branch shared by two colonies necessarily occurred before those on a branch private to one of those colonies); determine the relationship of all cells to one another; and apply phylodynamic models that were developed in the fields of population genetics and epidemiology to our data. These are explored below.

Figure 1.3. Inferences from the phylogeny. a. phylogeny of 140 single haematopoietic stem and progenitor cells showing the relationship between cell types. At each tip of the tree is a colony. Branches connect colonies to each other to form a family tree. Branch lengths are proportional to the number of somatic mutations: thus, a branch that is ancestral to two colonies and 100 substitutions long reflects the fact that these two colonies share 100 substitutions that are not present in any of the others. Symbols at the tips of the branches represent the phenotype of the cell. **b**, the same phylogeny as in **a**, but showing only the first 10 mutations of molecular time. **c**, the number of descendants of each node for the first 10 mutations of molecular time, used to estimate the embryonic mutation rate (Methods). **d**, phylodynamic inference of the population size trajectory of the stem cell pool reveals changes in the effective population size of stem cells over life based on the timing of coalescences (branch-points) in our observed phylogeny. Shading illustrates different credibility intervals. The y axis is shown in units of 'population size multiplied by generation time' (cell-years) because the same distribution of coalescences can be generated from a population of 10 times the size with 10 times as many generations.

a

CMP   MEP   GMP   BM HSC   PB HSC

b

33 lines-of-descent by 10 mutations of molecular time

c

Polytomies

Dichotomies

Number of nodes before 10 mutations of molecular time

d

Credibility interval: 95% 90% 80% 60% 40% 20%

Mean:

Conception

59 years old
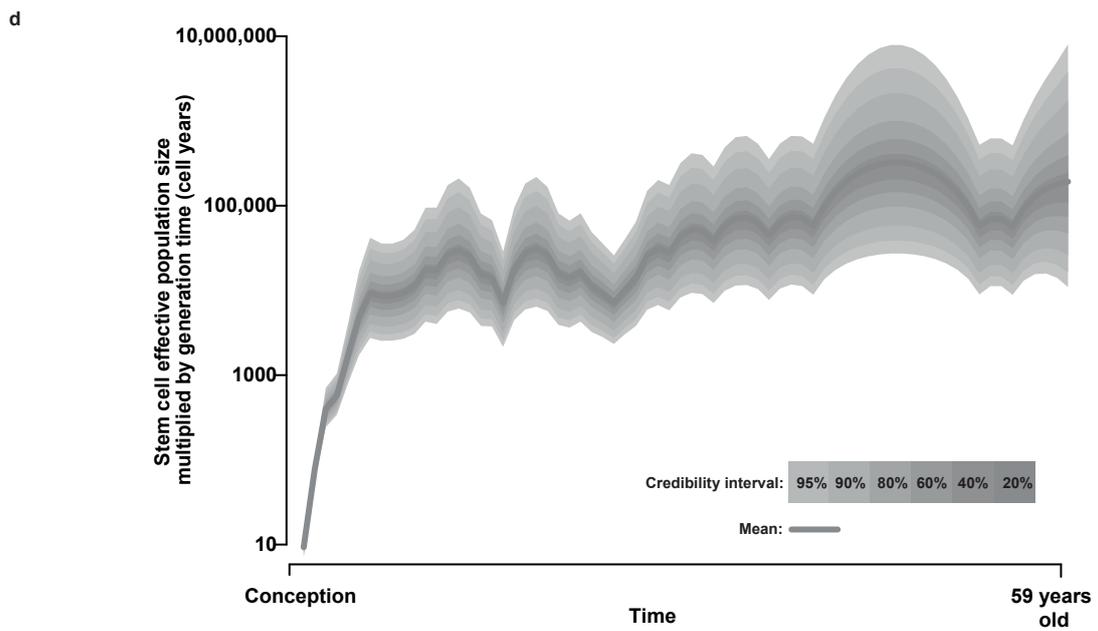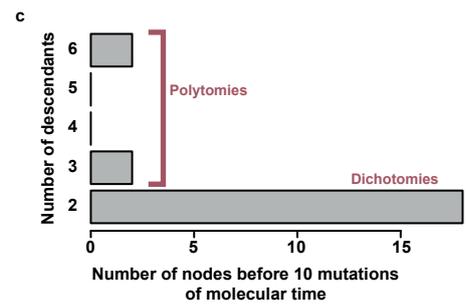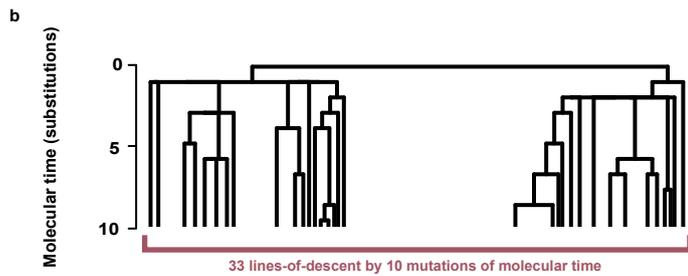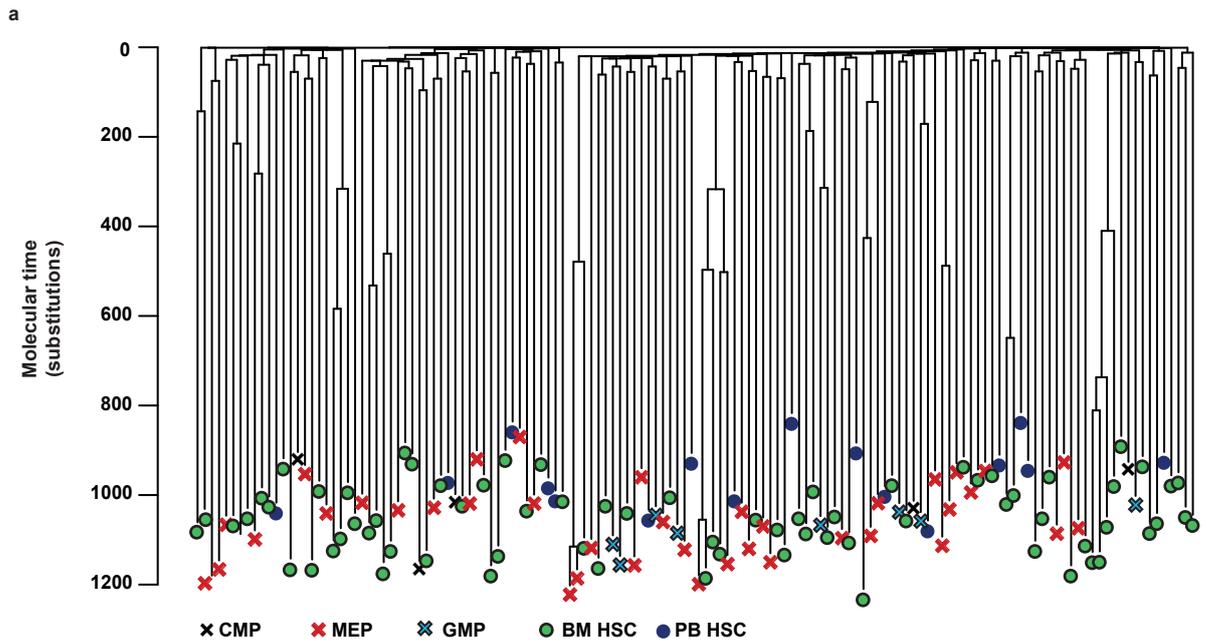
Time

Figure 1.4. Construction of the phylogeny using different methods. a, the phylogeny of cells, presented with p values next to every node, derived by bootstrapping the substitution matrix 1000 times, building a tree using SCITE for each replicate, and counting the proportion of the bootstrapped trees that support each node. **b-f,** phylogenies constructed using different datasets and methods. In each case the phylogeny was constructed using 100 bootstraps of the data, and the p value for each node shown underneath it. Branches are coloured by whether a branch ancestral to exactly the same descendants is also present in the SCITE tree, and are drawn with a thicker line if the branch is recovered in >=70% of bootstrap replicates. **b**, substitution and indel datasets combined, building the tree by maximum parsimony. **c,** substitution, indel, and neighbour joining datasets combined, building the tree by neighbour joining. **d,** substitutions, tree build by maximum parsimony. **e,** indels, tree built by maximum parsimony. **f**, Short tandem repeats, tree built by neighbour joining.

a
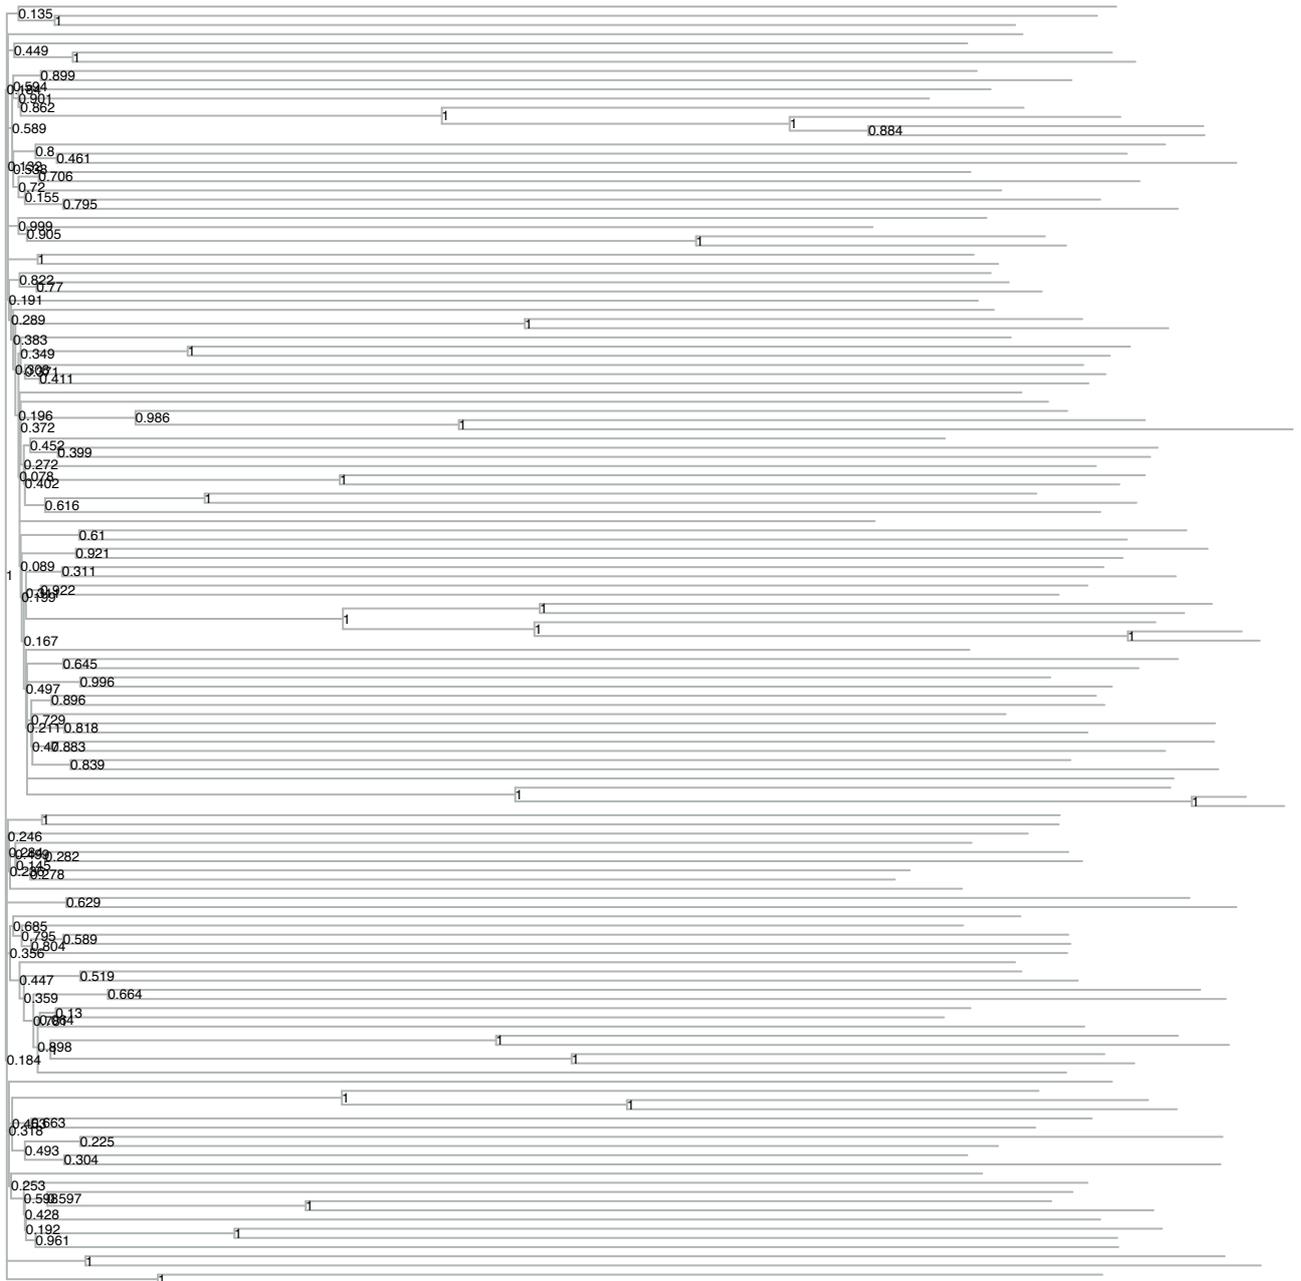
b
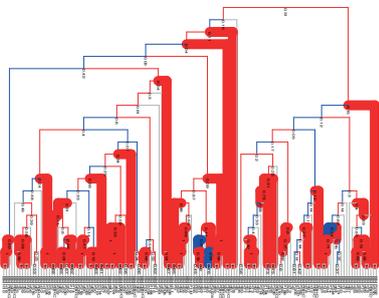
c
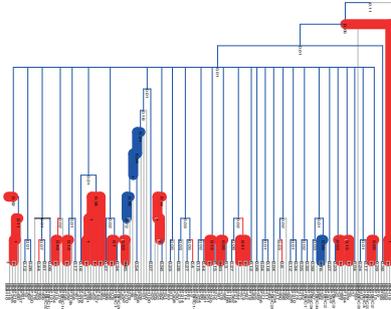
d
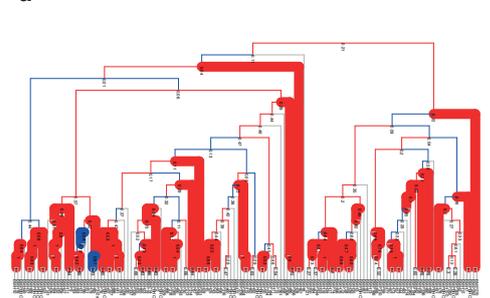
e
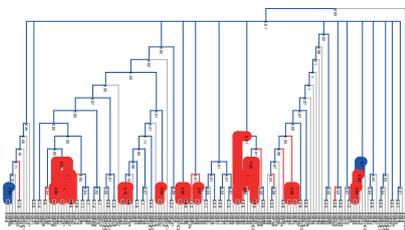
f
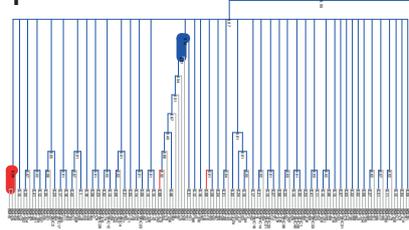
Present, high confidence
Absent, high confidence
Present, low confidence
Absent, low confidence
Uninformative (private)

R.2.d. Embryological insights

R.2.d.i. The ontogeny of blood

As one looks up the phylogeny, one is looking back in time into the somatic history of the haematopoietic system. The probability of capturing the most recent common ancestor of a population with a sample of size n for a randomly bifurcating tree is equal to (n-1)/(n+1); in our case, assuming that ours is a random sample of blood, 197/199=0.99. The earliest branch-point on our phylogeny therefore in all likelihood represents the cell division of the most recent common ancestor of blood. It partitions the tree into two uneven groups, one with one third of the cells (52/140) referred to as 'clade A' and the other with two thirds (88/140) referred to as 'clade B'. All cells in clade A all share one mutation ('mutation A') which is absent from clade B, and, conversely, those in clade B share a mutation which is absent from clade A ('mutation B'). We searched for these early embryonic somatic mutations in whole genome sequencing of the buccal swab sample. Strikingly, mutation A was found in the buccal epithelium at a clonal fraction of 1/3 (on 6 out of 30 reads) and mutation B at a clonal fraction of 2/3 (on 11 out of 33 reads). Mutation B was sequenced more deeply using a bait-set (discussed below), and found to occur on 124 out of 414 read (0.30%). It is possible that the buccal swab has some blood contamination, but it seems very unlikely that it should be made up exclusively of blood with no epithelial contribution at all. These data therefore indicate that the most recent common ancestor of blood is shared with the most recent common ancestor of buccal epithelium. Thus, at no point in embryogenesis is there a cell that is fated to produce all of blood and only blood. Furthermore, since blood is mesoderm-derived and buccal epithelium ectoderm-derived (Rothova et al. 2012), this common ancestor must precede gastrulation. Indeed, it seems likely that it existed very early, perhaps as one of the first cells in the embryo.

R.2.d.ii. Early embryonic mutation rates

Beyond this first branch-point, we observed a number of short branches, supported by 0-4 mutations (figure 1.3b). We were able to use the structure of the phylogeny to estimate the mutation

rate for the first 10 mutations of molecular time, presumably reflecting the first days or weeks of life. Firstly, the rate of branching in our tree informs on the upper bound of mutation rate per cell division. By 10 mutations of molecular time, we observe 33 lines-of-descent in our tree. Therefore, at least 33 cells must have existed in the embryo by this stage of molecular time. With doubling divisions, in the absence of cell death, it would take at least five generations to grow to this number of cells ($2^5 = 32$). Under this simplified scheme, the mutation rate is likely to be under 2 mutations per generation (10 mutations divided by 5 generations). Two principal caveats complicate this model: first, cell loss (either death or contribution to extraembryonic tissues exclusively, such that no descendants contribute to the adult); and, second, we may not have sampled every embryonic clone. Both cell loss and incomplete sampling of embryonic clones would mean that more generations need to have elapsed for us to obtain 33 lines-of-descent in our tree. Therefore, both would lower our estimate of the mutation rate per cell division. Thus, we can use an upper bound of 2 mutations per cell division.

The number of polytomies (or multifurcations) in this early embryonic tree also informs on the mutation rate per cell division. Polytomies in the tree can only occur if no mutations were acquired in a cell division. The lower the mutation rate, the more cell divisions are associated with no mutations and the more polytomies appear in the tree. Cell loss can make a true polytomy appear as a dichotomy, but not the other way around, and so increasing cell loss decreases the number of polytomies seen. The number of polytomies, therefore, also provides an upper bound: a combination of increased cell loss and decreased mutation rate would result in the same number of polytomies.

In the absence of cell loss, the following calculations allow us to estimate the mutation rate per cell division. In the first 10 mutations of molecular time of our phylogeny we observe 18 dichotomies, two trichotomies, and two hexachotomies (figure 1.3b,c). In the absence of cell loss, a trichotomy means that two cell divisions occurred, creating four branches, and one out of the four was associated with no mutations. Similarly, a hexachotomy means that five cell divisions occurred, creating 10 branches, and four of them were associated with no mutations. In total, the 32 cell divisions required to make 33 lines-of-descent created 64 branches, 10 of which were associated with no mutations, one for each of two trichotomies and four for each of two hexachotomies. From the Poisson distribution, the maximum likelihood of 10 branches with no mutations is 1.86 mutations per division (-log(10.64)), with a 99% confidence interval of 1.19 to

2.80. We can therefore conclude that in the absence of cell loss, mutation rate per cell division in the early embryo is likely to fall between 1.2 and 2 mutations per cell division, a rate that overlaps with estimates from human neural progenitor clones (Bae et al. 2018) and *de novo* germline mutations (Rahbari et al. 2016). With more cell loss, the mutation rate could be lower.

R.2.e. Phylogenetic relationships between different cell types

The mapping of phenotypes (as defined by surface markers and colony morphology) onto our phylogeny provides us with the opportunity to study clonal relationships. A statistical test for clustering on the phylogeny (Analysis of Molecular Variance, or AMOVA (Methods)) allows us to perform three comparisons: bone marrow-derived *versus* peripheral blood-derived stem cells; stem cells *versus* progenitors; and different kinds of progenitors *versus* one another.

First, no clustering of bone marrow-derived stem cells was observed relative to their circulating counterparts (p=0.14). This shows that there is no detectable geographic clustering of related cells in the marrow and, as far as we can tell, our sample from iliac crest is representative of the bone marrow throughout the body of this individual. This is consistent with two explanations: 1) a large polyclonal group of stem cells seeds the iliac crest marrow (and probably other marrow locations) around the time of birth, and no local clonal sweep occurs thereafter; or 2) stem cells must regularly recirculate and redistribute sufficiently often that the population in the iliac crest is a reasonably random sample of the whole stem cell pool. If neither of these scenarios were occurring, we would expect to see marrow-specific clades of the phylogenetic tree.

Second, the degree of clustering of progenitors relative to stem cells informs on the proportion of stem cells contributing to blood. If only a small proportion of stem cells were contributing to blood, and we had sequenced the entire haematopoietic phylogeny, we would see clades of progenitors clustering around their parent stem cells, and large numbers of stem cells with no nearby progenitor. Even if the parent stem cell were not sequenced, clustering of its progenitor daughters would be observed. In contrast, if all stem cells were contributing equally, stem and progenitor cells should be distributed around the tree. We were not able to detect clustering of stem cells relative to progenitors (p=0.12), indicating that more than a few clones are generating progenitors. The test that we use to show that progenitors are not clustered relative to

stem cells does not have an intuitive statistic, and it is hard to encapsulate the relationships on the tree with a few numbers. We can say, for example, that of the 51 progenitors, 37 are more closely related to another progenitor than to a stem cell. This or a more extreme value occurs 8% of the time in a randomisation test. However, this does not tell us when the common ancestor of the pair of cells existed, which is important for establishing whether currently active progenitors derive from a small number of stem cell ancestors. Perhaps a more useful statement is that the 51 progenitor colonies distribute across 47 different lines-of-descent extant at 100 mutations of molecular time, compared to 89 stem cells across 76 lines-of-descent. Thus, the progenitors we sequenced are not drawn from a substantially more restricted set of historic lines-of-descent than the stem cells.

Third, we were not able to demonstrate any clustering of different types of stem or progenitor cells on our tree. For example, MEPs were no more likely to be closely related to one another than they were to GMPs or CMPs (p=0.10), although the small numbers of GMPs and CMPs reduce the power of this analysis. These data are consistent with either a model whereby stem cell clones provide multilineage replenishment of progenitor cells with individual commitments, or a model with unilineage priming but such a large pool of ancestors that the cells of the same type that we have sampled were unlikely to descend from the same clone.

All of the above analyses have limited power. It is possible that by sequencing more stem and progenitor cells we would have been able to detect clustering of these various cell types. It is difficult to quantify the effect size that we would have been powered to detect, as this will depend on the size of the stem cell pool and the clonal dynamics of the tissue, both of which we are exploring here largely for the first time. These results should be revised by future studies.

Nevertheless, the observation that stem and progenitor cells are interspersed on the tree is useful to our analyses. Taken in conjunction with the relative paucity of recent branch-points, it implies that the phylogeny is dominated by events that occurred in stem cells. Progenitors can be thought of as random samples of stem cell clones and can therefore be treated in the same way as stem cells. Furthermore, since progenitors have short life-spans, if one were to retrace their life history they would rapidly coalesce with a stem cell. Branch-points that occur more than a few tens of mutations up the tree, and so more than a few years ago, necessarily represent symmetrical cell divisions where one stem cell has divided to give rise to two daughter stem cells. We know that both daughters are stem cells since we have sampled their progeny years to decades later. It

does not matter that some of the branch-points that we observe in our phylogeny we might have discovered through sequencing two progenitors, since our data indicate that we are sampling the descendants of two different stem cells. In analogy to this, in studies of mitochondrial inheritance, male samples can be treated as proxies for their mothers (Griffiths and Tavaré 1994). Therefore, for the analyses that follow, which rely upon the pattern of branch-points in our phylogeny, we can treat progenitors as stem cells, and, indeed, hereafter we refer to them as such.

R.2.f. Population size trajectory over life

While most branch-points are observed at the top of the phylogeny, a sizeable number also occur later. As explained above, these represent symmetrical cell divisions of adult stem cells. The pattern of branch-points throughout the tree is remarkably informative on how the size of the population of stem cells has changed over life. This has been extensively studied in population genetics and epidemiology, where, for example, the seasonal rises and falls of influenza virus populations can be reconstructed through an analysis of the branching pattern of the phylogeny (Lan et al. 2015; Karcher et al. 2017). These approaches are based upon Kingman's coalescent (Kingman 1982), which is discussed briefly above. Essentially, under a constant population size, the time intervals between coalescent events on a phylogeny should follow independent geometric distributions (with means determined by the number of lines of descent which traverse each inter-coalescent interval). Deviations from this statistical pattern of coalescences indicate changes in the size of the population. By comparing the pattern of coalescences before and after each time point, fluctuations in the population size can inferred.

The population size trajectory was computed using an established Bayesian method (Lan et al. 2015, Karcher et al. 2017) (figure 1.3d). We observe almost logarithmic growth early in life, for the first 100 mutations of molecular time, as one might expect given that the population must expand from one cell in the early embryo to an adult stem cell pool size. At this point, however, the population size stabilises. Assuming a constant mutation rate over life, 100 mutations of molecular time represents approximately 6 years of age. We cannot be quite this precise with great certainty, since the mutation rate may be higher in the embryo and perhaps in early childhood. If this were the case, the inflexion point would occur at a younger age. It would be possible to draw

a straight horizontal line through the 95% confidence interval on the population size trajectory from mutation 100 right until the end of life, such that we cannot reject the null hypothesis of a constant population size in adulthood. In humans, the number of immunophenotypic HSCs increases in life (Pang et al. 2011). Here, we are assaying the effective population size – that is to say, stem cells that are still able to reproduce – and so our results are a readout of stem cell *function*. It should be noted too that our individual is only 59 years old, whereas most studies of stem cell ageing look at older people.

The stability of the stem cell pool indicates a degree of homeostasis: the symmetric stem cell divisions – where one stem cell begets two others – that we observe as branch-points in our tree must be counterbalanced by death and/or differentiation of a number of cells equal to that created. It is unknown at which level this control is exercised. It could be provided by a niche of limited size beyond which stem cells are forced to die/differentiate (Zhang et al. 2003), by stochastic fate decisions that are balanced at the population level, as has been observed in epithelium (Alcolea et al. 2014), or even through a deterministic model.

The absolute size of the stem cell pool cannot be estimated through this approach without knowing the generation time of the population of stem cells, which is equivalent to the average time between symmetrical stem cell divisions. As explained in the introduction to this chapter, the same expected pattern of coalescences in a phylogeny is generated by a population that is 10 times as large as another going through 10 times as many generations: they are confounded. Very little is known about the generation time of stem cells in humans, and so we could not simply use an estimate from the literature. Instead, the next phase of our experiment allowed us to learn both at the same time.

R.3. Mutation quantification phase: estimating the number of human stem cells

R.3.a. Targeted sequencing of peripheral blood

In this phase of the experiment, we aimed to quantify the frequency in peripheral blood of the clonal markers that we had discovered by whole genome sequencing. All 129,582 mutations assigned to branches of the tree were potential clonal markers. To design a bait-set for and perform

high depth-targeted sequencing of all of these mutations would have been extremely costly. We therefore designed a bait to as many of the 8,676 mutations present in more than one colony as possible and aimed to target approximately 10 mutations per private branch of the tree. When selecting private mutations, we picked sites in the genome with the lowest sequencing error rate possible, based on a panel of ~1000 genomes (Methods). Not all mutations are suitable for bait design (because, for example, they fall in a repetitive region of the genome), resulting in a bait-set of 7,116 mutations, of which 6,317 were shared and 799 were private.

With this bait-set, we performed targeted sequencing of granulocytes taken at 4-month (mean coverage 776X), 9-month (mean coverage 4,669X), and 14-month (mean coverage 268X) time-points after the bone marrow aspirate. Different samples were sequenced with different levels of coverage as a range-finding exercise. In addition, a negative control of two cord bloods from unrelated individuals (kindly provided by Grace Collord) was sequenced (mean coverage 5,305X). An algorithm for counting mutant and wild type reads with stringent error correction, written and run by Robert Osborne, was used to generate mutation counts over every bait. Further error correction was performed statistically, using a Markov chain Monte Carlo generalised linear mixed model that used the site-specific error rate from the control DNA to estimate the true VAF in test samples, written by Peter Campbell (with minor amendments by me). Reassuringly, mutations on higher branches were at higher VAF than mutations on lower branches. 96.5% of mutations had no mutations on lower branches with a higher allele fraction. The VAFs of mutations were stable across the three time-points (figure 1.7). This is reassuring from a quality-control point of view, but also of biological interest, demonstrating clonal stability (at this level of the phylogeny, at least) over a 10-month period.

Most mutations further down the phylogeny were not detectable (figure 1.5). In contrast, mutations on the majority of embryonic and early childhood branches were detectable, as indicated by multiple black or red branches at the top of the phylogeny in figure 1.5. Every clade with detectable mutations represents a clone whose descendants are making blood nowadays. The picture we observe, therefore, is one of polyclonal haematopoiesis, with contributions from stem cells dispersed on the phylogeny of blood. A total of 47 non-nested clones have mutations that are detectable in the 9-month granulocytes (not all of these can be seen in figure 1.5, as a VAF cut-off is used to allow comparisons between different timepoints, but they can be seen in figure 1.8 where no cut-off is used). As with the Rodewald group's embryonic clonal marking strategy (Pei et al.

2017, section I.3.a.ii. of this chapter), the number of early clones that we observe actively contributing represents an extreme lower bound on the total number of actively contributing stem cells. Furthermore, if we make the assumption that stem cells make equal numbers of granulocytes, the distribution of VAFs itself suggests that thousands of stem cells are ancestral to today's granulocytes. If only 500 stem cells made blood, we should hardly ever see mutations at a VAF much below one in 1,000. Over some sites, the coverage is sufficiently high and the error rate in controls is sufficiently low that we can detect mutations at VAFs of the order of 1 in 3,000-4,000. We observe VAFs right down to our detection limit, suggesting that at least 2,000 stem cells are contributing to the granulocyte pool.

R.3.b. Capture-recapture approach for estimating stem cell number

We hypothesised that the actual number of stem cells might be estimated through a capture-recapture (also known as 'mark and recapture') approach that is typically used in ecology. The total number of animals in a population, N, can be estimated by marking n animals on one visit and recapturing K animals on a second visit. The number of recaptured animals that have been marked, k, tells us the size of the population through the formula: $N=(K*n)/k$. In our study, 'stem cells' can be substituted for 'animals', and 'mutations' for 'markers': by taking multiple blood samples from one time-point, lysing and sequencing them separately, and comparing the overlap between the mutant alleles found in one sample and those in another sample, one could begin to estimate the size of the population.

In our study, however, this is complicated by several factors. First, mutations are not discovered by sequencing of peripheral blood, but have rather been discovered previously by whole genome sequencing; the question is if they are recaptured in different peripheral blood samples. Second, we separated our blood sample into six subsamples, rather than two, for greater resolution. To adapt the analysis for these first two factors only requires minor amendments.

More challenging, however, is the fact that a mutation does not mark a single stem cell, but rather a clone of stem cells. For this to be most informative we need to know when the MRCA of the clone existed (i.e. when the clonal mutation occurred). Although we know on which branch a

Figure 1.5. Quantification of mutations by targeted sequencing. **a,** the phylogenetic tree of cells is shown as in Figure 1.3a, but information from targeted sequencing of peripheral blood granulocytes from the 9 month time-point is overlaid. This is shown more clearly in the inset (**b**), which zooms in on a portion of the tree. The underlying structure of the tree is shown in grey. On top are placed horizontal bars, one for every mutation in the bait-set for targeted sequencing. The bars are coloured according the proportion of cells in the sample that carry the mutations (obtained by multiplying the variant allele fraction for autosomal mutations by two), indicated in the colour scale. Undetectable mutations are coloured grey and shown as smaller bars. Mutations are assigned to a branch based on which colonies they are present in. We cannot know the ordering of mutations along a branch other than from the targeted sequencing of peripheral blood; mutations have therefore been spaced evenly along a branch according to their mean VAF from targeted sequencing of all granulocyte and lymphocyte time-points combined. Small fluctuations in the estimated VAF due to random sampling of mutant reads mean that sometimes a mutation might be at a higher allele fraction in one particular sample than the mutation placed above it; this explains why sometimes a mutation at a low allele fraction in a particular sample is placed higher up the tree than a mutation which is at lower allele fraction in that sample. A higher density of baits was designed for branches shared by more than one colony. On these branches the mutations are so close together that they can appear as one continuous bar. **c**, This schematic explains that the VAFs of mutations decline down branches because of undetected coalescences with stem cells that were not whole genome sequenced but that are producing granulocytes.

a

Molecular time (substitutions)

0
200
400
600
800
1000
1200

b

Mutation at high VAF in granulocytes

Mutation at low VAF in granulocytes

Mutation undetectable in granulocytes

Proportion of cells with mutation

| Up to 1/2500 | 1/2500 to 1/500 | 1/500 to 1/100 | 1/100 to 1/20 | 1/20 to 1/4 | 1/4 to 1/1 |

c

WGS HSC 1   WGS HSC 2   Missed HSC 1   Missed HSC 2
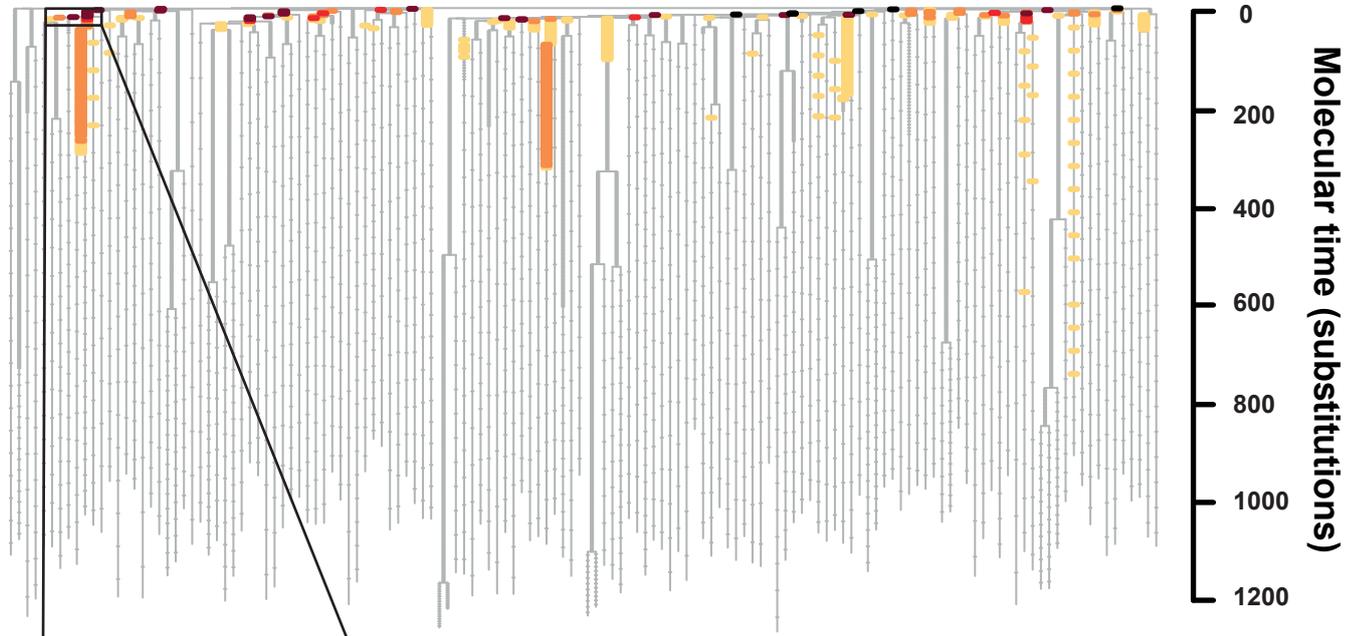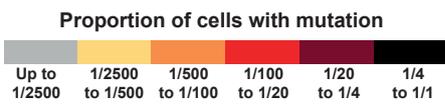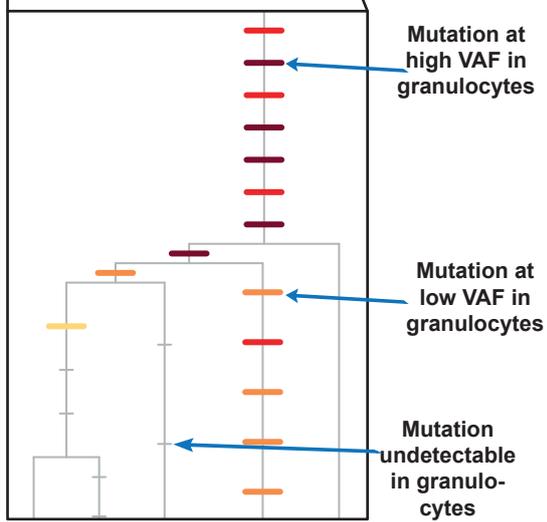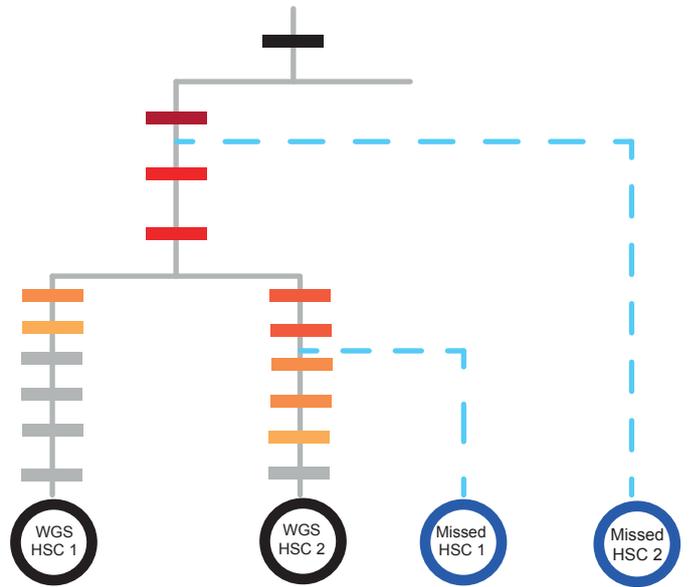
given mutation falls, we do not, from the whole genome sequencing alone, know the ordering of mutations along that branch. Targeted sequencing of peripheral blood informs on this: the VAF of mutations should decline down a branch, since they represent nested clones. To obtain a point estimate of their molecular time, therefore, we can spread mutations evenly down a branch in order of decreasing VAF. By performing a capture-recapture experiment on mutations from a given slice of molecular time, one could estimate the number of cells that existed at that time with extant descendants (i.e. the number of clones).

The rate of decline of VAFs down the tree is governed by the same coalescent process that determines the pattern of branchpoints in the phylogeny. Indeed, the VAFs decline because of undetected coalescences (figure 1.5c), and so this rate largely measures the same quantity as the rate of observed branching, albeit more sensitively. Recall that this defines the relationship between stem cell number and generation time but does not allow one to be calculated without knowledge of the other.

Extra information, however, comes from the mutations found at the very bottom of the tree, which occurred in the last few generations of stem cells: those that are ancestral to our granulocyte sample with few intermediate stem cells. The number of 'clones' at this level of the tree is essentially the same as the number of active stem cells. We had lysed and sequenced separately six different granulocyte samples from the same time point. Based on flow-sorting, each sample contained ~90,000 cells, and each position in the bait-set was sequenced in each sample at a mean depth of 800X.

Imagine the extreme case where only 10 stem cells make granulocytes at any one time. In this scenario, each of the six granulocyte samples will contain descendants from all of the 10 stem cells, and the VAFs of mutations in all six peripheral blood samples will be very similar to each other. In contrast, imagine a scenario in which one billion stem cells are making granulocytes at any one time. In this case, different samples of granulocytes will contain cells that descend from different ancestral stem cells: a sample of 90,000 granulocytes will not contain descendants from all stem cells. The VAFs of mutations will therefore differ more between different granulocyte samples than in the scenario of a small number of contributing stem cells. This is the information, in conjunction with the pattern of branchpoints in the tree and the rate of VAF decline, that we set out to use to estimate the total number of stem cells in the population.

R.3.c. Approximate Bayesian Computation to estimate stem cell numbers

Because of the complexity of using these multiple sources of information at the same time, it was not possible to solve this problem analytically. In future, it would be possible to apply a fully Bayesian approach, but because computation of the likelihood function involves integration of a large number of possible phylogenies, such an approach is impractically slow. In contrast, an Approximate Bayesian Computation (ABC) can be parallelised and run within the timescale available.

The concept of an ABC is simple: a model is simulated many times, varying the parameters of interest; for each simulation, summary statistics that capture the features of the data of interest are computed; simulated summary statistics are compared to the same statistics calculated on the observed data; and the values of the parameters of interest that produce summary statistics similar to the observed are deemed to be more likely than those that produce very different summary statistics (Tavaré et al. 1997; Beaumont et al. 2002; Bertorelle et al. 2010). The b simulations (where b is a number to be determined, but typically around 1,000) that are most similar to the observed are chosen and the values of the parameters in these simulations can be treated as a credibility interval. Finally, a regression step on a given parameter may be performed (Blum and Francois 2010, Csillery et al. 2012), regressing the value of the parameter against a measure of the dissimilarity between observed and expected summary statistics, thus finding the parameter value that would most closely approximate the observed data.

In our case, a simple model of haematopoiesis (explained below) was simulated thousands of times, varying the number of active stem cells and the generation time, and the whole experiment was repeated *in silico* on every simulation, from whole genome sequencing of the colonies, to design of the bait-set, to estimating the VAFs of mutations in peripheral blood. Such an approach allows us to recreate the uncertainties in our data: for example, we can pretend in the simulation that we do not know the true ordering of mutations, and rather reconstruct it by ordering mutations down a branch according to their estimated VAF (which need not be equal to their true VAF, since it is calculated from the binomial resampling of mutant and wild-type reads), just as in the real experiment. Similarly, we can use the observed sequencing depths from our data to generate a slight unevenness in coverage, making the simulation more realistic. Summary statistics

were chosen to include a measure of the pattern of branch-points on the tree and to reflect a capture recapture approach, by looking at how many mutations were shared between different blood samples. A more detailed description of each step of the method follows.

Please note that while I wrote the ABC code, I received extensive guidance from Kevin Dawson and Peter Campbell.

R.3.c.i. The model of haematopoiesis

R.3.c.i.1. Overview of model

Haematopoiesis is complex, and to model it requires a number of assumptions. We therefore opted for a simple model of haematopoiesis: individual stem cells within the population of total stem cells replicate stochastically over life and their clonal dynamics approximate neutral drift. Stem cells were simulated with a forward-in-time Wright-Fisher model, varying the number of stem cells and generation time. Because of uncertainty in the clonal dynamics of growth to adulthood, we commence our simulations at the point where the adult stem cell pool has reached a constant size, at 100 mutations of molecular time (figure 1.4d). In each cell division, each stem cell acquired a number of somatic mutations drawn from the Poisson distribution. The mean of the Poisson distribution is chosen such that the mean number of mutations per stem cell after the requisite number of generations in that simulation would equal the observed mean of 1,023.

It should be noted that we are using the term 'stem cell' to refer to cells that have self-renewed over the 59 years of life of our subject and are ancestral to circulating granulocytes. This population of cells may not translate directly to the standard definition of haematopoietic stem cells as those capable of long-term multi-lineage reconstitution when transplanted into immunocompromised mice. Nevertheless, our 'stem cells' meet the conceptual definition of a stem cell since they have self-renewed for so long, and in addition to producing granulocytes are likely also to have produced B lymphocytes (figure 1.8, discussed in section R.4.). We only concern ourselves with the stem cells that have given rise to cells actively making blood. In this experimental set-up, we are blind to long-term dormant stem cells (i.e., years of not making any

sort of long-lived multi-lineage progeny) and to stem cells that exclusively produce other mature blood cell types but no granulocytes.

R.3.c.i.2. Assumptions

First, we model stem cells with active contributions from their progeny as belonging to one population. There is evidence that there may, in fact, be different pools of stem cells cycling at different rates (Wilson et al. 2008), but it is currently unclear whether this applies to humans, how big each pool is, whether cells move from one pool to the other during homeostasis and if so at what rate, and what the relative contribution of each pool to circulating granulocytes is. We therefore chose to model the pool of stem cells actively contributing to granulopoiesis as one population. It is also important to note that while all stem cells share the same probability of dividing in a given window of time, our model in effect allows them to behave stochastically, as in previous models of haematopoiesis (Abkowitz et al. 1996, Catlin et al. 2011).

Second, we assume that the size of this stem cell pool is constant over the period of adulthood studied (figure 1.4s, Werner et al. 2015). We do not concern ourselves with the dynamics of how the population grew to be this size: all our summary statistics only use information from after 100 mutations of molecular time, when we see that the stem cell population size has stabilised. We thus have a population of stem cells of constant size that we allow to replicate stochastically over the course of adult life.

Third, we make the assumption that there is no selection in the stem cell pool. This is based upon our inability to detect selection in our experiment (discussed above) and the assumption that the vast majority of somatic mutations will not affect stem cell function. We thus have a population of stem cells that is replicating stochastically, resulting in a process of neutral drift.

Fourth, we assume that those stem cells that are making granulocytes are all making an approximately equal number of them. This assumption is likely to matter less with larger stem cell population sizes.

The final important assumption that we make is that stem cells accumulate somatic mutations at the same rate over life, drawing the number of mutations that each stem cell acquires at every generation from a Poisson distribution. This is justified by the linear accumulation of

somatic mutations over time reported by others (Welch et al. 2012) and the relatively narrow range of mutation burden that we observe from our 140 clonal whole genomes (figure 1.1c).

With this model, we simulate haematopoiesis hundreds of thousands of times with different values of the stem cell pool size and time between symmetrical stem cell divisions, drawing both from a uniform prior on the log scale (which minimal knowledge of the distribution of the data). We replicate the experiment on each simulation: we sample cells for whole genome sequencing, design a bait-set, sample granulocytes from our stem cell population, and count the number of mutant reads found in the granulocytes over each position in the bait-set.

R.3.c.i.3. Choice of priors

R.3.c.i.3.a. Prior on stem cell population size

The adult stem cell population size was drawn from a uniform distribution on the log scale between 1,096 (exp(7)) and 3,269,017 (exp(15)) stem cells. A minimum number of 1,096 stem cells was chosen because we knew from preliminary simulations and from the reasoning above that a smaller number of stem cells than 1,000 could not produce the low VAF mutations that we observe. The maximum number of stem cells was chosen because, firstly, it was at the limits of what was computationally feasible with the resources available (each simulation at this upper limit requires approximately 150 GB of memory), and, secondly, because it was two orders of magnitude higher than the proposed number of stem cells (Abkowitz et al. 2002).

R.3.c.i.3.b. Prior on generation time

The Wright-Fisher generation time (which is equivalent to the mean time between symmetrical cell divisions for one line-of-descent) was drawn from a uniform distribution on a log scale between 20 days (exp(3)) and 8,103 days (exp(9)) i.e. 22 years). The minimum time of 20 days was chosen because stem cells are reportedly relatively quiescent (Arai and Suda 2007; Orford and Scadden 2008). Not all of a stem cell's divisions need be symmetrical: a proportion is

100

likely to be asymmetrical, producing one daughter stem cell and one progenitor cell (Werner et al. 2015). We are blind to asymmetrical divisions. Therefore, if a cell is dividing asymmetrically in addition to symmetrical divisions on average every 20 days, it will be dividing significantly faster than every 20 days, which seemed unlikely given prior knowledge of stem cell quiescence. Furthermore, shorter times between cell divisions mean that more generations need to be simulated, which is computationally costly. The maximum time between symmetrical cell divisions of 22 years was chosen because it required a very small number of HSCs to create a phylogeny of the right shape, and such a small number was not compatible with the observed range of VAFs.

R.3.c.i.4. *In silico* recapitulation of our experiment

3.c.4.i.a. Mutation discovery phase

After simulating drift for the whole population of stem cells for the requisite number of generations, we choose 155 colonies for whole genome sequencing and construct a phylogeny from them. Of 198 colonies sequenced in our experiment, only 140 were clonal. These 140 were used to build the tree, and for all analyses except for the ABC. However, of the 58 polyclonal colonies, we could salvage 15 because they had a dominant clone and shared more than ten mutations with a clonal colony that was on our tree. We could therefore graft these 15 extra colonies onto the tree of 140 clones. This is helpful because it provides an extra time point on each branch onto which an extra colony is grafted. Mutations can then be classified as being shared with the polyclonal colony that has been grafted on, or absent from the polyclonal colony, thus providing additional information about the timing of the mutation. No mutations that were present only in polyclonal colonies (and not in clonal colonies) were used, as we could not be sure where to place them on the tree.

R.3.c.i.4.b. Mutation quantification phase

We design a bait-set for the simulated tree, using the same criteria as were used to design the real bait-set (Methods). We then simulate the sampling of peripheral blood granulocytes. We generate a sample of $M = 540,000$ granulocyte by sampling with replacement from the stem cell population. This simulated sample of 540,000 cells is then split into six sub-samples each of size 90,000 (to simulate the six sub-samples obtained from our volunteer at the nine-month time-point). We then simulate targeted sequencing of the mutations in the bait-set. In the observed data, there were 3,952 mutations in the bait-set that (after our duplicate removal and consensus calling step (methods)) were covered by at least 4,000 reads in the control cord blood DNA, but where no mutant reads were found in the cord blood. We therefore used these 3,952 mutations for analysis of the observed data, and also only use 3,952 mutations in the simulated bait-set. For every bait-set locus in every sub-sample, we randomly draw a sequencing depth from the empirical distribution of sequencing depths for the real targeted data from the nine month time-point. We sample the chosen number of reads over this locus from the granulocytes without replacement and count the number of reads that have the mutation.

Sequencing errors were included in the simulation as follows. The sequencing error rate was learnt from the control cord blood. For all 7,116 positions in the bait-set, the VAF in the cord blood was calculated. Where the VAF was zero (if there were no mutant reads), the VAF was set to 1/10,000, a value just below the rarest mutations that we could detect. For each of the 3,952 loci used in simulations, then, an error rate was drawn from these VAF distributions. The number of false positive reads was obtained by drawing from the binomial distribution, with parameter p equal to the randomly chosen error rate and parameter n equal to the sequencing depth. We also tested two other error models: one with no sequencing errors, and one with double the sequencing error rate observed in the cord blood controls. These made little difference to the median of the posterior distribution of the model but did affect its width. A separate false positive rate was included based on the estimated rate of homoplasy, assuming that every granulocyte has 2,000 mutations (double the number of mutations present in a stem cell, which seemed a reasonable upper bound for the number of additional mutations that a granulocyte could acquire) and that these are spread evenly across the genome.

R.3.c.ii. Extraction of summary statistics

We then extract summary statistics from the resulting simulated data set. There are two categories of summary statistics used. Only the first category of summary statistics was extracted for the first set of simulations, as explained below.

The first category of summary statistics is comprised of those designed to capture the shape of the phylogeny, referred to as lineages through time or LTT summary statistics. We divide the molecular time scale, from mutation 100 to mutation 800 (beyond which some branches in our observed phylogeny end), into bins of molecular time each 100 substitutions wide. For each mutation of molecular time in a given bin, we count the number of branches extant at that slice of time and take the mean across all hundred mutations in that bin.

The second category of summary statistics is made up of those designed to exploit the 'capture-recapture' aspect of our approach with targeted sequencing of multiple peripheral blood samples. We wanted statistics that would capture whether different granulocyte samples descended from the same stem cell population or not. As explained above, the less overlap in the granulocyte sub-samples, the larger the contributing stem cell population is likely to be. For a given mutant read threshold (from 1 to 6 mutant reads), and a given number of samples (from 0 to 6 samples) we count – out of the 3,952 mutations included – how many loci have this many or more mutant reads in precisely this many samples. For example, one summary statistic would be the number of mutations in the bait-set that are supported by at least three mutant reads in two samples. These summary statistics are recorded for every simulation, and the same summary statistics were calculated for the observed data.

R.3.c.iii. Details of ABC implementation

Two sets of simulations were run, referred to hereafter as ABC1 and ABC2, both using the same model, but with a different prior and extracting different summary statistics.

First, we generated 121,329 simulations drawing both the number of stem cells and the time between symmetrical stem cell divisions from a uniform prior on a log scale (figure 1.6a) and

extracted summary statistics that reflect the shape of the phylogeny (the LTT summary statistics explained above). This allowed us to identify a relationship between stem cell number and generation time, effectively defining a plausible diagonal band on the sample space of stem cell numbers plotted against generation time (figure 1.6b). Going backwards in time, the faster the rate of random drift, the more rapidly the number of lineages decreases. Thus, simulations that have too rapid a random drift rate (simulations with a small population size and short generation time) have LTT statistics that are too low for early bins of molecular time (figure 1.6o), and simulations that have too slow a random drift rate (with a large population size and long generation time) have LTT statistics that are too high for early bins of molecular time (figure 1.6p).

For ABC2, we ran another 80,762 simulations targeting this area of the sample space (figure 1.6c). For this set of simulations, both the LTT and peripheral blood-derived summary statistics were extracted.

3.c.iv. Comparison of simulations and observed data

Summary statistics were normalised, and for each simulation the Euclidean distance between the simulated and observed vector of summary statistics was calculated. For ABC1, we simply defined the region that contained the most similar 20% of simulations, which was then used as the prior for ABC2. Simulations outside this range can be shown to have a branching structure that does not resemble that of our phylogeny (figure 1.6o-p).

For ABC2, to maximise the accuracy of our model, we cross-validated both the number of simulations included in the acceptance region, and the weighting to give to the LTT statistics. For each of 1,000 cross validation samples, we drew one simulation to act as fake observed data and removed it from the pack of simulations. We then analysed the data as though our fake observed data were the true data. We took the n (where n is the number of accepted simulations) simulations that produced summary statistics that were most similar to our fake observed data, as determined by their Euclidean distance.

We then calculated a number of statistics (figure 1.6e-i). First, we plotted the accepted simulations on a graph of stem cell numbers *versus* generation time (both on a log scale) and drew an ellipse that contained 90% of the n points inside it. We then saw whether the true value of the

fake observed data fell inside this ellipse. The proportion of cross validation samples for which the fake observed data fell inside the ellipse is shown in figure 1.6e. We also measured: the mean area of the ellipse (figure 1.6f); the distance between the median number of stem cells of the accepted simulations and the fake observed data number of stem cells (figure 1.6g); the distance between the median generation time of the accepted simulations and the fake observed data generation time (figure 1.6h); and, finally, the distance between median of the posterior of a neural network regression run on the accepted simulations and the fake observed stem cell number (figure 1.6i). We chose an LTT weighting of 1 and an acceptance region of 1,000, since this resulted in an accurate prediction of the stem cell number from the neural network regression and a high proportion of the fake observed values falling in the ellipse, while keeping the size of the ellipse relatively small.

We then analysed the true observed data, using the error model that took VAFs from the observed control data, weighting the LTT summary statistics by 1 and choosing the 1,000 most similar simulations to fall in the accepted region. Of the best 1,000 simulations, 90% had more than 15,995 and less than 295,189 HSCs, and a generation time of more than two and less than 20 months (figure 1.6r).

For added precision on the number of stem cells, neural network regression was performed on these best 1,000 simulations using the R package *abc* (Csillery et al. 2012), to find the number of stem cells that minimised the distance between the observed and simulated summary statistics. The neural network regression was run using the default of one hidden layer with five units. As predictions from different neural networks can vary, thirty neural networks were run and the median provided. We found that the most likely number of stem cells ancestral to the sampled granulocytes was 97,000, with a 90% credibility interval of 45,000 to 215,000 (figure 1.6s).

To test the robustness of our analysis, we repeated it with the other two error models described and ignored summary statistics that used a mutant read cut-off of 1, since these would be most sensitive to incorrect modelling of the sequencing error rate. Both of these additional analyses resulted in a widening of the posterior distribution for the number of contributing stem cells but did not significantly change its location.

Figure 1.6. Approximate Bayesian Computation (ABC) to estimate HSC numbers. **a,** The joint prior distribution for HSC numbers and the generation time for the first ABC. **b**, The location in sample space of the 10% of simulations that produced ltt summary statistics most similar to the observed summary statistics. **c**, the joint prior distribution for the second ABC, in the area of sample space indicated to be plausible by the first set of simulations. **d**, The joint posterior distribution of the best 500 simulations from the second ABC. Letters n, o, and p on the plot indicate the position in sample space from which panels **n**, **o**, and **p** were drawn, respectively. **e-i,** cross-validation of the model to choose the number of accepted simulations and the weighting applied to the ltt summary statistics. **j**, for illustrative purposes, five simulations were sampled for each of three population sizes along the plausible diagonal of sample space indicated in panel **b**. One set of summary statistics is shown for these simulations in **k**. A red line indicates a simulation coming from the area of sample space indicated by a red point in **j**; *idem* for blue and green lines. The black dotted line indicates the observed values for these summary statistics. This set of summary statistics counts, for different numbers of samples (x axis), how many of the 3952 mutations considered (y axis) are in this many samples with two or more reads, using error model 1 (which simulates errors according to the error rate in control DNA). **l**, For each of the 1000 simulations that produce summary statistics most similar to the observed, the Euclidean distance from the observed (y axis) is plotted against the number of stem cells in that simulation (x axis). This information is used by the neural network regression to define the most likely value for the number of stem cells. The most similar values are seen at around 100,000 stem cells, which was the location of the median of the posterior distribution from neural network regression. **m**, the observed phylogeny, with branch points indicated by asterisks. **n-p,** phylogenies from simulations that occur at the points in sample space indicated in panel **d**. **n** represents a plausible simulation, since the pattern of branch points is not dissimilar from that seen in the observed phylogeny **m**. Simulations with smaller stem cell populations and faster stem cell turnover rates resulted in phylogenies where the stem cells are very closely related to each other (**o**), whereas those with larger populations and slower turnover result in phylogenies where the stem cells only share an embryonic common ancestor (**p**). **q**, the prior distribution for the number of stem cells contributing to granulocytes for the second ABC (i.e. the stem cell numbers for all 80,000 simulations). **r,** the distribution of stem cell numbers for the 1,000 simulations that produced summary statistics most similar to the observed summary statistics. **s,** the posterior distribution of a neural network regression run on these 1,000 simulations. The 90% credibility interval is quoted for the stem cell population in each of **q-s**.
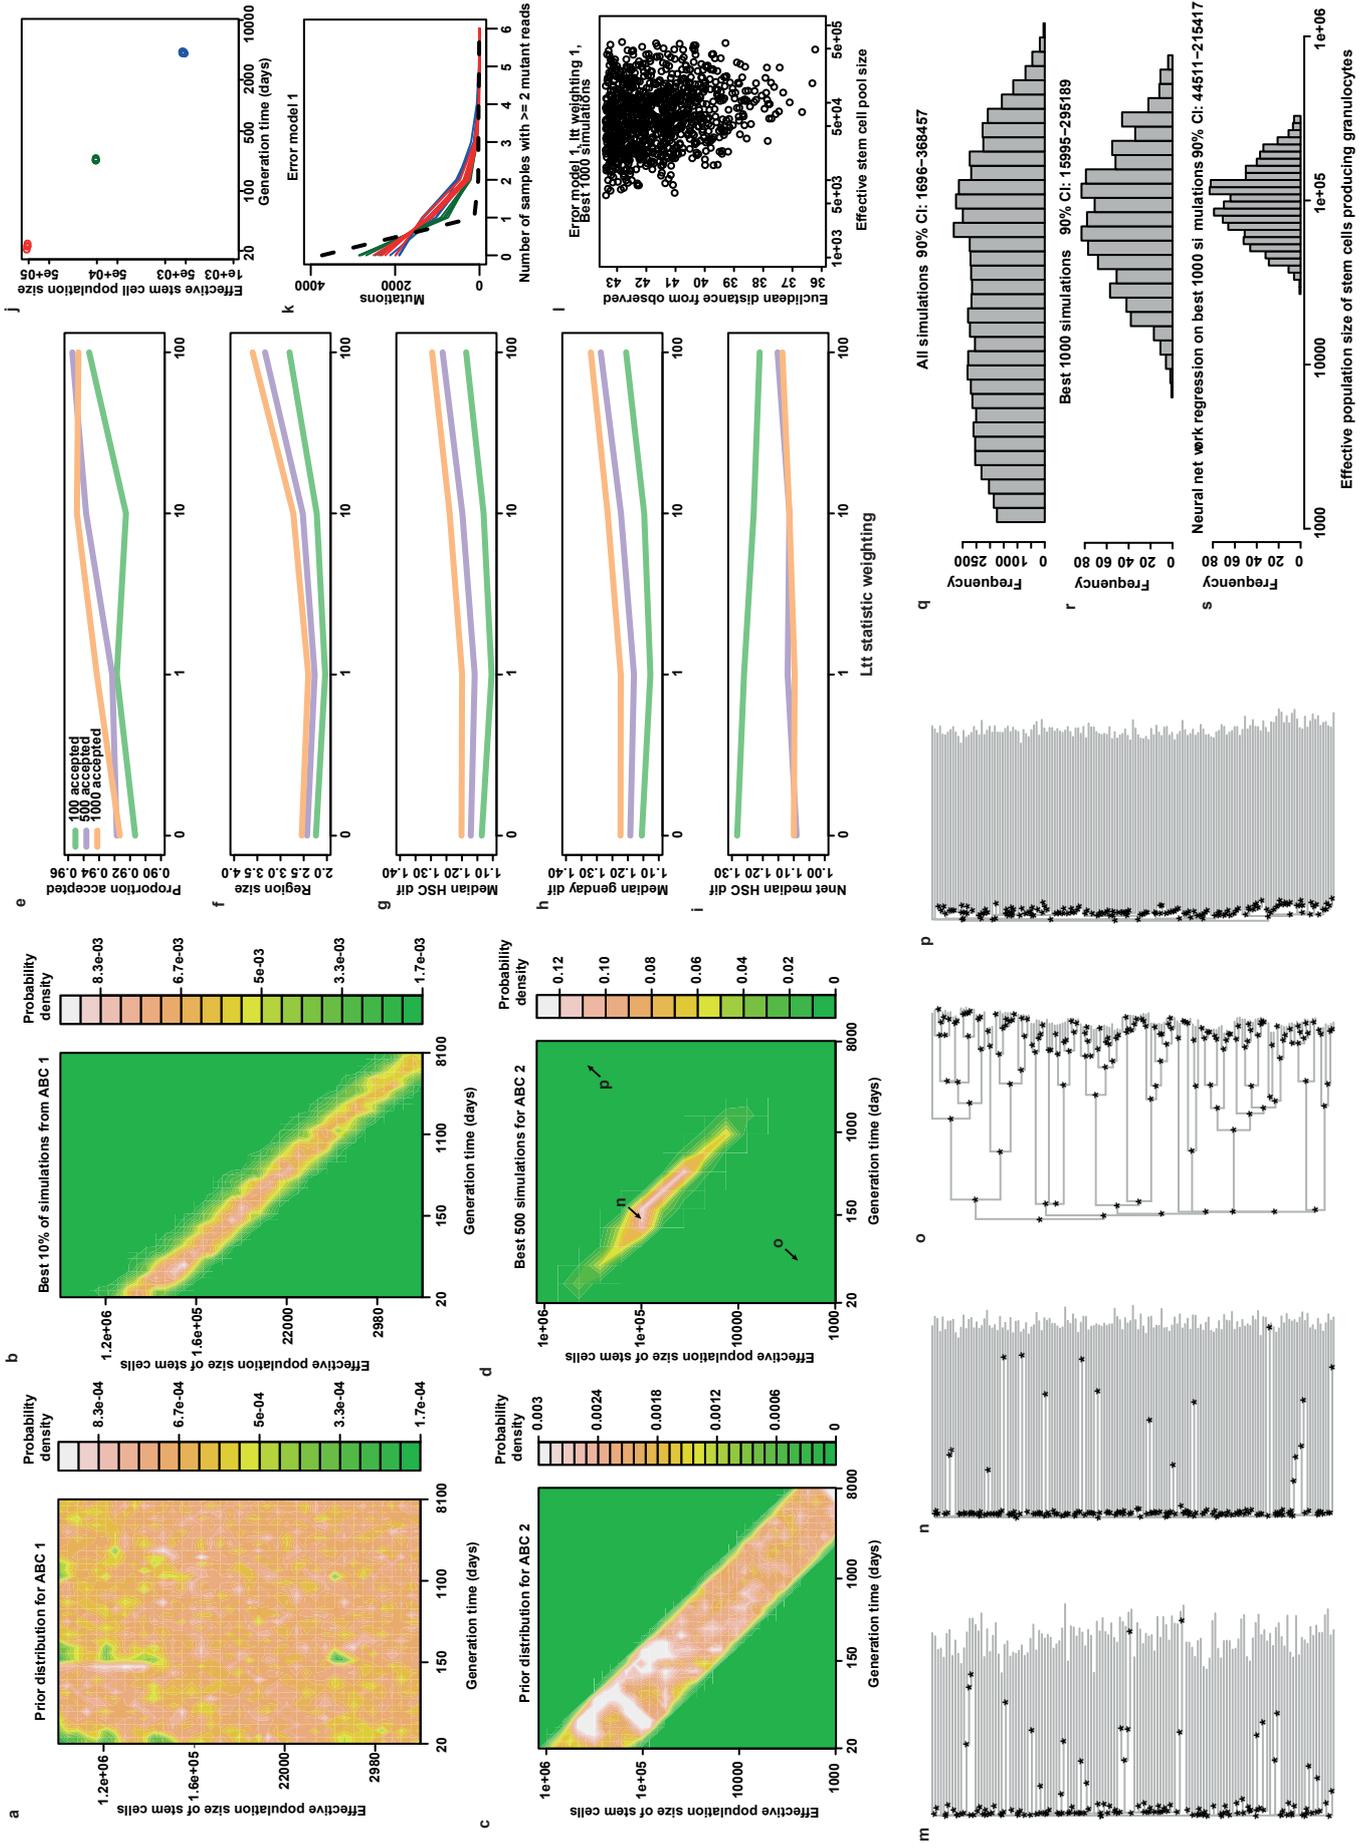
4. Clonal contributions to granulocytes and lymphocytes.

In order to examine lineage relationships at the clonal level, we performed targeted sequencing of B (mean coverage 2,830X) and T lymphocytes (mean coverage 3,298X), both from the 9-month time-point, using the same bait-set and error-correction method as for granulocytes (figures 1.6 and 1.8). As can be observed from the preponderance of black bars at the top of the phylogeny, most early embryonic branches contribute to all three cell types at the time of sampling. Thus, clones marked in the embryo, and perhaps early childhood, have multilineage potential, as has been observed in mice (Pei et al. 2017).

Beyond 100 mutations in molecular time, when the population size reached a plateau (figure 1.4d), 464 mutations distributed across 39 branches could be detected, of which 217 on 12 branches were detected in more than one cell type. Hardly any of these were detected in all three lineages. We do, however, observe a number of clones that make detectable amounts of both B lymphocytes and granulocytes. Sequencing coverage was equally good in T lymphocytes and in B lymphocytes, so the failure to detect these mutations in T lymphocytes is not a result of decreased sensitivity. At least five of these shared granulocyte-B cell branches stretch beyond the 100 mutations mark (figure 1.8) that we take to signify adult stem cell dynamics. Furthermore, the VAFs observed are similar between B cells and granulocytes, indicating a more recent common ancestor: the neutral drift to detectable clone sizes will have been gradual, with much of it occurring since the last detectable mutation. If the last detectable purple mutation in figure 1.8a represented the last common ancestor of granulocytes and B lymphocytes, both the B-restricted and granulocyte-restricted clones would have had to drift to similar sizes in the intervening time independently. It is more likely that they share a common ancestor well beyond the last detectable mutation found in both B cells and granulocytes. Thus, we demonstrate the presence of adult stem cell clones with both myeloid and lymphoid output in a human at homeostasis. Admittedly, the number of clones in which we have detected this multilineage output is small. However, there are no clones in which we would have been able to detect multilineage output in which we have not detected it (no branches beyond 100 mutations of molecular time that are exclusively dark red, blue, or green). This suggests that multilineage clonal output is not rare.

We cannot exclude smaller contributions of these granulocyte- and B cell-producing clones to the T lymphocyte pool on the order of less than one in three thousand T lymphocytes, which

would fall below our detection limit for lymphocytes. Nonetheless, since this is less than half the proportion that these clones are contributing to granulocytes and B lymphocytes, we can unequivocally show an imbalance. In contrast, we observe some branches where we can detect very rare mutations (allele fractions of less than 1/2,000) in granulocytes (pink branches in figure 1.8), which we did not have the sequencing depth to detect reliably in lymphocyte populations. It is quite plausible that these branches contribute to lymphocytes, too; we cannot, therefore, infer the presence of exclusively myeloid clones.

Only a handful of mutations were found to be shared by B and T lymphocytes and absent from granulocytes. Even if lymphoid-biased clones were frequent, however, we would not expect to find many of them, as the 'capture' phase of our experiment was biased towards myeloid lineages, since isolation strategies for human stem cells inevitably also capture more differentiated cells, predominantly myeloid progenitors (despite possessing the same combination of surface markers as true HSCs).

The observation of granulocyte-B cell restricted clones in adults would be reinforced by studying more cells from more people and with the ability to resolve rarer mutations. Even without these, our data already indicate that adult stem cell clones with both myeloid and lymphoid output are not infrequent in humans. In mice, tags have been found to be shared between granulocytes and B lymphocytes (Rodriquez-Fraticelli et al. 2018), but, unfortunately, T lymphocytes were not assessed. The observation of clones that produce detectable granulocyte and B lymphoid but not T lymphoid output could be explained in three ways. First, some individual stem cells may have a propensity to produce more of one mature cell type than another. This would have to be a somatically heritable trait, since we observe this behaviour at the level of clones that are ancestral to hundreds of stem cells. Such a finding of heritable stem cell heterogeneity is anticipated by the animal and *in vitro* studies discussed at the beginning of this chapter. Second, unequal clonal contributions to present-day blood samples might reflect the long life-span of T lymphocytes relative to granulocytes and some B cells: T lymphocytes could reveal the stem cell clones that dominated decades ago, while granulocytes and B lymphocytes show us more recent clone distributions. A third and related possibility is of a strong bottleneck in the T lymphocyte population imposed by thymic selection or infections, which would reduce the overlap between these populations. This could be resolved by including lymphocyte progenitors in the whole genome sequencing phase of the study, and by performing this experiment in children (since we

Figure 1.7. Quantification of mutations in different fractions of peripheral blood. a, b, targeted sequencing data represented as in figure 1.5 for all fractions of peripheral blood that were sequenced, showing only the first 350 mutations of molecular time, beyond which no mutations were detectable. To allow a better comparison between samples sequenced at different depths, a different detection threshold is used relative to figure 1.5. Data are shown prior to sequencing error correction using cord blood controls in the Bayesian generalised poisson mixed effects model (**a**) and after applying the error correction (**b**). **c,** correlations between the VAFs of all sequenced samples, shown on a log scale. Note that samples that were sequenced to lower depth cannot have VAFs as small as samples sequenced to higher depths.

a

Granulocytes, 4 months

Granulocytes, 9 months

Granulocytes, 14 months

B lymphocytes, 4 months

T lymphocytes, 4 months

Molecular time

b

Granulocytes, 4 months

Granulocytes, 9 months

Granulocytes, 14 months

B lymphocytes, 4 months

T lymphocytes, 4 months

Molecular time

Proportion of cells with mutation

Up to 1/1000	1/1000 to 1/300	1/300 to 1/100	1/100 to 1/30	1/30 to 1/10	1/10 to 1/3	1/3 to 1/1

c

Controls VAF

Granulocytes 9 months VAF

Granulocytes 4 months VAF

Granulocytes 14 months VAF

B cells 4 months VAF

T cells 4 months VAF

Figure 1.8. Comparison of allele fractions in granulocytes and lymphocytes. a, the phylogeny is depicted as in figure 1.5, with the underlying structure of the tree shown in grey, and horizontal bars drawn to represent every mutation in the bait-set. Here the colouring of mutations reflects which peripheral blood cell fractions they could be detected in, as indicated by the colour key. Granulocytes were sequenced at greater depth than lymphocytes and so mutations could be detected at lower allele fractions in granulocytes than in lymphocytes. Therefore, two colours are used for granulocytes: red for mutations only detected in granulocytes that were are at a sufficiently high allele fraction to have been found in lymphocytes should they have been present at a similar allele fraction in lymphocytes, and pink for mutations that were only detected in granulocytes, but were at such a low allele fraction ($<1/2000$ reads) that if they had been present in lymphocytes at this allele fraction they would not have been detected. Arrows indicate adult clones with multilineage output. G, granulocytes; G low VAF, granulocytes, allele fraction too low to be detected in lymphocytes; B, B lymphocytes; T, T lymphocytes. **b-f,** VAFs of all mutations on branches (indicated by arrows in **a**) with mutations beyond molecular time 100 that are detectable in granulocytes and B lymphocytes but not T lymphocytes.

a

Legend:
- G, B, T (black)
- G, B (purple)
- G, T (orange)
- B, T (teal)
- G (red)
- G low VAF (pink)
- B (blue)
- T (green)
- none (gray)

b

c

d

e

f
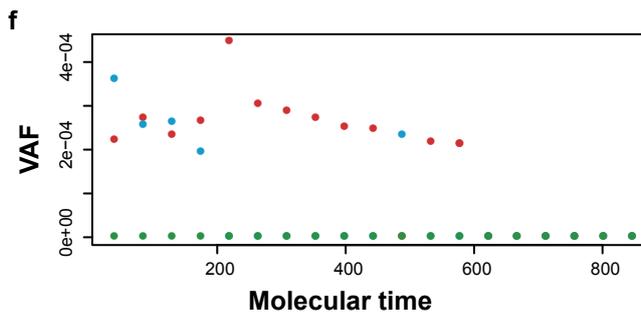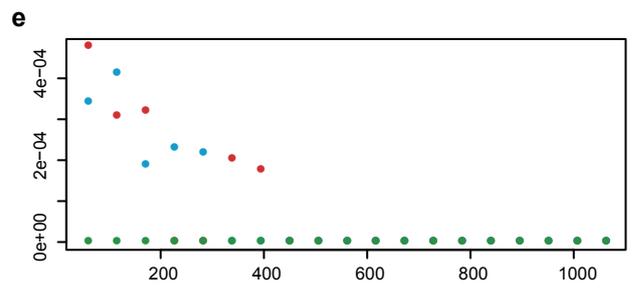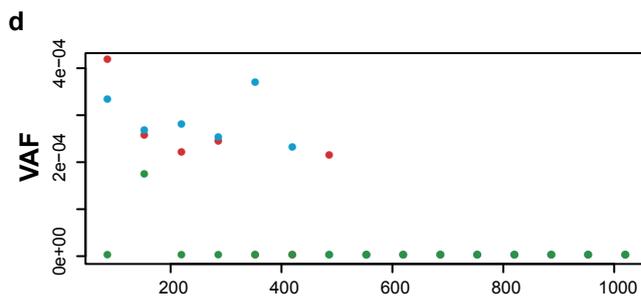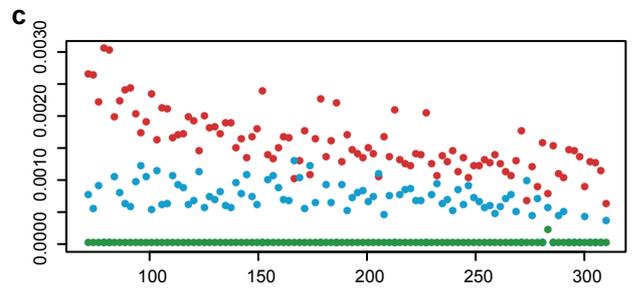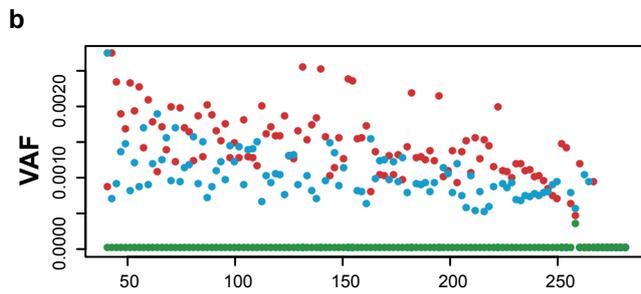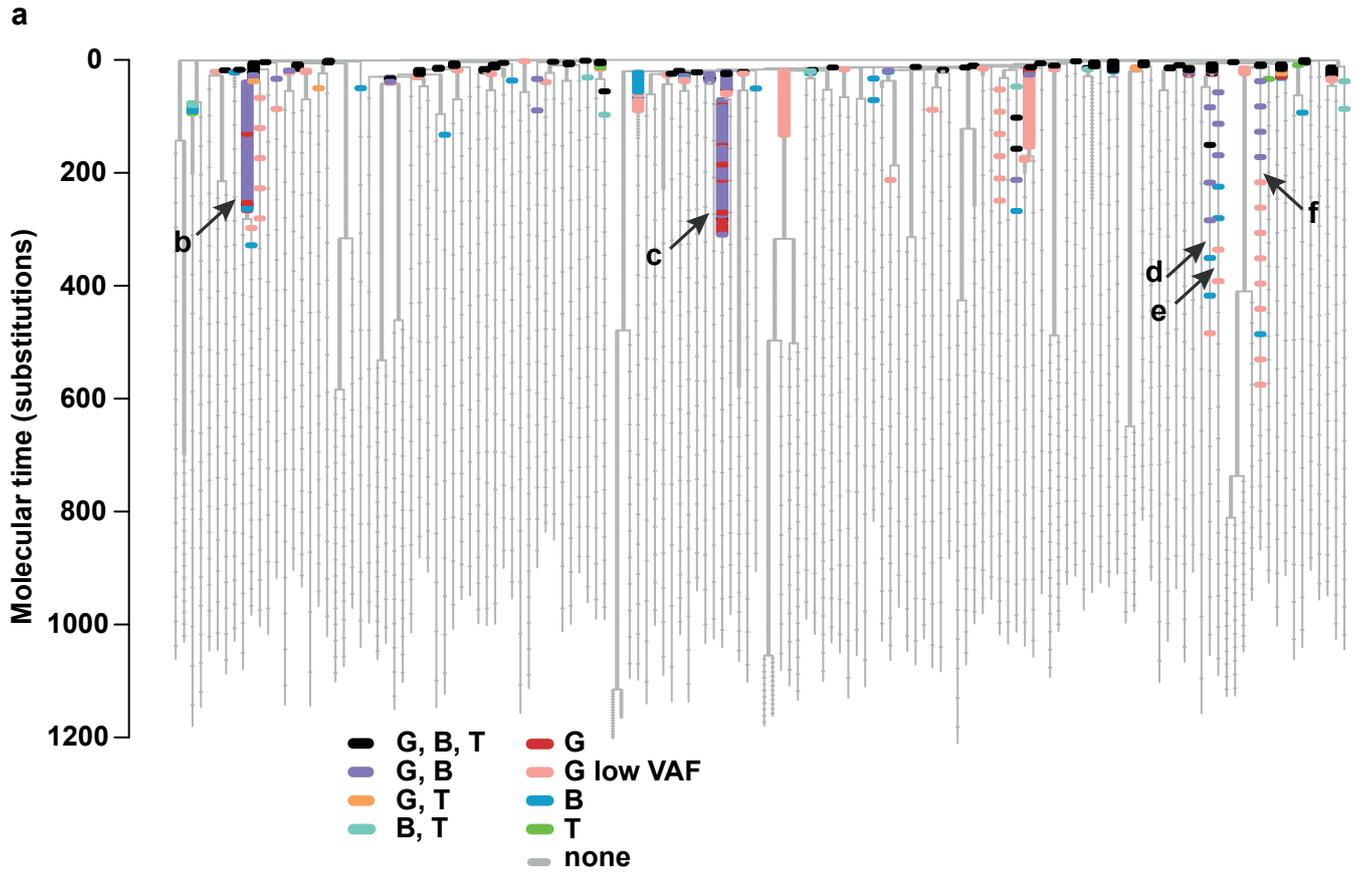
- Granulocytes (June) (red)
- B lymphocytes (blue)
- T lymphocytes (green)

were able to culture a reasonable number of stem cells from peripheral blood, a bone marrow aspirate would not be necessary).

Summary of results in this chapter

The number of haematopoietic stem cells, their clonal dynamics, and the clonal relationships between different cell types are all poorly understood in humans because of the difficulty of tracking clones in unperturbed people. We used spontaneously acquired somatic mutations to reconstruct lineage relationships among stem cells in normal human haematopoiesis. We sequenced the whole genomes of 140 colonies derived from single HSPCs from one healthy 59 year-old man. We identified the somatic mutations, reconstructed the phylogenetic relationships of the cells to one another, and sequenced at high depth bulk populations of granulocytes and lymphocytes for mutations that had been discovered by whole genome sequencing.

We were able to reconstruct cell divisions in the early embryo. Comparisons with buccal epithelium indicated that the most recent common ancestor of blood existed prior to gastrulation. The mutation rate in the early embryo is likely to be less than two mutations per mitosis. Averaging over the whole of life, the mutation burden and mutational signatures were consistent across different HSPCs. No positive or negative selection of mutations could be identified.

The trajectory of the number of haematopoietic stem cells over life could be inferred from the branching patterns in the phylogeny. The stem pool size increased rapidly during development and childhood and reached a stable plateau in adulthood. Using deep sequencing data, we inferred that the number of stem cells contributing actively to granulocytes at any given time is in the range 45,000 – 215,000.

Finally, we found that T and B lymphocytes have different clonal dynamics. We observed adult clones that produced detectable fractions of both granulocytes and B lymphocytes, but not T lymphocytes. Our data indicate a contribution of multipotent HSCs to B lymphopoiesis throughout life.

These findings are discussed in the Discussion chapter.

RESULTS CHAPTER 2

MUTATIONAL LANDSCAPE OF NORMAL COLON

Introduction to this chapter

I.1. Colonic stem cell numbers and clonal dynamics in health

I.1.a. Monoclonal origin of crypts

The human colon undergoes extraordinary amounts of cellular turnover. Its luminal surface area is approximately 3,300cm^2 and it is renewed every 3-4 days (Potten et al. 1992). 15 million invaginations of the epithelium (or crypts) (Boman and Huang 2008), each containing about 2,000 cells (Potten et al. 1992), form the regenerative unit of the colon and house four key functional cell types: enterocytes, goblet cells, crypt base secretory cells, and enteroendocrine cells. Rarer cell types include tuft, microfold, and cup cells.

All of these differentiated cell types derive ultimately from stem cells that sit at the base of the crypt. The propensity of cells at the base of mouse crypts[1] to phagocytose cellular debris allowed an early form of functional lineage tracing: when they were fed tritiated thymidine, phagosomes in crypt basal cells were rapidly labelled. The tritiated thymidine label could only later be detected in different mature cell types higher up the crypt, suggesting that basal cells were multipotent progenitors (Cheng and Leblond 1974).

Not only do all cells in the crypt have a stem cell as a recent ancestor, but they can all retrace their line of descent to the same ancestral stem cell. Crypts from mice, chimaeric for either the H2 antigen, detectable with a monoclonal antibody, or a carbohydrate polymorphism stained with a lectin, were clonal for a given marker (Ponder et al. 1985). Similarly, detection of a Y-chromosome-linked marker in XX-XY chimaeric mice showed the shared common origin of all

---

[1] Here, as in much of the discussion on stem cell biology that follows, the experiment was performed on small intestine rather than colon. The stem cell biology of the colon is thought to be similar, and when there are notable differences between the two organs I will draw attention to them.

cell types in a gastric gland (including neuroendocrine cells, which had not been demonstrated previously) (Thompson et al. 1990). In these chimaeras, however, large patches of epithelium shared a marker, and so the finding could have been a consequence of embryonic mosaicism rather than adult stem cell dynamics. That crypts derive from a single adult cell was confirmed using mice heterozygous for an inactivating mutation of a lectin-binding protein. ENU random mutagenesis resulted in loss of the wild type copy – and consequently the lectin stain – in a small proportion of cells, such that the progeny of a single cell could be traced (Winton et al. 1988). Treatment with ENU in adulthood resulted initially in ribbons of cells lacking staining emanating from one mutated cell, but after a chase period it only left wholly stained or unstained crypts (Winton and Ponder 1990). Similarly, treatment of mice with a mutagen resulted in sporadic loss of the X-linked G6PD biochemical marker in a small proportion of colonic crypts, but, importantly, in all the cells of affected crypts (Griffiths et al. 1988).

I.1.b. Neutral drift dynamics

The ancestral stem cell need not still be alive, and could have multiple self-renewing descendants (Potten and Loeffler 1990). More recent experiments have shown that a number of extant stem cells replace one another through a process of neutral drift, resulting, over time, in monoclonal conversion. By inducibly labelling less than 2% of cells with a marker at a defined point it was possible to determine that half of small intestinal crypt cells in a mouse were clonal within eight weeks, and, furthermore, that the trajectory to monoclonality was consistent with neutral dynamics (López-García et al. 2010). The same result was derived using a confetti system in mice, under the control of Lgr5, a stem cell marker (Barker et al. 2007). Initially multicoloured crypts became monochrome over a period of 1-6 months in a manner consistent with neutral drift (Snippert et al. 2010). In humans, the monoclonality of colonic crypts has been demonstrated by staining for loss of the cytochrome oxidase (CCO) protein, which is encoded by the mitochondrial genome and so – due to the higher mutation rate of the mitochondrial genome and the process of drift to homoplasmy such that only one mutation is necessary – is reasonably frequently inactivated (Baker et al. 2014). This process of neutral drift is made possible, in part, by the relative non-

116

quiescence of intestinal stem cells: bromodeoxyuridine measurements indicate that the cell cycle time at the base of human colonic crypts is of the order of 30 hours (Potten et al. 1992).

I.1.c. Quantification of stem cell numbers and the rate of drift

The number of stem cells and the rate of neutral drift are important parameters to understand cancer risk, both because stem cells are believed to be the cell of origin of cancers and because the number of stem cells per crypt affects the probability that a driver mutation will be able to colonise a crypt (see General Introduction). Furthermore, the time to monoclonality of a crypt is of technical interest in this dissertation. As described below, in our study we sequenced individual colonic crypts, which allowed us to detect mutations that were present in every cell in the crypt. Thus, in effect we recover the genome of the most recent common ancestor (MRCA) of the crypt. If this common ancestor existed a very long time ago, we might be significantly underestimating the mutation burden of stem cells at the time of resection.

In mice, inducible labelling strategies have begun to unpick these parameters. Kozar et al. used a continuous labelling approach to mark a small proportion of cells, effectively allowing the output of single cells to be monitored. An out of frame reporter, under a house-keeping promoter, was placed next to a CA[30] microsatellite tract such that slippage of this highly mutable stretch would infrequently place the reporter back in frame (Kozar et al. 2013). Both wholly- and partially-labelled crypts were observed. The former increased linearly with mouse age, while the latter remained constant. The proportion of each of these informs on the stem cell number and replacement rate (if the rate is known at which the reporter is activated), allowing the inference of approximately seven stem cells per crypt in mouse colon, with most crypts drifting to monoclonality within a few months (Lopez-Garcia 2010). Interestingly, fewer stem cells per crypt were observed in the small intestine, which may contribute to the decreased cancer incidence.

Estimates in humans are more controversial. Coalescent modelling of variation in methylation patterns in human crypts, proposed to act as a somatically heritable but mutable mark, suggested the presence of at least eight stem cells and that the time to the MRCA of the crypt in humans is between 15 and 40 years (Nicolas et al. 2007). Substantial uncertainty exists in these estimates: a small number of CpG sites were assayed in a small number of cells per crypt, and the

model relies on assumptions on the kinetics of methylation and demethylation as well as the shape of the phylogeny of how differentiated cells are related to the common ancestor of the crypt. Cytochrome oxidase (CCO) staining has also been used to estimate these parameters (Baker et al. 2014). The observation that a ribbon of CCO⁻ staining varied in width as it rose from the base of the crypt to the lumen led the authors to conjecture that these 'wiggles' might be a read-out of symmetrical stem cell divisions at the base. This led to the estimation of approximately six functional stem cells and a rapid stem cell replacement rate, with monoclonal conversion times approaching three weeks (calculated in Nicholson et al. 2018). Ingenious though this idea was, it does not fit particularly well with our understanding of transit times within the crypt. Given that it takes about one week for a cell to migrate from the crypt base to the gut lumen in humans, and cells at the crypt base divide every few days (and not all of these need be symmetric cell divisions (Kozar et al. 2013)), it is arguable whether a ribbon of CCO⁻ staining captures a long enough time period to assay neutral drift. Only the bottom half of crypts were examined, and sometimes over eight wiggles per crypt are reported, which might indicate that the wiggles are rather a result of the behaviour of cells in the transit-amplifying compartment or that the software used to detect them is overly sensitive. In addition, due to mitochondrial heteroplasmy and neutral drift of mitochondria within the cell, the amount of cellular CCO can, in theory, fall below detectable levels and then recover. Finally, the continuous clonal labelling approach used in mice by Kozar et al. was applied to humans by staining for loss of the mPAS protein due to spontaneous somatic mutations (Nicholson et al. 2018). The median time to monoclonal conversion was estimated to be 6.3 years, and the number of stem cells to be between five and 10. A final caveat should be added to all the models of neutral drift discussed so far, whether in mouse or human, in that they treat all stem cells as having an equal chance of survival. However, 3D intravital imaging of a confetti mouse has indicated that stem cells further from the centre of the crypt base are more likely to be lost by differentiation (Ritsma et al. 2014).

I.1.d. Crypt fission

Crypts themselves occasionally divide to produce two daughter crypts (a process termed crypt fission) throughout life. The relevance of crypt fission to our understanding of colorectal

cancer is that neoplasia occurs, initially, through a process of crypt fission that allows a driver mutation to extend beyond the borders of the crypt in which it arose. A low crypt fission rate in normal colon probably reduces the rate of clonal expansion of mutations that could contribute to malignant transformation, and so reduces the probability of a 'second hit'. An understanding of the dynamics and regulation of crypt fission is, therefore, of central importance in our model of the evolution to cancer.

In humans, patches of multiple crypts that are CCO⁻ are observed, and both their frequency and size increase with age (Greaves et al. 2006). Sequencing the CCO gene in the two arms of a CCO⁻ bifurcating crypt showed that they shared the same inactivating mutation, indicating that fission rather than fusion was occurring (Greaves et al. 2006). Modelling of the CCO⁻ patch size as a simple birth process allowed an estimate of a fission event every 36 years, while modelling of crypt fission rates based on protein stains results in an estimate of crypts dividing on average every 140 years (Nicholson et al. 2018). It is unclear that a simple birth process is appropriate, as that would result in an increase in crypt number over the course of life, which – in mice at least – does not seem to be the case (Bruens et al. 2017). Recently, *in vivo* imaging has provided evidence for crypt fusion events in the mouse small intestine (Bruens et al. 2017), which could serve to control crypt numbers. Crypt exhaustion may additionally occur. Nonetheless, simulations indicated that the inclusion of crypt fusion has a negligible effect on the estimation of fission rates (Nicholson et al. 2018). Presumably, the ability of a crypt to fission has evolved as a regenerative response to damage. Indeed, inflammatory conditions often increase the number of crypts seen in fission.

I.2. The genomics of colorectal cancer

Cancer is a late product of somatic evolution, representing the end-point of the adenoma-carcinoma sequence (General Introduction). There now follows a brief overview of the genomic landscape of colorectal cancers. Cancers are the winners of somatic evolution, and their features provide clues as to the forces that govern natural selection in the colon. Two features warrant discussion: the mutation burden of colorectal cancers, and the features and numbers of driver mutations per cancer.

I.2.a. Mutational processes in colorectal cancer

Colorectal adenocarcinomas are one of the most highly mutated cancer types. They are surpassed only by melanomas, lung cancers – which are associated with exposure to the potent mutagens of ultraviolet radiation and tobacco exposure, respectively – and oesophageal cancers (Alexandrov et al. 2018). Colorectal cancers can be separated into two groups based on their mutation burden: so-called 'hypermutators' account for ~15% of cancers with a median of ~30 mutations per million coding bases, while the other ~85% typically have ~3 mutations per million coding bases (Cancer Genome Atlas Network et al. 2012). These two groups are generally associated with different histopathology, age-, site-, and gender-specific incidence, clinical features, and patterns of driver mutations. As alluded to in the General Introduction, hypermutators, with their high burdens of point mutations and short insertions and deletions (indels), tend to have far fewer copy number alterations (Cancer Genome Atlas Network et al. 2012), which is suggestive of two evolutionary paths to colorectal cancer, both of which are facilitated by an increase in mutation rate.

The analysis of mutational signatures in colorectal cancer provides a window into the origin of these mutations. Deconvolution of 60 whole colorectal cancer genomes, analysed in conjunction with thousands of other tumours as part of the Pan Cancer Analysis of Whole Genomes, revealed the activity of myriad mutational processes across the cohort: 13 single base substitution (SBS), 10 doublet base substitution (DBS), and four insertion and deletion (ID) signatures (Alexandrov et al. 2018) (figure 2.1). Most colorectal cancers had three to five SBS signatures, two to three ID signatures, and four to five DBS signatures. Signatures can be divided into those that affected almost all cancers and sporadic signatures that affected only a subset of cancers.

Common signatures might represent processes that are active in normal colorectal stem cells or that are necessarily associated with the process of transformation; without sequencing normal tissues, they cannot be told apart. Signatures that are common in colorectal cancers include SBS1, SBS5, SBS18, DBS2, DBS4, DBS6, DBS9, ID1, and ID2. Both SBS1 and SBS5 are found in almost all cancers sequenced to date. SBS1 accounts for a median of ~3,000 mutations per colorectal cancer genome (Alexandrov et al. 2018). It is characterised by C to T mutations in an NCG context (the mutated base is underlined), and is thought be due to the hydrolytic deamination of 5-methylcytosine to uracil (Rideout et al. 1990). In replication, an A is paired with the uracil,

and in the following round of replication, a T is paired with that A, such that the mutation is fixed as a T. As such, this is a process that is likely to occur spontaneously in all cells, but whose rate could be increased in cancers if mutations are more likely to be fixed by DNA replication rather than repaired. Cancers that arise from less mitotically-active tissues tend to have lower rates of SBS1 and on this basis it has been proposed that SBS1 acts as a mitotic clock (Alexandrov et al. 2018, Alexandrov et al. 2015). SBS5 is also found in almost all tissues sequenced, but is of unknown cause; its relatively featureless trinucleotide profile provides few clues as to its aetiology. SBS18 is present in a large fraction of colorectal cancers, but not all. It should be noted, though, that the complex nature of cancer genomes, with multiple mutational processes with overlapping trinucleotide profiles active in a given cancer, makes the extraction of mutational processes relatively complicated. SBS18 may well truly be present in all cancers but not always be detected because some of its mutations could be misattributed to a different signature. SBS18 is characterised by C to A mutations and has been linked to the activity of reactive oxygen species attacking guanines to form 8-Oxoguanine, which, if not excised, can pair with an A (Viel et al. 2017). ID1 and ID2 are, respectively, single base insertions and deletions of a single T in a polyT tract, postulated to be due to replication slippage, and DBS2, DBS4, DBS6, and DBS9 are of unknown origin, although some have been observed in normal mouse cells and DBS2 and DBS4 have been noted to correlate with the age of cancer diagnosis (Alexandrov et al. 2018).

Sporadic mutational processes detected in the PCAWG cohort of 60 colorectal cancers include: SBS10a, SBS10b, SBS15, SBS17a, SBS17b, SBS28, SBS37, SBS44, SBS45, DBS5, DBS8, DBS10, DBS11, and ID14. Some have a known cause. For example, SBS15, SBS44, DBS7, DBS10, as well as a marked increase of ID1 and ID2, are associated with loss of DNA mismatch repair. Cases with these signatures are the 'hypermutators' described above; it is notable that hypermutation is the result of a strong increase in a small number of processes rather than a generalised increase in all processes. Mutations in the proof-reading domain of polymerase epsilon are associated with vast numbers of mutations due to SBS10a and SBS10b. A large number of sporadic mutation processes active in colorectal cancer, however, are still of unknown cause (Alexandrov et al. 2018).

Larger structural changes have also begun to be classified thoroughly. Most colorectal cancers have of the order of a hundred structural variants, most of which are complex events, large

Figure 2.1. Categories for signature decomposition, reproduced with permission from Alexandrov et al. (2018). An example signature is shown for each of single bases substitutions (**a**), doublet base substitutions (**b**), and small insertions and deletions (**c**), in order to show the categories into which every signature is separated. In figures to follow in this chapter the category labels are often removed due to space constraints, but all are plotted with the same order and colouring as here.

a

SBS4

b

DBS2

c

ID3

deletions, or tandem duplications (Li et al. 2017), although, as noted above, hypermutated cancers have very few or none.

Thus, the mutational processes associated with colorectal cancer are diverse and variable across tumours. From sequencing cancers alone, it is unclear whether this is a reflection of the process of transformation or of the diversity in mutational processes operative across normal colorectal epithelium.

I.2.b. Driver mutations in colorectal cancer

I.2.b.i. Features of driver mutations

A comprehensive discussion of all driver mutations in colorectal cancer is beyond the scope of this chapter. Here, rather, I attempt to sketch out succinctly two of the molecular pathways that are frequently deregulated, and whose nature informs our understanding of somatic evolution in the colon. I will not discuss driver mutations that are thought to act by increasing mutation rate, since the debate about whether an increased mutation rate is of importance has been covered in the General Introduction. Studies modelling the effect of driver mutations in normal tissues are discussed in section I.5. of this chapter.

The vast majority of colorectal cancers harbour mutations that de-regulate the Wnt signalling pathway. Wnt ligands are absent from the microenvironment of differentiated colonocytes. A complex formed of APC, AXIN, and GSK3B phosphorylates β-catenin, targeting it for degradation. At the crypt base, however, Wnt ligands are present, and these bind to the Wnt receptor on crypt stem cells, signalling to inhibit the degradation of β-catenin. β-catenin may then shuttle to the nucleus, where, it binds the TCF4 protein and turns it from transcriptional repression to transcriptional activation of a wide array of genes, including *CMYC* and *CCND1* (reviewed in Fearon 2011 and in Bienz and Clevers 2000). Inactivating mutations have been documented in colorectal cancers in *APC* (homozygous truncating mutations in 70-80% of sporadic tumours and heterozygous loss of function mutations inherited in the fully penetrant colorectal cancer syndrome Familial Adenomatous Polyposis) and *AXIN2* (truncating mutations that are frequently heterozygous (Segditsas and Tomlinson 2006)). Both result in a failure to degrade β-catenin.

Mutations are also observed in *TCF7L2* (which encodes TCF4). While it is counterintuitive that loss of TCF4 should promote colorectal cancer, it has been proposed that loss of the repressive effect of TCF4 on some genes is a large part of the effect of complexing with β-catenin (Fearon 2011). It seems that these mutations of the Wnt signalling pathway essentially revert cells to a stem cell-like phenotype: disruption of the β-catenin-TCF4 complex in colorectal cancer cell lines halted proliferation and increased the expression of genes associated with differentiation (van der Wettering et al. 2002). Kinzler and Vogelstein proposed that *APC* is a gatekeeper gene, which they defined as follows: 'gatekeeper genes are responsible for maintaining a constant cell number in renewing cell populations [...] Mutation of the gatekeeper leads to a permanent imbalance of cell division over cell death' (Kinzler and Vogelstein 1996).

The second pathway that warrants brief discussion is that of Ras signalling. This pathway mediates pro-proliferative signalling downstream of receptor tyrosine kinases such as EGFR. KRAS, BRAF, PIK3CA, and PTEN all form a part of this pathway and their genes are mutated in colorectal cancers at frequencies of 40%, 10% (mostly in hypermutators), 20%, and 10% respectively (Fearon 2011). Interestingly, *KRAS* mutations are frequently found in hyperplastic polyps with little chance of progressing to cancer, indicating that they might drive proliferation sufficiently powerfully to create a macroscopic lesion without necessarily setting a cell firmly on the road to cancer.

To simplify the features of the most frequent driver mutations in colorectal cancer, then, it seems that these act through a combination inhibiting differentiation and driving proliferation. A plethora of other driver mutations in colorectal cancer, however, have a broad range of effects.

I.2.b.ii. Numbers of driver mutations

The number of driver mutations needed for colorectal cancer is a primary determinant of cancer risk (Armitage and Doll 1954, Nordling 1953, General Introduction) and contextualises the observation of driver mutations in normal colon: if only two driver mutations were necessary to cause cancer we would be far more alarmed by the observation of a single driver mutation in normal colorectal epithelium than if 20 were required. A distinction should be drawn between the number of driver mutations *observed* in colorectal cancers, and the number of driver mutations

necessary to cause a cancer. Ongoing evolution within a cancer will increase the number of mutations found in a cancer over the bare minimum necessary. Both will be discussed here, but the latter quantity is the one of greatest interest to us.

Tabulating the driver mutations that were known in 1990, Fearon and Vogelstein estimated that most cancers have four to five driver mutations (Fearon and Vogelstein 1990). This was probably an underestimate, as only small parts of the genome were assayed and driver mutations have been discovered since. A more recent analysis of 52 colorectal cancer whole genomes that identified coding and non-coding driver mutations with a method that was largely based on prior knowledge found a mean of 7.4 driver point mutations and 2.5 driver genomic rearrangements per colorectal cancer (Sabarinathan et al. 2017). There was substantial variation, however, as almost 20% of the cancers had three or fewer drivers, and another ~20% had 10 or more. Approaches that do not rely on prior knowledge are perhaps more robust and comprehensive. Martincorena et al. developed a dNdS method to detect driver point mutations in the coding regions of cancer genomes. This method compares the ratio of nonsynonymous to synonymous mutations in genes, under the assumption that nonsynonymous mutations code for proteins and are therefore subject to selection, while synonymous mutations have no functional impact. Thus, an excess of nonsynonymous mutations indicates positive selection of the gene, while a depletion indicates negative selection. The trinucleotide composition of the gene, the mutational processes active in the tumour, and local mutation rate were taken into account. Using this method, a mean of 10 positively-selected coding point mutations per tumour was estimated from hundreds of colorectal cancer exomes (Martincorena et al. 2017). Perhaps a similar analysis, restricted to the tumours that are estimated to have transformed recently, might be informative for the minimum number of mutations necessary to cause cancer, although some driver events, including rearrangements, noncoding mutations, and epigenetic alterations would be missed.

Estimates of the minimum number of driver mutations necessary were first derived using cancer incidence data. It was observed that the incidence of cancer rose proportionately to four to six times the power of age, which was interpreted as indicating that five to seven rate-limiting events (often assumed to be driver mutations) were necessary to cause a cancer (Nordling 1953, Armitage and Doll 1954, General Introduction). This method assumes that driver mutations are independent events; if each driver mutation induces a clonal expansion or increases the mutation rate, and thus increases the probability of another mutation affecting one cell of the mutated clone,

126

the number of driver mutations could be lower. Realisation of this, along with the discovery of recessive oncogenesis (Knudson 1971) led to the elaboration of a multi-stage clonal expansion model. In this model, a number of mutations, each considered to be an independent rare event, must accumulate in a cell before it undergoes a clonal expansion, and then another high frequency event must occur (another mutation in any one of the many mutated cells). Fitting this model to colorectal cancer incidence data, trying different numbers of pre-clonal expansion mutations, showed that the most likely number of pre-clonal expansion mutations was just two (Luebeck and Moolgavkar 2002). These two mutations were posited to correspond each to the inactivation of one allele of the APC gene. Revealing though this is, assumptions are made that could be incorrect: the first mutation may already cause a small clonal expansion and neutral drift within the crypt is not taken into account.

Tomasetti et al. observed that patients with Lynch syndrome have a 114-fold risk of cancer over the general population, but that the microsatellite unstable tumours that these patients get only have eight times as many mutations as microsatellite stable sporadic tumours. From this they inferred that three driver mutations are sufficient to cause cancer (Tomasetti et al. 2014). This analysis ignores the fact that the mutation rate is only accelerated in some of the cells in the colon of Lynch syndrome patients, those that have lost the second copy of a mismatch repair gene. It would hold better for patients born with biallelic mismatch repair, who get colorectal cancer much younger than Lynch syndrome patients (Wimmer and Kratz 2010). Secondly, the analysis does not consider the different pathways to cancer taken by mismatch repair deficient *versus* proficient cancers; for instance, the former often have driver mutations in *TGFBRII* and *BAX*, which contain microsatellite tracts and so must be mutated disproportionately more rapidly than other driver mutations.

An orthogonal approach to estimate the number of driver mutations per patient is to induce different combinations of driver mutations in cancer models. For example, intestinal organoids (*in vitro* clonal expansions that recapitulate crypt organisation) engineered to have driver mutations in *APC, TP53, KRAS*, and *SMAD4* grew independently of niche factors and developed a morphology similar to organoids derived from invasive carcinomas (Drost et al. 2015).

In summary, then, we still do not know conclusively the minimum number of driver mutations needed for a colorectal cancer. The weight of the evidence, however, seems to point to a handful of mutations, probably between two and seven, and thus fewer than the observed counts.

I.3. Colorectal cancer incidence

Some features of colon cancer incidence deserve a brief mention. The age incidence of colorectal cancer has been discussed in the context of driver mutations. Three other features of cancer incidence, however, raise questions of colon biology.

Firstly, while in the West colonic cancer is a common disease, with a lifetime risk of about 5% (Cancer Research UK Bowel Cancer Incidence Statistics), other parts of the world have much lower rates. Comparing extremes, there is a 10-fold age-standardised incidence difference between central Africa and Oceania. While this may be partly genetic, countries in Eastern Europe or in Asia that have recently adopted a more Western diet have seen rapid increases in incidence. A number of risk factors have been identified, including smoking, alcohol, and eating processed meats, as well as composition of the microbiome, infections with *Fusobacterium spp.*, and inflammatory diseases like ulcerative colitis. Conversely, low-dose aspirin has been shown to have a protective effect (Brenner et al. 2014). All of these could presumably alter one or both of the mutation rate and selection pressures in the colon, but their precise mechanism of action is as yet incompletely understood.

Secondly, differences in cancer incidence between parts of the gut are intriguing. Most strikingly, in the UK, the incidence of adenocarcinoma in the large bowel is approximately 60-fold higher than in the small bowel (Cancer Research UK Bowel Cancer Incidence Statistics, Aparicio et al. 2014), despite their similar stem cell biology. I am not aware of a simple explanation for this from the perspective of comparative evolution. Relative to other primates the colon has reduced in size even more than the small intestine: the colon takes up about 20% of the gut, whereas in other large primates it is closer to 50% (Milton 1987). While the function of the human colon is largely limited to water reabsorption, in primates such as gorillas the colon plays a major role as a fermenting chamber. Short chain fatty acids produced as a result of bacterial fermentation of fibre provides over half of a gorilla's calories, compared to less than 10% in a human (Popovich et al. 1997). Combined with the geographic variability in colon cancer risk, we can speculate that dietary changes associated with ever more energy-rich foods might be responsible for the high incidence of colorectal cancer observed nowadays in the Western world.

Thirdly, even within the colon there are substantial differences in cancer incidence. In men, cancer incidence is approximately twice as high in the sigmoid colon as in the caecum (23% v. 12%), while in women it is only slightly higher (20% v. 17%) (Cancer Research UK Bowel Cancer Incidence Statistics). Clinical factors may play a role in these incidence rates, since left-sided cancers typically cause symptoms earlier, but the difference is nonetheless striking and, to my knowledge, unexplained. These differences could be caused by small differences in the rate of known mutational processes, different frequencies of activity of sporadic mutational processes, differences in stem cell dynamics (such as stem cell numbers per crypt), or other factors.

The sequencing of normal colonic and small intestinal stem cells can begin to resolve the role of somatic mutations in these curious discrepancies in cancer incidence.

I.4. Current understanding of somatic mutations in normal colon

I.4.a. Mutation rates

In recent years there has been a flurry of interest in the mutational processes in normal tissues, including in normal colon. The mutations in normal colon can be inferred from cancer genomes, which represent the sum of the mutations that occurred before and after transformation. Assuming a similar lag time across colorectal adenocarcinomas between the departure from normal mutational processes and resection, the number of mutations due to normal mutational processes should correlate with patient age. Examination of this relationship in colon revealed that only signature 1 correlated with age, with a rate of 23 (95% CI 19-28) mutations per year (Alexandrov et al. 2015).[2] Interestingly, the curve passes through the origin. An increase in the mutation rate of signature 1 during carcinogenesis would shift the curve up, whereas a lag between diagnosis and the time to the most recent common ancestor of the tumour in whom mutations can be called would shift the curve down. While a novel insight, the authors note that 'Peering through the "cracked lens" of cancer genomes may obscure or distort the estimates of clock-like mutation

---

[2] Please note that this is the COSMIC version of signature 1, which is not composed exclusively of C to T at CpG mutations but has some background in other contexts that resemble those of signature 5. Using the PCAWG SBS1, the number of mutations attributed to this signature would be lower, and signature 5 may be found to accumulate linearly.

rates of normal cells that are progenitors of the cancers' (Alexandrov et al. 2015). Sequencing the genomes of normal cells provides far greater resolution and allows the investigation of non-clock-like processes.

Normal organoids derived from 21 single colonic stem cells from six patients ranging in age from nine to 67 years old (although none was between the ages of 15 and 53) provided the first clear insight into the mutations in normal colonic stem cells (Blokzijl et al. 2016). Small intestinal organoids were derived as well and, remarkably, both had the same mutation rate of 36 mutations per year (95% CI 26.9-50.6 for colon and 25.8-43.6 for small bowel), which indicates that the model of cancer risk being mostly due to stem cell divisions (Tomasetti and Vogelstein 2015, General Introduction) does not explain the difference between cancer incidence in the large and small bowel. Interestingly, four out of 15 colonic organoids that could be assessed were found to have structural variants, including a complex translocation and a trisomy of chromosome 13, while small deletions were found in three out of 14 small intestinal organoids. Similarly, high density SNP arrays on individual human colonic crypts showed the presence of deletions and amplifications, which increased in prevalence with age, with detectable copy number changes in one in seven crypts from a 78 year-old. (Hsieh et al. 2013).

Blokzijl and colleagues found three single base substitution signatures to be operative in colonic organoids: signatures 1, 5, and 18. Numbers are not provided in the text, but judging from the figures, signature 1 accumulated at a rate of about 25 (95% confidence interval ~18-38),[3] signature 5 at about 10, and signature 18 at about five mutations per year. Similar numbers were found in the small bowel, with a little less signature 1 and more signature 5. The rate of accumulation for signature 18 is not significantly different from 0, and signature 18 was found to be enriched in sequential *in vitro* cultures, which led the authors to ascribe it to an oxidative process during organoid culture. For all signatures in colon the slope of the curve of mutation burden versus age cuts the y intercept near to, but slightly above, the origin, which hints at a period of transiently increased mutation rate. The authors remark on the lack of interindividual variability in mutational processes, but with only six individuals they were unlikely to capture mutational processes that are not ubiquitous. No driver mutations were found, but one disadvantage of the organoid culture system is that it can select against driver mutations in certain tissues; indeed, in the colon, wild-

---

[3] Again, this is the COSMIC version of signature 1. Using the PCAWG definition, probably fewer mutations would be attributed to SBS1 and more to SBS5.

type organoids will outcompete mutant ones unless niche factors are removed (van de Wetering et al. 2015). Despite its relatively limited power, this study is seminal in that clean whole genomes of single normal cells were seen for the first time. The discrepancy between the mutation burden of ~3,000 mutations in a colonic stem cell from a 60 year-old and the average mutation burden of 10,000-20,000 mutations in a non-hypermutated colorectal cancer demonstrates that cancers have an elevated mutation burden over their normal counterparts.

Nonetheless, until normal and tumour from the same people are studied, it remains theoretically possible that those people who get cancer have an elevated mutation rate all around their bowel. In an attempt to resolve this, organoids were derived from APCmin/+ mouse adenomatous crypts and normal crypts and exome sequenced (Lugli et al. 2017). The rate of acquisition of point mutations was found to be ~11 fold higher in adenomas, although small numbers of mutations were captured: only 71 mutations were seen in total across seven normal organoids and 15 tumour-derived organoids. The caveats of organoids (as discussed above) remain, and mouse intestinal organoids have been shown to have different mutational spectra to human ones, with an enrichment of C to A and fewer C to T mutations (Blokzijl et al. 2016, Behjati et al. 2014). Furthermore, while this indicates that the *people* who get cancer need not have a generally increased mutation rate, it is possible that the *cells* that become cancerous could, even prior to transformation, have had an increased mutation rate. A phylogenetic analysis of tumours and comparison to normal tissues from the same patients could resolve this.

An orthogonal approach to quantify mutation rates in normal tissues despite their polyclonality is to perform very deep and highly error-corrected sequencing, such that mutations in individual molecules of DNA can be called reliably (General Introduction). Analysis of normal colonic epithelium from 11 individuals showed an increase in mutation rate with age, reaching ~3,500 mutations per genome in people over the age of 40 (Hoang et al. 2016), which is similar to the number found by sequencing organoids. It should be noted that this assays mutations across all cells in the epithelium, some of which will have arisen during the process of differentiation, while organoids only report the mutations in stem cells. The similarity between these two estimates then indicates that there is no dramatic increase in mutation rates over the course of differentiation.

I.4.b. Driver mutations in normal colon

Relatively little is known about the frequency of driver mutations in normal human colon, largely due to difficulty in their detection. To my knowledge, two approaches have been used so far: PCR-based methods to detect specific mutations, which are only practical for assaying hotspots, and staining for tumour suppressor proteins. This detects homozygous mutations that result in a loss of expression rather than merely a loss of function.

KRAS hotspot mutations have been detected through PCR-based methods since 1998 in a number of studies, but most were non-quantitative or only semi-quantitative and frequently failed to detect *KRAS* mutations in normal mucosa (discussed in Parsons et al. 2010). Parsons and colleagues used allele-specific competitive blocker polymerase chain reaction (ACB-PCR) to quantify *KRAS* codon 12 GTT and GAT mutations in 89 samples of colonic mucosa (Parsons et al. 2010). Mutant *KRAS* was detected in all samples of normal mucosa, and it was estimated that 1 in 3,500 normal cells contained a KRAS codon 12 mutation. This is 60 times more frequent than the mutation is expected to occur by chance (Tomasetti et al. 2013), indicating positive selection. *KRAS* mutations were more frequent in the sigmoid colon, concordant with the observation that sigmoid tumours more frequently have *KRAS* mutations. Interestingly, the frequency of *KRAS* codon 12 GTT mutations was found to be higher in adenomas than in carcinomas, indicating that some *KRAS* mutations may promote the transition to malignancy more effectively than others. No correlation of variant allele fraction was observed with the patients' age, which may be a result of the relatively small age span covered (50 to 80 years old). A more recent study using targeted sequencing on a larger cohort of patients validated the presence of *KRAS* mutations in normal colon (Nicholson et al. 2018). Similarly, ACB-PCR investigation of *PIK3CA* found that the H1047R mutation was above the detection threshold of 1×10^{-5} in 20 out of 20 normal samples, whereas the E545K mutation was not detected in any of the 20 samples (Parsons et al. 2017).

Recently, staining for four putative tumour suppressor proteins located on the X chromosome (such that loss of one allele was sufficient to inactivate all copies of the protein) was performed for 186 patients across an age range (Nicholson et al. 2018). One of these, STAG2, was found to be lost in most patients at a mean frequency of about one in 7,000 crypts in a 60 year-old. The mechanisms by which these driver mutations might colonise colorectal epithelium are discussed below.

I.5. Quantitative insights into the earliest stages of colorectal cancer evolution

In order to form a tumour, driver mutations that occur in colonic stem cells must be able first to sweep through the crypt and second to spread beyond it. Quantitative analyses of the effects of driver mutations in mouse models, and more recently in humans, have begun to elucidate how this occurs.

I.5.a. Driver mutations skew the odds of stem cell competition

Under a model of neutral drift, neutral mutations that arise in a single stem cell have a probability of becoming fixed that is inversely proportional to the number of stem cells per crypt. A mutation that decreases the probability that the cell in which it occurs will be lost from the crypt is more likely to be able to go on to form a tumour. Mouse models indicate that this is a property of common colorectal cancer driver mutations in *Apc* and *Kras* (Vermeulen et al. 2013, Snippert et al. 2014). Vermeulen et al. (2013) induced driver mutations and a coloured tag in mice at infrequent levels, such that only one cell in a crypt would be mutated. Quantifying the growth of the labelled cells allowed a model of the benefit that the driver mutation conferred. *Kras* G12D mutant stem cells outcompeted their wild-type neighbours 80% of the time. Snippert et al. (2014) induced both *Kras* G12D and the confetti reporter at a low frequency in Lgr5 positive cells, and similarly found that the mutant cells had a higher chance of sweeping through a crypt. An EdU pulse showed that *Kras* mutant cells were cycling faster, consistent with our understanding of the Ras pathway in driving proliferation.

Vermeulen et al. also studied *Apc* and *P53*. Interestingly, *Apc* +/- cells outcompeted *Apc* +/+ cells, but *Apc* -/- cells outcompeted *Apc* +/- cells. This shows that mutations of tumour suppressor gene alleles are not necessarily independent events, as they are frequently modelled to be. The selective advantage of *Apc* mutations within the crypt is consistent with its suggested role in controlling the balance stem cell self-renewal and differentiation (section I.4.b.). As discussed in the General Introduction, *P53* mutations were only found to be advantageous when colitis was

induced. The concept of a context-specific driver mutation is both fascinating and daunting in that it adds another layer of complexity to the construction of quantitative models of cancer.

In humans, comparing STAG2 to a neutral mark showed that the proportion of crypts that had lost the expression of STAG2 in all cells was increased relative to the proportion of crypts where STAG2 was lost in only a fraction of cells, indicating more rapid clonal sweeps (Nicholson et al. 2018). STAG2 loss was estimated to increase the probability of replacing a wild type neighbour from 0.5 to 0.99.

I.5.b. Driver mutations increase the rate of crypt fission

Crypt fission is rare physiologically (Baker et al. 2014, Nicholson et al. 2018), so unless a driver mutation can increase this rate, it is likely to remain entombed in its own crypt. It seems that as well as giving a selective advantage within the crypt, canonical driver mutations also promote clonal expansion beyond the crypt.

Snippert et al.'s (2014) multi-coloured labelling method allowed them to analyse the dynamics of crypts that had been fully colonised by *Kras*. In the presence of *Kras* mutations, adjacent crypts were more likely to be the same colour, which indicated an excess of crypt fission events. It was estimated that *Kras* G12D increased the rate of crypt fission 30-fold. In humans, the discrepancy between the rate at which *KRAS* mutations should occur by chance and their allele fraction in bulk epithelium indicates that they must increase the rate of crypt fission approximately 10-fold (Nicholson et al. 2018). Similarly, STAG2-negative patches of epithelium tended to be larger than patches of epithelium negative for neutral marks. Modelling the growth of these patches with age showed that 0.7% of crypts with neutral marks fissioned per year, whereas 2.15% of crypts with STAG2 loss fissioned per year.

Theoretically, early driver mutations need not cause crypt fission. In ulcerative colitis (a risk factor for colorectal cancer), crypt fission rates are increased, presumably as part of a wounding response, and so driver mutations could hitch-hike out of the crypt; once the clone was big enough, a second driver mutation that did allow a disruption of the tissue architecture would be more likely to strike the clone. Nonetheless, it seems probable that most colorectal cancer drivers that tend to occur early in the adenoma-carcinoma sequence will promote crypt fission as

it would provide a strong fitness advantage. Many known tumour suppressor genes seem to have evolved at the time when our ancestors became multicellular (Domazet-Loso and Tautz 2010), and it has been suggested that they served the purpose of limiting the selfish behaviour of cells in metazoa. With this in mind it is perhaps not so surprising that their deregulation results in an atavism that involves concomitant proliferation and disruption of tissue architecture.

Thus, we begin to be able to describe in a quantitative manner how driver mutations can colonise the colorectal epithelium. Much remains unanswered, however. To list but a few questions: what are the actual mechanisms behind a competitive advantage within the crypt, and what governs crypt fission? Are mutations that allow cells to outcompete their wild type neighbours necessarily advantageous to cancer? What are the effects of combinations of driver mutations? And what other driver mutations lurk in normal colonic epithelium?

Results

R.1. Study design

We aimed to investigate the landscape of somatic mutations in normal colon. A small number of normal colonic organoids had previously been sequenced (Blokzijl et al. 2016, section I.4.a), which indicated the activity of only three mutational signatures in normal colon and little variation in between different samples. We designed an experiment to explore the variety in mutation burden, mutational processes, and frequency of driver mutations in the normal colonic mucosa between different people and between different crypts within one person. We set out to exploit the stem cell architecture of the colon as a clonal unit by laser capture microdissection of single crypts, followed by sequencing. The advantages of laser capture microdissection over organoids (which were used by Blokzijl and colleagues (2016)) are the following: the method is more easily scaled to analysing hundreds of samples; spatial information on the location of the crypts is retained, allowing the investigation of processes such as crypt fission; there is no selection of crypts in culture, allowing an unbiased discovery of driver mutations; and no mutations are acquired *in vitro*. On the other hand, one downside of bulk sequencing whole crypts is that only

mutations in the most recent common ancestor of the crypt are called. The time to the most recent common ancestor of the crypt is likely to be within the decade before resection (section I.1.c.). This should be borne in mind in analyses of mutation burden, but is not problematic for the discovery of mutational signatures.

R.2. Development of a protocol to sequence individual crypts

At the time when this experiment was begun (Autumn 2015), 500ng of DNA were required for Illumina to guarantee sequencing success, and sequencing was rarely performed with less than 200ng. Colonic crypts, each with ~2,000 cells (~12ng of DNA) of which only a fraction are obtained in a given section, were far below this threshold. We developed a pipeline to allow the sequencing of single colonic crypts. Peter Ellis developed a more sensitive library preparation method (Methods), while I, with advice from Robert Osborne, optimised the thickness of sections, the choice of fixative, the staining protocol, and the lysis method. All experiments were performed on fresh frozen colonic tissue, initially from a mouse, and later from a human.

The thickness of sections was chosen to be the largest possible that still allowed the dissection of single crypts. In very thick sections, if crypts are visualised longitudinally it may be that a fragment of another crypt is hidden behind the back wall of the crypt that is being dissected, which would result in a polyclonal sample. The spacing between crypts may vary depending on factors such as mucosal oedema. Images of *en face* crypt sections from a number of mucosal samples, however, showed that 30 micron sections rarely resulted in capturing one crypt behind another. The staining regimen was chosen to be the simplest that still allowed crypts to be visualised, the rationale being that any unnecessary chemicals might damage DNA. Furthermore, as these experiments require very long days, any time that can be saved is valuable. Crypts were therefore stained only with Gill's haematoxylin and no eosin. With 30 micron sections and staining with haematoxylin, crypts could clearly be seen as clonal units (figure 2.2a). The images are much less attractive than those in textbooks due to a combination of the thickness of the section (pathology sections are often only 4 microns thick), the use of fresh frozen tissue (whereas pathology sections are formalin-fixed and paraffin-embedded), the absence of a coverslip (since

this cannot be used in laser capture microdissection), and the fact that sections were only stained with haematoxylin (rather than haematoxylin and eosin as in standard pathology).

The best fixative and lysis methods were evaluated jointly by the quantification of libraries made by Peter Ellis. Four fixatives were tested: acetone, paraformaldehyde, methanol, and ethanol. Three different lysis methods were assayed: alkaline lysis, protease lysis, and chaotropic lysis (RLT). Below are the results for half a plate testing these combinations. From these results and repeat experiments that confirmed them, fixation with methanol and protease lysis were chosen.

	7	8	9	10	11	12
A	blank	blank	blank	blank	blank	blank
B	Alk EtOH 1 crypt	Alk PFA 1 crypt	Prot EtOH 1 crypt	Prot PFA 1 crypt	RLT EtOH 1 crypt	RLT PFA 1 crypt
C	Alk EtOH 1 crypt	Alk PFA 1 crypt	Prot EtOH 1 crypt	Prot PFA 1 crypt	RLT EtOH 1 crypt	RLT PFA 1 crypt
D	Alk EtOH 1 crypt	Alk PFA 1 crypt	Prot EtOH 1 crypt	Prot PFA 1 crypt	RLT EtOH 1 crypt	RLT PFA 1 crypt
E	Alk ace 1 crypt	Alk MeOH 1 crypt	Prot ace 1 crypt	Prot MeOH 1 crypt	RLT ace 1 crypt	RLT MeOH 1 crypt
F	Alk ace 1 crypt	Alk MeOH 1 crypt	Prot ace 1 crypt	Prot MeOH 1 crypt	RLT ace 1 crypt	RLT MeOH 1 crypt
G	Alk ace 1 crypt	Alk MeOH 1 crypt	Prot ace 1 crypt	Prot MeOH 1 crypt	RLT ace 1 crypt	RLT MeOH 1 crypt
H	blank	blank	blank	blank	blank	blank

	7	8	9	10	11	12
A	0.076	0.12	0.113	0.12	0.14	0.167
B	0.233	0.114	0.12	1.71	0.125	0.168
C	0.086	0.092	0.098	2.55	0.139	0.133
D	0.079	0.08	1.504	0.198	0.121	0.152
E	0.144	0.204	0.693	28.895	0.118	0.152
F	0.073	0.248	11.444	18.849	0.121	0.135
G	0.156	0.096	19.788	6.247	0.12	0.173
H	0.107	0.1	0.11	0.111	0.125	0.127

Table 2.1. Quantification of libraries to test fixation and lysis condition. The top panel shows the layout of this half of the plate. The bottom panel shows library preparation results, in ng/ml. Alk, alkaline lysis; Prot, protease lysis; RLT, chaotropic lysis; EtOH, ethanol fixation; ace, acetone fixation; PFA, paraformaldehyde fixation; MeOH, methanol fixation.

With these sample preparation methods and the library construction protocol developed by Peter Ellis, 11 single colonic crypts from one sample were sequenced at 1-2x coverage each. Even coverage across the genome was observed, and sequencing metrics were acceptable. Pooling the crypts allowed known germline mutations in this patient to be recovered. Sequencing at higher coverage (~15X per crypt) allowed somatic mutations to be called using our standard algorithms

Figure 2.2. Laser capture microdissection of crypts. a, a representative image of a section of colonic tissue, with a magnified inset showing the section before and after dissection of a crypt. **b-c,** the coverage of crypts that underwent whole genome (**b**) and targeted (**c**) sequencing. Crypts are ordered by their mean depth (shown below), and for each crypt the proportion that is covered by a certain read depth is shown as a stacked barplot. **d-e**, their respective VAF (which is half of the clonal fraction). **f-h**, the distribution of coverage over exonic regions of putative colorectal cancer driver genes, from combining both whole genome and targeted data. **f,** the mean coverage across all samples of each gene. **g**, the number of crypts in which each gene was covered by an average of >7 reads, and **h**, the number of crypts in which each gene was covered by an average of >9 reads. Please note that the ordering of genes in each figure is different.

a

b

c

d Median VAF WGS crypts

e Median VAF targeted crypts

f

Mean coverage over all crypts

g

Number of crypts with average coverage over gene >7 reads

h

Number of crypts with average coverage over gene >9 reads
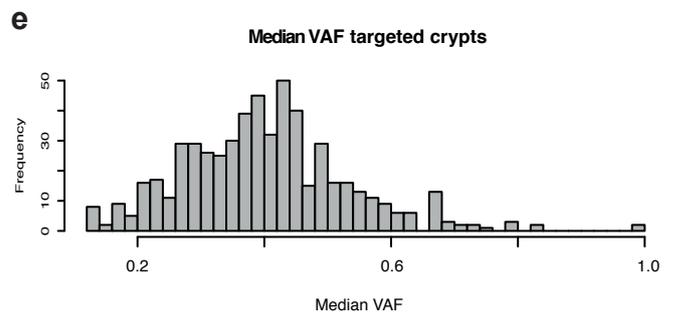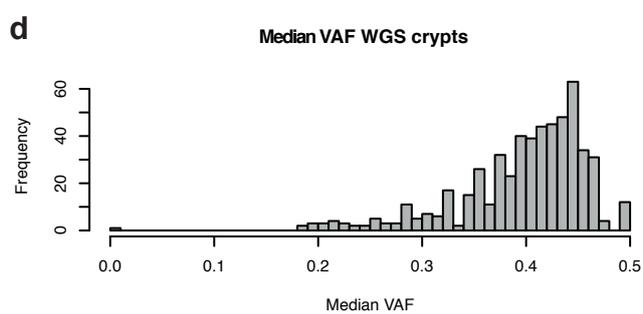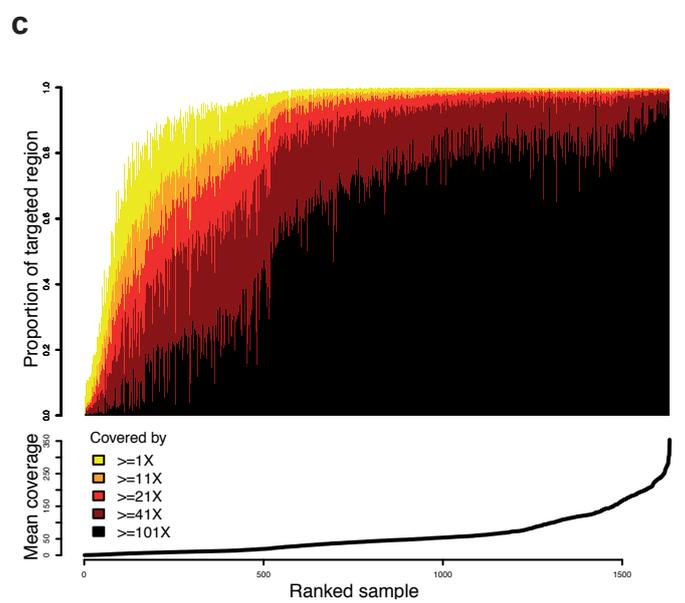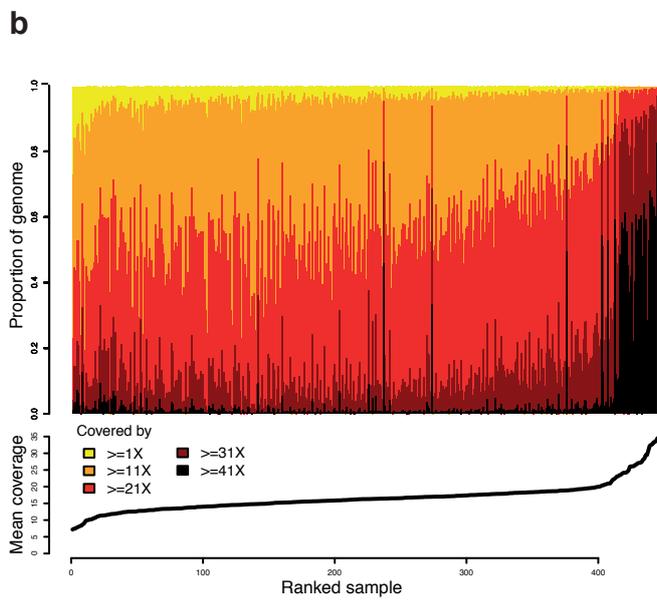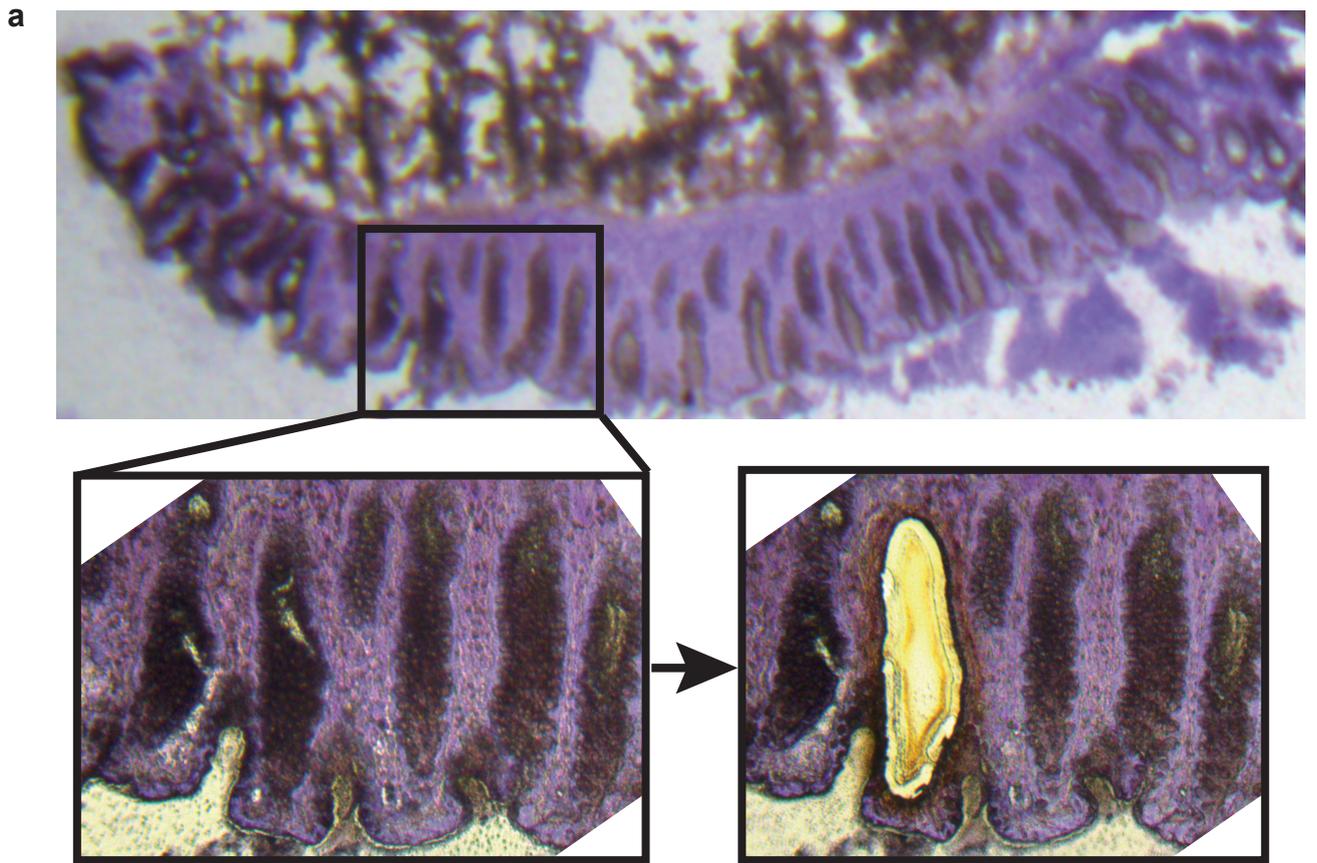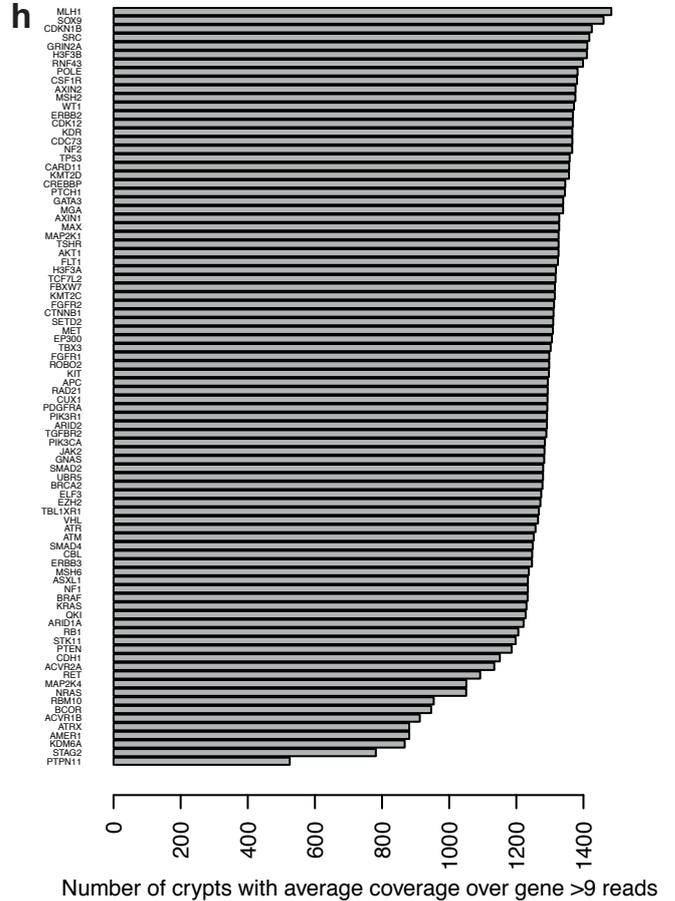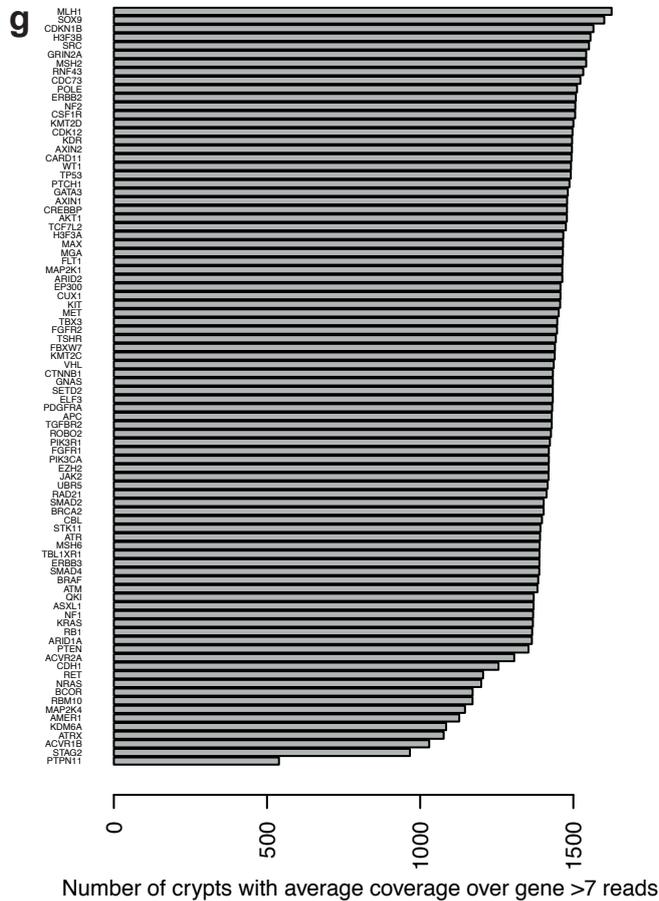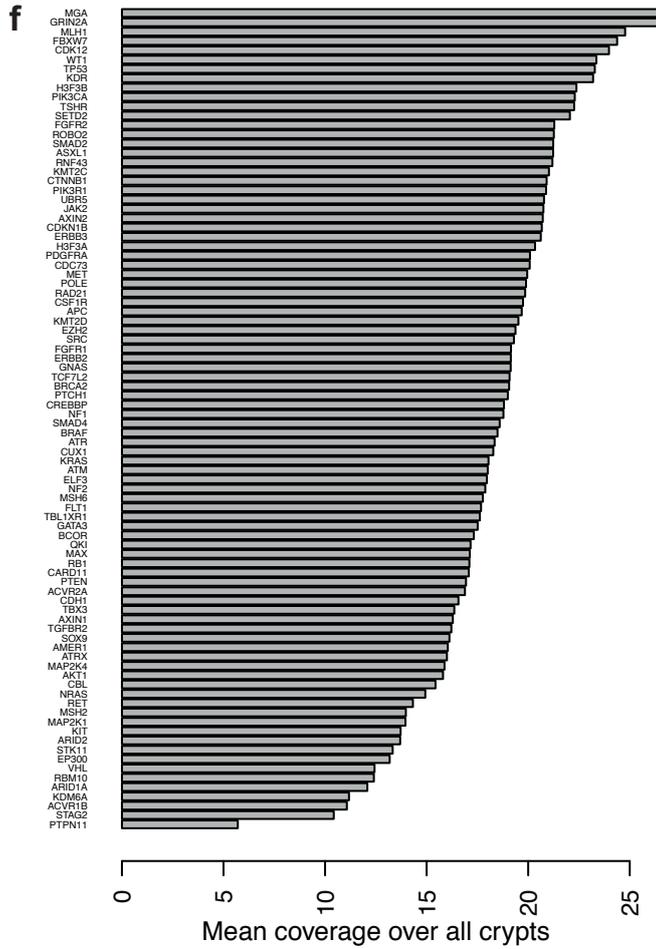
that have been developed for cancer genomes. Subsequently, bioinformatic filters were developed by Mathijs Sanders to remove artefacts that are due to this library preparation method (Methods).

It should be noted that some wells are empty because the crypt does not fall into the well. Electrostatic attractions often cause a dissected segment to stick to the underside of the slide. Visual inspection of adjacent wells never showed a crypt that had gone into the wrong well. In general, 30-40% of dissected crypts resulted in libraries with over 5ng/ul, which was chosen as the threshold to proceed to sequencing for most experiments.

R.3. Samples

Samples were obtained from four cohorts in order to cover as broad an age span as possible (Methods). 42 patients aged 11 to 78, 27 of whom had no diagnosis of colorectal disease and 15 of whom had been found to have a colorectal adenocarcinoma, were investigated. Wherever possible, biopsies from the caecum, transverse, and sigmoid colon were taken, as well as terminal ileum in a subset of cases. From these samples I dissected >5,000 crypts, of which 2,035 were sequenced: 571 were whole genome sequenced at ~15X coverage (figure 2.2b), and 1,464 underwent targeted sequencing using a bait-set of known cancer genes. Inspection of the allele fractions from the whole genomes showed that most crypts were 80-90% clonal (figure 2.2d), with some contamination which is likely to be stromal. The clonality of crypts that underwent targeted sequencing should be the same, but its assessment is less accurate as few mutations are called per genome (figure 2.2e). Targeted sequencing is more sensitive to low amounts of input DNA because of the additional step of bait hybridisation. Targeted sequencing on such small quantities of DNA was at the limit of current technical capabilities and variable coverage was achieved (figure 2.2c). CaVEMan (the algorithm used to call substitutions) requires – in most cases – three reads to call a mutation in a diploid genome. With a depth of eight reads, a clonal sample will achieve this many mutant reads 85% of the time (based on the binomial distribution). We therefore considered that only sites covered by at least eight reads could be genotyped with acceptable accuracy. Pindel (the algorithm used to call small insertions and deletions), however, requires four mutant reads to call a mutation. With a depth of 10 reads, a clonal sample will achieve this many mutant reads 83% of the time, and so we only considered that we could accurately genotype indels where sites were

covered by at least 10 reads. Because of the different coverage requirements for calling substitutions vs small indels, we estimate that we were adequately powered to call substitutions in 1,403 and indels in 1,046 of all crypts (genomes and targeted combined).

R.4. Driver mutations

Two approaches were taken to discover driver mutations in normal crypts: first a modelling method to detect positive or negative selection of genes, and second, manual annotation based on prior knowledge.

A dNdS approach (section I.2.b.ii., Methods, Universal patterns) was used to detect positive or negative selection. Two separate analyses were performed: a genome-wide analysis using only the whole genomes, and an analysis restricted to 90 putative colorectal cancer driver genes (Appendix B) that were included in our bait-set, using combined genomes and targeted sequences. In both analyses, the 95% CI for the global dNdS spanned 1, which indicates that the vast majority of the mutations in our dataset are selectively neutral. In the genome-wide analysis, no genes were significantly mutated. In the analysis of 90 genes, however, there was evidence of positive selection of two genes: *AXIN2* (three truncating mutations, adjusted p value 0.004), and *STAG2* (two truncating mutations, adjusted p value 0.038).

AXIN2 is a negative regulator of the WNT signalling pathway (section I.2.b.i.). *AXIN2* is inactivated in 2.3% of colorectal adenocarcinomas and smaller proportions of other cancer types (Forbes et al. 2017). One of the three *AXIN2* nonsense mutations was present in two adjacent crypts that were closely related genetically, sharing 1,606 SBS1 mutations; the *AXIN2* mutation must have occurred in their common ancestor which then underwent crypt fission. In one sister crypt, but not the other, the *AXIN2* mutation was rendered homozygous by copy number neutral loss of heterozygosity of chromosome 17q (figure 2.3a-d). This suggests that while loss of one copy of *AXIN2* already confers a growth advantage and may have contributed to the crypt fission, loss of a second copy could provide a further advantage and aid the expansion of the mutant clone, as has been shown for other tumour suppressors in mouse models (section I.5.a.). This provides evidence for ongoing clonal evolution in normal colon.

STAG2 is a component of the cohesin complex, which has roles in sister chromatid cohesion, DNA repair, and regulation of gene expression and chromatin structure (Hill et al. 2016). Although *STAG2* loss has been associated with aneuploidy in solid tumours (Solomon et al. 2011), this is not always the case (Hill et al. 2016, Balbas-Martinez et al. 2013, Taylor et al. 2014), and we do not observe this here. It is inactivated in 0.9% of colonic adenocarcinomas, and more frequently in other tumour types (Forbes et al. 2017). In our dataset, both *STAG2* nonsense mutations occurred in men, so no wild type copies of this X chromosome recessive cancer gene would remain in these cells. STAG2 loss has previously been shown to confer a proliferative advantage in human colon (section I.5.).

Figure 2.3. Driver mutations in normal colon. a-d, driver mutation in *AXIN2*. **a,** a section (after dissection) in which an inactivating *AXIN2* mutation was found. Red dots represent crypts with the *AXIN2* mutation. Blue dots represent crypts that could be assessed and were found not to have the mutation. Crypts without dots failed sequencing and could not be assessed. **b**, the two crypts with the *AXIN2* mutations prior to dissection did not appear different to any other crypts. **c**, copy neutral loss of heterozygosity (CNN-LOH) of one of the crypts over the *AXIN2* locus. The copy number state (y axis) for every chromosome is shown, with one allele coloured red and the other green. **d**, Jbrowse image of reads supporting the *AXIN2* mutations in each of the crypts. The mutation is coloured red. 25 out of 29 reads support the mutation in the crypt that has CNN-LOH; the four reads that do not are presumably the result of stromal contamination. **e**, putative driver missense mutations in oncogene hotspots. The number of substitutions catalogued in COSMIC (Forbes et al. 2017) are shown on the y axis at each position along the gene, with the mutations observed in our cohort highlighted.

a

b

c

d

PD37449e_lo0014 BWA alignment

PD37449e_lo0015 BWA alignment

e

PIK3CA

E542K

R38H

FBXW7

R505C

R658Q

ERBB2

R678Q

V842I

T862A

ERBB3

R475W

R667L
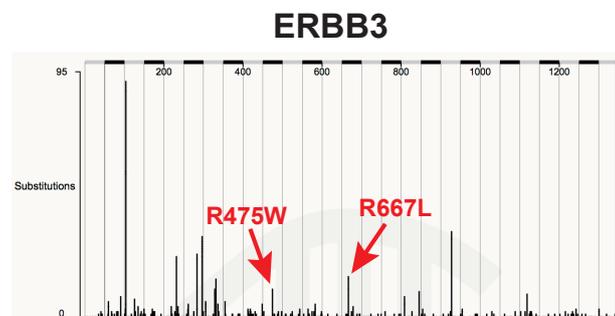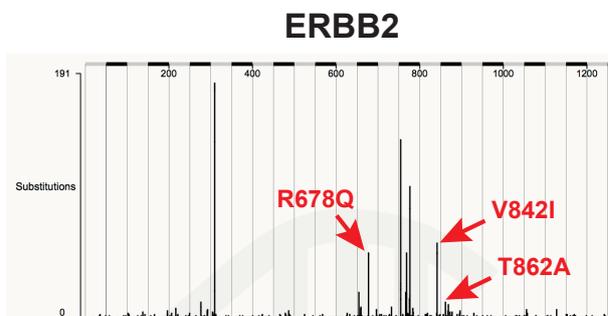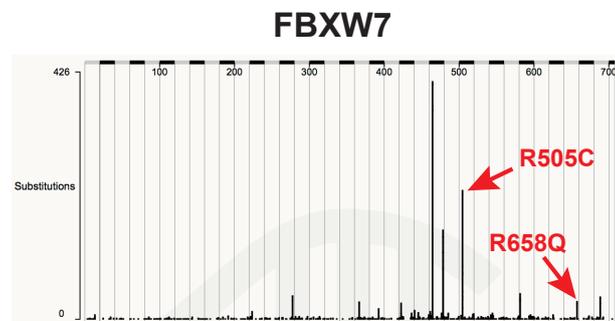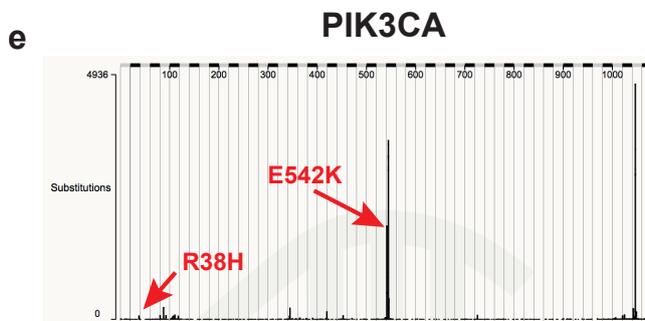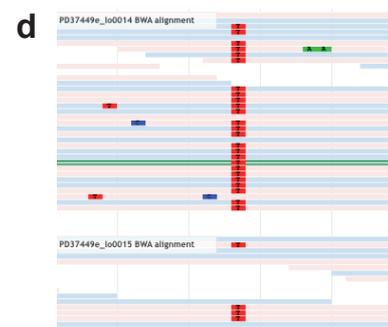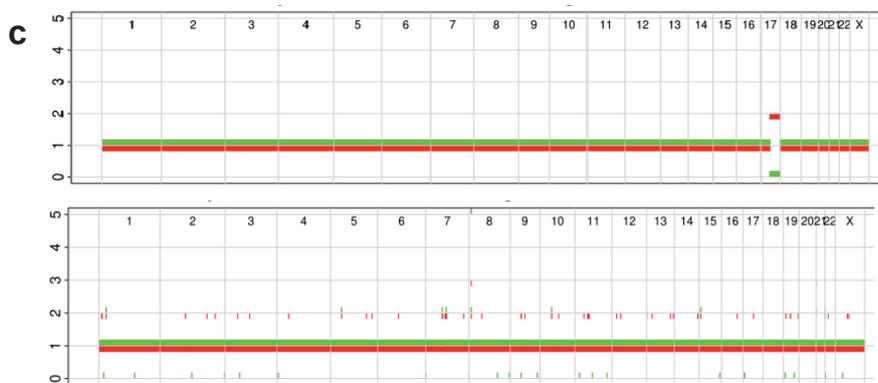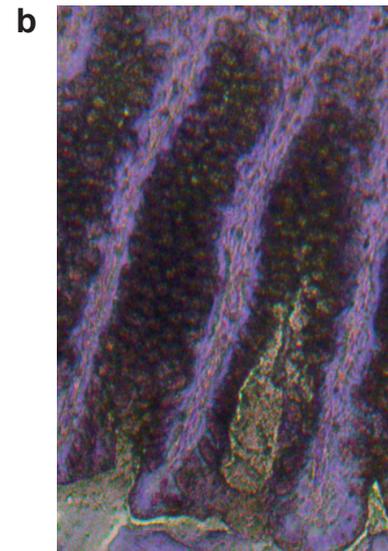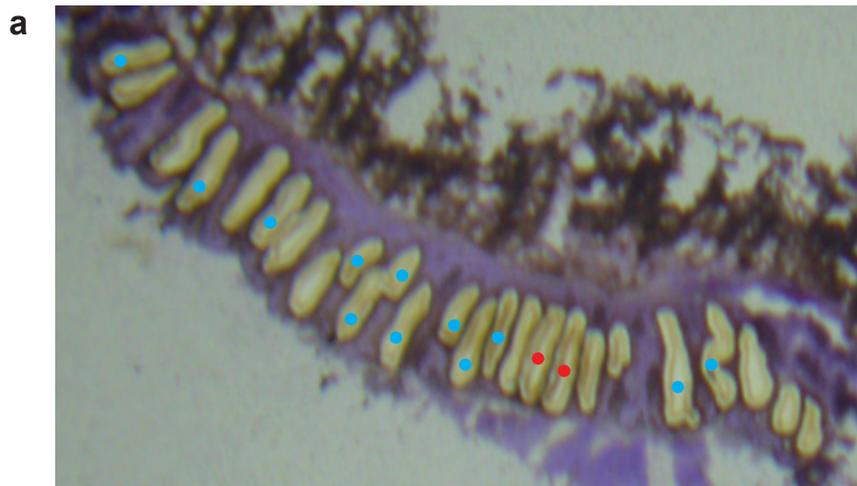
Additional potential driver mutations were identified by manual curation based on the known cancer genes in colorectal cancer and their distinctive patterns of mutation. Nine canonical missense hotspot mutations in the dominantly acting cancer genes *PIK3CA* (E542K, R38H (a minor hotspot)), *ERBB2* (R678Q, V842I, T862A), *ERBB3* (R475W, R667L), and *FBXW7* (R505C, a major hotspot, and R658Q, a minor hotspot) were observed (figure 2.3e). Given the specificity of these mutation hotspots, the majority of these are likely to be driver mutations and confer some growth advantage. As with *AXIN2*, the *PIK3CA* E542K mutation was also in two adjacent crypts which shared 2,516 SBS1 mutations and had 93 and 208 private SBS1 mutations, implying a recent crypt fission event. Indeed, these two crypts shared more mutations than any other pair of crypts in our dataset.

Finally, a series of heterozygous truncating mutations in the recessive colorectal cancer genes *ARID2*, *ATM* (two mutations), *ATR*, *BRCA2*, *CDK12* (two mutations), *CDKN1B*, *RNF43* (two mutations), *TBL1XR1*, and *TP53* were found. They were not associated with loss of heterozygosity, and no crypt had more than one driver mutation. It is likely that some of these did not confer any selective advantage. Nonetheless, in mouse colon heterozygous nonsense mutations of the *Apc* tumour suppressor gene can confer a selective advantage (Vermeulen et al. 2013), and indeed the *AXIN2* mutations for which we have compelling evidence of driver function were mostly heterozygous. Even if not currently advantageous, these mutations could set the scene for future clonal expansions, through loss of the remaining wild type copy or a change in microenvironment as has been observed for *P53* (General Introduction, section I.5.a. of this chapter).

On the basis of these findings, treating the *AXIN2*, *STAG2*, and dominant cancer gene hotspot mutations as drivers (i.e. all but the heterozygous mutations in recessive cancer genes) we estimate that at least 1% of normal colorectal crypts in a 50-60 year old carries a driver mutation. We are underpowered to detect a change in driver frequency with age. Since there are ~15 million crypts in the colon, ~150,000 crypts carry a driver. In the over 70s, ~40% of people have an adenoma on colonoscopy (Corley et al. 2013) and ~5% of people develop colorectal cancer over their lifetime (Cancer Research UK bowel cancer incidence statistics), and some of these may arise from more recently-acquired driver mutations. Therefore, only an extremely small proportion of these crypts with driver mutations becomes a macroscopically detectable adenoma (< 1/375,000) or carcinoma (< 1/3,000,000) within the following few decades.

146

Given that they are so common in colorectal cancer and absent from our dataset, we might conjecture that crypts with driver mutations in *APC* or *KRAS* as a first hit might have a higher chance of progressing. Using the mutation rate observed in our data, we can estimate that even if they were selectively neutral, 12,000 crypts from a 50-60 year old would have inactivated one copy of *APC*, and five crypts would have both copies inactivated. As even heterozygous *Apc* mutations confer a selective advantage in mouse models (Vermeulen et al. 2013), the true frequency is likely to be higher. PCR-based analysis of bulk epithelium has shown that 1 in 3,500 epithelial cells bears *KRAS* G12D (Parsons et al. 2010), which would indicate a few thousand crypts with this mutation per colon. Thus, even for these mutations, the probability of progression to cancer must remain low.

Comparison of the frequency of particular cancer gene mutations between normal epithelium and colorectal cancers extends the adenoma-carcinoma sequence and informs on the properties of driver mutations. Mutations reported in 260 cancers (Cancer Genome Atlas Research Network et al. 2012) were annotated for driver mutations using the same criteria as for the manual annotation of driver mutations in normal tissues. The pattern of driver mutations is different between cancer and normal tissues (p=0.003 by randomization test, figure 2.4). In colorectal cancer, mutations in *APC*, *KRAS* and *TP53* are common, accounting for 56% of base substitution and indel drivers but are comparatively rare among normal crypts with driver mutations (1 out 14 drivers). By contrast, mutations in, for example, *ERBB2* and *ERBB3* are relatively common in normal crypts with drivers (5/14) but rare in colorectal cancer (7/631). It is, therefore, not simply that the genes found in cancers are found at lower frequencies in normal tissues, but rather some account for a greater proportion of driver mutations in cancer than they do in normal tissues.

A mutation may be enriched in a cancer for four reasons: (1) the mutation itself has itself promoted progression of the lesion; (2) the mutation provides a selective advantage only in a neoplastic microenvironment; (3) the mutation occurs more frequently in the cancer due to an increased mutation rate; or (4), the mutation itself increases the mutation rate and so the probability of one of its descendants acquiring a mutation with properties (1) or (2). Mutations that provide a selective advantage in normal tissues but do not confer any of the above properties will be found at equal frequencies in cancer and in normal tissue.

Figure 2.4. Comparison of driver frequency in normal colon and colorectal cancer. The frequency of driver mutations in colorectal cancer is derived using data from Cancer Genome Atlas Network (2012). **a**, the proportion of crypts or cancers with driver mutations in each gene found in either of the two groups. **b**, the proportion of driver mutations in each gene in normal and cancer.

a

Normal crypts   **Cancer**

Significantly enriched in cancer
Non−significantly enriched in cancer
Significantly depleted in cancer
Non−significantly depleted in cancer

Proportion mutant   Proportion mutant

b

Normal drivers   **Cancer drivers**
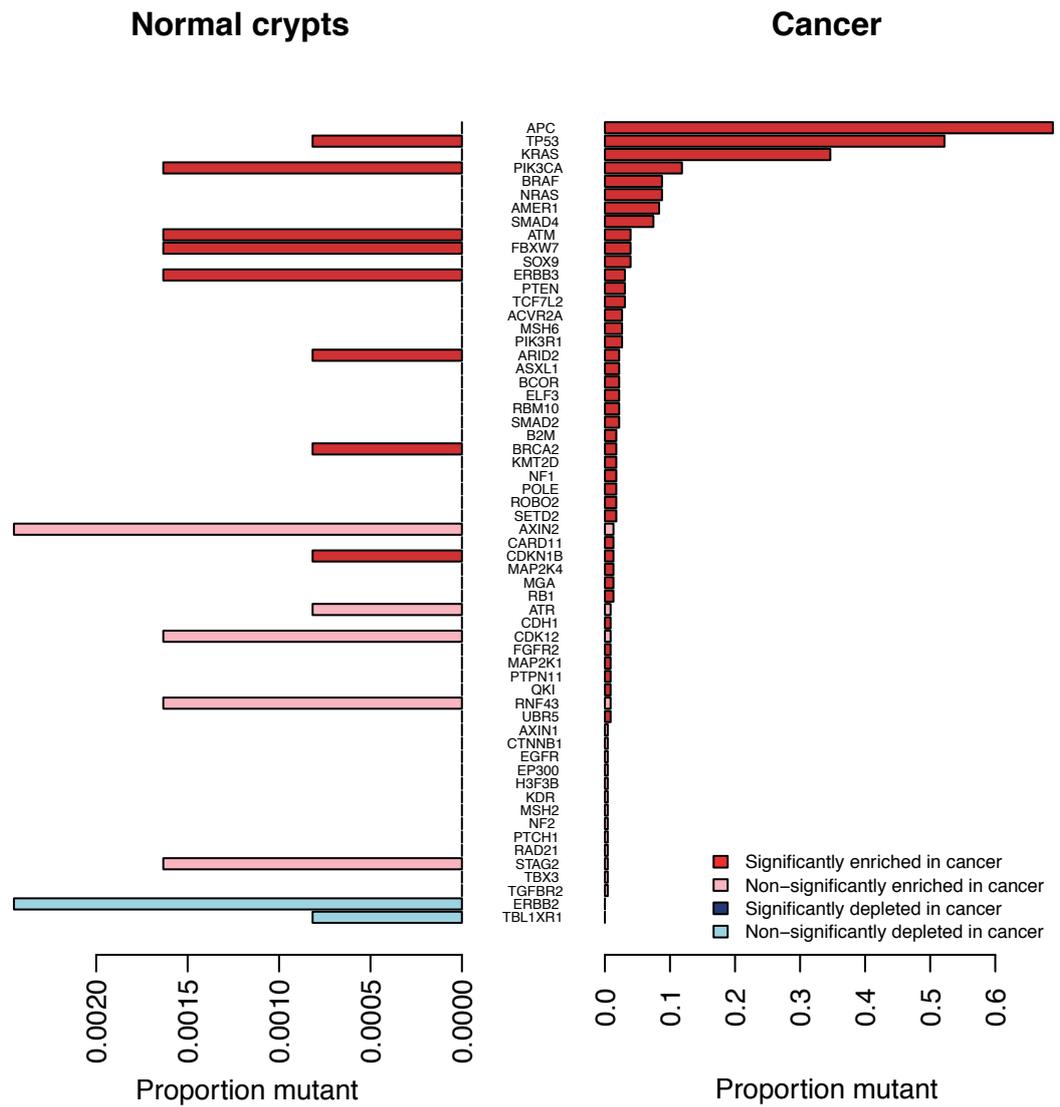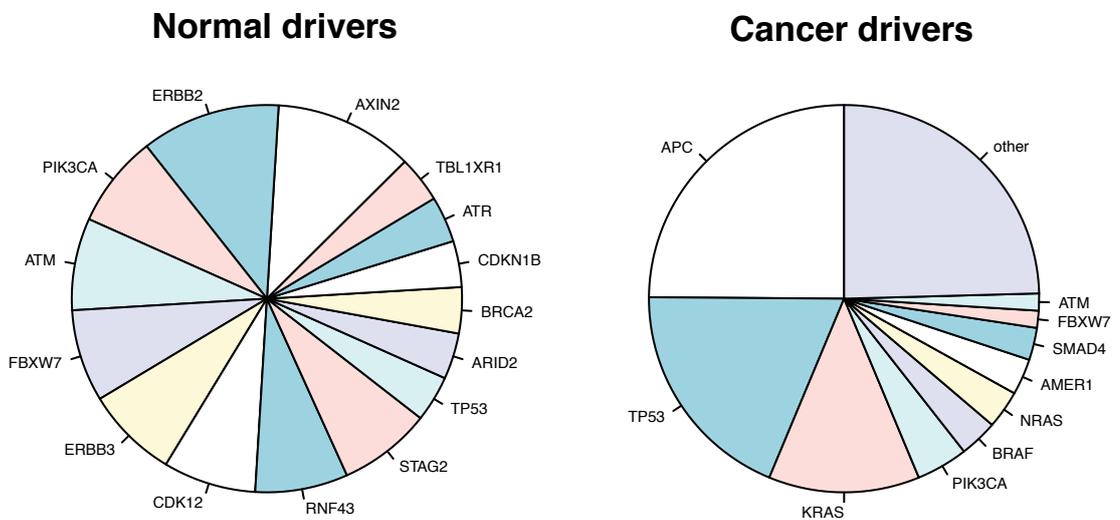
Thus, *APC*, *KRAS,* and *TP53* mutations, which are present orders of magnitude more frequently in cancer than in normal tissues (figure 2.4a) will confer properties (1)-(4), while mutations in *ERBB2* or *STAG2*, which are not significantly enriched in cancers, may merely provide a selective advantage in normal tissues without promoting cancer development. The distinction must therefore be drawn between mutations that are under positive selection and those that actually promote cancer development, although there will be significant overlap between the two categories.

Uneven coverage over the genome presents some difficulties when estimating the frequency of driver mutations. Ideally, every locus in every gene would be covered by many reads in every crypt, and the true frequency of driver mutations would need no correction. Second best, all genes would be covered equally well, but all genes would not be callable in a certain proportion of crypts. The frequency of driver mutations could then be estimated by dividing the number of drivers by the number of crypts in which they were callable. Third best, coverage would be even within each gene, but some genes would be better covered than others. If there were a sufficient frequency of drivers in each gene, we could estimate the true frequency of drivers per gene. Let us imagine that gene A achieves sufficient coverage to call mutations in 1,000 crypts and 10 drivers are found in it, and gene B achieves sufficient coverage to call mutations in 2,000 crypts and 100 drivers are found in it, we would say that the 1% of crypts have gene A drivers and 5% of crypts have gene B drivers. Assuming that these were occurring in different crypts, we would conclude that 6% of crypts have driver mutations.

The first challenge of our data is that the driver frequency is so low that we cannot estimate a per-gene driver frequency. All we can do to derive a meaningful estimate is to pool our driver mutations. Thus, if we can detect one gene A and one gene B driver mutation in our cohort, and the mean number of crypts in which we can call mutations accurately over all base pairs is 1,500, our best estimate is that two in 1,500 crypts has a driver mutation.

The second challenge is that coverage may fluctuate even within a gene. Some portions of a gene may be well covered in 1,000 crypts, and others in 2,000 crypts. All we can do here is treat different *parts of a gene* with different coverage like the different *genes* in the section above; that is to say, to take an average of the number of crypts in which they are covered. The approach that we took, therefore, was to calculate, for the average exonic base pair in our 90 cancer genes, the number of crypts in which that base pair was covered by >=8 reads (for substitutions) and by >=10 reads (for indels). 64% of all bases in the targeted panel across all crypts are covered by >=8 reads,

150

which equates to a number of callable bases equivalent to having sequenced ~1,400 crypts with perfect coverage over every base in every crypt. This average number of crypts in which all base pairs achieve good coverage becomes the denominator for calculating the driver mutation frequency (with the number of drivers observed in the dataset as the numerator). A similar approach can be taken with indels.

Note that for this reason of uneven coverage, this approach is less suitable to estimating the frequency of mutations in a given gene. Particular driver mutations, may be under-represented in our cohort. If one part of gene A is covered by a very low number of crypts, and that it is this part where most driver mutations occur, we will underestimate the frequency of driver mutations in gene A. Similarly, other driver mutations may be over-represented. This should be borne in mind when considering figure 2.4. The true frequencies of driver mutations in these genes may, in time, reveal themselves to be different to those that we have estimated from imperfect data here. We nonetheless present this figure as a preliminary indication of the landscape of driver mutations in normal colon.

Our estimate of 1% uses a global correction, on the assumption that under-representation and over-representation will even itself out when estimating the total frequency of driver mutations in the whole dataset. Without prior knowledge of which are under-represented and which are over-represented, using a mean is a valid approach. We stress the highly simplified nature of this approach. It is our resort because the frequency of driver mutations in our dataset is so low. The value that we derive of 1% of colonic crypts bearing a driver should be taken as a first ballpark estimate to guide further investigation. Further studies of larger numbers of crypts will be required to achieve greater accuracy.

R.5. Mutational processes and rates in the colons of different people

There was substantial variation in mutation burdens between individual crypts, ranging from 1,508 to 15,329 for individuals in their sixties, which was not obviously attributable to technical factors. To explore the biological basis of this variation we extracted mutational signatures from the whole genomes and estimated their contribution to the mutation burden of each crypt.

We first derived phylogenies of how crypts from each patient were related to one another and assigned every mutation to a branch of a phylogeny. This allowed us to treat every branch as a sample in signature extraction. This has the dual advantage of avoiding double counting of mutations in signature extraction (as mutations shared between two crypts are only counted once with this method, while they would be counted twice if every crypt were treated as a sample), and of allowing us to time mutational signatures over the course of life, since mutations shared by two samples must have occurred before mutations that are private to one of them.

The mutational signatures that are extracted from an analysis are dependent on the samples that went into it. If all samples have perfectly correlated contributions of different processes, these will only be extracted as one signature. A cohort of normal genomes from a single tissue runs this risk. We wanted to be able to frame our signature extraction results in the context of previous work in cancers, in order to allow comparisons with different studies. We therefore performed a signature extraction using a hierarchical Dirichlet process (HDP) (Roberts et al. 2015, Roberts et al. 2018), providing the algorithm with the catalogue of mutational signatures extracted from the Pan Cancer Analysis of Whole Genomes (PCAWG) (Alexandrov et al. 2018). This allows simultaneous discovery of new signatures and matching to known ones. Nine single base substitution (SBS), two doublet base substitution (DBS), and five indel (ID) signatures were discovered (figure 2.5). Despite pre-conditioning, signatures that were perfectly correlated in all samples were still amalgamated. This occurred, for example, with signatures 1, 5, and 18. Therefore, expectation maximisation was used to deconvolute all HDP signatures into known PCAWG signatures. If a signature reconstituted from the components that expectation maximisation extracted (only including PCAWG signatures that accounted for at least 10% of mutations in each sample to avoid over-fitting) had a cosine similarity to the HDP signature of more than 0.95, the signature is hereafter presented as its expectation maximisation deconvolution (Methods). Three HDP signatures met these criteria: the HDP SBS1 signature was deconvoluted into a mixture of PCAWG SBS1, PCAWG SBS5, and PCAWG SBS18; the HDP DBSN1 was deconvoluted in PCAWG DBS2, PCAWG DBS4, PCAWG DBS6, PCAWG DBS9, and PCAWG DBS11; and the HDP IDN1 was deconvoluted into PCAWG ID1, PCAWG ID2, and PCAWG ID5 (figure 2.6). To test the robustness of this signature analysis, other signature extraction methods were used: HDP with no pre-conditioning, the non-negative matrix factorisation (NNMF) method

used by Blokzijl and colleagues (2016), and a version of the NNMF algorithm used by Alexandrov and colleagues (Alexandrov et al. 2018). These all produced comparable results (figure 2.8).

Following expectation maximisation deconvolution, we found nine SBS, six DBS, and five ID signatures in our dataset (figure 2.7). Of these, 14 were known and six were novel (denoted by an "N" in their nomenclature, e.g. "SBSN1"). The signatures can be divided into those that are common and those that are rare depending on whether they are present in more or less than 85% of crypts. The common signatures are: SBS1, SBS5, SBS18, DBS2, DBS4, DBS6, DBS9, DBS11, ID1, ID2, and ID5. The rare signatures are: SBS2, SBS13, four novel SBS signatures (SBSN1 – SBSN4), DBS8, and two novel ID signatures (IDN2 and IDN3). The correlation with age of every signature is shown in figure 2.10.

The mutational signatures extracted are dependent both on the process causing mutations and on the trinucleotide composition of the genome. As variable coverage was achieved across the genome of crypts (figure 2.2), it was theoretically possible that some of our lower coverage crypts may be altering the profile of the signatures that were extracted. Mutational signatures were historically extracted from exome data (Alexandrov et al. 2013), and when they were updated to include genomes (Alexandrov et al. 2018), mutational processes largely remained the same. Those that changed in profile did so as a result of including additional samples with very simple trinucleotide profiles that allowed the NMF algorithm to draw out the salient features of the signature more clearly, rather than because of including other parts of the genome with a different trinucleotide composition. This would suggest that the trinucleotide profile of parts of the genome that are less well covered in some crypts would have less of an effect on the signatures extracted than one might expect. So as to detect changes in mutational signature composition that might be due to coverage, we ordered crypts by increasing coverage, and plotted the proportional contribution of different signatures to them (figure 2.9a). No obvious systematic differences with coverage are observed, and all novel signatures are seen in crypts with both good and bad coverage. Second, we compared the raw trinucleotide profiles of a representative selection of six of our lowest-depth crypts (all with average depth <10X), with the trinucleotide profile of six normal colonic organoids (see section R.7.), all of which were sequenced at 30-40X. There are no obvious differences between the trinucleotide profiles of these high- and low-coverage samples (figure 2.9b).

Figure 2.5. Signatures of mutational processes in normal colon. Signatures extracted by HDP are shown, with the trinucleotide context of a sample that contains a large proportion of the relevant signature shown underneath. Signatures are presented as in figure 2.1.

SBS1

Sample rich in SBS1

SBS5

Sample rich in SBS5

SBS18

Sample rich in SBS18

SBS2

Sample rich in SBS2

SBS13

Sample rich in SBS13

SBSN1

Sample rich in SBSN1

SBSN2

Sample rich in SBSN2

SBSN3

Sample rich in SBSN3

SBSN4

Sample rich in SBSN4

DBS8

Sample rich in DBS8

DBSN1

Sample rich in DBSN1

ID1

Sample rich in ID1

ID2

Sample rich in ID2

IDN1

Sample rich in IDN1

IDN2

Sample rich in IDN2

IDN3

Sample rich in IDN3

Figure 2.6. Expectation maximisation (EM) decomposition of HDP signatures. Three signatures were decomposed by EM. For each signature, the original HDP version in shown on the top left, the PCAWG signatures that are deemed by EM to contribute at least 10% of mutations to it on the right, and the reconstituted signature built by combining the PCAWG signatures on the bottom left. The cosine similarity of the reconstituted signature to the original is shown in the title to the reconstituted signature plot.

Original hdp SBS1

PCAWG SBS1 accounts for 0.39

PCAWG SBS5 accounts for 0.48

Reconstituted SBS1, cosine similarity to original: 0.99

PCAWG SBS18 accounts for 0.14

Original hdp DBSN1

PCAWG DBS2 accounts for 0.221

PCAWG DBS4 accounts for 0.235

PCAWG DBS6 accounts for 0.126

PCAWG DBS9 accounts for 0.176

Reconstituted, cosine similarity to original: 0.962

PCAWG DBS11 accounts for 0.144

Original hdp IDN1

PCAWG ID1 accounts for 0.284

PCAWG ID2 accounts for 0.144

Reconstituted, cosine similarity to original: 0.963

PCAWG ID5 accounts for 0.547

Figure 2.7. Final set of signatures in normal colon, following EM decomposition. These are the set that are used in analyses.

SBS1

SBS5

SBS18

SBS2

SBS13

SBSN1

SBSN2

SBSN3

SBSN4

DBS2

DBS4

DBS6

DBS9

DBS11

DBS8

ID1

ID2

ID5

IDN2

IDN3
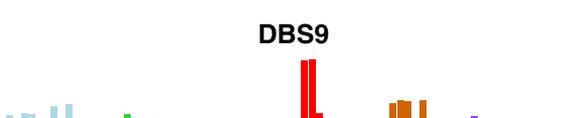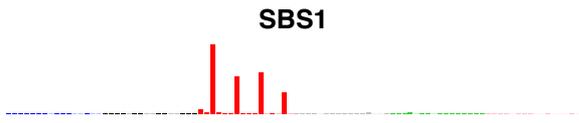
Figure 2.8. Validation of single base substitution signatures. Other methods of signature extraction were run to test the robustness of signature decomposition. **a**, HDP without pre-conditioning on PCAWG. **b**, In-house NNMF without pre-conditioning on PCAWG. **c**, NNMF implemented by the MutationalPatterns R package (Blokzijl et al. 2016).

a

hdp sig1 hdp sig2 hdp sig3

hdp sig4 hdp sig5 hdp sig6

hdp sig7 hdp sig8 hdp sig9

b

nmf sig1 nmf sig2 nmf sig3

nmf sig4 nmf sig5 nmf sig6

nmf sig7 nmf sig8 nmf sig9

c

mutpat sig1 mutpat sig2 mutpat sig3

mutpat sig4 mutpat sig5 mutpat sig6

mutpat sig7 mutpat sig8 mutpat sig9
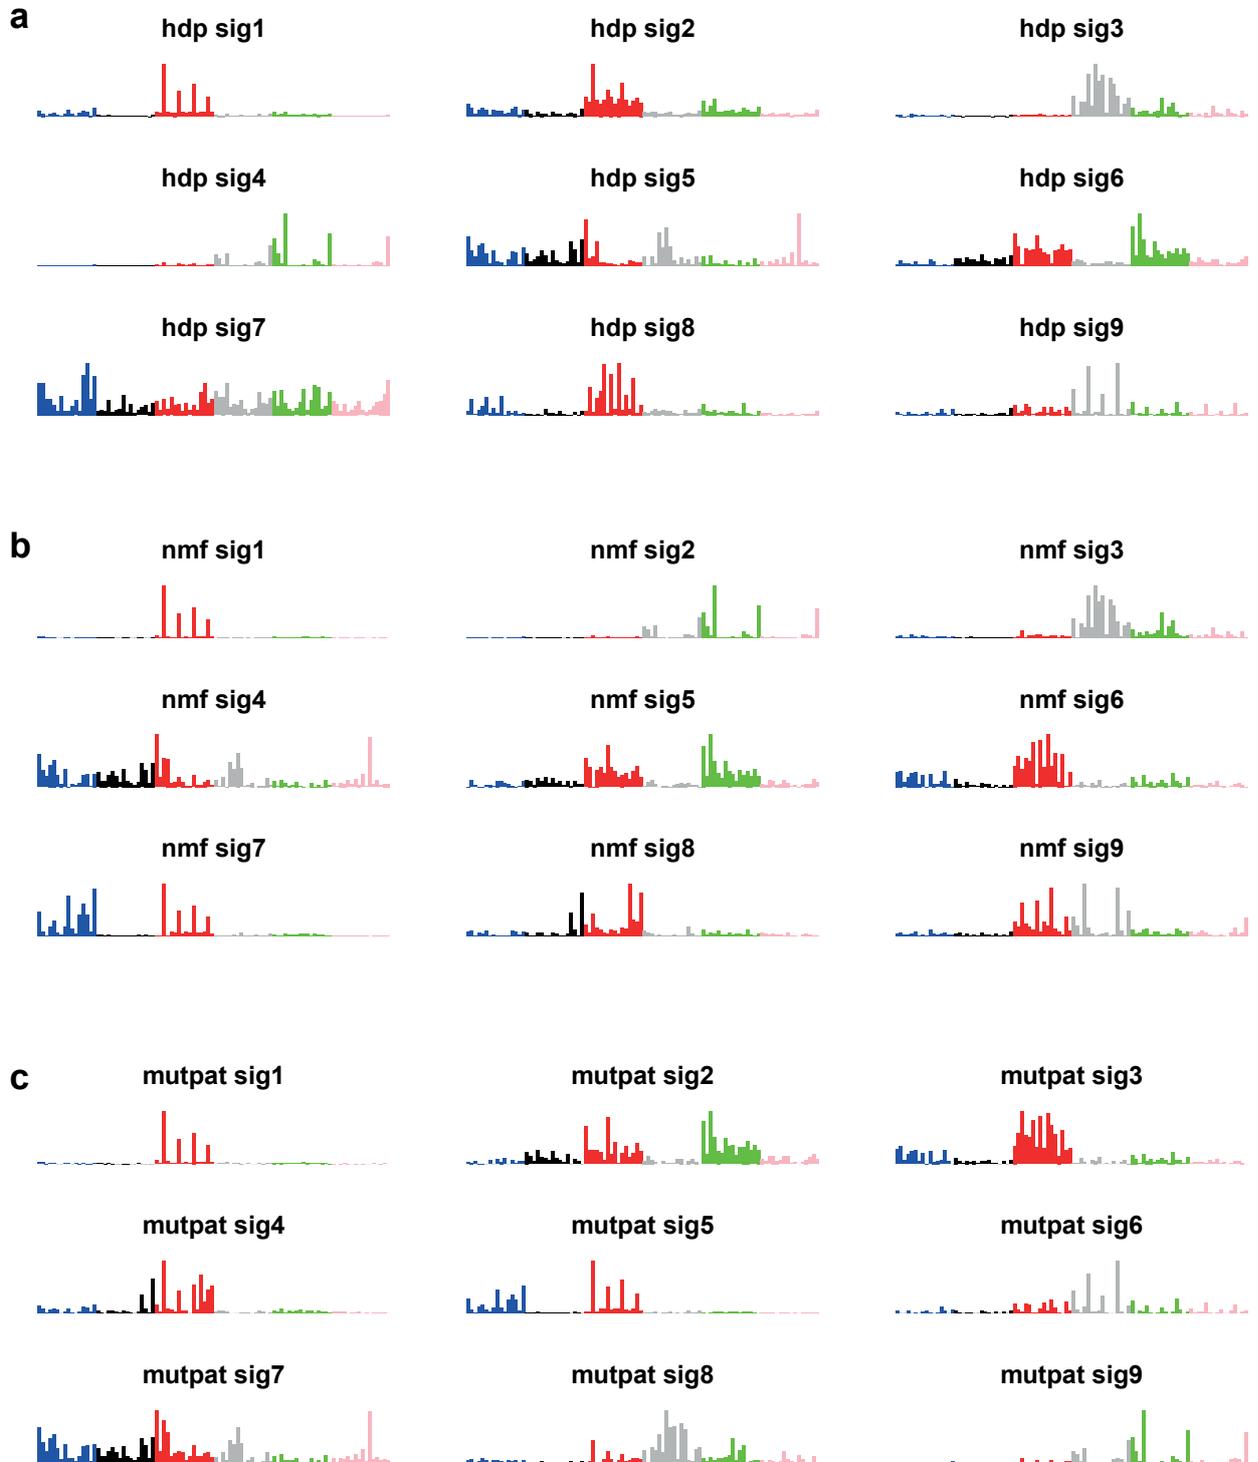
Figure 2.9. Mutational signatures with coverage. a, crypts are ranked by their median coverage, with coverage increasing from left to right. The proportional contribution of each signature is presented a stacked barplot. **b,** the trinucleotide profile of six crypts with high coverage are boxed in red, and the trinucleotide profile of six crypts with low coverage are boxed in blue. The high coverage samples are normal colonic organoids (see section R7) all sequenced at >30X, whereas the low coverage samples are laser capture microdissected crypts all with coverage <10X.

a
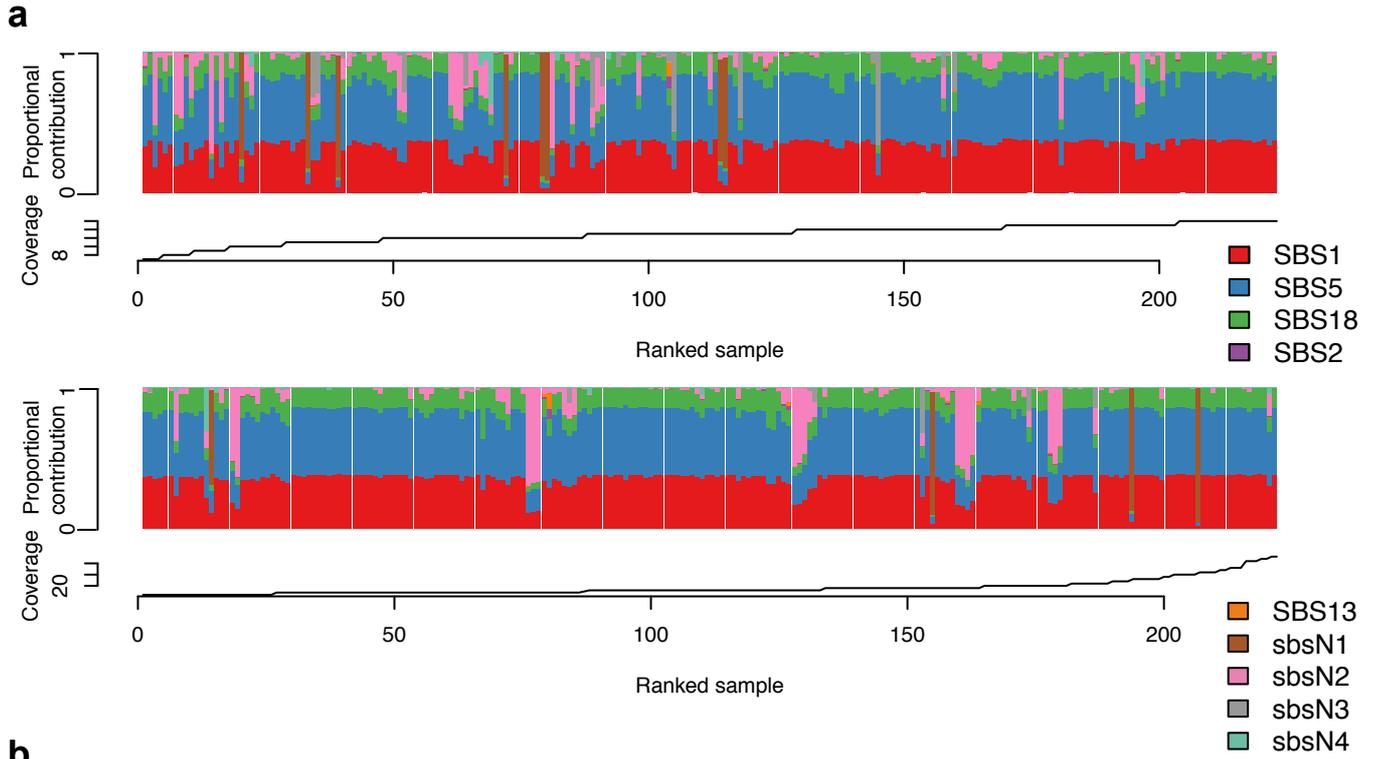
b
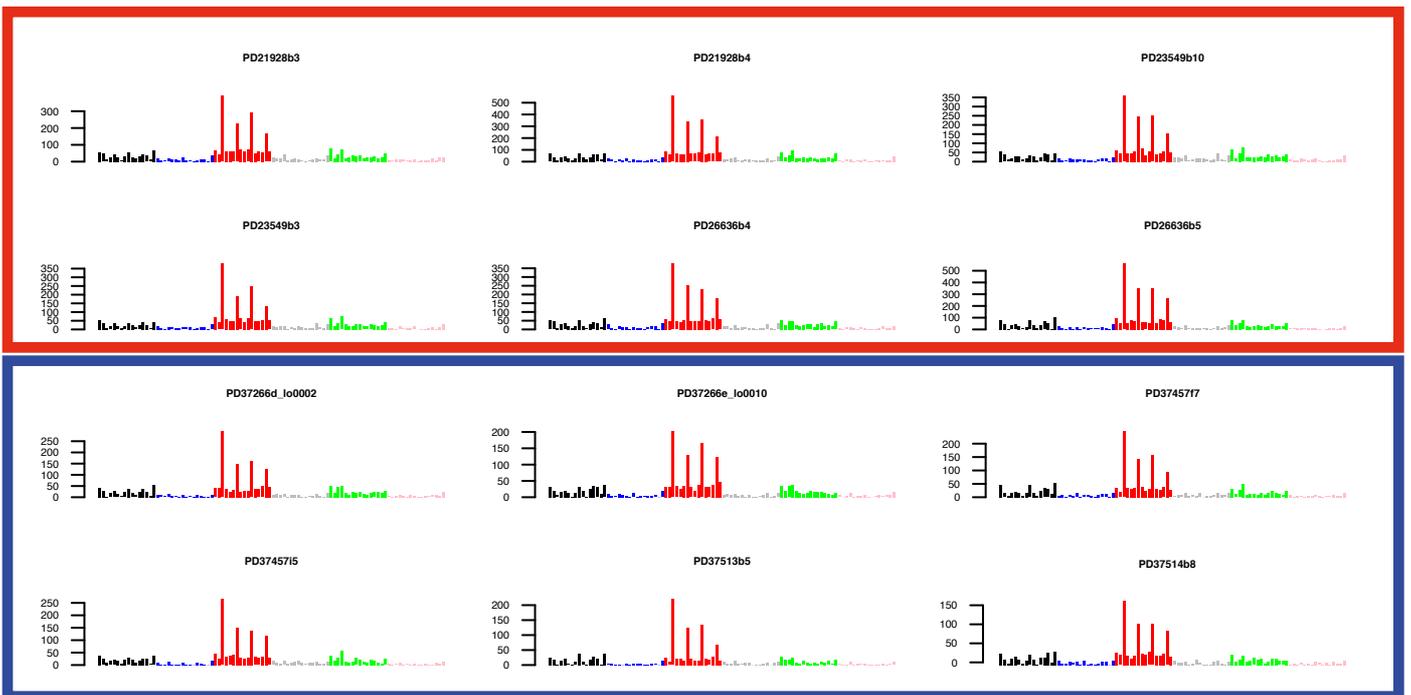
Figure 2.10. Mutation burden v. age for every signature. For each signature, the total number of mutations in every crypt attributed to each signature is plotted against the age of the patient from whom the crypt derives. Points are coloured by the anatomical location from of the crypt. Mutation burden is adjusted for the callable proportion of the genome (Methods section 12).

total   sbs.SBS1   sbs.SBS5   sbs.SBS18

sbs.SBS2   sbs.SBS13   sbs.N1   sbs.N2

sbs.N3   sbs.N4   dbs.DBS2   dbs.DBS4

dbs.DBS6   dbs.DBS9   dbs.DBS11   dbs.DBS8

id.ID1   id.ID2   id.ID5   id.N2

id.N3   svs   chromamps   loh

Age (years)   Age (years)   Age (years)   Age (years)

Mutations

- Right
- Transverse
- Left
- Ileum
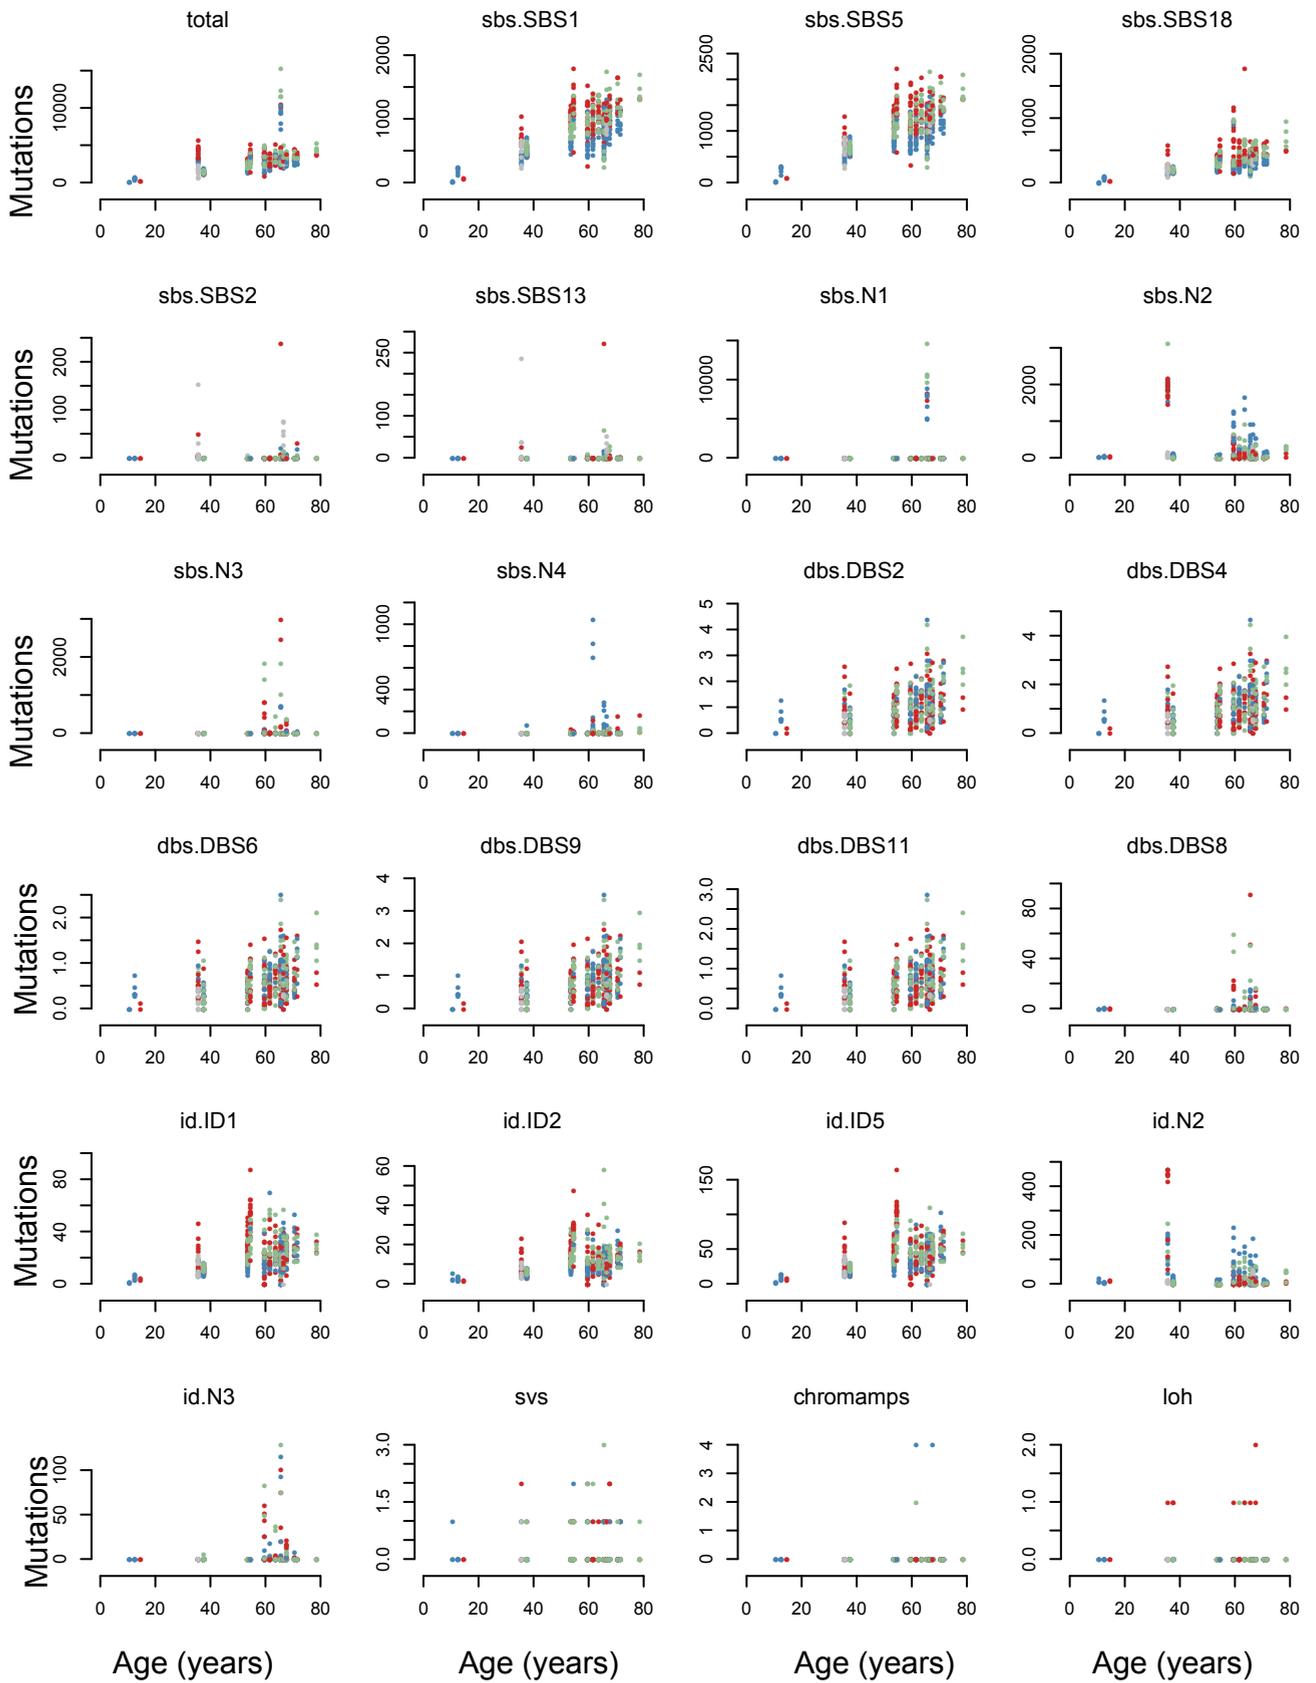
R.5.a. Common mutational signatures

All three of SBS1, SBS5, and SBS18 were found in all crypts from all patients and correlated linearly with their age, indicating that the underlying mutational processes are continuously active in all individuals throughout life. There was, however, substantial variation in mutation burdens for each of these signatures between crypts from the same individual. These were partly explained by differences between sites in the colon: for all three signatures, mutations accumulated at the fastest rate in right colon (ascending and caecum), then transverse colon, and slowest in left colon (caecum and sigmoid). Terminal ileum accumulated mutations at a similar rate to the left colon (figure 2.11). Even taking site differences into account, however, there was still substantial spread of mutation burden, and in the case of SBS18 one strong outlier. For example, the mean ratio of SBS1 burden between the most and least mutated crypts from the same site of one person is 1:1.3, and the mean ratio between the SBS1 mutation burden of crypts from the same site in two people in their 60s is 1:1.4. The cause of this variability is unknown, but possible explanations include differences in cell division rate, exposure to locally acting mutagens, and variability in the time to the most recent common ancestor of each crypt. Interestingly, bromodeoxyuridine staining of colonic crypts showed substantial inter- and intrapatient variability in cell division rates (Potten et al. 1992). For SBS1 and SBS5, which appear to accumulate in a clock-like manner, the x axis intercept is of approximately five to 10 years of age. With a linear mutation rate the x axis intercept would represent the time to the most recent common ancestor of the crypt, and this is similar to previous estimates of the duration of monoclonal conversion in humans (section I.1.c.).

DBS 2, 4, 6, 8, and 11 were all extracted by HDP as one mutational processes that was present in almost all crypts and deconvoluted into these signatures by expectation maximization. These processes must therefore be tightly correlated in normal colonic crypts. As numbers of doublet base substitutions are small, we are underpowered to detect differences between sites.

Similarly, ID1, ID2, and ID5 were all extracted as one ubiquitous signature by HDP, and deconvoluted into their constituents by expectation maximization. They correlate with the age of diagnosis of the patient, which is consistent with their proposed aetiology in replication slippage, and show the same ordering of mutation burden between sites as SBS1, SBS5, and SBS18 (figure 2.11).

166

Thus, the signatures that we found to be common have in some shape or form been described before. The size of the cohort and the low complexity of normal genomes relative to cancerous ones allow us to gain novel insights.

R.5.b. Rare mutational signatures

Many mutational processes, both known and novel, only affected subsets of crypts. SBS2 and SBS13 are known signatures predominantly characterised by C>T and C>G mutations respectively at T$\underline{C}$A and T$\underline{C}$T trinucleotides. They are thought to be due to activity of the APOBEC family of cytidine deaminases and often occur together (Alexandrov et al. 2013, Roberts et al. 2013). The C>T mutations may be a result of cytidine deamination, while the C>G and C>A (and possibly some C>T) mutations may be the errors of translesion polymerases following excision of uracils by repair machinery (Helleday et al. 2014). SBS2 and SBS13 were clearly observed in a colonic crypt from one individual and in an ileal crypt from another, each with >100 mutations of SBS2/13; smaller contributions may be present in other crypts. Therefore, APOBEC DNA-editing of the human genome occurs in gastrointestinal stem cells, to our knowledge the first time that it has been shown in normal cells (beyond the physiological role of AID, an APOBEC family member, in lymphocytes). The wider sequence context of the mutations suggests that APOBEC3A is the major contributing enzyme (Chan et al. 2015).

A novel signature (described as novel here, as it was unknown at the time that we found it, although it has since been observed in oral squamous cell carcinomas (Boot et al. 2018)), SBSN2, characterised predominantly by T>C mutations at A$\underline{T}$A, A$\underline{T}$T, and T$\underline{T}$T, and T>G mutations at T$\underline{T}$T was detected in 30% of crypts. In the most affected crypts, it accounted for 3,000 mutations, doubling the mutation burden. SBSN2 exhibits strong transcription strand bias, with 2.5 times as many T>A mutations occurring on the untranscribed as on the transcribed strand. Transcription strand bias is often due to transcription coupled nucleotide excision repair (TC-NER) acting on DNA with covalently bound bulky adducts and distortion of the helical structure. Assuming that this is the case, damage to adenine by a carcinogen may underlie SBSN2. SBSN2 exhibited a highly variable mutational burden between individuals and between crypts from the same individual that was not attributable to age. Examination of the branches of phylogenies in which

Figure 2.11. Linear modelling of signature accumulation. For signatures that appeared to show a linear accumulation with age, the mutation rate per site was determined using mixed models, with age and site as fixed effects, and patient as a random effect. Confidence intervals were determined by bootstrapping. Mutation burden is adjusted for the callable proportion of the genome (Methods section 12).

sbs.SBS1

Ileum: 12.7 (CI95 10.5–14.7)
Right: 16.8 (CI95 15.4–18.3)
Transverse: 16 (CI95 14.7–17.6)
Left: 12.7 (CI95 11.3–14.4)

sbs.SBS5

Ileum: 16.4 (CI95 13.7–18.9)
Right: 21.3 (CI95 19.3–23.2)
Transverse: 20.4 (CI95 18.3–22.2)
Left: 16.4 (CI95 14.3–18.3)

sbs.SBS18

Ileum: 5.13 (CI95 3.77–6.56)
Right: 7.48 (CI95 6.37–8.54)
Transverse: 6.99 (CI95 5.9–8.03)
Left: 5.38 (CI95 4.29–6.49)

id.ID1

Ileum: 0.218 (CI95 0.105–0.344)
Right: 0.378 (CI95 0.293–0.467)
Transverse: 0.372 (CI95 0.287–0.46)
Left: 0.31 (CI95 0.223–0.405)

id.ID2

Ileum: 0.0885 (CI95 0.0123–0.162)
Right: 0.171 (CI95 0.119–0.222)
Transverse: 0.18 (CI95 0.127–0.231)
Left: 0.124 (CI95 0.0732–0.18)

id.ID5

Ileum: 0.402 (CI95 0.203–0.615)
Right: 0.675 (CI95 0.523–0.834)
Transverse: 0.68 (CI95 0.525–0.84)
Left: 0.547 (CI95 0.396–0.704)

SBSN2 occurred showed that whenever SBSN2 could be timed, it always occurred early in life (figure 2.12 f, h, z, aa, am, ao). Assuming a constant mutation rate of SBS1 (as indicated by the age correlation observed), SBSN2 occurs in the first ten years of life. It cannot be the result of one brief mutagenic burst, since it is found in sequential branches (figure 2.12ao). In addition, SBSN2 clustered spatially. In a thirty-six year-old patient, crypts from all around the colon were affected, but not crypts from the ileum. In other patients, crypts from the left colon were much more affected than crypts from other parts of the colon. This was not due simply to the sharing of mutations between crypts. SBSN2 correlates with IDN2, another novel signature characterised by deletion of single T in a run of Ts with a mode length of four (compare ochre bars in ID trees and pink bars in SBS trees in figure 2.12). The initiating events for this relatively common mutational process are unknown, but our data indicate an extrinsic, locally-acting mutagen that affects children. Many causes are possible, including diet, infections, and microbiome composition.

SBSN3 was predominantly characterised by C>T substitutions at A$\underline{C}$A, T>A at C$\underline{T}$N, and T>G at G$\underline{T}$G. It was present in five individuals and in these in a subset of crypts (figure 2.12 e, aa, af, ai, aj). Like SBSN2, SBSN3 was active early in life in the two patients in whom we could time it (figure 2.12 aa, ai), and even when mutations were not shared we could detect spatial clustering. The cause of SBSN3 is unknown, but here again we have evidence of an early, locally-acting process. SBSN3 correlated with DBS8 and IDN3 (compare dark grey in SBS trees, dark grey in DBS trees, and green in ID trees). DBS8 is composed of AC > CA and AC > CT mutations, thus representing some of the same set of base changes as SBSN3. DBS8 has been reported in rare hypermutated cancers with no obvious cause for their hypermutation (Alexandrov et al. 2018). IDN3 is dominated by the deletion of a single T with no other Ts surrounding it.

All crypts from the left, right, and transverse colon of a 66 year-old man were dominated by a mean of 8,567 mutations due to a new signature, SBSN1. This signature is characterised by T>A substitutions with a transcriptional strand bias that is again consistent with damage to adenines. The mutation burden in his colorectal epithelium was three- to five-fold higher than expected for his age, and thus by extrapolation equivalent to that of a 200-300 year-old. This individual had a rich and unusual clinical history: initially diagnosed with a large anaplastic lymphoma in 1994 and treated with six cycles of CHOP (cyclophosphamide, doxorubicin, vincristine, prednisolone), the diagnosis was subsequently revised to Hodgkin's lymphoma, and in 2002 he was treated with three cycles of Chl-VPP / PABlOE (chlorambucil, vinblastine,

procarbazine and prednisolone, alternating with prednisolone, doxorubicin, bleomycin, vincristine and etoposide). In 2011 he had a positive faecal occult blood test, for which he had a colonoscopy (when the biopsies that we used were taken) and a caecal adenocarcinoma was found. Two years later this gentleman died from a diffuse large B cell lymphoma.

SBSN1 bears a strong resemblance to SBS25 (cosine similarity 0.9), and they share the same transcription strand bias. Signature 25 has previously been detected only in two lymphoma cell lines derived from lymphomas from patients who had received chemotherapy. The case history of one of these patients is available (Wolf et al. 1996): he too had been treated with a cocktail of chemotherapies, with overlap of several drugs (cyclophosphamide, doxorubicin, vincristine/vinblastine, prednisolone, procarbazine, bleomycin). He had also been treated with an experimental ricin-coupled anti-CD25 immunotoxin. Signature 25 was not detected in any of 23,829 cancers in PCAWG and so represents a very rare mutational process. The vast majority of PCAWG cancers were primaries that had not been exposed to chemotherapy, and so it seems likely that chemotherapy is responsible, either through the action of one cytotoxic agent, a combination of them, or their interaction with a germline polymorphism. Cyclophosphamide has been observed to cause T>A mutations in the chicken DT40 cell line (Szilkrist et al. 2016); further work testing different chemotherapies on human cells would be required to identify the causal agent. To our knowledge, this is the first report of the mutagenic consequence of chemotherapy in normal human cells *in vivo*.

To determine whether this process could have played a part in the development of this gentleman's adenocarcinoma, we obtained the biopsy of his tumour taken during the same colonoscopy as the normal samples. The tumour was necrotic and individual crypts could not be distinguished. We therefore bulk-sequenced it. This complicates comparisons of mutation burden between the tumour and normal sample, since the time to the most recent common ancestor of the tumour and the normal samples need not be the same. The tumour was not closely related to any of the individual crypts. As with most tumours, it had an excess of mutational processes not seen in the normal crypts, including single base substitutions, small insertions and deletions, and copy number changes. The T>A mutational process was present in the clonal peak of mutations, but not subclonally, indicating that exposure to chemotherapy predated the last clonal sweep of the tumour; indeed, it is quite plausible that it could have occurred while the tumour was still a normal crypt. After adjusting for copy number changes, a similar number of T>A mutations were present

Figure 2.12. Crypt phylogenies. For every patient, the phylogeny of crypts is shown three times: on top, with branch lengths proportional to the number of single base substitutions; in the middle, with branch lengths proportional to the number of doublet base substitutions; on the bottom, with branch lengths proportional to the number of small insertions and deletions. Scale bars are shown on the right-hand side. A stacked barplot of the mutational signatures that contribute to each branch is overlaid over every branch. Please note that the ordering of signatures along a given branch is just for visualisation purposes: we cannot distinguish the timing of different signatures along a branch. In most cases, crypts from the same individual are distant from one another, and so we would not expect to see late branching events.

PIK3CA
E542K

Mutations

s t u v w x y

Mutations

Mutations

z aa ab ac ad ae af ag ah ai aj

Mutations

Mutations

Insuffi-
cient
cover-
age

Mutations

AXIN2 loss
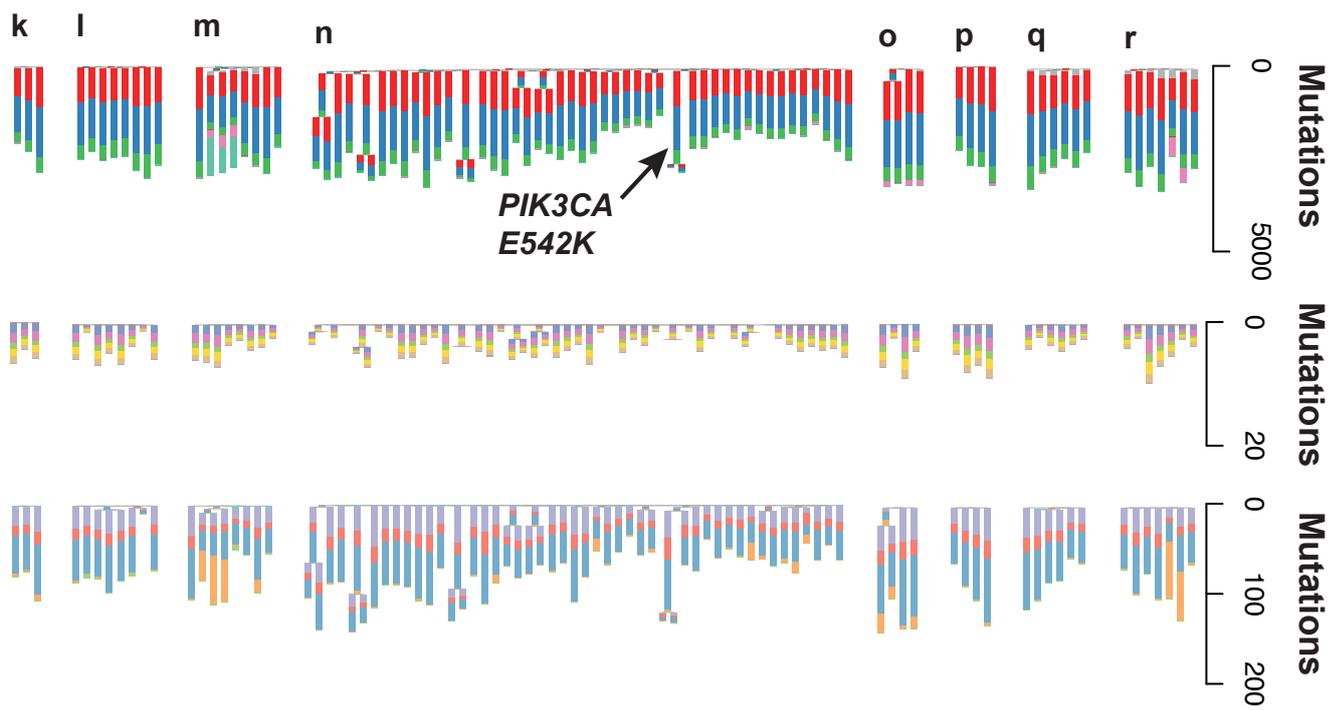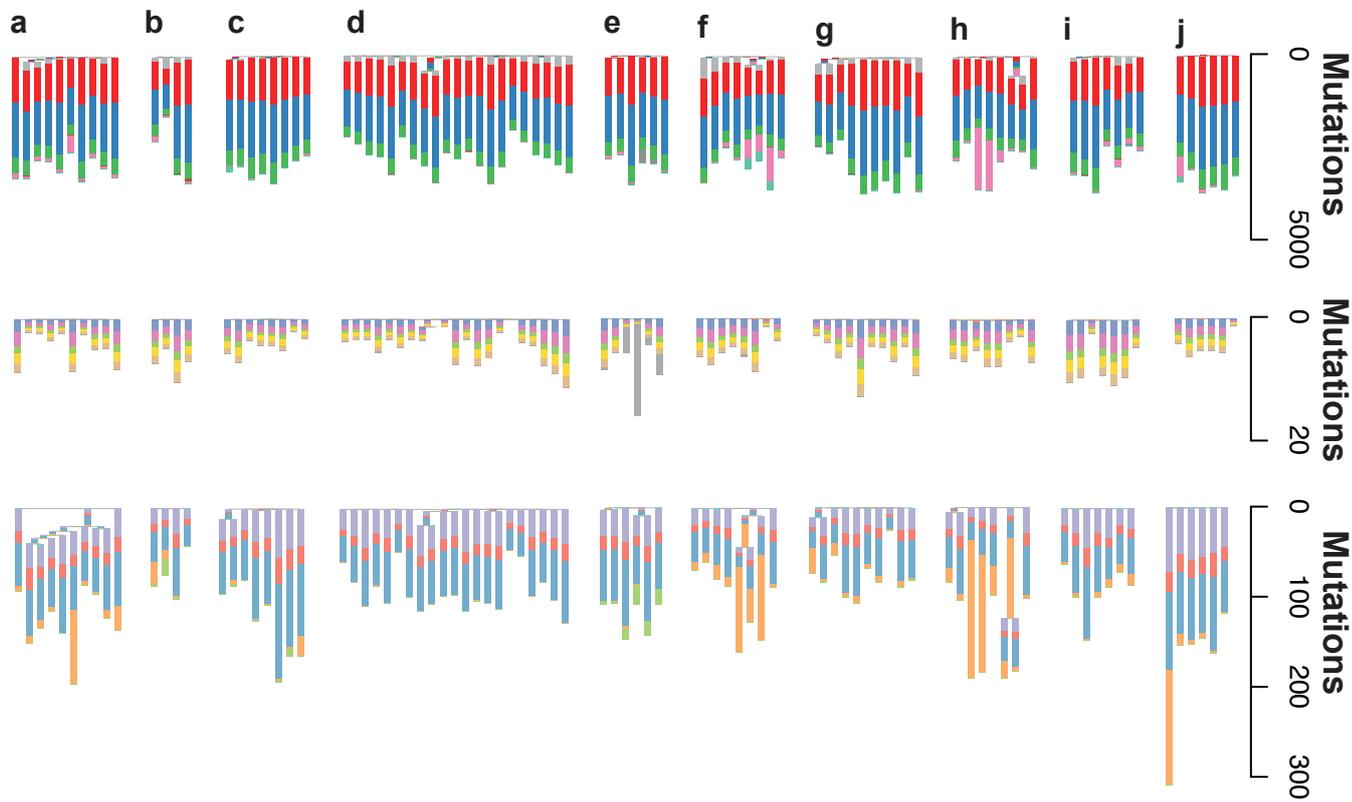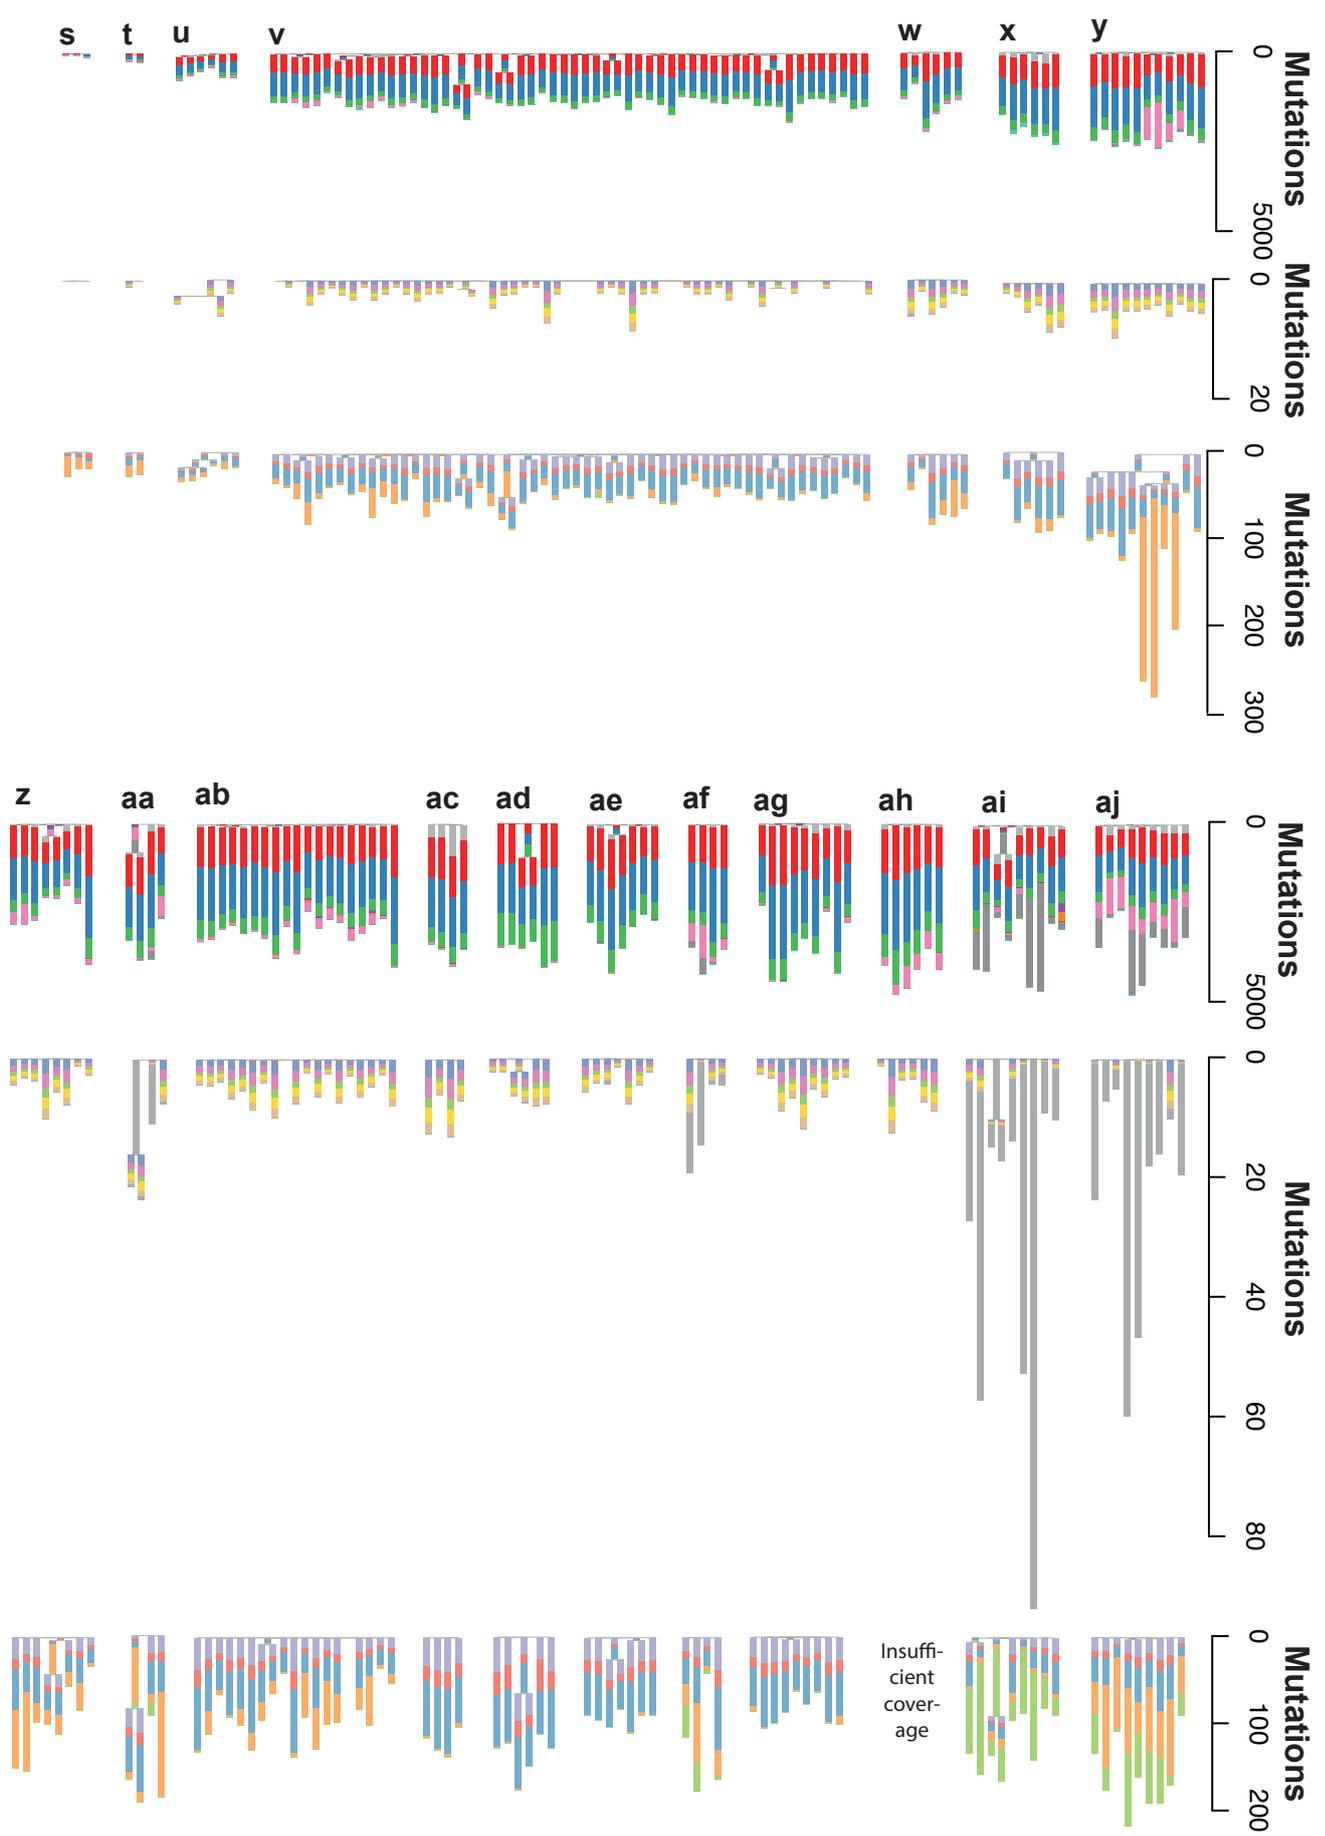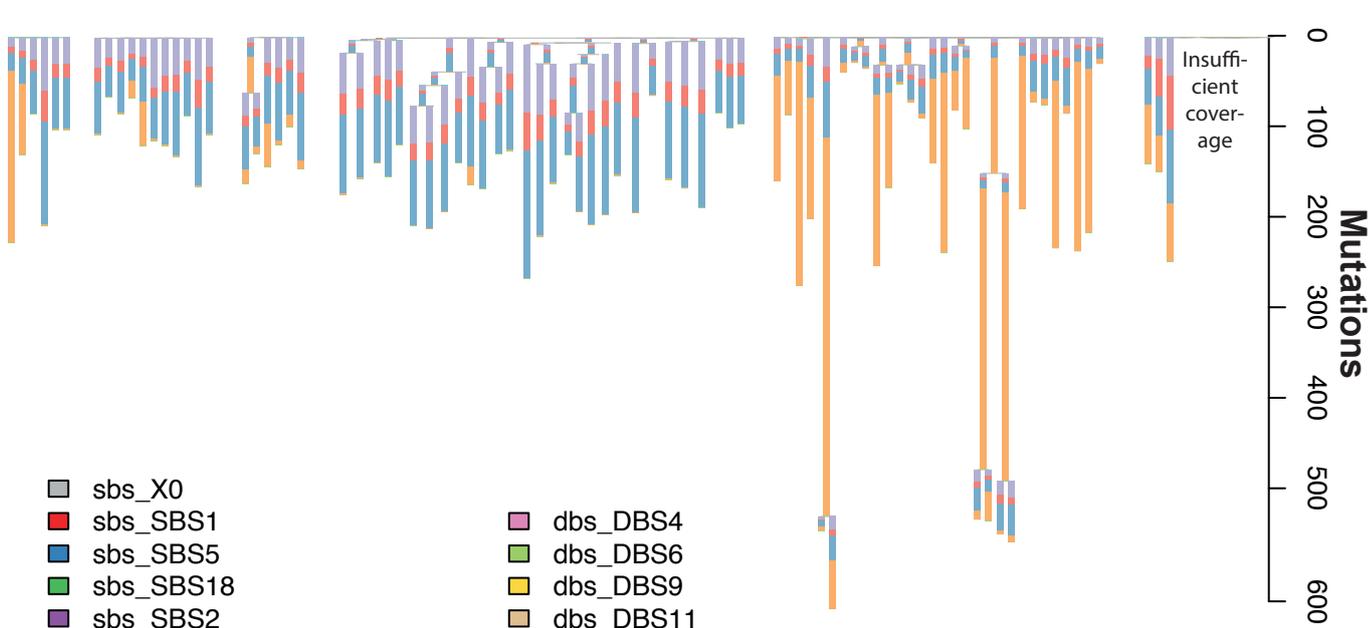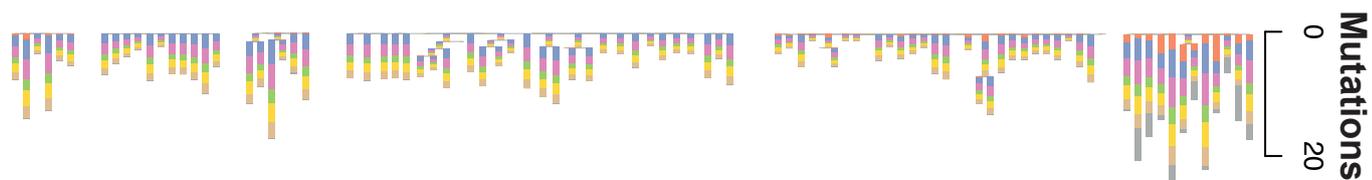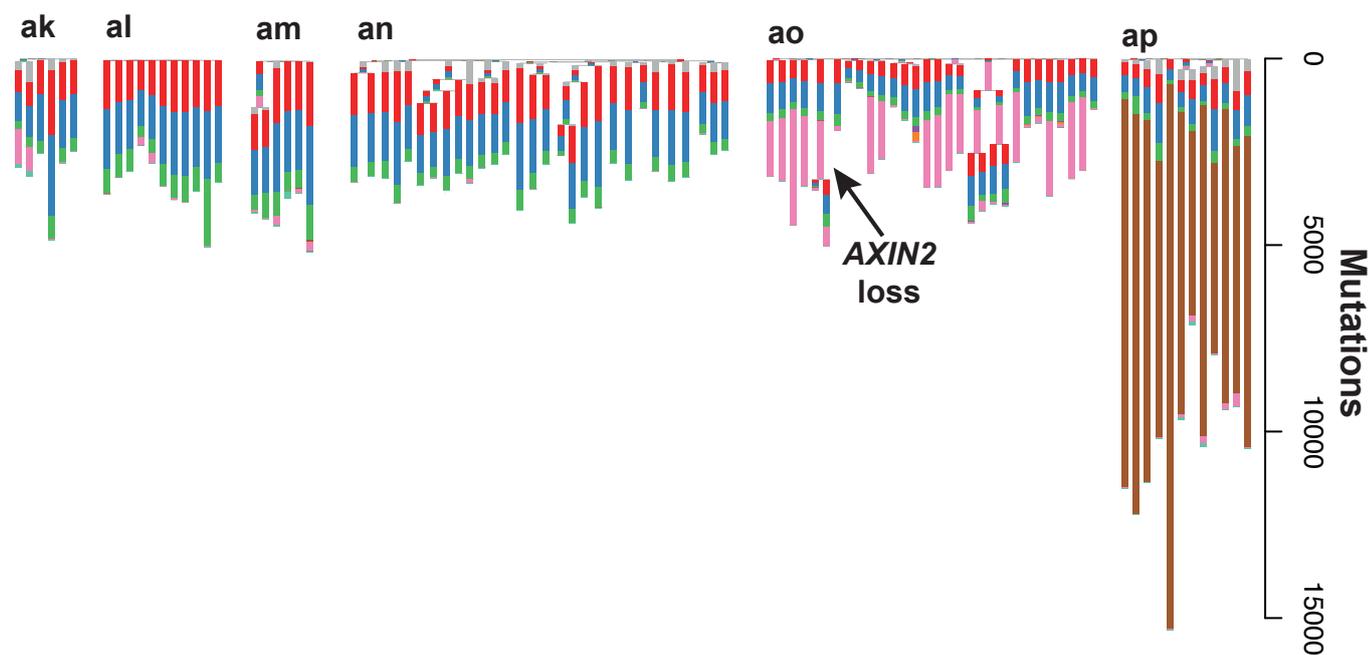
Legend:
- sbs_X0
- sbs_SBS1
- sbs_SBS5
- sbs_SBS18
- sbs_SBS2
- sbs_SBS13
- sbs_N1
- sbs_N2
- sbs_N3
- sbs_N4
- dbs_X0
- dbs_DBS2
- dbs_DBS4
- dbs_DBS6
- dbs_DBS9
- dbs_DBS11
- dbs_DBS8
- id_X0
- id_ID1
- id_ID2
- id_ID5
- id_N2
- id_N3

Insufficient coverage

in the trunk of the tumour to the normal crypts. The tumour was driven by a clonal *KRAS* G12D hotspot substitution and two inactivating mutations (one substitution and one indel) in *APC*. These were not due to T>A mutations, and so the mutagenic activity of this chemotherapy cannot be held directly responsible for the development of the tumour. It is possible, however, that the selection pressure of chemotherapy could have provided the conditions for a pre-malignant clone to expand and progress.

R.5.c. Large scale genomic rearrangements in normal colon

Colorectal cancers frequently bear a large number of genomic rearrangements (Li et al. 2017), and rearrangements were found in four out of 15 organoids (Blokzijl et al. 2016). In our larger cohort, we sought evidence of large copy number changes in 449 crypts that had >10X coverage and >0.3 VAF sufficient for us to call copy number changes accurately. In stark contrast to colorectal cancers, 82% of evaluable crypts had no large-scale genomic rearrangements. Remarkably, however, five crypts had seven whole chromosome copy number increases affecting the same three chromosomes. We observed: amplification of both copies of chromosome 3; trisomy 3 and trisomy 9; trisomy 7; amplification of both copies of both chromosome 7 and of the X chromosome; and amplifications of both copies of both chromosome 7 and chromosome 9 (figure 2.13a). We also observed an amplification of the X chromosome. All amplifications increased the copy number of the chromosome by one or two copies. We did not observe any chromosomal deletions. While regions of chromosome 7 are often amplified in colorectal cancer (Cancer Genome Atlas Network et al. 2012), we are not aware of frequent amplifications of chromosomes 3 and 9. In addition, we found large regions of copy neutral loss of heterozygosity in 12 crypts, affecting 1p, 6p, 7p, 8q, 9q, 10q (twice), 17p, 17q, 18q, 21q, 22q, and the X chromosome (e.g. figure 2.3c).

Five of these copy number changes could be timed reliably by using their allele fractions to assign mutations to a copy number state. The count of mutations at each copy state can be used to estimate when the copy number change occurred. This is because mutations that occur before a chromosomal gain will be on two copies, whereas those that occurred after it will be on one copy. Timing these seven changes showed that they all occurred between the ages of 20 and 51 (figure

2.13b). Two gains in the same crypt were timed independently and found to occur at the same time, suggesting that they occurred as one event. This analysis requires the assumption of a constant mutation rate per length of DNA over life, which seems reasonable given that the signatures that are responsible for the majority of mutations in normal colorectal genomes correlate linearly with the age of the patients from whom they are derived. Trinucleotide plots of the mutations assigned to each copy number state were inspected and showed no dramatic change in mutational profile.

Structural variants were detected by abnormally mapping reads. We observed 48 deletions, 18 tandem duplications, four translocations, and two inversions. All were private to a single crypt, except for one which was shared between two adjacent crypts from one patient, which share a very distant common ancestor: this must reflect a deletion that occurred in the embryo or in early childhood.

R.6. Comparison of mutational signatures in our cohort and PCAWG colorectal cancers

The total mutation burden in most crypts that we sequenced is of the order of 3,000 mutations per genome, substantially less than the mutation burden of 10,000-20,000 mutations per genome in even non-hypermutated colorectal adenocarcinomas. We sought to explain the source of the mutational excess in cancers by comparing the number of mutations due to each mutational signature in PCAWG colorectal adenocarcinomas and our normal tissues (figure 2.14). A number of complicating factors should be borne in mind: tumours and normal crypts have a different time to their MRCA; patients were of different ages, although in both cohorts most samples come from patients in their 50s and 60s; they were sequenced on different platforms and mutations called with different filters; and signatures were extracted separately (although the same set of reference signatures was used).

As anticipated, the total mutation burden in colorectal cancers is always higher than in normal crypts, with the exception of those from the patient who had been exposed to chemotherapy. The mutation burden of the near-ubiquitous signatures SBS1, SBS5, SBS18, DBS2, DBS4, DBS6, DBS9, ID1, and ID2 was higher in tumour samples than in normal samples, suggesting an acceleration of normal mutational processes. Some rare mutational processes were

Figure 2.13. Copy number changes in normal colon. a, whole chromosome amplifications in five crypts. The copy number state (y axis) for each allele, one coloured red, and one coloured green, is shown. Chromosomes are labelled along the top of the graph. **b,** timing of copy number changes throughout life. Vertical bars represent 95% confidence intervals determined by bootstrapping.

a

b
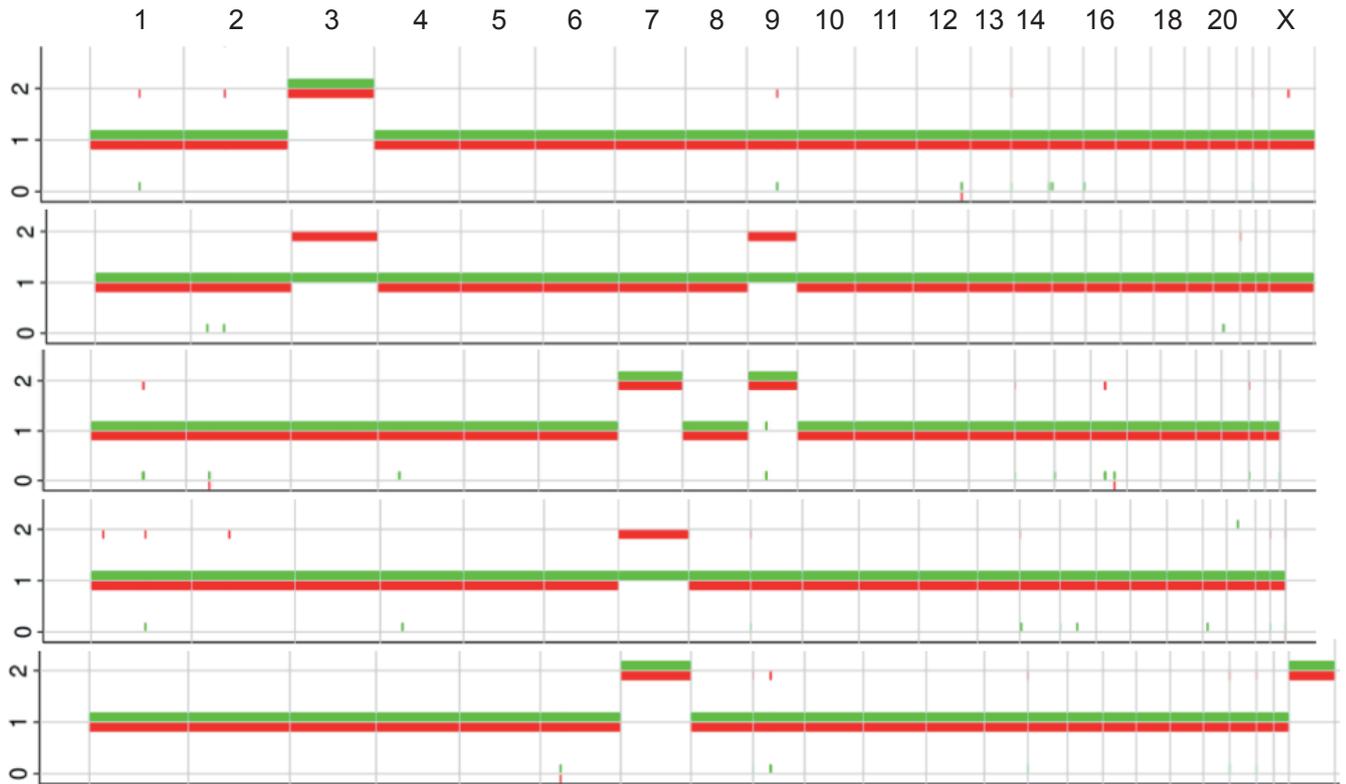
Timing of copy number changes
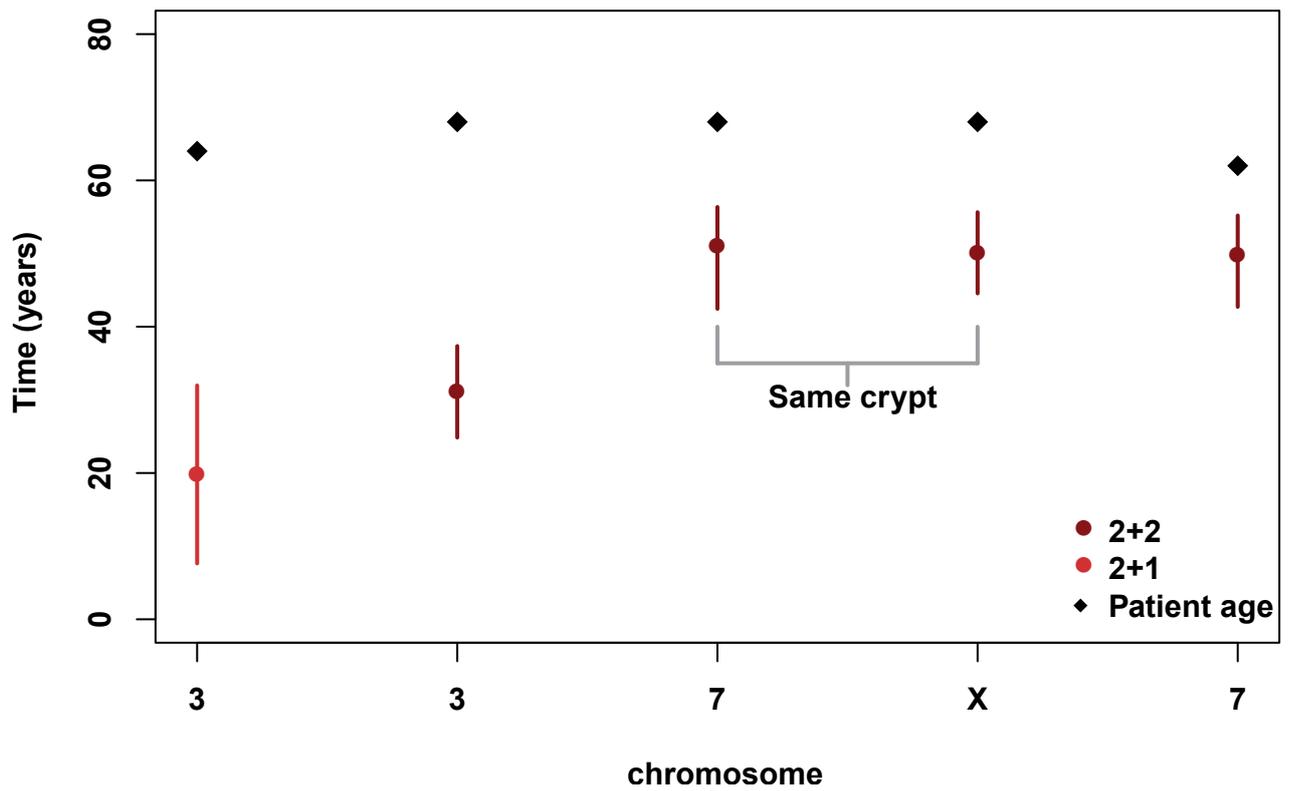
Figure 2.14. Mutation burden of each signature in normal crypts versus cancer. For every signature, the (mutation burden+1) of every sample is shown on the y axis on a log scale. Normal colon and cancer samples are ordered within their groups. Colorectal adenocarcinoma signature attributions and burden are from Alexandrov et al. (2018).
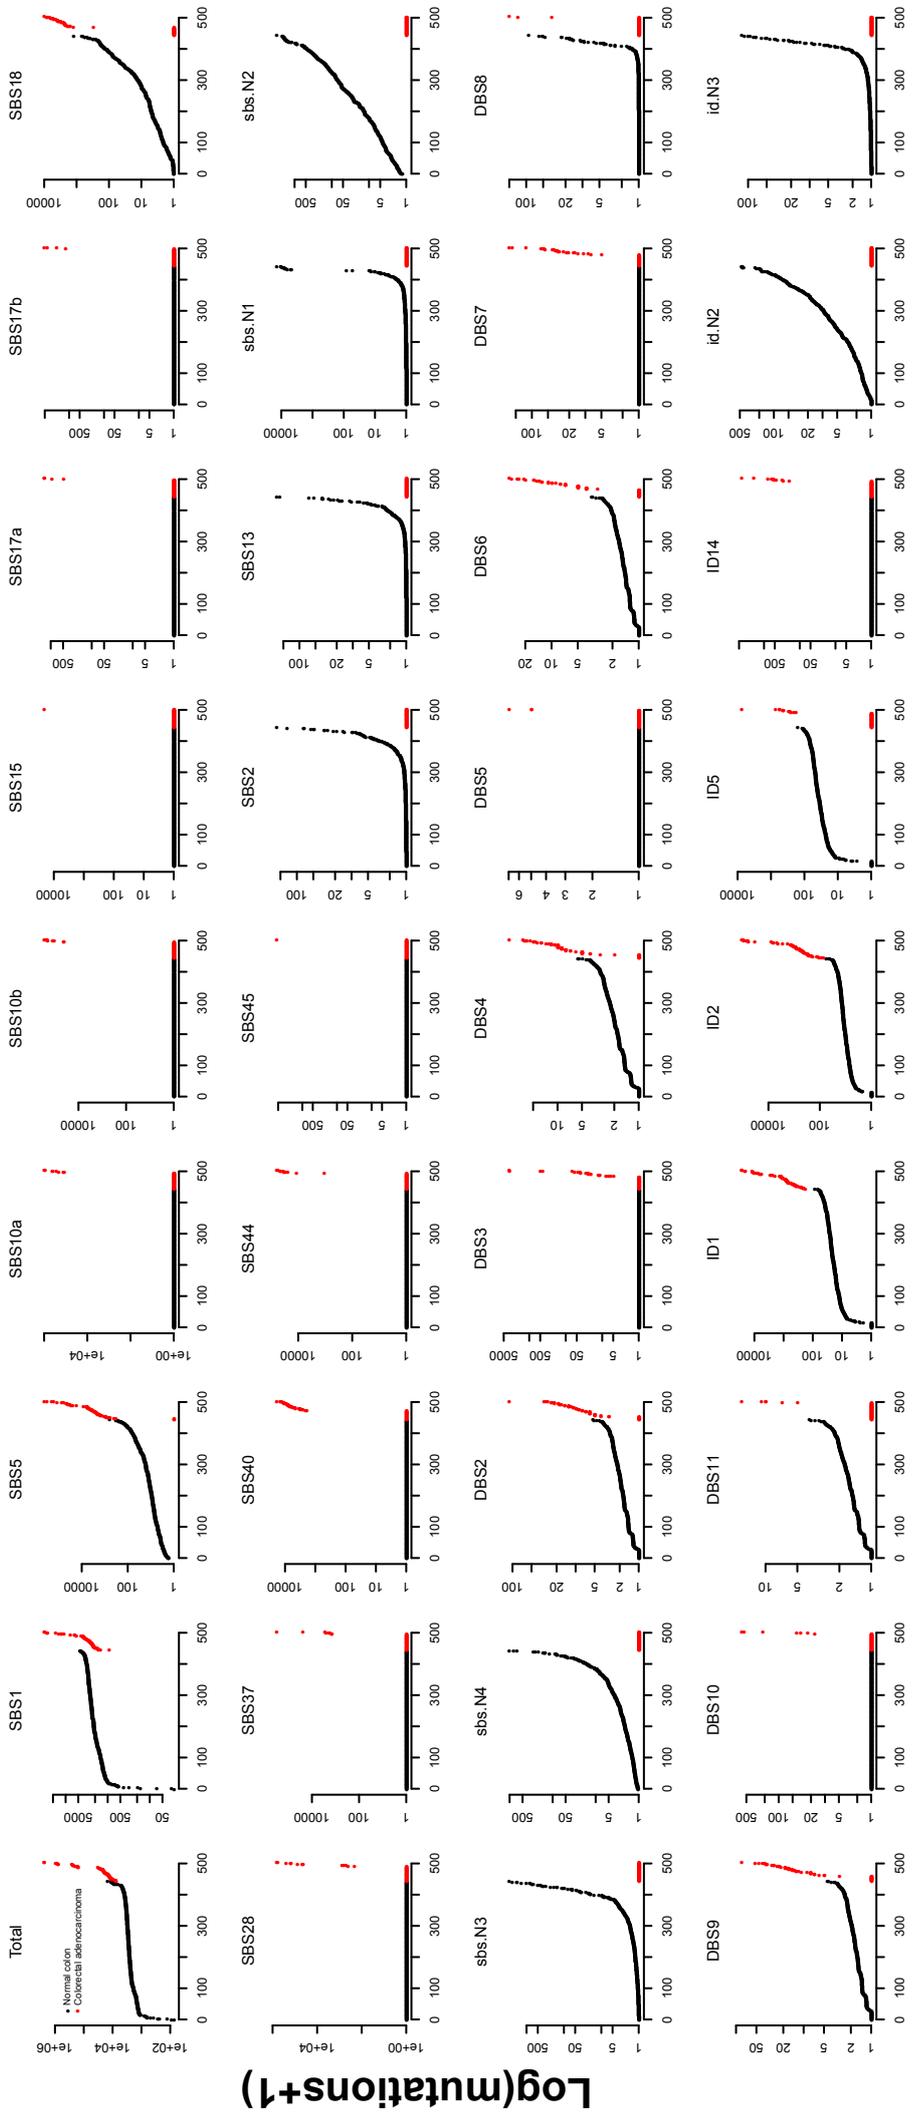
Ranked sample

Log(mutations+1)

only found in tumours, such as SBS10a or SBS10b and SBS17. Perhaps if a larger number of normal crypts had been sequenced these mutational processes would eventually have been found in normal tissues too; nonetheless, sporadic mutational processes do seem to be enriched in cancers relative to normal tissue. Conversely, all the novel signatures were (by definition) found only in normal tissues. No mutational processes were found in all tumours and only in tumours. Thus, there is no process that is specific to and intrinsically linked with malignancy.

R.7. Mutational processes and rates in the progression from normal to cancer

The additional mutational processes detected in cancers could have occurred during the process of transformation or have preceded it: it could be that rare colonic cells with a naturally higher mutation burden are more likely to become cancerous. Furthermore, as explained above, the comparison of bulk tumour samples with normal colonic crypts is complicated by the difference in the time to the most recent common ancestor of each sample. We reasoned that the former difficulty could be resolved by reconstructing the mutational life history of the tumour through multi-region sequencing, and the latter by sequencing single cells derived from tumour and normal epithelium. We set out to achieve this by analysing organoids derived from three patients undergoing a resection for colorectal cancer. For each patient, organoids were derived from single cells from a tumour and from individual normal crypts five centimetres away from the tumour. From the first patient we whole genome sequenced four normal and seven tumour organoids, from the second five normal and 11 tumour organoids, and from the third four normal and eight tumour organoids. This study was set up by Sam Behjati and Sophie Roerink and organoids were generated by the Clevers group at the Hubrecht Institute.

R.7.a. Comparison of mutational processes in single cells from cancer and normal colon

For all three patients, all tumour organoids had many more mutations than any of the normal organoids. Subcloning of organoids showed that the mutation burden acquired *in vitro* was

minimal, as had been shown by others (Blokzijl et al. 2016). Patient 1 had a mismatch repair deficient tumour. Methylation profiling (analysed by Matthew Young) showed that this was due to hypermethylation of the *MLH1* promoter. This patient had a mean of 3,792 substitutions in normal clones but 72,398 substitutions in tumour clones. Patients 2 and 3, who had mismatch repair proficient tumours, had means of 3,172 and 3,621 substitutions in their normal organoids, and of 22,291 and 14,209 substitutions, respectively, in their tumour-derived organoids (figure 2.15a). Similarly, the mean burden of indels in normal organoids for patients 1-3 was 227, 130, and 167, while the mean indel burden in their tumour organoids was 27,893, 1,485, and 2,021. There was a mean of one structural variant in normal organoids, compared to a mean of 71, 176, and 67 from the tumours of patients 1-3. Thus, as one would expect from the comparison of mutation burden between normal colonic crypts and cancer genomes, the mutation burden is increased in cancer relative to normal tissue even within the same patient. Tumour organoids were derived from single cells, while normal organoids were derived from single crypts. More than one stem cell from a crypt might have contributed to the organoid, and so the most recent common ancestor of normal organoids may predate the resection, while that of tumour organoids will not. Nonetheless, the difference is likely to be less than a decade and cannot explain the discrepancy in mutation burden between tumour and normal organoids.

To investigate the origin of the increased burden of mutations in cancer relative to normal, signatures of mutational processes were extracted for substitutions by Ludmil Alexandrov using NNMF. Signatures are referred to here by their numbers in COSMIC (Forbes et al. 2017).[4] All the mutations in normal clones were attributed to signatures 1 and 5. Additional mutational signatures were found in the tumours (figure 2.15). These included: signatures associated with mismatch repair deficiency in patient 1 (signatures 6, 20, and 26); signature 17 in patients 1 and 2, which is of unknown cause and is found in a minority of colorectal cancers; signature 18 in patients 2 and 3, which is thought to be associated with reactive oxygen species; and a novel signature (figure 2.15b), found in only six out of eight cancer-derived organoids in patient 3. 100 mutations that were likely to be due to this novel signature were validated by capillary sequencing (data not shown).

---

[4] Please note that signature profiles here reflect the TCGA and cosmic versions, which means the number of mutations attributed to each signature cannot be compared directly to analyses that used the PCAWG catalogue.

Figure 2.15. Mutation burden of colonic organoids. a, barplots of the mutation burden in every organoid. For substitutions, mutations are broken down by signature, according to the COSMIC catalogue (NB not the PCAWG catalogue), whereas indels are classified as insertion or deletion, and structural variants as inversion, deletion, tandem-duplication, or translocation. For each patient the burden in normal and tumour organoids are shown. **b**, the trinucleotide spectrum of the novel signature (yellow in **a**).

While the majority of the discrepancy between tumour and normal could be attributed to additional mutational signatures, tumour cells also had an increased burden of signatures 1 and 5 relative to the normal cells. Given signature 1's proposed correlation with mitoses, this suggests that cancer cells have undergone more cell divisions. Indeed, increased rates of Ki67 staining, a marker of proliferation, are observed in the resected cancer specimens (staining performed by collaborators; data not shown). However, the increase in signature 1 is not sufficient to explain the discrepancy in the number of mutations between tumour and normal. The fact that cancer cells have additional mutational signatures that are not seen in normal cells shows that additional mutational processes are operative, which must mean an increase in mutation rate per cell division. Thus, cancer cells, even in MMR proficient patients, show an increase in mutation rate per cell division. The cause of this increased mutation rate per cell division warrants further investigation. Some cancer-specific signatures, as in mismatch-repair deficiency, may be a result from the loss of a repair process that occurs in normal cells. Others may be the result of new mutagenic exposures. For example, as signature 18 has been associated with oxidative damage (Viel et al. 2017), its acceleration in all three tumours relative to normal tempts one to speculate that it might be a result of the change in metabolism associated with tumour growth (Warburg et al. 1927).

R.7.b. Reconstructing the mutational histories of tumours with phylogenies

Phylogenies were constructed for the samples derived from each patient (Methods), and mutations were assigned to each branch. Signatures were extracted treating every branch as an independent sample (Methods), which allowed us to see the mutational processes operative at different stages of tumour evolution (figure 2.16). For patient 2, in the trunk of whose tumour a whole genome duplication was present, mutations in the trunk were further separated according to whether they occurred before or after the whole genome duplication (Methods).

In all patients, the trunk of the tumour contained a greater proportion of signatures 1 and 5 than later branches, suggesting that the additional mutational signatures in cancer cells were not present through their whole life history. This was clearest in patient 2. While signature 1 accounted for approximately 60% of the mutations before the whole genome duplication, afterwards it only accounted for 30%, with signature 17 becoming much more prominent. In patient 3, a subclonal

signature was observed (yellow in figure 2.16c), indicating ongoing diversification of mutational processes through tumour evolution. This could be a result of a heritable change (such as loss of another repair protein) or may reflect the particular microenvironment of different tumour regions, since the phylogeny covaries with the spatial structure of the tumour. Likewise, rearrangements and small indels increase markedly in rate beyond the trunk of the tumour (figure 2.17 and 2.18). In patient 3, the rate of rearrangement acquisition can be seen to vary subclonally too, as the clade of the tree that had acquired *TP53* mutations had far more rearrangements than the clade that had not.

The fact that the genomic landscape of the trunk of the tumour phylogenies resembles that of normal colorectal stem cells indicates that there was nothing particularly special about the mutational processes of the cell-of-origin of cancers. It suggests that cancers derive from cells that were exposed to normal mutational processes for much of life until an event, such as one or more driver mutations, caused a change in the active mutational processes. In patient 1, the trunk of the tumour contains almost no signatures of mismatch repair deficiency-associated processes (although there was an increase in indels), despite the fact that all organoids have hypermethylated the *MLH1* promoter, and so this event presumably did occur in the trunk of the tumour. This suggests that the acquisition of aberrant mutational processes was rapidly followed by intratumoral growth and diversification. It could be that both are a result of the *BRAF* V600E mutation, given that Ras pathway mutations are associated with growth (see introduction to this chapter) and that this particular mutation frequently coincides with mismatch repair deficiency (Fearon 2011).

To understand the transition from normal to cancer, we attempted to time the onset of the cancer-specific mutational signatures (Methods). This analysis requires several assumptions, as we can only time mutations relative to signature 1 and cannot relate them directly to real time. As there is more signature 1 in tumour samples than in normal samples, it cannot be used as a real-time clock in the tumour, although we use it in this way in the normal organoids. To estimate the onset of cancer-specific mutational signatures, we assumed that the rate of signature 1 started increasing at approximately the same time as the onset of new mutational signatures. This may be valid if both were associated with the acquisition of driver mutations. In all three patients, signature 18 seemed to be the first cancer-specific mutational signature to become operative. Using the ratio of signature 1 to signature 18 in the next branch after our first timepoint (branches after the most recent common ancestor (MRCA) in patients 1 and 3, and the branch between the WGD and the

Figure 2.16. Phylogenies of colonic organoids. For each patient, the anatomical location in the colon of the tumour is shown, with the phylogeny underneath. White circles represent organoids derived from normal tissue, while filled coloured circles represent organoids derived from tumour, with the colour matching the colour of the biopsy site in the schematic above the phylogeny. Phylogenies are depicted as in figure 2.12 with branch lengths proportional to numbers of mutations, and signature contributions overlaid as stacked barplots.

a. Patient 1

b. Patient 2

c. Patient 3

1000 mutations

homozyg ACVR2A p.K437fs*5
PIK3CA p.E81K
BRAF p.V600E
MLH-1 hypermethylation

APC p.K311fs*25
APC p.S1321fs*3
homozyg. TP53 essential splice site

WGD

KRAS p. A146T
APC p.H427fs*1
APC p.L1488fs*41

homozyg. TP53
p.N235_Y236delNY

PTEN p.W111*
PTEN p.R335*

RNF43 p.G659fs*41
RNF43 p.R225fs*194

1000 mutations

signature 1
signature 5
signature 17
signature 18
signature 6
signature 20
signature 26
new signature
indels
rearrangements

Figure 2.17. Phylogenies of colonic organoids showing indel burden. Phylogenies are shown as in figure 2.16, but showing indels only as figure 2.16 was dominated by substitutions.

P1

homozyg. ACVR2A p.K437fs*5
PIK3CA p.E81K
BRAF p.V600E
MLH-1 hypermethylation

4000 mutations

PTEN p.W111*
PTEN p.R335*

RNF43 p.G659fs*41

RNF43 p.R225fs*194

N.1 N.2 N.3 N.4

T1.1 T1.4 T3.3 T3.4 T4.3 T4.4 T2.4

P2

APC p.K311fs*25
APC p.S1321fs*3
homozyg.TP53 essential splice site

WGD

N.1 N.2 N.3 N.4

N.2_sc

T1.1 T1.3 T6.2 T6.6 T2.5 T5.1 T5.4 T4.2 T4.3

T6.2 subclor

P3

KRAS p. A146T
APC p.H427fs*1
APC p.L1488fs*41

homozyg.TP53 ess. splice

N.1 N.2 N.3 N.4

T1.1 T1.4 T2.1 T2.2 T3.4 T3.2 T4.1 T4.2

1000 mutations   insertion
                      deletion

Figure 2.18. Phylogenies of colonic organoids showing rearrangement burden. Phylogenies are shown as in figure 2.16, but showing rearrangements only as figure 2.16 was dominated by substitutions.

P1

N.1 N.2 N.3
N.4

homozyg.AVCR2A p.K437fs*5
PIK3CA p.E81K
BRAF p.V600E
MLH-1 hypermethylation

PTEN p.W111*
PTEN p.R335*

RNF43 p.G659fs*41
RNF43 p.R225fs*194

T2.4

T3.3

T3.4

T1.4

T4.3   T4.4

T1.1

P2

N.1
N.2
N.3
N.4

N.2 subclone

APC p.K311fs*25
APC p.S1321fs*3
homozyg.TP53 ess splice site

T6.2

T6.2 subclone

T6.6

T2.5

T1.1

T5.1

T5.4

T4.2

T4.3

T1.3

P3

N.1 N.2 N.3 N.4

KRAS p. A146T
APC p.H427fs*1
APC p.L1488fs*41

homozyg.
TP53 p.N235_Y236delNY

T4.1

T4.2

T1.1

T1.4

T2.1

T2.2

T3.2

T3.4

10 mutations

deletion
inversion
tandem duplication
translocation

MRCA in patient 2) we calculated the number of signature 1 mutations present in the cell when signature 18 started. Assuming a linear acquisition of signature 1 in the trunk of the tumour at the same rate as in the normal samples, we could translate this signature 1 burden into the age at which signature 18 began. As the ratio of other signatures to signature 1 increases down the tree it is likely that that the rate of acquisition of signature 18 relative to signature 1 increases over time. If this is the case, the estimates that we obtain can be seen as a lower bound of age. Following this line of reasoning, in patient 1 (aged 67) signature 18 began in the last 24 years, in patient 2 (aged 68) in the last 20 years, and in patient 3 (aged 56) in the last 22 years.

Summary of results in this chapter

Our understanding of the process of somatic mutation, and its consequences, in normal cells remains rudimentary. This has largely been because of the difficulties of detecting somatic mutations in normal cells. A protocol to sequence individual colonic crypts isolated by laser capture microdissection was derived and applied to sequence hundreds of crypts from 42 different individuals. We characterised the genomic landscape of these normal cells, quantifying their mutation burden, the signatures of the mutational processes that have affected them, and the spectrum of driver mutations found in the colons of healthy people.

Putative driver mutations were found to occur in approximately 1% of crypts, and so number in the hundreds of thousands in the colons of middle-aged healthy individuals. The spectrum of driver mutations was different to that observed in cancers, with some of the most common colorectal cancer driver mutations absent from our data. Caution should be exercised with these estimates, however, as we are limited by a low frequency of driver mutations and variable coverage within and between genes; further studies (guided by this initial estimate) will be required to establish the frequency of driver mutations more accurately.

Striking variability in somatic mutation rates in normal cells was found between different people and between different cells within the same person. This was due both to variation in the number of mutations due to ubiquitous mutational processes and to the presence of sporadic mutational processes that only affected certain individuals or certain crypts within one individual. One novel signature, which quintupled the normal mutation burden, could be linked to a previous

194

exposure to chemotherapy. Others were of unknown aetiology, but accounted for thousands of mutations in focal patches of the colons of children. APOBEC mutations, which have never previously been reported in normal cells, were found in a small subset of crypts.

The genomes of colorectal cancers show mutational signatures that were not found even in a large cohort of normal cells. Those mutational processes that were found in all normal cells are universally accelerated in cancer. Phylogenetic analysis of the genomes of cancer cell-derived organoids reveals that the cell ancestral to the cancer was for most of its life subject to the same mutational processes as its neighbours that did not transform. In the process of transformation to cancer, this clone both accelerated normal mutational processes and acquired novel ones.

These findings will be discussed in the Discussion chapter.

DISCUSSION

In this dissertation, I have presented two case studies of somatic evolution. One investigates clonal dynamics, and the other somatic mutations, using blood and colon respectively as the model tissues. Here, I will discuss how this work informs our understanding of somatic evolution, drawing out future perspectives for how their extension, and the development of the field more broadly, might change the way in which we manage diseases that stem from natural selection within a multicellular organism.

1. Discussion of findings

1.a. The clonal dynamics of blood

Our study is the first to investigate haematopoiesis in an unperturbed healthy human using spontaneous somatic mutations. Our method is analogous to lineage tracing experiments in animal models, but we were able to use thousands of nested markers, detecting clones that arose at different times throughout life. We provide proof-of-concept that somatic mutations are a powerful tool to dissect the stem cell dynamics of normal tissues, as they have been used to study population dynamics and cancer evolution.

There is no previous *ab initio* estimate of the number of blood stem cells in a human in steady state. The value that we derive of 97,000 (90% CI 45,000−215,000) stem cells contributing to granulopoiesis will require validation by studying a large number of individuals. A more precise estimate could be obtained by whole genome sequencing a larger number of cells and, probably more importantly, developing a method of error corrected sequencing that would enable us to detect rarer mutations with high confidence. Methods related to duplex sequencing (Schmitt et al. 2012, Kennedy 2014, Hoang et al. 2016) may provide such an opportunity.

Quantification of the number of haematopoietic stem cells is an important biological insight and is also crucial to our understanding of the development of malignancy. As stem cells are believed to be at risk of transformation (General Introduction) each one may represent a potential cancer. In their discussion of cancer risk, Tomasetti and Vogelstein estimate a number of blood stem cells that is two orders of magnitude greater than ours (Tomasetti and Vogelstein 2015, General Introduction). Other factors, such as microenvironmental changes (General Introduction), are likely to be involved in determining the probabilities of mutant clones arising and outcompeting their neighbours, but our estimate provides a step on the road to a quantitative model of cancer development.

It remains to be shown precisely how the cells that we define as stem cells based on their long-term self-renewal *in vivo* map onto those at risk of malignant transformation. Our observation of granulocyte-B cell restricted clones in adults indicates that the cells that we are counting have both lymphoid and myeloid potential. Our study could be expanded to all phenotypically distinct cell types, resolving the clonal contribution of stem cells to each. We would, in theory, be able to build a quantitative map of the differentiation hierarchy of blood, calculating the number of stem cells responsible for maintaining each cell type. The gamut of haematological malignancies could be investigated, revealing the clone that gave rise to the neoplasm, and the lineage restriction of the line-of-descent that predates transformation. As clonal haematopoiesis is a common feature of ageing it could be followed prospectively, tracking its emergence from a normal clone through longitudinal blood draws. The method could further be applied to bone marrow failure syndromes, inflammatory conditions, and ageing.

1.b. Somatic mutations in normal colon

1.b.i. Driver mutations

We investigated the landscape of somatic mutations in hundreds of normal cells from 42 people. Due to the large number of crypts investigated, we were able to capture infrequent events, such as driver mutations. Despite sequencing over 1,000 crypts, driver mutations are so rare that assaying even more crypts will be necessary to develop an accurate estimate of the frequency of

each individual mutation in normal colon. Studies in mice and humans have begun to quantify the selective advantage that mutations confer in competition within the crypt and between crypts (Results Chapter 2, section I.5.). The laser capture microdissection method developed here would allow these studies to be extended to all driver mutations in humans. Ultimately, the selective advantage of every mutation, in different genetic and microenvironmental contexts, would allow us to chart all the possible paths to cancer and their respective probabilities, indicating the best points to intervene. It might be that early 'clonal purging' is the best approach. For example, the presence of *PIK3CA* mutations in normal colorectal mucosa, combined with a recent successful trial of a PIK3CA inhibitor in patients with an overgrowth syndrome as a result of an early embryonic somatic mosaic *PIK3CA* mutation (Venot et al. 2018), indicates that it would be possible to intervene pharmacologically in somatic evolution in the normal colon. The *ERBB2* and *ERBB3* mutations that we found may also represent druggable targets.

Caution should be exercised, however, as not all mutations that confer a selective advantage in normal mucosa need promote cancer development. Though *STAG2* loss confers a selective advantage in normal colon (Results Chapter 2 section I.5. and R.4.), it is rare in colon cancers. Similarly, in the skin and oesophagus, *NOTCH* family mutations are, if anything, more frequent in normal epithelia than in the cancers that derive from them (Martincorena et al. 2015, Inigo Martincorena personal communication). It seems, therefore, that some mutations provide a selective advantage without conferring the hallmarks of cancer (Hanahan and Weinberg 2000). In the colon, where the progression to malignancy is well described, a distinction could be drawn between potentially innocuous mutations and those that actually cause cancer by repeating the experiment along the adenoma carcinoma sequence.

Although the frequency of driver mutations *per colon* is high, the frequency of driver mutations *per crypt* is low, certainly in comparison to skin (30% of cells, (Martincorena et al. 2015)), oesophagus (40% of cells, Inigo Martincorena personal communication). This is despite the fact that the incidence of colon cancer is similar to that of oesophageal (5% and 1% lifetime risk respectively). We speculate that the difference in driver frequencies between these tissues is due in part to the modular structure of glandular epithelium, which determines the strength of selection within the crypt and prevents clonal expansions beyond it (General Introduction). Indeed, it seems that some colorectal cancer driver mutations promote crypt fission (Results Chapter 2 section I.5.b.), allowing mutant stem cells to escape the constraint of the gland. Two of the 14

driver mutations that we observed were associated with crypt fission events. Recently, somatic driver mutations were found at a high frequency in endometrial glands (Suda et al. 2018). *PIK3CA* mutations were reported in 20% of glands, although not all of these need be driver mutations. Endometrium may represent an unusual glandular tissue in that it is shed and regenerated on a monthly basis during reproductive life. This might provide opportunities for glands harbouring mutant stem cells to fission and colonise the epithelium more effectively than in the colon. The larger size of endometrial glands, if coupled to a larger number of stem cells per gland, may also provide greater opportunities for stem cells that have acquired driver mutations to colonise the gland. Further work will be required to establish the frequency of driver mutations in other glandular tissues and their basis in terms of the rate of acquisition of somatic mutations, their selective advantage, and the constraints on clonal expansion.

1.b.ii. Mutational signatures

We observe a surprising diversity of mutational processes active across the normal colons of different people. Many of these processes had previously been reported, but the size of our cohort and the low complexity of normal genomes allowed us to investigate how their rates vary across and within different people and reveal differences between sites in the colon. Intriguingly, rates of common mutational processes were higher on the right side of the colon than on the left, which is the opposite trend to cancer incidence. We also confirm the finding that mutation rates are similar between the colon and the ileum (Blokzijl et al. 2016), despite the former developing cancers more frequently. Part of this discrepancy may be due to the number of stem cells per crypt and differences in selection pressures.

We found signatures that have not been detected in the PCAWG analysis of tens of thousands of cancers (Alexandrov et al. 2018). These novel signatures may not have been described in cancers simply because they are specific to the colon and fewer colorectal cancer genomes have been sequenced than we have sequenced normal crypts. Alternatively, these processes may be present in cancer genomes, but obscured by the panoply of additional mutational processes. Some novel signatures have an explanation, such as SBSN1 which is found in a patient who had been exposed to a cocktail of chemotherapies. It is sobering that chemotherapy could

cause so many mutations all around the colon. Nevertheless, during this patient's colonoscopy, which occurred 17 years after his first chemotherapy, only one adenocarcinoma and one adenoma were found, which is not a particularly rare finding in a 66 year-old man. Furthermore, although the increase in mutations is likely to have affected many other tissues too, this individual survived for 19 years after his first chemotherapy. Some signatures were of unknown cause, such as SBSN2 and SBSN3. These seem to be active early in life, and their features suggest that they are caused by extrinsic mutational processes. They warrant further investigation as they may represent preventable exposures that account for thousands of somatic mutations in a high proportion of people.

A comparison of the mutational processes in tumours and normal tissues reveals that the process of transformation is associated with both an increase in the rate of normal mutational processes and the acquisition of additional ones, even in non-hypermutated cancers. Ostensibly, they allow an accelerated evolution of the cancer and contribute to the subclonal diversity within the tumour that allows some cells to survive chemotherapy. The origin of many of these aberrant mutational processes is unknown; one can hope that the identification of their aetiologies might indicate ways to intervene in tumour evolution.

2. A comparison of the somatic evolution of blood and colon

A brief comparison of blood and colon highlights the diversity of the forces that govern somatic evolution across the human body.

Stem cell numbers appear to be quite different between the two tissues. Assuming seven functional stem cells per crypt and 15 million crypts per colon, there are approximately 100 million stem cells per human colon. This is 1,000 times larger than our estimate of the number of active blood stem cells. This is a striking difference given that the number of mature blood cells produced per day is larger than the number of colonocytes. Presumably, the optimal number of stem cells per tissue has been determined over the course of species evolution in part by balancing the requirement to produce large numbers of differentiated cells against the risk of malignancy. Why, then, should there be so many more colorectal stem cells than blood stem cells? We speculate that

the colon can have far more stem cells than blood without a substantially increased cancer risk[5] because of the organisation of colonic stem cells into crypts, which reduces the effect of mutations that confer a selective advantage (General introduction, section 3.c.). We might predict from this that most mutations that are known to cause colorectal cancers confer a stronger selective advantage than those that are known to cause blood cancers, since even a weak driver mutation can, over time, sweep through the relatively unstructured blood stem cell pool. The presence of common and indolent blood clones that rarely (or, at least, slowly) progress to malignancy supports this view. Recent evidence indicates that clonal haematopoiesis is likely to be responsible for age-related morbidities beyond cancer (General Introduction, section 4). The very liquidity of blood and the recirculation of blood stem cells, which must limit architectural controls on clonal competition, can perhaps be considered to be partly responsible for these diseases.

Mutation burdens and processes are different across the two tissues, as a colonic stem cell has approximately three times as many mutations as a blood stem cell of the same age. Although only 140 blood stem cells from one individual were sequenced, it seems that there is more variability in the mutational processes and mutation burden across the colon than in blood. Some of the sporadic mutational processes in the colon are likely to be caused by mutagens, a likely corollary of the exposure of intestinal stem cells to the luminal contents of microflora and dietary carcinogens. Blood stem cells, in contrast, are maintained in a more homogeneous environment.

Both variation and selection are therefore different in blood and colon, and are likely to vary across all tissues. Indeed, cancers from different organs have a different incidence, repertoire of mutational signatures, and predilection for particular driver mutations. Every tissue will have to be studied in order to build up a holistic view of somatic evolution across the human body.

3. Perspective on managing the diseases of somatic evolution

Somatic evolution is inevitable in a multicellular organism with a long lifespan and imperfect DNA replication and repair. Nonetheless, an understanding of the forces of variation

---

[5] Colorectal cancer incidence is approximately four times that of all leukaemias combined (Cancer Research UK), but this excess is probably largely a result of recent dietary changes (chapter 2, section I.3.), and the disease mostly affects the elderly. Throughout the course of much of human evolution it seems probable that colorectal cancers were not a significantly more frequent cause of premature death than leukaemias.

and selection in each tissue provides clues as to how we can manage the diseases that stem from it.

First, the description of mutational signatures across tissues and the identification of their causes may provide opportunities to decrease the mutation rate. Extrinsic mutagens may be avoided, and protective drugs may be discovered. For example, low-dose aspirin decreases the risk of bowel cancer (Brenner et al. 2014), and it has been suggested that this effect of non-steroidals on cancer incidence may be through reducing the mutation rate (Kostadinov et al. 2013). Studying the mutational signatures of normal colonic crypts of those who have and have not taken aspirin would reveal whether or not this was the case. Some mutational processes, such as polymerase slippage in normal DNA replication, will always, however, be unavoidable. Measuring a person's normal mutation rate may allow patient stratification: those with higher mutation rates might receive more regular screening tests than those with lower ones.

Second, clonal dynamics might be monitored. It has recently been shown that it is possible to predict which patients with clonal haematopoiesis are at high risk of progression to acute myeloid leukaemia (Abelson et al. 2018). By developing a good understanding of healthy clonal dynamics, it will be possible to detect a departure from normality and quantify the risk of progression to malignancy or another disease. This would be easiest in blood, but one can imagine that a similar approach might be possible for other tissues, such as through circulating DNA (Wan et al. 2017) or other methods of randomly sampling epithelium (Lao-Sirieix and Fitzgerald 2012). This could inform risk stratification and treatment.

Finally, as discussed above, judicious 'clonal purging' of mutations associated with a high risk of progression might remove clones early, before they can grow large enough that they contain sufficient genetic diversity to be resistant to drugs.

4. Conclusion

> It is interesting to contemplate an entangled bank, clothed with many plants of many kinds, with birds singing on the bushes, with various insects flitting about, and with worms crawling through the damp earth, and to reflect that these elaborately constructed forms, so different from each other, and dependent on each other in so complex a manner have all been produced by laws acting around us. These laws, taken in the largest sense, being Growth with Reproduction; Inheritance which is almost implied by reproduction; Variability from the indirect and direct action of the external conditions of life, and from use and disuse; a Ratio of Increase so high as to lead to a Struggle for Life, and as a consequence to Natural Selection, entailing Divergence of Character and the Extinction of less-improved forms.
>
> (Charles Darwin, *On the Origin of Species*)

A slide of cancer histology presents a comparable richness and diversity to Darwin's entangled bank, but on a different scale: cells acquire mutations, creating diversity, and compete with one other. Cancers emerge as a result of this struggle for existence. Darwinian theory can be applied to somatic tissues: this requires the integration of the fields of cancer genomics, evolutionary biology, population genetics, and cell and developmental biology. It is to be hoped that such an approach offers opportunities to extend our understanding of health and disease.

BIBLIOGRAPHY

Abelson, S. et al. (2018). Prediction of acute myeloid leukaemia risk in healthy individuals. Nature. *559(7714)*, 400-404.

Abkowitz, J. L., et al. (1996). Evidence that hematopoiesis may be a stochastic process in vivo. Nat. Med. *2*, 190–197.

Abkowitz, J.L. et al. (2000). In vivo kinetics of murine hemopoietic stem cells. Blood. *96(10)*, 3399-405.

Abkowitz, J.L. et al. (2002). Evidence that the number of hematopoietic stem cells per animal is conserved in mammals. Blood. *100(7)*, 2665-2667.

Alcolea, M.P. et al. (2014). Differentiation imbalance in single oesophageal progenitor cells causes clonal immortalization and field change. Nat. Cell Biol. *16(6)*, 615-622.

Alexandrov, L.B., et al. (2013). Signatures of mutational processes in human cancer. Nature. *500*, 415–421.

Alexandrov. L.B. et al. (2013b). Deciphering signatures of mutational processes operative in human cancer. Cell Rep. *3(1)*, 246-59

Alexandrov, L.B. and Stratton, M.R. (2014). Mutational signatures: the patterns of somatic mutations hidden in cancer genomes. Curr. Opin. Genet. Dev. *24*, 52-60.

Alexandrov, L.B. et al. (2015). Clock-like mutational processes in human somatic cells. Nat Genet. *47(12)*, 1402-1407.

Alexandrov, L.B. et al. (2018). The repertoire of mutational signatures in human cancer. Preprint at: https://www.biorxiv.org/content/early/2018/05/15/322859.

Aparicio, T. et al. (2014). Small bowel adenocarcinoma: epidemiology, risk factors, diagnosis and treatment. Dig Liver Dis. *46(2)*, 97-104.

Arai, F. and Suda, T. (2007). Maintenance of quiescent hematopoietic stem cells in the osteoblastic niche. Ann N Y Acad Sci. *1106*, 41-53.

Armitage, P., and Doll, R. (1954). The age distribution of cancer and a multistage theory of carcinogenesis. Br. J. Cancer *8*, 1–12.

Bae, T. et al. (2018). Different mutational rates and mechanisms in human cells at pregastrulation and neurogenesis. Science *359*, 550–555.

Baker, A.M. et al. (2014). Quantification of crypt and stem cell evolution in the normal and neoplastic human colon. Cell Rep. *8(4)*, 940-7.

Balbás-Martínez, C., et al. (2013). Recurrent inactivation of STAG2 in bladder cancer is not associated with aneuploidy. Nat Genet. *45(12)*, 1464-1469.

Barker, N. et al. (2007). Identification of stem cells in small intestine and colon by marker gene Lgr5. Nature. *449(7165)*, 1003-1007.

Barker, N. et al. (2009). Crypt stem cells as the cells-of-origin of intestinal cancer. Nature. *457(7229)*, 608-611.

Becker, A.L., McCulloch, E.A., and Till, J.E. (1963). Cytological demonstration of the clonal nature of spleen colonies derived from transplanted mouse marrow cells. Nature. *197*, 452-454.

Beaumont, M. A., Zhang, W., and Balding, D. J. (2002). Approximate Bayesian computation in population genetics. Genetics, *162*, 2025-2035.

Behjati, S. et al. (2014). Genome sequencing of normal cells reveals developmental lineages and mutational processes. Nature. *513(7518)*, 422-425.

Bertorelle, G. et al. (2010). ABC as a flexible framework to estimate demography over space and time: some cons, many pros. Molecular Ecology. *19*, 2609-2625.

Biasco, L. et al. (2016). In vivo tracking of human hematopoiesis reveals patterns of clonal dynamics during early and steady-state reconstitution phases. Cell Stem Cell. *19*, 107-119.

Bienz, M. and Clevers, H. (2000). Linking colorectal cancer to Wnt signaling. Cell. *103(2)*, 311-320.

Bignold, L.P., Coghlan, B.L.D., and Jersmann, H.P.A. (2006). Hansemann, Boveri, chromosomes and the gametogenesis-related theories of tumours. Cell Biology International *30* 640-644

Bird C. et al. (2011) Detecting and measuring genetic differentiation. Crustacean Issues: Phylogeography and Population Genetics in Crustacea. *19*, 31-55.

Bizzozero, G. (1868). Sulla funzione ematopoetica del midollo delle ossa. Comunicazione preventiva. Gazz. Med. Ital. Lombardia *28*, 381-382.

Blokzijl, F. et al. (2016). Tissue-specific mutation accumulation in human adult stem cells during life. Nature *538*, 260–264.

Blum, M. G. and Francois, O. (2010). Non-linear regression models for Approximate Bayesian Computation. Statistics and Computing. *20(1)*, 63-73.

Boman, B.M. and Huang, E. (2008). Human Colon Cancer Stem Cells: A New Paradigm in Gastrointestinal Oncology. J. Clin. Oncol. *26*, 2828-2838.

Bonnet, D. and Dick, J.E. (1997). Human acute myeloid leukemia is organized as a hierarchy that originates from a primitive hematopoietic cell. Nat Med. *3(7)*, 730-7.

Boot, A. et al. (2018). Mutational signature analysis of Asian OSCCs reveals novel mutational signature with exceptional sequence context specificity. Preprint at: https://www.biorxiv.org/content/early/2018/07/19/368753.1.

Boveri, T. (1914). Zur Frage der Entstehung Maligner Tumoren (Gustav Fischer, Jena); English translation The Origin of Malignant Tumors by Boveri, M. (Williams and Wilkins, Baltimore, 1929, 1914).

Bozic, I. et al. (2010). Accumulation of driver and passenger mutations during tumor progression. Proc. Natl. Acad. Sci. *107(43)*, 18545–18550.

Breivik, J. (2001). Don't stop for repairs in a war zone: Darwinian evolution unites genes and environment in cancer development. Proc. Natl. Acad. Sci. *98(10)*, 5379-5381.

Brenner, H., Kloor, M., and Pox, C.P. (2014). Colorectal cancer. Lancet. *383(9927)*, 1490-1502.

Brown, S. et al. (2017). Correction of aberrant growth preserves tissue homeostasis. Nature. *548(7667)*, 334-337.

Bruens, L. et al. (2017). In Vivo Imaging Reveals Existence of Crypt Fission and Fusion in Adult Mouse Intestine. Gastroenterology. *153(3)*, 674-677.

Bryder, D., Rossi, D.J., and Weissman, I.L. (2006). Hematopoietic stem cells: the paradigmatic tissue-specific stem cell. Am. J. Pathol. *169(2)*, 338-346.

Buczacki, S.J. et al. (2013). Intestinal label-retaining cells are secretory precursors expressing Lgr5. Nature. *495(7439),* 65-69.

Buels, R. et al. (2016). Jbrowse: a dynamic web platform for genome visualization and analysis. Genome Biology. Doi: 10.1186/s13059-016-0924-1.

Buescher E.S. et al. (1985). Use of an X-linked human neutrophil marker to estimate timing of lyonization and size of the dividing stem cell pool. J Clin Invest. *76(4)*, 1581-4.

Busch, K. et al. (2015). Fundamental properties of unperturbed haematopoiesis from stem cells in vivo. Nature *518*, 542-546.

Busque, L. et al. (2012). Recurrent somatic TET2 mutations in normal elderly individuals with clonal hematopoiesis. Nat Genet. *44(11)*, 1179-81.

Bystrykh, L.V. et al. (2012). Counting stem cells: methodological constraints. Nature methods. *9(6)*, 567-574.

Cabezas-Wallscheid, N. et al. (2017). Vitamin A-Retinoic Acid Signaling Regulates Hematopoietic Stem Cell Dormancy. Cell. *169*, 807–823.

Cairns, J. (1975). Mutation selection and the natural history of cancer. Nature *255*, 197–200.

Calabrese and Shibata (2010). A simple algebraic cancer equation: calculating how cancers may arise with normal mutation rates. BMC Cancer, *10(3)*, https://doi.org/10.1186/1471-2407-10-3.

Campbell, P.J. et al. (2017). Pan-cancer analysis of whole genomes. Preprint at: https://www.biorxiv.org/content/early/2017/07/12/162784.

Cancer Genome Atlas Network et al. (2012). Comprehensive molecular characterization of human colon and rectal cancer. Nature *487*, 330-337.

Cancer Genome Atlas Research Network et al. (2013). Genomic and epigenomic landscapes of adult de novo acute myeloid leukemia. N. Engl. J. Med. *368(22)*, 2059-2074.

Cancer Research UK, Bowel Cancer Incidence Statistics, https://www.cancerresearchuk.org/health-professional/cancer-statistics/statistics-by-cancer-type/bowel-cancer/incidence#heading-Seven (Accessed August 2018).

Cartier, N. et al. (2009). Hematopoietic stem cell gene therapy with a lentiviral vector in X-linked adrenoleukodystrophy. Science. *326*, 818–823.

Catlin, S. N., et al. (2011). The replication rate of human hematopoietic stem cells in vivo. Blood. *117*, 4460-4466.

Cavazzana-Calvo, M. et al. (2010). Transfusion independence and HMGA2 activation after gene therapy of human β-thalassaemia. Nature. *467(7313)*, 318-322.

Challen G.A. et al. (2010). Distinct hematopoietic stem cell subtypes are differentially regulated by TGF-beta1. Cell Stem Cell. *6(3)*, 265-78.

Chan, K. et al. (2015). An APOBEC3A hypermutation signature is distinguishable from the signature of background mutagenesis by APOBEC3B in human cancers. Nat Genet. *47(9)*, 1067-72.

Cheng, H. and Leblond, C.P. (1974). Origin, differentiation and renewal of the four main epithelial cell types in the mouse small intestine. V. Unitarian Theory of the origin of the four epithelial cell types. Am J Anat. *141*, 537-561.

Choi, K. et al. (1998). A common precursor for hematopoietic and endothelial cells. Development. *125*, 725-32.

Christmann, M. et al. (2016). Adaptive upregulation of DNA repair genes following benzo(a)pyrene diol epoxide protects against cell death at the expense of mutations. Nucleic Acids Res. *44(22)*, 10727-10743.

Corley, D.A. et al. (2013). Variation of adenoma prevalence by age, sex, race, and colon location in a large population: implications for screening and quality programs. Clin Gastroenterol Hepatol. *11(2)*, 172-180.

Cooper, B. (2011). The origins of bone marrow as the seedbed of our blood: from antiquity to the time of Osler. Proc (Bayl. Univ. Med. Cent.). *24(2)*, 115-118.

Copley, M.R., Beer, P.A., and Eaves, C.J. (2012). Hematopoietic stem cell heterogeneity takes center stage. Cell Stem Cell. *10(6),* 690-697.

Copley, M.R. et al. (2013). The Lin28b-let-7-Hmga2 axis determines the higher self-renewal potential of fetal haematopoietic stem cells. Nat. Cell Biol. *15(8)*, 916-925.

Csillery, K., Francois, O., and Blum, M. G. (2012). abc: an R package for approximate Bayesian computation (ABC). Methods in ecology and evolution. *3(3)*, 475-479.

Darwin, C. (1859) On the origin of species by means of natural selection, or the preservation of favoured races in the struggle for life. London: John Murray.

DeGregori, J. (2017). Connecting cancer to its causes requires incorporation of effects on tissue microenvironments. Cancer Res. doi: 10.1158/0008-5472.CAN-17-1207

Dentro, S.C. et al. (2018). Portraits of genetic intra-tumour heterogeneity and subclonal selection across cancer types. Preprint at: https://www.biorxiv.org/content/early/2018/05/12/312041

Dick, J.E. et al. (1985). Introduction of a selectable gene into primitive stem cells capable of long-term reconstitution of the hemopoietic system of W/Wv mice. Cell. *42(1)*, 71-9.

Dieterlen-Lievre, F. (1975). On the origin of haemopoietic stem cells in the avian embryo: an experimental approach. J Embryol Exp Morphol. *33(3)*, 607-19.

Ding, L. et al. (2010), Genome remodelling in a basal-like breast cancer metastasis and xenograft. Nature. *464(7291)*, 999-1005.

Ding, L. et al. (2012), Clonal evolution in relapsed acute myeloid leukaemia revealed by whole-genome sequencing. Nature. *481(7382)*, 506-10.

Domazet-Loso, T., and Tautz, D. (2010). Phylostratigraphic tracking of cancer genes suggests a link to the emergence of multicellularity in metazoa. BMC Biol. *8*, 66

Doulatov S et al. (2010). Revised map of the human progenitor hierarchy shows the origin of macrophages and dendritic cells in early lymphoid development. Nat. Immunol. *7*, 585-593.

Doulatov, S. et al. (2012). Hematopoiesis: a human perspective. Cell Stem Cell. *10(2)*, 120-136.

Drost, J. et al. (2015). Sequential cancer mutations in cultured human intestinal stem cells. Nature. *521(7550)*, 43-7.

Dykstra, B. et al. (2007). Long-term propagation of distinct hematopoietic differentiation programs in vivo. Cell Stem Cell. *1(2)*, 218-29.

Dzierzak, E. and Speck, N.A. (2008). Of lineage and legacy: the development of mammalian hematopoietic stem cells. Nat. Immunol. *9(2)*, 129-136.

Eaves, C.J. (2015). Hematopoietic stem cells: concepts, definitions, and the new reality. Blood. *125(17)*, 2605-2613.

Excoffier L. et al. (1992). Analysis of molecular variance inferred from metric distances among DNA haplotypes: application to human mitochondrial DNA restriction data. Genetics. *131*, 479-491.

Fearon, E. R. & Vogelstein, B. (1990). A genetic model for colorectal tumorigenesis. Cell. *61*, 759–767.

Fearon, E.R. (2011). Molecular genetics of colorectal cancer. Annu. Rev. Pathol. Mech. Dis. *6*, 479–507.

Felsenstein, J. (1989). PHYLIP — Phylogeny Inference Package (Version 3.2). Cladistics 5, 973 164–166.

Fialkow, P.J. (1974). Human genetic markers as tracers of tumour histogenesis. J. Clin. Pathol. Suppl. (R. Coll. Pathol.). *7*, 11–15.

Fialkow, P.J. et al. (1978). Chronic myelocytic leukemia. Origin of some lymphocytes from leukemic stem cells. J. Clin. Invest. *62(4)*, 815-823.

Fischer, A. et al. (2013). EMu: probabilistic inference of mutational processes and their localization in the cancer genome. Genome Biol. 14.4, R39.

Fisher, R.A. (1922). On the dominance ratio. Proc. R. Soc. Edinburg *42*, 321–431.

Forbes S.A. et al. (2017). COSMIC: somatic cancer genetics at high-resolution. Nucleic Acids Res. *45*, D777-D783.

Ford, C.E. et al. (1956). Cytological identification of radiation-chimaeras. Nature. *177*, 452-454.

Ford, A.M. et al. (1993). In utero rearrangements in the trithorax- related oncogene in infant leukaemias. Nature *363*, 358–360.

Ford, A.M. et al. (2009). The *TEL-AML1* leukemia fusion gene dysregulates the TGF-β pathway in early B lineage progenitor cells. J. Clin. Invest. *119(4)*, 826-836.
Foudi, A. et al. (2009). Analysis of histone 2B-GFP retention reveals slowly cycling hematopoietic stem cells. Nat Biotechnol. *27(1)*, 84-90.

Foulds, L. (1958). The natural history of cancer. J. Chronic Dis. *8(1)*, 2-37.

Friend, S. H. et al. (1986). A human DNA segment with properties of the gene that predisposes to retinoblastoma and osteosarcoma. Nature. *323*, 643–646.

Gallagher, J. 2nd Jan. 2015. Most cancer types 'just bad luck'. BBC News Website. Available at: https://www.bbc.co.uk/news/health-30641833

Garinis, G.A. et al. (2008). DNA damage and ageing: new-age ideas for an age-old problem. Nat Cell Biol. *10(11)*, 1241-7.

Genovese, G. et al. (2014). Clonal hematopoiesis and blood-cancer risk inferred from blood DNA sequence. N. Engl. J. Med. *371*, 2477–2487.

Gerstung, M. et al. (2014). Subclonal variant calling with multiple samples and prior knowledge. Bioinformatics *30*, 1198–1204.

Gerstung, M. et al. (2017). The evolutionary history of 2,658 cancers. Preprint at: https://www.biorxiv.org/content/early/2018/09/12/161562

Gladstein, K. (1978). The characteristic values and vectors for a class of stochastic matrices arising in genetics. SIAM Journal of Applied Mathematics. *34*, 630-642

Griffiths D.F. et al. (1988). Demonstration of somatic mutation and colonic crypt clonality by X-linked enzyme histochemistry. Nature. *333*, 461-463.

Griffiths, R.C. and Tavaré S. (1994). Sampling theory for neutral alleles in a varying environment. Philos Trans R Soc Lond B Biol Sci. *344(1310)*, 403-10.

Greaves, L.C. et al. (2006). Mitochondrial DNA mutations are established in human colonic stem cells, and mutated clones expand by crypt fission. Proc Natl Acad Sci. *103(3)*, 714-719.

Greaves, M. (1993). A natural history for pediatric acute leukemia. Blood. *82(4)*, 1043-51.

Greaves, M. et al. (2003). Leukemia in twins: lessons in natural history. Blood. *102*, 2321–2333.

Greaves, M. (2018). A causal mechanism for childhood acute lymphoblastic leukaemia. Nat. Rev. Cancer. *18(8)*, 471-484.

Greaves, M. and Maley, C.C. (2012). Clonal evolution in cancer. Nature. *481*, 306-313.

Hacein-Bey-Abina, S. et al. (2003). A serious adverse event after successful gene therapy for X-linked severe combined immunodeficiency. N. Engl. J. Med. *348(3)*, 255-256.

Haeckel, E. (1868). Natürliche Schöpfungsgeschichte. Berlin: Georg Reimer.

Hanahan, D. and Weinberg, R.A. (2000). The hallmarks of cancer. Cell. *100(1), 57-70.*

Haradhvala, N.J. et al. (2016). Mutational strand asymmetries in cancer genomes reveal mechanisms of DNa damage and repair. Cell *164*, 538–549.

Helleday, T., Eshtad, S., and Nik-Zainal, S. (2014). Mechanisms underlying mutational signatures in human cancers. Nat Rev Genet. *15(9)*, 585-98.

Henshaw, P.S. et al. (1947). The carcinogenic effect of pile radiations. Cancer res. *7(1)*, 48.
Ho, T. T. et al. (2017). Autophagy maintains the metabolism and function of young and old stem cells. Nature. *543***,** 205–210.

Hill, V.K., Kim, J.S., and Waldman, T. (2016). Cohesin mutations in human cancer. Biochim Biophys Acta. *1866(1)*, 1-11.

Hoang, M.L. et al. (2016). Genome-wide quantification of rare somatic mutations in normal human tissues using massively parallel sequencing. Proc. Natl. Acad. Sci. *113(35)*, 9846-9851.

Hoffbrand, A.V., Catovsky, D., Tuddenham, E.G.D., Green, A.R. editors. 2011. Postgraduate haematology. 6th Edition. Chichester (UK): Wiley-Blackwell.

Hsieh, J.C. et al. (2013). Large chromosome deletions, duplications, and gene conversion events accumulate with age in normal human colon crypts. Aging Cell. *12(2)*, 269-279.

Huber, T.L. et al. (2004). Haemangioblast commitment is initiated in the primitive streak of the mouse embryo. Nature. *432(7017),* 625-630.

Huntly B.J.P., et al. (2004). MOZ-TIF2, but not BCR-ABL, confers properties of leukaemic stem cells to committed murine haematopoietic progenitors. Cancer Cell. *6*, 586-595

Ikehata, H. and Ono, T. (2011). The mechanisms of UV mutagenesis. J. Radiat. Res. *52(2)*, 115-125.

Ito, K. and Suda, T. (2014). Metabolic requirements for the maintenance of self-renewing stem cells. Nat. Rev. Mol. Cell Biol. *15(4)*, 243-256.

Jahn K, Kuipers J & Beerenwinkel N. (2016). Tree inference for single-cell data. Genome Biol. *17*, 86.

Jaiswal, S. et al. (2014). Age-related clonal hematopoiesis associated with adverse outcomes. N. Engl. J. Med. *371*, 2488–2498.

Jaiswal, S. et al. (2017). Clonal haematopoiesis and risk of atherosclerotic cardiovascular disease. N. Engl. J. Med. *377*, 111-121.

Jones, D. et al. (2016). cgpCaVEManWrapper: Simple execution of CaVEMan in order to detect somatic single nucleotide variants in NGS data. Curr. Protoc. Bioinformatics. *56*, 15.10.1–15.10.18.

Karcher, M. et al. (2017). phylodyn: an R package for phylodynamic simulation and inference. Mol. Ecol. Resour. *17(1)*, 96-100.

Kay, H.E.M. (1965). How many cell-generations? Lancet. *2*, 418.

Keller, G. et al. (1985). Expression of a foreign gene in myeloid and lymphoid cells derived from multipotent haematopoietic precursors. Nature. *318(6042)*, 149-154.

Kennedy, S.R. et al. (2014). Detecting ultralow-frequency mutations by Duplex Sequencing. Nat Protoc. *9(11),* 2586-2606.

Kent, D.G. et al. (2009). Prospective isolation and molecular characterization of hematopoietic stem cells with durable self-renewal potential. Blood. *113(25)*, 6342-6350.

Kim, H.J. et al. (2000). Many multipotential gene-marked progenitor or stem cell clones contribute to hematopoiesis in nonhuman primates. Blood. *96(1)*, 1-8.

Kim, J. et al. (2016). Somatic ERCC2 mutations are associated with a distinct genomic signature in urothelial tumors. Nat. Genet. *48.6*, 600–606.

Kim, S. et al. (2014). Dynamics of HSPC repopulation in nonhuman primates revealed by a decade-long clonal-tracking study. Cell Stem Cell. *14(4)*, 473-475.

Kingman, J. F. C. (1982). On the genealogy of large populations. J. Appl. Probab. *19A*, 27-43.

Kinzler, K.W. and Vogelstein, B. (1996). Lessons from hereditary colorectal cancer. Cell. *87(2)*, 159-170.

Knudson, A. G. (1971). Mutation and cancer: statistical study of retinoblastoma. Proc. Natl Acad. Sci. *68*, 820–823.

Knudson, A.G. (2001). Two genetic hits (more or less) to cancer. Nat. Rev. Cancer. *1(2)*, 157-62.

Koelle, S.J. et al. (2017). Quantitative stability of hematopoietic stem and progenitor cell clonal output in rhesus macaques receiving transplants. Blood. *129(11)*, 1448-1457.

Kostadinov, R.L. et al. (2013). NSAIDs modulate clonal evolution in Barrett's esophagus. PLoS Genet. 9(6):e1003553.

Kozar, S. et al. (2013). Continuous clonal labeling reveals small numbers of functional stem cells in intestinal crypts and adenomas. Cell Stem Cell. *13(5)*, 626-633.

Krontiris, T.G. and Cooper, G.M. (1981). Transforming activity of human tumor DNAs. Proc. Natl Acad. Sci. *78(2)*, 1181–1184.

Lan, S. et al. (2015). An efficient Bayesian inference framework for coalescent-based nonparametric phylodynamics. Bioinformatics. *31*, 3282–3289.

Lansdorp, P.M. (2008). Telomeres, stem cells, and hematology. Blood. *111(4)*, 1759-1766.

Lao-Sirieix, P., and Fitzgerald R.C. (2012). Screening for oesophageal cancer. Nat Rev Clin Oncol. *9(5)*, 278-87.

Larochelle, A. et al. (1996). Identification of primitive human hematopoietic cells capable of repopulating NOD/SCID mouse bone marrow: implications for gene therapy. Nat. Med. *2(12)*, 1329-1337.

Laurenti, E. and Gottgens, B. (2018). From haematopoietic stem cells to complex differentiation landscapes. *553(7689)*, 418-426.

Lee-Six, H. et al. (2018). Population dynamics of normal human blood inferred from somatic mutations. Nature. https://doi.org/10.1038/s41586-018-0497-0

Leedham et al. (2009). Clonality, founder mutations, and field cancerization in human ulcerative colitis-associated aeoplasia. Gastroenterology. *136*, 542–550.

Lemischka, I.R. et al. (1986). Developmental potential and dynamic behavior of hematopoietic stem cells. Cell. *45(6)*, 917-927.

Ley, T.J. et al. (2008). DNA sequencing of cytogenetically normal acute myeloid leukemia genome. Nature. *456*, 66-72.

Li, Y. et al. (2017). Patterns of structural variation in human cancer. Preprint at: https://www.biorxiv.org/content/biorxiv/early/2017/08/27/181339.full.pdf

Lodato, M.A. et al. (2018). Aging and neurodegeneration are associated with increased mutations in single human neurons. Science *359*, 555–559.

Loeb, L.A. (1991). Mutator phenotype may be required for multistage carcinogenesis. Cancer Res. *51(12)*, 3075-3079.

Lopez-Garcia, C. et al. (2010). Intestinal stem cell replacement follows a pattern of neutral drift. Science. *330*, 822-825.

Lorenz, E. (1951). Modification of irradiation injury in mice and guinea pigs by bone marrow injections. J. Natl. Cancer. *12*, 197-201.

Luebeck, E.G. and Moolgavkar, S.H. (2002). Multistage carcinogenesis and the incidence of colorectal cancer. Proc Natl Acad Sci. *99(23)*, 15095-15100.

Lugli, N. et al. (2017). Enhanced Rate of Acquisition of Point Mutations in Mouse Intestinal Adenomas Compared to Normal Tissue. Cell Rep. *19(11)*, 2185-2192.

Martin, P.J. et al. (1980). Involvement of the B-lymphoid system in chronic myelogenous leukaemia. Nature. *287(5777)*, 49-50.

Martin, G.S. (1970). Rous Sarcoma Virus: a function required for maintenance of the transformed state. Nature *227*, 1021-1023.

Martincorena, I., et al. (2015). High burden and pervasive positive selection of somatic mutations in normal human skin. Science *348*, 880–886.

Martincorena, I. et al. (2017). Universal patterns of selection in cancer and somatic tissues. Cell. *171*, 1029–1041.

Mckerrell, T. et al. (2015). Leukemia-associated somatic mutations drive distinct patterns of age-related clonal hemopoiesis. Cell Rep. *10(8)*, 1239-1245.

Mckerrell, T., and Vassiliou. G.S. (2015). Aging as a driver of leukaemogenesis. Sci. Trans. Med. *7*, 306fs38

Mikkola, H.K. and Orkin, S.H. (2006). The journey of developing hematopoietic stem cells. Development. *133(19)*, 3733-3744.

Milton, K. (1987). Primate diets and gut morphology: implications for hominid evolution. In: Food and Evolution: Toward a Theory of Food Habits, eds. Harris M, Ross EB; Temple University Press, Philadelphia, pp. 93-115.

Mitchell, T.J. et al. (2018). Timing the landmark events in the evolution of clear cell renal cell cancer: TRACERx Renal. Cell. *173(3)*, 611-623.

Moran, P. A. P. (1958). Random processes in genetics. Proceedings of the Cambridge Philosophical Society, 54(1), 60-71.

Morley, A.A. (1995). The somatic mutation theory of ageing. Mutat. Res. *338*, 19–23.

Muller, H.J. (1927). Artificial Transmutation of the Gene. Science *66*, 84–87.

Muller-Sieburg, C.E. et al. (2004). Myeloid-biased hematopoietic stem cells have extensive self-renewal capacity but generate diminished lymphoid progeny with impaired IL-7 responsiveness. Blood, 103(11), 4111-4118.

Muller-Sieburg, C.E. et al. (2012). Stem cell heterogeneity: implications for aging and regenerative medicine. Blood. *119(17)*, 3900-3907.

Nangalia, J. and Green, A.R. (2018). Myeloproliferative neoplasms: from origins to outcomes. Blood. *130(23)*, 2475-2483.

Neumann, E. (1868). Über die Bedeutung des Knochenmarkes für die Blutbildung. Central med. Wissensch. 6, 689.

Ng, A.W.T. et al. (2017). Aristolochic acids and their derivatives are widely implicated in liver cancers in Taiwan and throughout Asia. Sci. Transl. Med. *9(412)*, eaan6446.

Nicholson, A. et al. (2018). Fixation and spread of somatic mutations in adult human colonic epithelium. Cell Stem Cell. *22*, 909-918.

Nicolas, P. et al. (2007). The Stem Cell Population of the Human Colon Crypt: Analysis via Methylation Patterns. PLoS Comput Biol 3(3): e28.

Nik-Zainal, S. et al. (2012). Mutational processes molding the genomes of 21 breast cancers. Cell *149(5)*, 979–993.

Nocturne, G. et al. (2013). Germline and somatic variations of TNFAIP3 in lymphoma complicating primary Sjogren's syndrome. Blood. *122(25)*, 4068-4076.

Nordling, C.O. (1953). A new theory on cancer-inducing mechanism. Br. J. Cancer *7*, 68–72.

Nowell, P.C. and Hungerford, D.A. (1960). Chromosome studies on normal and leukemic human leukocytes. J. Natl. Cancer Inst. *25*, 85–109.

Nowell, P.C. et al. (1956). Growth and continued fraction of rat marrow cells in X-radiated mice. Cancer Res. *16*, 258-261.

Nowell, P.C. (1976). The clonal evolution of tumor cell populations. Science *194*, 23–28.

Orford, K.W. and Scadden, D.T. (2008). Deconstructing stem cell self-renewal: genetic insights into cell-cycle regulation. Nat Rev Genet. *9(2)*, 115-128.

Ortmann, C.A. et al. (2015). Effect of mutation order on myeloproliferative neoplasms. N. Engl. J. Med. *372(7)*, 601-612.

Osawa, M. et al. Long-term lymphohematopoietic reconstitution by a single CD34-low/negative hematopoietic stem cell. Science. *273(5272)*, 242-245.

Owen, R.D. (1945). Immunogenetic consequences of vascular anastomoses between bovine twins. Science. *102*, 400-401.

Pang, W.W. et al. (2011). Human bone marrow hematopoietic stem cells are increased in frequency and myeloid-biased with age. Procl. Natl. Acad. Sci. *108(50)*, 20012-20017.

Pappenheim, A. (1896). Ueber Entwickelung und Ausbildung der Erythroblasten. Virchows Arch. *145*, 587-643.

Parsons, B.L. et al. (2010). ACB-PCR quantification of K-RAS codon 12 GAT and GTT mutant fraction in colon tumor and non-tumor tissue. Cancer Invest. *28(4)*, 364-75.

Parsons, B.L., McKim, K.L., and Myers, M.B. (2017). Variation in organ-specific PIK3CA and KRAS mutant levels in normal human tissues correlates with mutation prevalence in corresponding carcinomas. Environ Mol Mutagen. *58(7)*, 466-476

Paul, F. et al. (2015). Transcriptional Heterogeneity and Lineage Commitment in Myeloid Progenitors. Cell. *163(7)*, 1663-77.

Pei, W. et al. (2017). Polylox barcoding reveals haematopoietic stem cell fates realized in vivo. Nature. *548*, 456-460.

Perié, L. et al. The Branching Point in Erythro-Myeloid Differentiation. Cell. *163(7)*, 1655-62.

Petljak, M. and Alexandrov, L.B. (2016). Understanding mutagenesis through delineation of mutational signatures in human cancer. Carcinogenesis *37(6)*, 531-540.

Pleasance, E.D. et al. (2010a). A comprehensive catalogue of somatic mutations from a human cancer genome. Nature *463*, 191–196.

Pleasance, E.D. et al. (2010b). A small-cell lung cancer genome with complex signatures of tobacco exposure. Nature *463*, 184–190.

Ponder, B.A., et al. (1985). Derivation of mouse intestinal crypts from single progenitor cells. Nature. *313*, 689-691.

Ponder, B.A. and Wilkinson, M.M. (1986). Direct examination of the clonality of carcinogen-induced colonic epithelial dysplasia in chimeric mice. J. Natl. Cancer Inst. *77(4)*, 967-976.

Popovich, D.G. et al. (1997). The western lowland gorilla diet has implications for the health of humans and other hominoids. J. Nutr. *127(10)*, 2000-2005.

Potten, C.S. et al. (1992). Measurement of in vivo proliferation in human colorectal mucosa using bromodeoxyuridine. Gut. *33*, 71-78.

Potten, C.S. and Loeffler, M. (1990). Stem cells: attributes, cycles, spirals, pitfalls and uncertainties. Lessons for and from the crypt. Development. *110(4)*, 1001-1020.

Rahbari, R. et al. (2016). Timing, rates and spectra of human germline mutation. Nat. Genet. *48*, 126–133.

Raine, K.M. et al. (2015). cgpPindel: identifying somatically ac- quired insertion and deletion events from paired end sequencing. Curr. Protoc. Bioinformatics *52*, 15.17.11–15.17.12.

Reddy, E.P. et al. (1982). A point mutation is responsible for the acquisition of transforming properties by the T24 human bladder carcinoma oncogene. Nature. *300(5888)*, 149-152.

Rideout, W.M. et al. (1990). 5-Methylcytosine as an endogenous mutagen in the human LDL receptor and p53 genes. Science. *249(4974)*, 1288-90.

Ritsma, L. et al. (2014). Intestinal crypt homeostasis revealed at single-stem-cell level by in vivo live imaging. *507(7492)*, 362-365.

Roberts, N.D. 2018. Patterns of somatic genome rearrangement in human cancer [PhD thesis]. [Univ Cambridge, (UK)]: Wellcome Trust Sanger Institute.

Roberts, N.D. (2015). R pkg for Hierarchical Dirichlet Process. https://github.com/nicolaroberts/hdp. Accessed May 2018.

Roberts, S.A. et al. (2013). An APOBEC cytidine deaminase mutagenesis pattern is widespread in human cancers. Nat Genet. *45(9),* 970-6.

Rodriguez-Fraticelli, A.E. et al. (2018). Clonal analysis of lineage fate in native haematopoiesis. Nature. *553(7687)*, 212-216.

Rosales, R.A. et al. (2016). signeR: An empirical Bayesian approach to mutational signature discovery. Bioinformatics. doi: 10.1093/bioinformatics/btw572.

Rothova, M. et al. (2012). Lineage tracing of the endoderm during oral development. Dev Dyn. *241(7)*,1183-1191.

Rozhok, A.L. et al. (2015). A Critical Examination of the "Bad Luck" Explanation of Cancer Risk. Cancer Prev. Res. *8(9)*, 762-764.

218

Rozhok, A.L. and DeGregori, J. (2015). Toward an evolutionary model of cancer: Considering the mechanisms that govern the fate of somatic mutations. Proc. Natl. Acad. Sci. *112(29)*, 8914-8921.

Sabarinathan, R. et al. (2017). The whole-genome panorama of cancer drivers. Preprint at: https://www.biorxiv.org/content/early/2017/09/20/190330**.**

Sabin, F.R. (1920). Studies on the origin of blood vessels and of red blood corpuscles as seen in the living blastoderm of chicks during the second day of incubation. Carnegie Inst. Wash. Pub. n°272, Contrib. Embryol. *9*, 214-262.

Sandberg, A.A. and Hossfeld, D.K. (1970). Chromosomal abnormalities in human neoplasia. Annu. Rev. Med. *21*, 379.

Sanjuan-Pla, A. et al. Platelet-biased stem cells reside at the apex of the haematopoietic stem-cell hierarchy. Nature. *502(7470)*, 232-236.

Sawai, C. M. et al. (2016). Hematopoietic stem cells are the major source of multilineage hematopoiesis in adult animals. Immunity 45, 597-609.

Scheinin, I, et al. (2014). DNA copy number analysis of fresh and formalin-fixed speciments by shallow whole-genome sequencing with identification and exclusion of problematic regions in the genome assembly. Genome Research. *24*, 2022-2032.

Schmitt, M.W. et al. (2012). Detection of ultra-rare mutations by next-generation sequencing. Proc Natl Acad Sci. 109(36):14508-14513.

Schroeder, T. (2010). Hematopoietic stem cell heterogeneity: subtypes, not unpredictable behavior. Cell Stem Cell. *6(3)*, 203-207.

Schwitalla, S. et al. (2013). Intestinal tumorigenesis initiated by dedifferentiation and acquisition of stem-cell-like properties. Cell. *152*, 25-38.

Segditsas, S. and Tomlinson I. (2006). Colorectal cancer and genetic alterations in the Wnt pathway. Oncogene. *25(57)*, 7531-7537.

Shih, C. et al. (1981). Transforming genes of carcinomas and neuroblastomas introduced into mouse fibroblasts. Nature. *290(5803)*, 261-264.

Signer, R. A. J. et al. (2014). Haematopoietic stem cells require a highly regulated protein synthesis rate. Nature. *509*, 49–54.

Siminovitch, L., McCulloch, E.A., Till, J.E. (1963). The distribution of colony-forming cells among spleen colonies. J Cell Physiol. *62*, 327-336.

Simsek, T. et al. (2010). The Distinct Metabolic Profile of Hematopoietic Stem Cells Reflects Their Location in a Hypoxic Niche. Cell Stem Cell. *7*, 380-390.

Snippert, H.J. et al. (2010). Intestinal crypt homeostasis results from neutral competition between symmetrically dividing Lgr5 stem cells. Cell. *143*, 134-144.

Snippert, H.J. et al. (2013). Biased competition between Lgr5 intestinal stem cells driven by oncogenic mutation induces clonal expansion. EMBO Rep. *15(1)*, 62-9.

Solomon, D.A. et al. (2011). Mutational inactivation of STAG2 causes aneuploidy in human cancer. Science. *333(6045)*, 1039-43.

Sottoriva, A. et al. (2015). A big bang model of human colorectal tumor growth. Nat. Genet. *47(3)*, 209-216.

Spangrude, G.L. et al. (1988). Purification and characterization of mouse hematopoietic stem cells. Science. *241*, 58-62.

Stehelin, D. et al. (1976). DNA related to the transforming gene(s) of avian sarcoma viruses is present in normal avian DNA. Nature. *260(5547)*, 170-3.

Stratton, M.R., Campbell, P.J., and Futreal, P.A. (2009). The cancer genome. Nature *458*, 719–724.

Suda, K. et al. (2018). Clonal Expansion and Diversification of Cancer-Associated Mutations in Endometriosis and Normal Endometrium. Cell Rep. *24(7)*, 1777-1789.

Sugarbaker, J.P. Gunderson, L.L., and Wittes, R.E. 1985. Colorectal cancer. In Cancer: Principles and Practices of Oncology. Philadelphia: J. B. Lipincott.

Sun, J. et al. (2014). Clonal dynamics of native haematopoiesis. Nature *514*, 322–327.

Szilard, L. (1959). On the nature of the aging process. Proc. Natl. Acad. Sci. *45(1)*, 30-45.

Szilkriszt, B. et al. (2016). A comprehensive survey of the mutagenic impact of common cancer cytotoxics. Genome Biol. *17*, 99.

Tarabichi et al. (2017). Neutral tumor evolution? Preprint at: https://www.biorxiv.org/content/early/2017/06/30/158006

Tavaré, S. et al. (1997) Inferring coalescence times from DNA sequence data. Genetics. *145(2)*, 505-518.

Taylor, C.F., et al. (2014). Frequent inactivating mutations of STAG2 in bladder cancer are associated with low tumour grade and stage and inversely related to chromosomal copy number changes. Hum Mol Genet. *23(8)*, 1964-1974.

Teh, Y. W. et al. (2006). Hierarchical Dirichlet Processes. J. Am. Stat. Assoc. *101(476)*, 1566–1581.

Tetteh, P.W. et al. (2016). Replacement of Lost Lgr5-Positive Stem Cells through Plasticity of Their Enterocyte-Lineage Daughters. Cell Stem Cell. *18(2)*, 203-13.

Tischler, G. and Leonard, S. (2014). biobambam: tools for read pair collation based algorithms on BAM files. Source Code for Biology and Medicine. *9*, 13.

Thompson, M. et al. (1990). Gastric endocrine cells share a clonal origin with other gut cell lineages. Development. *110(2)*, 477-81.

Tomasetti, C. and Vogelstein, B. (2015). Cancer etiology. Variation in cancer risk among tissues can be explained by the number of stem cell divisions. Science. *347(6217)*, 78-81.

Tomasetti, C. et al. (2014). Only three driver gene mutations are required for the development of lung and colorectal cancers. Proc. Natl. Acad. Sci. *112(1),* 118-123.

Tomasetti, C. et al. (2015). Only three driver gene mutations are required for the development of lung and colorectal cancers. Proc. Natl. Acad. Sci. USA *112*, 118–123.

Tomasetti, C. and Vogelstein, B. (2017). Stem cell divisions, somatic mutations, cancer etiology, and cancer prevention. Science. *355(6331)*, 1330-1334.

Tomlinson, I.P. et al. (1996). The mutation rate and cancer. Proc. Natl. Acad. Sci. *93(25)*, 14800-14803.

Tsao, J.L. et al. (1999). Colorectal adenoma and cancer divergence. Evidence of multilineage progression. Am J Pathol. *154(6)*, 1815-24.

van de Wetering, M. et al. (2002). The beta-catenin/TCF-4 complex imposes a crypt progenitor phenotype on colorectal cancer cells. Cell. *111(2)*, 241-250.

van de Wetering, M. et al. (2015). Prospective derivation of a living organoid biobank of colorectal cancer patients. Cell. *161*, 933–945.

Van Loo, P. et al. (2010). Allele-specific copy number analysis of tumors. Proc. Natl. Acad. Sci. USA *107*, 16910–16915.

Venot, Q. et al. (2018). Targeted therapy in patients with PIK3CA-related overgrowth syndrome. Nature. *558(7711)*, 540-546.

Vermeulen, L. et al. (2013). Defining stem cell dynamics in models of intestinal tumor initiation. Science. *342(6161)*, 995-8.

Viel, A. et al. (2017). A Specific Mutational Signature Associated with DNA 8-Oxoguanine Persistence in MUTYH-defective Colorectal Cancer. EBioMedicine. *20*, 39-49.

Visser, J.W. et al. (1981). Characterization and enrichment of murine hemopoietic stem cells by fluorescence activated cell sorting. Exp Hematol. *9(6)*, 644-55.

Visser, J.W. et al. (1984). Isolation of murine pluripotent hemopoietic stem cells. J. Exp. Med. *159(6)*, 1576-1590.

Vogelstein, B. et al. (2013). Cancer genome landscapes. Science *339*, 1546–1558.

von Hansemann, D. (1980). Über asymmetrische Zellteilung in Epithelkrebsen und deren biologische Bedeutung. Virchow's Arch. Path. Anat. *119*, 299–326.

Wan, J.C.M. et al. (2017). Liquid biopsies come of age: towards implementation of circulating tumour DNA. Nat. Rev. Cancer. *17.4*, 223–238.

Wang, J.C. et al. (1997). Primitive human hematopoietic cells are enriched in cord blood compared with adult bone marrow or mobilized peripheral blood as measured by the quantitative in vivo SCID-repopulating cell assay. Blood. *89(11)*, 3919-24.

Wang, K. et al. (2010). ANNOVAR: functional annotation of genetic variants from high-throughput sequencing data. Nucleic Acids Res. *38(16)*, 598-609.

Wang, Y. and Navin, N.E. (2015). Advances and applications of single-cell sequencing technologies. Mol. Cell. *58(4)*, 598-609.

Warburg O., Wind F. and Negelein E. (1927) The metabolism of tumors in the body. J. Gen. Physiol. *8*, 519–530.

Warr, M. R. et al. (2013). FOXO3A directs a protective autophagy program in haematopoietic stem cells. Nature. *494***,** 323–32.

Weaver, J.M.J. et al. (2014). Ordering of mutations in preinvasive disease stages of esophageal carcinogenesis. Nat Genet. *46(8)*, 837-843.

Welch, J.S. et al. (2012). The origin and evolution of mutations in acute myeloid leukemia. Cell *150*, 264–278.

Weissman, I.L. and Shizuru, J.A. (2008). The origins of the identification and isolation of hematopoietic stem cells, and their capability to induce donor-specific transplantation tolerance and treat autoimmune diseases. Blood. 112(9), 3543-3553.

Werner, B. et al. (2015). Reconstructing the in vivo dynamics of hematopoietic stem cells from telomere length distributions. eLife 4, e08687.

Whitlock, M.C. (2003). Fixation Probability and Time in Subdivided Populations. Genetics *164*, 767–779.

Willems, T. et al. (2017). Genome-wide profiling of heritable and de novo STR variations. Nat Methods. *14*, 590-592.

Williams, M.J. et al. (2016). Identification of neutral tumor evolution across cancer types. Nat Genet. *48(3)*, 238-244.

Wilson, A. et al. (2008). Hematopoietic stem cells reversibly switch from dormancy to self-renewal during homeostasis and repair. Cell. *135(6)*, 1118-1129.

Wimmer, K. and Kratz, C.P. (2010). Constitutional mismatch repair-deficiency syndrome. Haematologica. *95(5)*, 699-701.

Winton, D.J. et al. (1988). A clonal marker induced by mutation in mouse intestinal epithelium. Nature. *333(6172)*, 463-6.

Winton, D.J., and Ponder, B.A. (1990). Stem-cell organization in mouse small intestine. Proc. Biol. Sci. *241*, 13-18.

Wolf, J. et al. (1996). Peripheral blood mononuclear cells of a patient with advanced Hodgkin's lymphoma give rise to permanently growing Hodgkin-Reed Sternberg cells. Blood. *87(8)*, 3418-3428.

Wright, S. (1931). Evolution in Mendelian populations. Genetics *16*, 97–159.

Wu, A.M. et al. (1968). Cytological evidence for a relationship between normal hemotopoietic colony-forming cells and cells of the lymphoid system. J. Exp. Med. *127*, 455-464.

Wu, S. et al. (2016). Substantial contribution of extrinsic risk factors to cancer development. Nature. *529(7584)*, 43-47.

Xie, M. et al. (2014). Age-related mutations associated with clonal hematopoietic expansion and malignancies. Nat. Med. *20*, 1472–1478.

Yates, L.R. and Campbell, P.J. (2012). Evolution of the cancer genome. Nat Rev Genet. *13(11)*, 795-806.

Yoshizato, T. et al. (2015). Somatic Mutations and Clonal Hematopoiesis in Aplastic Anemia. N Engl J Med. *373(1)*, 35-47

Young, A. L., et al. (2016). Clonal haematopoiesis harbouring AML-associated mutations is ubiquitous in healthy adults. Nat. Commun. *7*, 12484.

Zafar, H. et al. (2017). SiFit: inferring tumor trees from single-cell sequencing data under finite-sites models. Genome Biology. *18*, 178.

Zhang, J. et al. (2003). Identification of the haematopoietic stem cell niche and control of the niche size. Nature. *425(6960)*, 836-841.

Zink, F. et al. (2017). Clonal hematopoiesis, with and without candidate driver mutations, is common in the elderly. Blood. *130*, 742–752.

ABBREVIATIONS AND ACRONYMS

ABC	Approximate Bayesian computation
CMP	Common myeloid progenitor
CRC	Colorectal cancer
DBS	Doublet base substitution
GMP	Granulocyte macrophage progenitor
HDP	Hierarchical Dirichlet process
HPC	Haematopoietic progenitor cell
HSC	Haematopoietic stem cell
HSPC	Haematopoietic stem or progenitor cell
ID	Insertion or deletion (indel)
LCM	Laser capture microdissection
LOH	Loss of heterozygosity
MEP	Megakaryocyte erythroid progenitor
MRCA	Most recent common ancestor
NNMF	Non-negative matrix factorisation
PCAWG	Pan cancer analysis of whole genomes
SBS	Single base substitution
SV	Structural variant
TCGA	The Cancer Genome Atlas
VAF	Variant allele fraction
WGS	Whole genome sequencing

APPENDIX A

MYELOID LEUKAEMIA AND CLONAL HAEMATOPOIESIS MUTATION HOTSPOT LOCI INCLUDED IN BAIT-SET FOR TARGETED SEQUENCING OF PERIPHERAL BLOOD

Gene	Mutation	Position
NRAS	Q61	1:115256530-115256531
NRAS	G12	1:115258748-115258748
CD58	G210	1:117078587-117078587
GNB1	K57	1:1747229-1747229
MPL	W515	1:43815009-43815009
CBL	Y371	11:119148892-119148892
CBL	C381	11:119148922-119148922
CBL	R420	11:119149251-119149251
PTPN11	A72	12:112888199-112888200
PTPN11	G503	12:112926888-112926888
KRAS	G12	12:25398285-25398285
FLT3	D835	13:28592642-28592642
FLT3	N676	13:28602341-28602341
IDH2	R172	15:90631839-90631839
IDH2	R140	15:90631935-90631935
CREBBP	R1446	16:3788618-3788618
STAT3	D661	17:40474420-40474420
STAT3	Y640	17:40474482-40474482
PPM1D	S468	17:58740498-58740498
PPM1D	Q520	17:58740655-58740655
CD79B	Y196	17:62006799-62006799
SRSF2	P95	17:74732960-74732960
TP53	R273C	17:7577121-7577121
TP53	R248Q	17:7577539-7577540
TP53	R175G	17:7578407-7578407
TP53	H168R	17:7578428-7578428
RPS15	P131	19:1440415-1440415
JAK3	R657	19:17945970-17945970
JAK3	A572	19:17948010-17948010
JAK3	M511	19:17949108-17949108

CNOT3	E20	19:54646887-54646887
CNOT3	E70	19:54647435-54647435
SF3B1	K700	2:198266834-198266834
SF3B1	K666	2:198267361-198267361
IDH1	R132	2:209113113-209113113
DNMT3A	P904	2:25457177-25457177
DNMT3A	R882	2:25457243-25457243
DNMT3A	W860	2:25458595-25458595
DNMT3A	R771	2:25463182-25463182
DNMT3A	S770	2:25463184-25463184
DNMT3A	F752	2:25463239-25463239
DNMT3A	R749	2:25463248-25463248
DNMT3A	R736	2:25463287-25463287
DNMT3A	F732	2:25463299-25463299
DNMT3A	R729	2:25463308-25463308
DNMT3A	S714	2:25463541-25463541
DNMT3A	I705	2:25463569-25463569
DNMT3A	G699	2:25463587-25463587
DNMT3A	G685	2:25464460-25464460
DNMT3A	R635	2:25466800-25466800
DNMT3A	R598	2:25467083-25467083
DNMT3A	L547	2:25467436-25467436
DNMT3A	G543	2:25467449-25467449
DNMT3A	P385	2:25469614-25469614
DNMT3A	W330	2:25470484-25470484
DNMT3A	R326	2:25470498-25470498
DNMT3A	R320	2:25470516-25470516
DNMT3A	P307	2:25470554-25470554
ASXL1	R404	20:31021211-31021211
ASXL1	R417	20:31021250-31021250
ASXL1	Y591	20:31022288-31022288
ASXL1	G642	20:31022439-31022439
ASXL1	R693	20:31022592-31022592
GNAS	R201	20:57484422-57484422
U2AF1	Q157	21:44514777-44514777
U2AF1	S34	21:44524456-44524456
SF3A1	E75	22:30742471-30742471
EP300	D1399	22:41565529-41565529
STAG2	R305	23:123184056-123184056
BCORL1	R784	23:129149098-129149098

PHF6	R274	23:133549137-133549137
STAG1	A116	3:136287655-136287655
PIK3CA	H1047	3:178952086-178952086
MYD88	L265	3:38182641-38182641
TET2	R550	4:106156747-106156747
TET2	Q916	4:106157845-106157845
TET2	W954	4:106157961-106157961
TET2	R1261	4:106164914-106164914
TET2	G1288	4:106180835-106180835
TET2	Y1337	4:106182972-106182972
TET2	A1381	4:106190863-106190863
TET2	R1465	4:106193931-106193931
TET2	S1486	4:106193995-106193995
TET2	I1873	4:106197285-106197285
FBXW7	R505	4:153247289-153247289
FBXW7	R479	4:153247367-153247367
FBXW7	R465	4:153249385-153249385
KIT	V559	4:55593610-55593610
KIT	D816	4:55599322-55599322
NPM1	L287	5:170837543-170837545
BRAF	V600	7:140453137-140453137
EZH2	Y646	7:148508728-148508728
CARD11	R1092	7:2946461-2946461
CARD11	E626	7:2963931-2963931
CARD11	D401	7:2976811-2976811
CARD11	D230	7:2979559-2979559
RAD21	A480	8:117864217-117864217
JAK2	V617	9:5073772-5073772

APPENDIX B

COLORECTAL CANCER GENES INCLUDED IN BAIT-SET FOR TARGETED SEQUENCING OF COLONIC CRYPTS

Gene
ACVR1B
ACVR2A
APC
ARID1A
ARID2
ASXL1
ATM
ATR
ATRX
AXIN2
BCOR
BRAF
BRCA2
CARD11
CBL
CDC73
CDH1
CDK12
CREBBP
CTNNB1
EGFR
ELF3
EP300
ERBB2
ERBB3
EZH2
AMER1
FBXW7
FGFR1
FGFR2

FLT1
GATA3
GNAS
H3F3A
H3F3B
JAK2
KDM6A
KDR
KIT
KRAS
MET
MGA
MLH1
KMT2D
KMT2C
MSH2
MSH6
NF1
NF2
NRAS
PDGFRA
PIK3CA
PIK3R1
POLE
PTCH1
PTEN
PTPN11
RB1
RBM10
RET
RNF43
SETD2
SMAD2
SMAD4
SOX9
STAG2
STK11
TBX3
TCF7L2
TGFBR2

TP53
TSHR
VHL
WT1
AKT1
AXIN1
B2M
CDKN1B
CSF1R
CUX1
GRIN2A
MAP2K1
MAP2K4
MAX
QKI
RAD21
ROBO2
SRC
TBL1XR1
UBR5

APPENDIX C

CHARACTERISTICS OF INFORMATIVE CRYPTS

crypt	site	patient	sex	age	VAF	Depth	cancer.pt
HLS_1C_30_B5	Right	HLS	female	60	0.44	12	No
HLS_1C_30_D5	Right	HLS	female	60	0.41	10	No
HLS_1C_30_G5	Right	HLS	female	60	0.43	12	No
HLS_1C_30_H5	Right	HLS	female	60	0.45	9	No
HLS_2C_30_D6	Right	HLS	female	60	0.46	8	No
HLS_2C_30_E6	Right	HLS	female	60	0.45	9	No
PD28690bx_2_a5	Right	PD28690	male	79	0.4	17	No
PD28690bx_2_d5	Right	PD28690	male	79	0.44	18	No
PD28690cb_2_g1	Transverse	PD28690	male	79	0.38	17	No
PD28690cb_2_g3	Transverse	PD28690	male	79	0.44	16	No
PD28690cb_2_h1	Transverse	PD28690	male	79	0.33	17	No
PD28690cb_2_h3	Transverse	PD28690	male	79	0.42	19	No
PD34199a_c17	Right	PD34199	male	67	0.42	19	No
PD34199a_c22	Right	PD34199	male	67	0.32	16	No
PD34199a_c27	Right	PD34199	male	67	0.39	17	No
PD34199a_c28	Right	PD34199	male	67	0.31	18	No
PD34199a_c3	Right	PD34199	male	67	0.38	19	No
PD34199a_c31	Right	PD34199	male	67	0.36	18	No
PD34199a_c5	Right	PD34199	male	67	0.33	19	No
PD34199a_c7	Right	PD34199	male	67	0.38	19	No
PD34199a_s13	Left	PD34199	male	67	0.4	18	No
PD34199a_s16	Left	PD34199	male	67	0.32	18	No
PD34199a_s2	Left	PD34199	male	67	0.38	18	No
PD34199a_s27	Left	PD34199	male	67	0.32	17	No
PD34199a_s31	Left	PD34199	male	67	0.4	16	No
PD34199a_s4	Left	PD34199	male	67	0.38	18	No
PD34199a_t14	Transverse	PD34199	male	67	0.36	20	No
PD34199a_t16	Transverse	PD34199	male	67	0.42	17	No
PD34199a_t18	Transverse	PD34199	male	67	0.41	18	No
PD34199a_t22	Transverse	PD34199	male	67	0.37	19	No
PD34199a_t24	Transverse	PD34199	male	67	0.42	17	No
PD34199a_t28	Transverse	PD34199	male	67	0.35	17	No
PD34199a_t31	Transverse	PD34199	male	67	0.4	18	No

PD34200a_c10	Right	PD34200	male	54	0.43	17	No
PD34200a_c14	Right	PD34200	male	54	0.47	14	No
PD34200a_c15	Right	PD34200	male	54	0.46	16	No
PD34200a_c17	Right	PD34200	male	54	0.46	17	No
PD34200a_c19	Right	PD34200	male	54	0.43	16	No
PD34200a_c23	Right	PD34200	male	54	0.45	17	No
PD34200a_c24	Right	PD34200	male	54	0.46	18	No
PD34200a_c26	Right	PD34200	male	54	0.42	19	No
PD34200a_c31	Right	PD34200	male	54	0.44	19	No
PD34200a_c32	Right	PD34200	male	54	0.44	18	No
PD34200a_c33	Right	PD34200	male	54	0.47	17.5	No
PD34200a_c35	Right	PD34200	male	54	0.46	15	No
PD34200a_c37	Right	PD34200	male	54	0.45	17	No
PD34200a_c41	Right	PD34200	male	54	0.39	17	No
PD34200a_c42	Right	PD34200	male	54	0.47	17	No
PD34200a_c43	Right	PD34200	male	54	0.44	17	No
PD34200a_c47	Right	PD34200	male	54	0.43	17	No
PD34200a_c5	Right	PD34200	male	54	0.44	17	No
PD34200a_c7	Right	PD34200	male	54	0.42	16	No
PD34200a_c8	Right	PD34200	male	54	0.42	15	No
PD34200a_i21	Ileum	PD34200	male	54	0.38	16	No
PD34200a_i22	Ileum	PD34200	male	54	0.43	15	No
PD34200a_i33	Ileum	PD34200	male	54	0.46	16	No
PD34200a_i36	Ileum	PD34200	male	54	0.43	17	No
PD34200a_i54	Ileum	PD34200	male	54	0.3	17	No
PD34200a_s10	Left	PD34200	male	54	0.41	17	No
PD34200a_s17	Left	PD34200	male	54	0.35	21	No
PD34200a_s18	Left	PD34200	male	54	0.33	20	No
PD34200a_s2	Left	PD34200	male	54	0.36	16	No
PD34200a_s20	Left	PD34200	male	54	0.38	18	No
PD34200a_s22	Left	PD34200	male	54	0.35	19	No
PD34200a_s27	Left	PD34200	male	54	0.33	19	No
PD34200a_s29	Left	PD34200	male	54	0.38	20	No
PD34200a_s3	Left	PD34200	male	54	0.4	18	No
PD34200a_s4	Left	PD34200	male	54	0.37	17	No
PD34200a_s8	Left	PD34200	male	54	0.41	18	No
PD34200a_t12	Transverse	PD34200	male	54	0.41	15	No
PD34200a_t14	Transverse	PD34200	male	54	0.41	20	No
PD34200a_t15	Transverse	PD34200	male	54	0.42	20	No

PD34200a_t16	Transverse	PD34200	male	54	0.43	18	No
PD34200a_t20	Transverse	PD34200	male	54	0.33	21	No
PD34200a_t23	Transverse	PD34200	male	54	0.41	20	No
PD34200a_t27	Transverse	PD34200	male	54	0.41	17	No
PD34200a_t28	Transverse	PD34200	male	54	0.38	19	No
PD34200a_t29	Transverse	PD34200	male	54	0.39	18	No
PD34200a_t3	Transverse	PD34200	male	54	0.4	18	No
PD34200a_t33	Transverse	PD34200	male	54	0.42	21	No
PD34200a_t37	Transverse	PD34200	male	54	0.39	21	No
PD34200a_t6	Transverse	PD34200	male	54	0.41	18	No
PD34201a_c1	Right	PD34201	male	38	0.42	15	No
PD34201a_c10	Right	PD34201	male	38	0.4	17	No
PD34201a_c18	Right	PD34201	male	38	0.41	16	No
PD34201a_c20	Right	PD34201	male	38	0.4	15	No
PD34201a_c22	Right	PD34201	male	38	0.4	18	No
PD34201a_c25	Right	PD34201	male	38	0.41	17	No
PD34201a_c27	Right	PD34201	male	38	0.35	16	No
PD34201a_c32	Right	PD34201	male	38	0.36	16	No
PD34201a_c35	Right	PD34201	male	38	0.41	16	No
PD34201a_c38	Right	PD34201	male	38	0.39	15	No
PD34201a_c4	Right	PD34201	male	38	0.37	17	No
PD34201a_c40	Right	PD34201	male	38	0.38	18	No
PD34201a_c41	Right	PD34201	male	38	0.35	16	No
PD34201a_c42	Right	PD34201	male	38	0.42	16	No
PD34201a_c44	Right	PD34201	male	38	0.44	17	No
PD34201a_c48	Right	PD34201	male	38	0.36	18	No
PD34201a_s1	Left	PD34201	male	38	0.36	15	No
PD34201a_s10	Left	PD34201	male	38	0.36	15	No
PD34201a_s11	Left	PD34201	male	38	0.35	15	No
PD34201a_s15	Left	PD34201	male	38	0.36	19	No
PD34201a_s18	Left	PD34201	male	38	0.41	17	No
PD34201a_s2	Left	PD34201	male	38	0.4	17	No
PD34201a_s20	Left	PD34201	male	38	0.38	16	No
PD34201a_s22	Left	PD34201	male	38	0.32	19	No
PD34201a_s24	Left	PD34201	male	38	0.36	17	No
PD34201a_s28	Left	PD34201	male	38	0.39	14	No
PD34201a_s31	Left	PD34201	male	38	0.39	24	No
PD34201a_s36	Left	PD34201	male	38	0.38	17	No
PD34201a_s4	Left	PD34201	male	38	0.36	18	No

PD34201a_s42	Left	PD34201	male	38	0.45	16	No
PD34201a_s43	Left	PD34201	male	38	0.39	18	No
PD34201a_s44	Left	PD34201	male	38	0.37	22	No
PD34201a_s45	Left	PD34201	male	38	0.39	20	No
PD34201a_s47	Left	PD34201	male	38	0.42	19	No
PD34201a_s6	Left	PD34201	male	38	0.43	17	No
PD34201a_s7	Left	PD34201	male	38	0.39	20	No
PD34201a_t18	Transverse	PD34201	male	38	0.42	17	No
PD34201a_t19	Transverse	PD34201	male	38	0.46	15	No
PD34201a_t20	Transverse	PD34201	male	38	0.43	16	No
PD34201a_t21	Transverse	PD34201	male	38	0.41	18	No
PD34201a_t22	Transverse	PD34201	male	38	0.43	18	No
PD34201a_t23	Transverse	PD34201	male	38	0.41	18	No
PD34201a_t28	Transverse	PD34201	male	38	0.42	19	No
PD34201a_t3	Transverse	PD34201	male	38	0.45	17	No
PD34201a_t30	Transverse	PD34201	male	38	0.41	19	No
PD34201a_t31	Transverse	PD34201	male	38	0.43	17	No
PD34201a_t32	Transverse	PD34201	male	38	0.4	14	No
PD34201a_t34	Transverse	PD34201	male	38	0.41	19	No
PD34201a_t35	Transverse	PD34201	male	38	0.43	18	No
PD34201a_t39	Transverse	PD34201	male	38	0.4	21	No
PD34201a_t41	Transverse	PD34201	male	38	0.4	18	No
PD34201a_t42	Transverse	PD34201	male	38	0.43	17	No
PD34201a_t44	Transverse	PD34201	male	38	0.42	19	No
PD34201a_t45	Transverse	PD34201	male	38	0.4	18	No
PD34201a_t46	Transverse	PD34201	male	38	0.41	19	No
PD34201a_t8	Transverse	PD34201	male	38	0.44	17	No
PD34202a_s12	Left	PD34202	male	13	0.35	18	No
PD34202a_s15	Left	PD34202	male	13	0.31	20	No
PD34202a_s27	Left	PD34202	male	13	0.36	12	No
PD34202a_s41	Left	PD34202	male	13	0.38	13	No
PD34202a_s43	Left	PD34202	male	13	0.3	17	No
PD34202a_s9	Left	PD34202	male	13	0.38	16	No
PD36719a_lo005	Left	OO13	male	60	0.42	12	No
PD36719a_lo008	Left	OO13	male	60	0.38	15	No
PD36719a_lo018	Right	OO13	male	60	0.42	20	No
PD36719a_lo022	Right	OO13	male	60	0.42	21	No
PD36719a_lo034	Transverse	OO13	male	60	0.47	13	No
PD36719a_lo048	Transverse	OO13	male	60	0.37	19	No

PD36719a_lo049	Left	OO13	male	60	0.385	13	No
PD36719a_lo066	Right	OO13	male	60	0.45	13	No
PD36719a_lo072	Right	OO13	male	60	0.44	12	No
PD36813aa	Transverse	OO94	female	66	0.46	12	No
PD36813aa11	Transverse	OO94	female	66	0.45	14	No
PD36813aa12	Transverse	OO94	female	66	0.36	17	No
PD36813aa16	Transverse	OO94	female	66	0.35	16	No
PD36813aa2	Transverse	OO94	female	66	0.4	14	No
PD36813aa7	Transverse	OO94	female	66	0.33	15	No
PD36813ab10	Transverse	OO99	male	66	0.39	18	No
PD36813ab4	Transverse	OO99	male	66	0.41	15	No
PD36813ab5	Transverse	OO99	male	66	0.47	15	No
PD36813ac15	Transverse	O251	female	71	0.43	18	No
PD36813ac18	Transverse	O251	female	71	0.39	18	No
PD36813ac2	Left	O251	female	71	0.43	31	No
PD36813ac4	Left	O251	female	71	0.43	28	No
PD36813ac7	Left	O251	female	71	0.4	31	No
PD36813ad17	Transverse	O189	female	62	0.44	22	No
PD36813ad2	Left	O189	female	62	0.38	33	No
PD36813ad8	Transverse	O189	female	62	0.44	17	No
PD36813ad9	Transverse	O189	female	62	0.43	18	No
PD36813ae	Transverse	O208	female	68	0.36	15	No
PD36813ae3	Transverse	O208	female	68	0.38	16	No
PD36813ae7	Transverse	O208	female	68	0.37	17	No
PD36813b4	Right	O251	female	71	0.45	12	No
PD36813b5	Right	O251	female	71	0.46	13	No
PD36813b6	Right	O251	female	71	0.44	13	No
PD36813c	Right	O208	female	68	0.44	13	No
PD36813c2	Right	O208	female	68	0.45	16	No
PD36813c3	Right	O208	female	68	0.45	13	No
PD36813d3	Right	O189	female	62	0.47	14	No
PD36813d5	Right	O189	female	62	0.47	14	No
PD36813e	Right	O187	female	60	0.47	14	No
PD36813e5	Right	O187	female	60	0.43	14	No
PD36813e6	Right	O187	female	60	0.47	14	No
PD36813g2	Right	O162	male	72	0.39	15	No
PD36813g4	Right	O162	male	72	0.42	18	No
PD36813h	Right	O132	female	67	0.47	14	No
PD36813h3	Right	O132	female	67	0.45	14	No

PD36813h5	Right	O132	female	67	0.47	15	No
PD36813i	Right	OO99	male	66	0.45	16	No
PD36813i4	Right	OO99	male	66	0.44	13	No
PD36813i6	Right	OO99	male	66	0.48	14	No
PD36813j	Right	OO94	female	66	0.46	12	No
PD36813j3	Right	OO94	female	66	0.44	12	No
PD36813j6	Right	OO94	female	66	0.5	11	No
PD36813k	Right	OO82	female	64	0.42	16	No
PD36813k3	Right	OO82	female	64	0.41	17	No
PD36813k5	Right	OO82	female	64	0.44	16	No
PD36813l	Right	OO35	male	66	0.41	17	No
PD36813l7	Right	OO35	male	66	0.46	13	No
PD36813l8	Right	OO35	male	66	0.47	15	No
PD36813m3	Right	OO23	male	66	0.46	14	No
PD36813m5	Right	OO23	male	66	0.44	15	No
PD36813m8	Right	OO23	male	66	0.33	17	No
PD36813n	Left	OO23	male	66	0.44	16	No
PD36813o2	Left	OO82	female	64	0.3	17	No
PD36813o9	Left	OO82	female	64	0.4	17	No
PD36813p	Left	OO94	female	66	0.43	17	No
PD36813q11	Left	OO99	male	66	0.35	17	No
PD36813q14	Left	OO99	male	66	0.38	17	No
PD36813r12	Left	O132	female	67	0.46	13	No
PD36813r14	Left	O132	female	67	0.44	16	No
PD36813s12	Left	O162	male	72	0.31	19	No
PD36813t20	Left	O187	female	60	0.39	15	No
PD36813u2	Transverse	O132	female	67	0.43	20	No
PD36813v4	Transverse	O162	male	72	0.32	21	No
PD36813x15	Transverse	O187	female	60	0.44	18	No
PD36813x16	Transverse	O187	female	60	0.46	18	No
PD36813y13	Transverse	OO39	male	68	0.42	16	No
PD36813y15	Transverse	OO39	male	68	0.47	16	No
PD36813y6	Transverse	OO39	male	68	0.47	14	No
PD36813y7	Transverse	OO39	male	68	0.5	12	No
PD36813z15	Transverse	OO82	female	64	0.42	16	No
PD36813z4	Transverse	OO82	female	64	0.4	15	No
PD36813z5	Transverse	OO82	female	64	0.43	15	No
PD36813z6	Transverse	OO82	female	64	0.45	15	No
PD36813z7	Transverse	OO82	female	64	0.35	15	No

PD36813z8	Transverse	OO82	female	64	0.33	16	No
PD36814b2	Right	OOO7	male	64	0.38	11	Yes
PD36814b7	Right	OOO7	male	64	0.46	15	Yes
PD36814c2	Right	OOO8	female	64	0.45	13	Yes
PD36814c3	Right	OOO8	female	64	0.46	13	Yes
PD36814c4	Right	OOO8	female	64	0.45	11	Yes
PD36814d5	Right	OO10	female	60	0.45	13	Yes
PD36814d6	Right	OO10	female	60	0.43	13	Yes
PD36814d7	Right	OO10	female	60	0.45	11	Yes
PD36814d8	Right	OO10	female	60	0.46	13	Yes
PD36814e5	Right	O174	female	66	0.41	14	Yes
PD36814e6	Right	O174	female	66	0.4	14	Yes
PD36814e7	Right	O174	female	66	0.45	13	Yes
PD36814e8	Right	O174	female	66	0.45	13	Yes
PD36814f4	Right	O176	female	68	0.46	13	Yes
PD36814f8	Right	O176	female	68	0.45	14	Yes
PD36814g4	Right	O193	female	71	0.43	12	Yes
PD36814g7	Right	O193	female	71	0.47	14	Yes
PD36814g8	Right	O193	female	71	0.45	14	Yes
PD36814h2	Right	O225	male	66	0.5	9	Yes
PD36814h4	Right	O225	male	66	0.45	12	Yes
PD36814i	Right	O232	male	62	0.47	13	Yes
PD36814i3	Right	O232	male	62	0.45	12	Yes
PD36814i4	Right	O232	male	62	0.46	13	Yes
PD36814k2	Right	O244	male	68	0.45	16	Yes
PD36814k6	Right	O244	male	68	0.47	14	Yes
PD36814k7	Right	O244	male	68	0.47	14	Yes
PD36814k8	Right	O244	male	68	0.5	14	Yes
PD36814l2	Right	O333	female	68	0.5	13	Yes
PD36814l4	Right	O333	female	68	0.47	10	Yes
PD36814l7	Right	O333	female	68	0.47	13	Yes
PD36814l8	Right	O333	female	68	0.5	13	Yes
PD36814m5	Right	O337	male	66	0.45	10	Yes
PD36814m8	Right	O337	male	66	0.45	11	Yes
PD37226b5	Left	OOO8	female	64	0.4	12	Yes
PD37226b6	Left	OOO8	female	64	0.39	12	Yes
PD37226c5	Left	OO10	female	60	0.45	10	Yes
PD37226c6	Left	OO10	female	60	0.46	13	Yes
PD37226c7	Left	OO10	female	60	0.45	15	Yes

PD37226c8	Left	OO10	female	60	0.44	12	Yes
PD37226d4	Left	O174	female	66	0.36	12	Yes
PD37226d5	Left	O174	female	66	0.33	11	Yes
PD37226d6	Left	O174	female	66	0.38	11	Yes
PD37226e4	Left	O176	female	68	0.39	12	Yes
PD37226e5	Left	O176	female	68	0.42	12	Yes
PD37226f5	Left	O193	female	71	0.42	11	Yes
PD37226f6	Left	O193	female	71	0.38	11	Yes
PD37226f7	Left	O193	female	71	0.33	14	Yes
PD37226g5	Left	O225	male	66	0.41	10	Yes
PD37226g6	Left	O225	male	66	0.44	12	Yes
PD37226g7	Left	O225	male	66	0.42	12	Yes
PD37226h4	Left	O232	male	62	0.47	12	Yes
PD37226h7	Left	O232	male	62	0.43	16	Yes
PD37226h8	Left	O232	male	62	0.38	15	Yes
PD37226i5	Left	O333	female	68	0.45	11	Yes
PD37226i6	Left	O333	female	68	0.42	12	Yes
PD37226i7	Left	O333	female	68	0.43	13	Yes
PD37226i8	Left	O333	female	68	0.38	12	Yes
PD37226j4	Left	O337	male	66	0.31	12	Yes
PD37226j5	Left	O337	male	66	0.38	13	Yes
PD37226j6	Left	O337	male	66	0.45	11	Yes
PD37226j7	Left	O337	male	66	0.47	13	Yes
PD37226j8	Left	O337	male	66	0.38	16	Yes
PD37226k2	Left	O340	male	72	0.4	15	Yes
PD37226k5	Left	O340	male	72	0.37	15	Yes
PD37226k6	Left	O340	male	72	0.38	14	Yes
PD37226k8	Left	O340	male	72	0.37	17	Yes
PD37226l5	Left	O395	male	60	0.36	14	Yes
PD37226l6	Left	O395	male	60	0.42	14	Yes
PD37226l7	Left	O395	male	60	0.4	12	Yes
PD37226l8	Left	O395	male	60	0.4	14	Yes
PD37226m5	Left	O401	male	62	0.41	13	Yes
PD37226m7	Left	O401	male	62	0.42	14	Yes
PD37226m8	Left	O401	male	62	0.4	16	Yes
PD37226o4	Transverse	O337	male	66	0.36	26	Yes
PD37226o5	Transverse	O337	male	66	0.44	23	Yes
PD37226p7	Left	O337	male	66	0.48	19	Yes
PD37265b	Transverse	OOO7	male	64	0.47	14	Yes

PD37265b7	Transverse	OOO7	male	64	0.5	11	Yes
PD37265c2	Transverse	OOO8	female	64	0.42	11	Yes
PD37265c3	Transverse	OOO8	female	64	0.43	13	Yes
PD37265c5	Transverse	OOO8	female	64	0.45	16	Yes
PD37265d	Transverse	OO10	female	60	0.44	14	Yes
PD37265d5	Transverse	OO10	female	60	0.44	12	Yes
PD37265d8	Transverse	OO10	female	60	0.5	11	Yes
PD37265e	Transverse	O174	female	66	0.43	14	Yes
PD37265e5	Transverse	O174	female	66	0.45	13	Yes
PD37265f	Transverse	O176	female	68	0.4	13	Yes
PD37265f2	Transverse	O176	female	68	0.45	13	Yes
PD37265f4	Transverse	O176	female	68	0.46	12	Yes
PD37265g6	Transverse	O193	female	71	0.42	14	Yes
PD37265g7	Transverse	O193	female	71	0.38	14	Yes
PD37265g8	Transverse	O193	female	71	0.42	12	Yes
PD37265h5	Transverse	O225	male	66	0.41	14	Yes
PD37265h6	Transverse	O225	male	66	0.44	14	Yes
PD37265h8	Transverse	O225	male	66	0.45	13	Yes
PD37265i2	Transverse	O232	male	62	0.46	10	Yes
PD37265i7	Transverse	O232	male	62	0.43	12	Yes
PD37265k	Transverse	O244	male	68	0.46	10	Yes
PD37265k2	Transverse	O244	male	68	0.45	10	Yes
PD37265k3	Transverse	O244	male	68	0.43	11	Yes
PD37265l2	Transverse	O333	female	68	0.46	12	Yes
PD37265l6	Transverse	O333	female	68	0.46	11	Yes
PD37265m3	Transverse	O337	male	66	0.37	12	Yes
PD37265m4	Transverse	O337	male	66	0.42	12	Yes
PD37266b_lo0004	Left	PD37266	female	67	0.35	20	No
PD37266b_lo0005	Left	PD37266	female	67	0.42	19	No
PD37266b_lo0006	Left	PD37266	female	67	0.38	18	No
PD37266b_lo0007	Left	PD37266	female	67	0.31	20	No
PD37266b_lo0008	Left	PD37266	female	67	0.39	22	No
PD37266b_lo0017	Left	PD37266	female	67	0.46	16	No
PD37266b_lo0023	Left	PD37266	female	67	0.38	13	No
PD37266c_lo0010	Ileum	PD37266	female	67	0.32	18	No
PD37266c_lo0012	Ileum	PD37266	female	67	0.41	19	No
PD37266c_lo0014	Ileum	PD37266	female	67	0.36	18	No
PD37266c_lo0016	Ileum	PD37266	female	67	0.43	19	No
PD37266c_lo0018	Ileum	PD37266	female	67	0.47	13	No

PD37266c_lo0023	Ileum	PD37266	female	67	0.47	7	No
PD37266d_lo0002	Right	PD37266	female	67	0.5	8	No
PD37266d_lo0010	Right	PD37266	female	67	0.47	11	No
PD37266d_lo0012	Right	PD37266	female	67	0.4	10	No
PD37266d_lo0014	Right	PD37266	female	67	0.43	11	No
PD37266e_lo0008	Transverse	PD37266	female	67	0.45	19	No
PD37266e_lo0010	Transverse	PD37266	female	67	0.46	7	No
PD37449c_lo0001	Transverse	PD37449	female	36	0.46	17	No
PD37449c_lo0009	Transverse	PD37449	female	36	0.44	18	No
PD37449c_lo0014	Transverse	PD37449	female	36	0.45	19	No
PD37449c_lo0017	Transverse	PD37449	female	36	0.43	16	No
PD37449d_lo0001	Left	PD37449	female	36	0.46	17	No
PD37449d_lo0002	Left	PD37449	female	36	0.44	17	No
PD37449d_lo0003	Left	PD37449	female	36	0.42	19	No
PD37449d_lo0006	Left	PD37449	female	36	0.36	19	No
PD37449d_lo0011	Left	PD37449	female	36	0.43	12	No
PD37449d_lo0012	Left	PD37449	female	36	0.39	9	No
PD37449d_lo0017	Left	PD37449	female	36	0.44	16	No
PD37449e_lo0001	Right	PD37449	female	36	0.45	8	No
PD37449e_lo0002	Right	PD37449	female	36	0.44	18	No
PD37449e_lo0003	Right	PD37449	female	36	0.44	19	No
PD37449e_lo0004	Right	PD37449	female	36	0.41	20	No
PD37449e_lo0006	Right	PD37449	female	36	0.46	7	No
PD37449e_lo0008	Right	PD37449	female	36	0.45	20	No
PD37449e_lo0010	Right	PD37449	female	36	0.42	18	No
PD37449e_lo0014	Right	PD37449	female	36	0.44	20	No
PD37449e_lo0015	Right	PD37449	female	36	0.4	8	No
PD37449f_lo0002	Ileum	PD37449	female	36	0.45	17	No
PD37449f_lo0006	Ileum	PD37449	female	36	0.38	17	No
PD37449f_lo0008	Ileum	PD37449	female	36	0.35	13	No
PD37449f_lo0010	Ileum	PD37449	female	36	0.39	17	No
PD37449f_lo0011	Ileum	PD37449	female	36	0.42	13	No
PD37449f_lo0016	Ileum	PD37449	female	36	0.44	16	No
PD37449f_lo0017	Ileum	PD37449	female	36	0.42	16	No
PD37449f_lo0019	Ileum	PD37449	female	36	0.43	16	No
PD37449f_lo0022	Ileum	PD37449	female	36	0.44	19	No
PD37449f_lo0023	Ileum	PD37449	female	36	0.43	15	No
PD37457b4	Right	O395	male	60	0.31	10	Yes
PD37457c4	Right	O401	male	62	0.46	12	Yes

PD37457e6	Right	O340	male	72	0.45	14	Yes
PD37457f4	Transverse	O340	male	72	0.43	13	Yes
PD37457f7	Transverse	O340	male	72	0.4	9	Yes
PD37457g5	Transverse	O395	male	60	0.45	13	Yes
PD37457h2	Transverse	O401	male	62	0.4	10	Yes
PD37457h5	Transverse	O401	male	62	0.5	12	Yes
PD37457i4	Transverse	O438	female	62	0.47	13	Yes
PD37457i5	Transverse	O438	female	62	0.5	8	Yes
PD37457j7	Left	O438	female	62	0.33	17	Yes
PD37457j8	Left	O438	female	62	0.47	12	Yes
PD37509b8	Right	300	female	15	0.48	17	No
PD37509b9	Right	300	female	15	0.47	15	No
PD37510b4	Left	312	female	11	0.42	8	No
PD37510b7	Left	312	female	11	0.4	9	No
PD37510b8	Left	312	female	11	0.37	12	No
PD37512a11	Left	O311	male	68	0.44	18	No
PD37512a4	Transverse	O311	male	68	0.405	18	No
PD37512a5	Transverse	O311	male	68	0.44	17	No
PD37512b7	Right	O311	male	68	0.47	12	No
PD37513a13	Left	O325	male	62	0.33	14	No
PD37513a2	Transverse	O325	male	62	0.45	16	No
PD37513a20	Left	O325	male	62	0.45	14	No
PD37513a21	Left	O325	male	62	0.45	14	No
PD37513a3	Transverse	O325	male	62	0.44	17	No
PD37513b5	Right	O325	male	62	0.43	7	No
PD37513b7	Right	O325	male	62	0.44	12	No
PD37513b8	Right	O325	male	62	0.45	14	No
PD37514a	Transverse	O342	male	62	0.4	17	No
PD37514a2	Transverse	O342	male	62	0.43	16	No
PD37514a3	Transverse	O342	male	62	0.41	16	No
PD37514b3	Right	O342	male	62	0.47	11	No
PD37514b7	Right	O342	male	62	0.41	15	No
PD37514b8	Right	O342	male	62	0.45	9	No
PD37590b_lo0010	Right	PD37590	Female	55	0.45	25	No
PD37590b_lo0015	Right	PD37590	Female	55	0.43	23	No
PD37590b_lo0018	Right	PD37590	Female	55	0.43	32	No
PD37590b_lo0030	Right	PD37590	Female	55	0.34	28	No
PD37590b_lo0033	Right	PD37590	Female	55	0.45	16	No
PD37590b_lo0041	Right	PD37590	Female	55	0.41	22	No

PD37590b_lo0051	Right	PD37590	Female	55	0.42	25	No
PD37590b_lo0052	Right	PD37590	Female	55	0.36	25	No
PD37590b_lo0053	Right	PD37590	Female	55	0.44	23	No
PD37590b_lo0058	Right	PD37590	Female	55	0.41	24	No
PD37590b_lo0074	Right	PD37590	Female	55	0.45	25	No
PD37590b_lo0075	Right	PD37590	Female	55	0.42	31	No
PD37590b_lo0090	Right	PD37590	Female	55	0.43	33	No
PD37590b_lo0092	Right	PD37590	Female	55	0.45	15	No
PD37590b_lo0107	Right	PD37590	Female	55	0.39	32	No
PD37590b_lo0116	Transverse	PD37590	Female	55	0.43	23	No
PD37590b_lo0117	Transverse	PD37590	Female	55	0.42	18	No
PD37590c_lo0001	Transverse	PD37590	Female	55	0.41	27	No
PD37590c_lo0005	Transverse	PD37590	Female	55	0.41	27	No
PD37590c_lo0009	Transverse	PD37590	Female	55	0.43	20	No
PD37590c_lo0019	Right	PD37590	Female	55	0.46	25	No
PD37590c_lo0021	Transverse	PD37590	Female	55	0.45	26	No
PD37590c_lo0025	Transverse	PD37590	Female	55	0.42	23	No
PD37590c_lo0032	Transverse	PD37590	Female	55	0.44	21	No
PD37590c_lo0034	Transverse	PD37590	Female	55	0.35	28	No
PD37590c_lo0040	Transverse	PD37590	Female	55	0.41	26	No
PD37590c_lo0048	Transverse	PD37590	Female	55	0.41	23	No
PD37590d_lo0012	Left	PD37590	Female	55	0.4	26	No