Quantifying the deterrent effect of anti-cartel enforcement[*]

Stephen Davies[†], Franco Mariuzzo[‡], and Peter L. Ormosi[§]

Abstract

This paper presents a rare attempt to quantify the deterrent effect of anti-cartel policy. It develops a conceptual framework, which establishes the sort of information necessary for such quantification. This is then illustrated and calibrated by drawing upon existing literatures and using evidence from legal cartels to approximate what would be observed absent policy. Measuring impact by the proportion of all potential harm that is deterred, our best estimate is two thirds and, even on conservative assumptions, at least half of all harms (or seven times the detected harm) is deterred.

Keywords: anti-competitive harm, cartels, detection, deterrence, Monte Carlo simulation, selection bias

JEL Classification codes: H11, K21, L44

[†]Corresponding author. School of Economics and Centre for Competition Policy, University of East Anglia, Norwich Research Park, NR4 7TJ, Norwich, United Kingdom, email: s.w.davies@uea.ac.uk

[‡]School of Economics and Centre for Competition Policy, University of East Anglia, Norwich Research Park, NR4 7TJ, Norwich, United Kingdom, email: f.mariuzzo@uea.ac.uk

[§]Norwich Business School and Centre for Competition Policy, University of East Anglia, Norwich Research Park, NR4 7TJ, Norwich, United Kingdom, email: p.ormosi@uea.ac.uk.

I Introduction

With the growth of anti-cartel enforcement throughout the world, increasing efforts are now devoted to estimating its impact on economic welfare. Most evaluations by the competition authorities (CAs) are very positive: although estimates are fairly rough and ready, they suggest that the benefits to consumers from anti-cartel enforcement more than outweigh the CA's costs, usually by an order of magnitude. For example, the European Commission (EC) estimates that in 2014 consumers benefited from the removal of overcharges by cartels that it detected and prohibited to the tune of €1.69 billion.[1]

However such estimates are based only on the cases that the CA busts - we call this the 'observed harm' – and this raises the obvious question: what about the harm that goes unobserved, because it is deterred? After all a major function of any law is to deter antisocial behaviour, and this is hopefully also true for competition policy. Both academics and the CAs themselves have long argued that the magnitude of deterred harm probably far exceeds the harm removed by direct intervention, but this hunch has never been confirmed: indeed, we know of no direct attempts to do so. In an oft-cited quotation the DoJ concedes that "*We firmly believe that deterrence is perhaps the single most important ultimate outcome of the [Antitrust] Division's work. We are just as sure that it presents the most significant measurement challenges. . .*" (Nelson and Sun, 2002, p. 49). More recently, a European consortium of CAs reports: "*We are not aware of any methodologies used by the Competition Authorities to assess, let alone estimate, the deterrence effects of their work and to assess its macroeconomic impact*".[2]

---

[1](European Commission DG Competition, 2015, p. 6). The US Department of Justice also estimates the benefits of its cartel enforcement to be $293 million for 2015 (DoJ Antitrust Division Congressional Submission, FY 2015, Performance Budget), and the UK's Competition and Market Authority (2016) reports an annual average of £73.6 million from cartel enforcement (including monopoly abuse) for 2013-16. The estimates of all three authorities relate to cartel cases prohibited in a given year. Their estimates are all based on very conservative assumptions, and for that reason, it is suggested that they are lower bound estimates. Note that the relative magnitudes are not comparable because the assumptions differ in just how conservative the three authorities are.

[2]"*Looking beyond the direct effects of the work of Competition Authorities: Deterrence and*

However, there is a further related issue, which is also typically neglected in the evaluation literature: what about a second type of unobserved case - those which the CA fails to detect, even though they involve anti-competitive harm? In effect, these represent a foregone opportunity, or even failure of policy; again, this is never quantified, although the magnitudes involved could also be substantial.

This suggests that a more encompassing approach to evaluation would be to ask first "how much potential for anti-competitive harm is out there in an economy"? Then, to assess the success of policy, further ask "how much of that potential harm is avoided and/or removed by the presence of anti-cartel laws and the activities of competition authorities"? The first question is reminiscent of an old literature, provoked by Harberger (1954), which attempted to quantify the "social costs of monopoly". Our approach is undoubtedly ambitious, but the potential value of any practicable methodology to incorporate deterrence within policy evaluation is self-evident – not only for assessing the impact of competition law and its agencies, but also for informing prioritisation of resources within the authorities.

We begin by setting out a conceptual framework to establish what sort of information would be needed to answer these questions. This is set up as a simple sampling problem in which there is a population of 'potential' cartels, all of which would occur in the absence of anti-cartel law. A sample is drawn - these are the set of detected cartels - and we derive an estimator for making deductions from the sample harm about the magnitudes of the unsampled (deterred or undetected) harms. This identifies three types of information required: (a) the aggregate deterrence and detection probabilities (which determine the aggregate sample size); (b) how these probabilities vary between cartels depending on their overcharge (the sampling differentials); and (c) the extent of cartel heterogeneity in overcharge. None of this information is obviously available. The objective of our work is to demonstrate how one could combine existing and published estimates and Monte Carlo simulations to retrieve this information.

We review existing cartel literatures for guidance on calibrating this framework. On some parts, we can be fairly confident in narrowing down the possible assumptions: for example, previous empirical literatures point to an aggregate probability of detection within the range (0.1, 0.33); and the distribution of cartel harm (overcharge) is likely to be positively skewed. On other parts, conventional cartel theory provides guidance: notably, on how the deterrence and detection rates are likely to vary across cartels with overcharge. On the other hand, we have very little hard evidence on the likely magnitude of the aggregate probability of deterrence.

Inevitably, estimation within this or any other framework will entail significant uncertainty and, for this reason, we opt for Monte Carlo experiments, with a strong emphasis on the lower bounds of the estimates, focussing more on what outcomes can be ruled out, than attaching undue confidence to specific point estimates. Two sets of experiments are presented. The first set is based on fairly "general" weak assumptions; for example, both aggregate deterrence and detection probabilities are allowed to lie anywhere within wide ranges, and asymmetry within the population distribution is allowed to lie anywhere between perfect homogeneity and extreme heterogeneity. This plays a largely ground-clearing role.

The second set of experiments is invested with more structure in its assumptions, and we explore how far this narrows down the range of possible outcomes. Here, assumptions are based on direct observation of the real world distribution of cartel overcharge gathered originally by Connor (2014). This database is wide ranging, both over time and geographically, and although most of the cartels were illegal, a sizeable minority were legal (occurring in jurisdictions without prohibition, or where exemptions were granted). The overcharge distribution amongst these legal cartels would be observed in the absence of competition law and agencies - the counterfactual population of potential cartels. This adds more structure not only on the population distribution, but also, by comparing the legal and illegal cartel overcharge distributions, Bos et al. (2017) provide guidance on which types of cartel are less likely to be observed (detected) once cartels become il-

4

legal. While two important sources of uncertainty remain: about the aggregates detection and deterrence rates, and how far the observed sampling differentials reflect differences in detection or deterrence, this extra structure does indeed narrow down the range of feasible outcomes, especially at the lower bound. Our three main results, all based on the most conservative estimates are that: detected harm is only the tip of the iceberg; at least half of all potential harm is deterred, and at least twice as much harm is undetected as is detected.

This paper contributes to, and draws on, various previous literatures. First and most obviously, it aims to contribute a missing link to the evaluation of the impact of competition policy. This has attracted a voluminous literature, in a variety of forms: ex-ante versus ex-post evaluation; evaluation of individual cases, or broad areas of policy, or aggregate competition policy. The quantitative methods used have also varied widely, including event studies, qualitative surveys, and causal inference studies (see Ormosi (2012) for summaries). However, all these works are constrained by the same limitation: driven by data availability, inferences are made on the basis only of the sample of cases in which the CA has intervened, and fail to take account of unobserved cases, which are deterred or undetected. Of course, this has long been noted by many others, including Werden (2008), (Geroski, 2004, p. 3) and (Baker, 2003, p. 40). More specifically within the evaluation literature, our paper is closest to the work of the CAs themselves in evaluating the impact of competition policy in aggregate, and/or the broad areas of anti-trust and merger enforcement (see footnote 1 above). An important theme running through this policy literature is the use of conservative assumptions which give rise to lower bound estimates, given that one purpose of evaluation is accountability to stakeholders, this is clearly advisable. We also focus on conservative lower bounds - values of deterrence and undetection that can we rule out.

Turning to the more academic literature, empirical studies on the rate of detection provide a useful input into the current paper. The early work of Bryant and Eckard (1991) suggests that in a given year only 13-17% of cartels are detected, and Combe et al. (2008) confirms a similar magnitude for Europe using

5

essentially the same method. Ormosi (2014) proposes an alternative method that draws an analogy with capture-recapture analysis in ecological science and finds detection rates to be in the same ballpark. Lande and Connor (2012) provide an exhaustive survey of these and other estimates of the rate of detection, and their findings are used to calibrate our experiments below.

While the deterrent effect of cartel law has attracted a considerable literature, much of it is only of indirect interest for present purposes.[3] Amongst the agencies, the UK and Dutch competition authorities have commissioned various research projects. These are survey studies, based on interviewing competition practitioners, lawyers and companies and asking them directly for their opinions. The most widely cited statistic to emanate from these studies (Deloitte (2007) for the OFT), is that, for each cartel detected by the CA, there were at least another five that were deterred by competition law.[4] This is sometimes referred to as the deterrence multiplier. Below, we refer to 'the simple multiplier approach', which we employ as a yardstick in our own simulations. Clearly, any estimates based on speculative opinions in answer to hypothetical questions should be treated with considerable caution, and we do not use any survey estimates directly in our own simulations.

In the more recent academic literature a number of papers are related to our own in that they explore theoretically and/or empirically how information on detected cartels might be used to draw inferences about other cartels which are unobserved. Miller (2009) develops a model of cartel formation over time and shows empirically that the introduction of the 1993 leniency programme strengthened deterrence. Theoretically, Harrington and Chang (2009) establish the conditions under which information on the number of detected cartels, and their durations, can support inferences about how the number of undetected cartels is changing.

---

[3]See for example the literatures on optimal fines: Elzinga and Breit (1973), Landes (1983), Kobayashi (2001), Ginsburg and Wright (2010), Werden et al. (2012). See also Katsoulacos and Ulph (2013) on whether revenue or profit should be used as the basis for calculating the optimal fine.

[4]Subsequent studies using similar survey methodologies confirm that such multipliers are likely to be significant, although the precise magnitudes vary considerably: a larger follow-up UK survey by London Economics (2011) reports much higher (1:28) deterrence rates.

Subsequently, Harrington and Wei (2017) focus on the relationship between the durations of detected and undetected cartels, and establish empirically the sort of selection bias that might be involved. All three papers have made valuable and imaginative contributions on how inferences can be drawn about an inherently unobservable population from a sample of observed cartels, and this is also our objective. However, none is designed to quantifying the magnitude of deterrence, as measured by overcharge: the latter two address non-detection rather than deterrence and duration rather than overcharge, and the first identifies only changes over time, rather than the magnitude, of harm.

Finally, two other very recent papers are closely related to this one. One, Bos et al. (2017) is complementary to ours in that it shows, both theoretically and empirically, that cartel deterrence will be highest for very low and very high mark-up cartels. This result feeds directly into the assumptions employed in this current paper on the sampling differentials, as will be see in Section 3.2. The other, Katsoulacos et al. (2016), shares the same conceptual objective as ours, but its approach is very different in at least two respects. First it abstracts from any heterogeneity amongst cartels, and thereby sidesteps the key issue of selection bias. Second, its main focus is on the rarely observed phenomenon of recidivist cartels.[5]

The remainder of the paper is structured as follows. Section 2 introduces the framework, couched as a simple sampling problem; this yields estimators for the unknown population harm, including deterred and undetected harms, in terms of sample-detected harms. It identifies three building blocks: the aggregate rates of detection and deterrence, the heterogeneity of cartels within the potential population, and the sampling differentials, which reflect differences in the likelihood of formation and detection between cartels with different levels of harm. Section 3 draws on existing literatures for our assumptions about these building blocks. The framework is then operationalised in Section 4 using Monte Carlo experiments to

---

[5]We believe this emphasis is unwarranted because it confuses the fairly frequently observed occurrence of individual firms being involved in more than one cartel, with the very rare possibility of a cartel reforming after having been detected.

generate estimates of the likely relative magnitudes of deterred, undetected and detected harms. Section 5 assesses the likely robustness of our results and Section 6 concludes.

II Framework

II.1 Preliminary notation

Denote the total population of all potential cartels in a given economy at a given point in time (say year) by N. This represents all harmful cases, which would occur in the absence of competition policy.[6] With probability ω some are deterred by the existence of competition law and enforcement. Within the undeterred subset, some cartels are discovered and prohibited by the CA, but others go undetected. The probability that a cartel is detected is σ.

In general, σ and ω can be time variant and interdependent. For example, there may be a dynamic causal relation running from detection, σ, to future deterrence, ω: success in detection by the CA might deter some firms from attempting to contravene the law in the future. Moreover, it is also possible that a more effective competition authority performs well both in terms of deterrence and detection, therefore ω and σ would be contemporaneously correlated as well. In our experiments below, we allow for the two parameters to be either positively correlated or uncorrelated.

Suppose the CA detects a sample N^S of the N cartel population, and estimates the magnitude of harm they cause, whose value is denoted by H^S. We refer to H^S below synonymously as sample or detected harm. It follows that the expected number of detected cartels is:

---

[6]This counterfactual is over-simplification to the extent that some cartels might be deterred by private enforcement, import liberalisation, deregulation and other non-competition policies. Note that we also abstract from the possibility of welfare-enhancing cartels, and Type 2 errors.

$$N^S = \underbrace{(1-\omega)\sigma}_{\equiv\lambda}\, N, \tag{1}$$

where λ is the proportionate sample size. This provides two preliminary insights. First, detected cases are probably a small proportion of the potential population. As an illustrative example, if $\sigma = 0.2$ and $\omega = 0.5$, only 10% of all potential cases are detected. Second, a head count of the number of detected cartels tells us little about the efficacy of the CA: detecting 10% of the population could indicate a very poor CA, with low detection and no deterrence ($\sigma = 0.1$ and $\omega = 0$), or a strong CA, which deters most cartels and detects all those which are undeterred ($\sigma = 1$ and $\omega = 0.9$). Both points are illustrated in the simulations that follow.

There are similar implications for the magnitudes of harm, rather than the number of cases. If sample cases are a random draw of the population, then detected harm is related to total potential population harm, H, by:

$$H^S = \lambda H, \tag{2}$$

which can be rewritten as

$$H = \frac{H^S}{\lambda}. \tag{3}$$

The magnitudes of deterred (H^{DR}) and undetected (H^{UT}) harms can be derived similarly, leading to:

Remark 1 *If the cartels detected by the CA are a random sample of the population, then population harm is a simple multiple of detected harm: $H/H^S = 1/\lambda$. The multiple for deterred harm is $H^{DR}/H^S = \omega/\lambda$, and for undetected harm $(H^{UT}/H^S) = (1-\sigma)/\sigma$. We refer to those as "the simple multipliers".*

We argue below that the assumption of a random sample is implausible, and our aim is to establish and calibrate alternative estimators, which address the classic selection bias faced by most studies of policy evaluation, and indeed for any empirical research in industrial organisation more widely defined.[7]

---

[7]For example, conventional wisdoms drawn from the empirical literature are that, typically,

II.2 The population distribution

We proceed by framing the task as a sampling problem in which the purpose is to estimate the magnitude of total harm (H) for a population, from which a sample is taken, where the sample harm (H^S) is known, from the CAs estimates for the cartels that it has detected. For analytical simplicity, we employ a stylised trichotomy in which the population is broken down into three, not necessarily equal, segments: *low*, *medium* and *high* harm cases (L, M, H). The lower tail of least harmful cases accounts for a proportion p_L of the population, and proportion h_L of population harm; the upper tail for p_H and h_H, and the middle segment for $p_M = (1 - p_H - p_L)$ and $h_M = (1 - h_H - h_L)$.

The population distribution is depicted in Figure 1 using a traditional Lorenz curve, in which the potential cartels are ranked in ascending order of harm along the horizontal axis, and the curve shows the cumulated proportion of aggregate harm (normalized to 1) they account for.[8]

As a purely presentational device below, we employ the special case of a *symmetric Lorenz curve*. This is symmetric either side of the downward sloping main dotted diagonal. Then, for any lower tail p_L accounting for harm h_L there is an upper tail of size $p_H = h_L$, which accounts for harm $h_H = p_L$, (and of course $p_M = h_M$), e.g. if the 30% least harmful cases account for 10% of harm, the 10% most harmful cartels account for 30% of harm. In the figure, this implies $AD = BC$. Lorenz-symmetry is analytically convenient and should offer a reasonable first approximation for many positively skewed distributions. It is satisfied for example by the lognormal (Aithchison and Brown, 1957, p.113). Below, we

---

cartels last for 5-8 years, involve 4-8 members and overcharge in the region of 20% (see Levenstein and Suslow (2006), Davies and De (2013)). All of these estimates are based on cartels which have been detected and do not necessarily apply to undetected or deterred cartels.

[8]Both axes are normalised to the interval [0,1]. The 45° diagonal shows the symmetric benchmark where all cartels are equally harmful; and the curve must be concave to the diagonal because cases are ranked in ascending order of harm. The distance of the curve below the diagonal reflects the degree of asymmetry in the size distribution. The Lorenz curve is non-parametric, but often proves helpful in the analysis of positively skewed distributions. Traditionally, 'size' might be personal income in studies of income distribution or firm size in studies of industrial concentration; see Lorenz (1905), or Gastwirth (1972).

Figure 1: Lorenz curve of potential cartel harm

employ it for preliminary illustrative purposes, but none of our substantive results require this assumption.

II.3 Selection bias with differential sampling

Suppose that sampling is random within each segment, but not necessarily across segments, then there is a straightforward estimator for the magnitude of population harm (H).[9] If the segments are sampled in potentially different proportions, λ_i $(i = L, M, H)$, overall proportionate sample size of the population is given by:

$$\frac{N^S}{N} \equiv \lambda = \lambda_L p_L + \lambda_M p_M + \lambda_H p_H, \tag{4}$$

and since $p_L + p_M + p_H = 1$, we can substitute p_M and have

---

[9]With this simple assumption, we effectively approximate continuous relationships between the sampling rate and case harm with a simple three-step function. Although analytically crude, this is sufficiently flexible for present purposes - the relative magnitudes of the λ_i can capture monotonicity or not, and concavity, linearity or convexity. Of course, for a given continuous relationship, the magnitudes of the λ_i are sensitive to how broadly defined are the sizes of the different segments (p_i).

$$\lambda = \lambda_M + (\lambda_L - \lambda_M)p_L + (\lambda_H - \lambda_M)p_H. \tag{5}$$

Similarly, proportionate sample harm is:

$$H^S/H = \lambda_M + (\lambda_L - \lambda_M)h_L + (\lambda_H - \lambda_M)h_H. \tag{6}$$

Subtracting 4 from 6, and with simple manipulation,

$$H = \frac{H^S/\lambda}{1 - \frac{(\lambda_L - \lambda_M)}{\lambda}(p_L - h_L) + \frac{(\lambda_H - \lambda_M)}{\lambda}(h_H - p_H)}. \tag{7}$$

Lemma 1 *With random sampling, the simple multiplier, $1/\lambda$, is an unbiased estimator of total population harm. But with differential sampling, population harm also depends on (i) the sampling differentials and (ii) population asymmetry.*

Proof. *From inspection, with random sampling across segments ($\lambda_L = \lambda_M = \lambda_H$), $H = H^S/\lambda$. But, more generally, $H \gtreqless H^S/\lambda$ as $(\lambda_L - \lambda_M)(p_L - h_L) - (\lambda_H - \lambda_M)(h_H - p_H) \gtreqless 0$.* ∎

The intuition is fairly obvious. Defining the sampling differentials as $\frac{(\lambda_L - \lambda_M)}{\lambda}$ and $\frac{(\lambda_H - \lambda_M)}{\lambda}$, then for a given sample proportion λ, and estimate of sample harm H^S, the implied population harm is larger the greater is sampling of the lower tail, $\frac{(\lambda_L - \lambda_M)}{\lambda}$, and the smaller is sampling of the upper tail, $\frac{(\lambda_H - \lambda_M)}{\lambda}$, i.e. the more that low harm cartels are featured disproportionately in samples, and vice-versa. Population asymmetry is captured by the magnitudes of $(p_L - h_L)$ and $(h_H - p_H)$, and the larger these are, the greater is the impact of sampling differentials on population harm.

This intuition can be brought into sharper focus by referring back to the Lorenz curve, in which $(h_H - p_H) = BC$ and $(p_L - h_L) = AD$. Thus

$$H \gtreqless H^S/\lambda \text{ as } (AD/BC) \gtreqless (\lambda_M - \lambda_H)/(\lambda_M - \lambda_L) \tag{8}$$

The magnitudes of AD and BC reflect the asymmetry of the distribution, i.e.

the distance of the Lorenz curve from the diagonal, but they also depend on the precise form of the underlying population distribution.

Proposition 1 *For a population distribution that generates symmetric Lorenz curves, for example the lognormal, H will deviate by more from the simple multiplier (H^S/λ) the larger are the sampling differentials and the more asymmetric is the population distribution.*

Proof. *For a symmetric concave Lorenz curve, $AD = BC$, i.e. $p_L - h_L = h_H - p_H$, and Eq. (7) simplifies to:*

$$H = \frac{H^S/\lambda}{1 + \frac{(\lambda_H - \lambda_L)}{\lambda}(h_H - p_H)}. \tag{9}$$

Then $H \gtreqless H^S/\lambda$ as $\lambda_H \lesseqgtr \lambda_L$; and for given $(\lambda_H - \lambda_L)$, the absolute difference is increasing with $(h_H - p_H)$, where $(h_H - p_H)$ is greater for more asymmetric distributions.[10] ∎

We assume Lorenz-symmetry in the general experiments below, but merely for ease of illustration.[11] This assumption is not necessary for our framework, and is not used in the evidence-based experiments on which our main conclusions are based.

---

[10]To avoid confusion note that we refer to two, different, types of asymmetry. A distribution is said to be asymmetric if it is skewed, in which case the degree of asymmetry refers to the degree of skewness. A Lorenz curve is asymmetric if it is not symmetric either side of the downward sloping diagonal. In the current context, all potential population distributions are likely to be asymmetric, but they may or may not exhibit asymmetric Lorenz curves. For example, the lognormal distribution is asymmetric but has symmetric Lorenz curves.

[11]For asymmetric Lorenz curves, the Lorenz asymmetry coefficient (LAC) (Damgaard and Weiner (2000)) enters as an additional term, where $LAC \gtreqless 1$ as the upper tail contributes relatively more (less) than the lower tail to overall asymmetry.

II.4 The estimators for population harm and deterred and undetected harms

This proposition relates to the relationship between population harm and sample harm. We can interpret it in the present context in two ways. First, to find total potential population harm in terms of detected and deterred sample harm, recall that $\lambda_i = (1 - \omega_i)\sigma_i$ and Eq.(7) becomes:

$$H = \frac{H^S/(1 - \omega)\sigma}{1 - \frac{(1-\omega_L)\sigma_L - (1-\omega_M)\sigma_M}{(1-\omega)\sigma}(p_L - h_L) + \frac{(1-\omega_H)\sigma_H - (1-\omega_M)\sigma_M}{(1-\omega)\sigma}(h_H - p_H)}. \quad (10)$$

Second, it can also be used to provide an estimator for total undeterred harm (H^{UR}) if we treat the detected cases as a sample drawn from the population of undeterred cartels. In that case, the sample proportion in segment i is $\lambda_i = \sigma_i$ and population harm refers to the total harm of all undeterred cartels, H^{UR}. Deterred harm (H^{DR}) and undetected harm (H^{UT}) can then be backed out as simple residuals:

$$H^{DR} = H - H^{UR} \quad (11)$$

and

$$H^{UT} = H^{UR} - H^S \quad (12)$$

Equations (10)-(12) form the basis for the simulations below. For this we require information on:

1. the aggregate sampling proportion (λ) and its two constituents: $1 - \omega$ and σ;

2. the sampling differentials $\frac{(\lambda_M - \lambda_L)}{\lambda}$ and $\frac{(\lambda_M - \lambda_H)}{\lambda}$, and their constituents (the segments ω_i and σ_i);

14

3. the asymmetry of the population distribution, as represented by $(p_L - h_L)$ and $(h_H - p_H)$.

To calibrate the parameters, we turn to existing empirical and theoretical literatures. This provides some broad guidance on the likely magnitudes of ω and σ, and some ordinal constraints on the sampling differentials.

III Assumptions on deterrence and detection

III.1 The aggregate parameters

In the following simulations, the probabilities of deterrence and detection ω and σ are modelled as random variables, using the existing empirical literature for evidence of plausible ranges. There is a large literature attempting to estimate the cartel detection rate. Lande and Connor (2012) report the results of 25 previous econometric studies, yielding estimates of σ within the range 0.1 to 0.33. We therefore specify σ as a random variable drawn from a distribution within those supports. For ω, there is much less to go on from the previous literature, and this parameter is allowed to take any possible value between 0 and 1, except the very extremes. To add a perspective, note that the midpoints of these two ranges, $\omega = 0.5$ and $\sigma = 0.215$, imply a deterrence multiplier of approximately 5:1 - i.e. for every detected case there are 5 others deterred.[12] This is consistent with the most widely cited estimates in the existing literature, from the Deloitte (2007) survey described earlier, although when interpreting results below, we focus less on midpoints, and more on conservative lower bounds.

---

[12]The proportions of detected and deterred cartels are $(1 - \omega)\sigma$, and ω respectively, and if the latter is 5 times the former, $\omega = 5(1 - \omega)\sigma$. This is satisfied approximately at the midpoints of our assumed distribution, $\sigma = 0.215$ and $\omega = 0.5$.

III.2 The sampling differentials

For the segment parameters, the ω_i and σ_i, and thus the sampling differentials, we draw on the existing theoretical literature to impose qualitative constraints on relative magnitudes. Again, this is for both sets of experiments, but later we will impose more structure on the relative magnitudes of segment parameters using real world evidence.

The relative magnitudes of the segment sampling proportions, the λ_i, depend on how the probabilities of deterrence and detection vary with cartel harm. Here, we employ insights from the existing cartel literature on deterrence and detection. For this purpose, we interpret cartel overcharge as synonymous with harm, which is an acceptable approximation given (Bolotova, 2009, p. 338) empirical finding of no significant correlation between overcharge and cartel turnover.

We believe the existing literature supports two ordinal assumptions.

Assumption 1 *The probability of cartel detection is non-decreasing with cartel harm:* $\sigma_L \leq \sigma_M \leq \sigma_H$.

Anti-cartel enforcement has two arms: leniency applications and 'ex-officio' detection (our short-hand for non-leniency cases provoked by consumer complaints or the CA's own monitoring independent of leniency). Our detection rate, σ, can be thought of as the weighted average of these two forms of detection. Most of the recent theoretical and experimental literatures on leniency, such as Bigoni et al. (2012), Harrington Jr and Chang (2015), Fonseca and Normann (2012), and Jensen and Sørgard (2016), suggest that the probability of leniency applications increases with harm. This is because in markets where cartels cause large harm there is also a larger incentive to deviate, and less stable cartels are more likely to apply for leniency. On ex-officio detection, following Block et al. (1981) model, the prevailing assumption is that the probability of detection increases with the cartel markup, because higher prices are more likely to raise customer suspicion and complaints, and/or to be spotted by the CA. This is also recognised explicitly in much of the cartel screening literature.

Assumption 2 *The probability of deterrence is higher for low harm and high harm than for medium harm cartels:* $\omega_L \geqslant \omega_M \leqslant \omega_H$.

Bos et al. (2017) provide both theoretical and empirical support for this assumption.[13] Using a fairly generic model of cartel formation, they show that low-overcharge cartels are more likely to be deterred from forming because the expected penalty makes them unstable and unprofitable. In addition, high-overcharge cartels are less likely to be observed because effective anti-cartel enforcement will tend to destabilise collusion and force colluding firms to limit their price increases.

They also provide empirical evidence, which is consistent with these assumptions. Using data from Connor's (2014) extensive historical dataset on cartel overcharge, they extract the distributions of (%) overcharge ($(p - mc)/p$) for cartels observed under regimes when cartels were illegal and those where they were legal (1107 and 390 observations respectively). These are reproduced here as Figure 2. This shows that relatively fewer low and high overcharges cartels are observed in regimes where cartels are illegal. They show this to be statistically significant, using quantile regression and propensity score methods to control for various potentially confounding factors, such as date and location of cartel. Assuming that the legal distribution reveals what would be observed in a world without policy, and therefore without deterrence, the figure shows that the combination of deterrence and non-detection leads to fewer cartels being observed in the tails when cartels are illegal. This is consistent with Assumptions 1 and 2.

In all simulations below, magnitudes for the segment parameters ω_i and σ_i are drawn subject to these assumptions, and in what we shall term the evidence-based experiments, additional structure is imposed using further evidence from the above comparison of legal and illegal cartels.

---

[13]A key finding of Bos et al. (2017) is that enforcement tightens the incentive compatibility constraints of high-harm cartels, which means that these cartels will lower their prices in order to deter cheating. Such an outcome would imply higher deterrence for high-harm cartels and lower for medium harm (for medium harm deterrence could even be negative, if composition deterrence reduces the markup of high harm cartels, which then become medium harm). This however does not affect our assumption. For our purposes what matters is that $\omega_H \geq \omega_M$.

Figure 2: Kernel density curves for cartel harm in legal and illegal cartels

IV   Experiments

Our empirical work involves simulating the harm by calibrating the parameters of Equation (10) based on the above Assumptions 1 and 2. It is done in two parts. The first part plays a preparatory role by establishing what sort of broad magnitudes might be involved, if one wishes to avoid imposing too much additional structure beyond these assumptions. This is referred to as general experiments. In the second part we impose more structure on the parameters, using evidence derived from the comparison of real world legal and illegal cartels.

IV.1   General Experiments

As explained above, we treat the aggregate parameters as random variables, both drawn from uniform distributions: $\omega \sim U(0.1, 0.9)$, and $\sigma \sim U(0.1, 0.33)$. Thus for aggregate deterrence we rule out only the most extreme possibilities that deterrence is either almost entirely ineffective or entirely effective. For detection, we

allow the parameter to lie anywhere in the range of values previously estimated in the detection literature.[14] In most of the simulations reported below, it is assumed that the two parameters are perfectly correlated - we believe this to be the natural default assumption, namely, a strong detection record for the CA is probably the strongest form of deterrence.[15]

Given the draws for aggregate ω and σ, a vector of values for σ_L, σ_M, and σ_H is drawn subject to the constraints: $\sigma = p_L\sigma_L + p_M\sigma_M + p_H\sigma_H$ and $\sigma_L \leq \sigma_M \leq \sigma_H$ (Assumption 1);[16] and a vector of ω_L, ω_M, ω_H is drawn subject to the constraints $\omega = p_L\omega_L + p_M\omega_M + p_H\omega_H$ and $\omega_L \geq \omega_M \leq \omega_H$ (Assumption 2).[17]

The population distribution is assumed to be lognormal, i.e. $ln(H) \sim N(\mu, s^2)$. The lognormal is a plausible candidate for this role - it typically provides a good description of many real-life business size related distributions, but nevertheless allows parsimony in the calibrations: we can vary the degree of population asymmetry across experiments with just one parameter, the variance of the distribution s^2. Here, we report results with 3 alternative values for $s = 1, 1.5$ or 0. A standard deviation $s = 1$ generates a fairly asymmetric distribution, in which the most harmful third of cartels account for nearly three quarters of total harm, and the least harmful third account for only 8%. Alternatively if $s = 0$, there is no cartel heterogenity, and if $s = 1.5$ there is even more asymmetry; the top tail accounts for the large majority, 86%, of total harm. We define arbitrarily the three segments to be of equal size, $p_{L,M,H} = 1/3$. The 'General' column of Table 1 summarises these assumptions.

---

[14]A variety of papers offer both theoretical and empirical results on the relationship between detection and deterrence, which forms part of our discussion of policy implications. Block et al. (1981) is the seminal reference on whether, and how, enforcement influences subsequent firm behaviour; other examples include Feinberg (1980), Block and Feinstein (1986), Brenner (2009), and Zhou (2011).

[15]In further experiments, we have replicated all the simulations reported below assuming alternatively, that the two parameters are uncorrelated. In fact all results are robust to this alternative assumption. For instance, see Section 4.3, Figure 3.

[16]The constraint is satisfied by ranking the drawn values in ascending order, and assigning values to σ_L, σ_M, and σ_H, respectively.

[17]The constraint is satisfied by assigning the lowest of the three values to ω_M, and allocating the other two values with a fair draw to ω_L and ω_H.

Table 1: Calibrations: parameter assumptions

	General	Evidence based
Aggregate deterrence	$\omega \sim U(0.1, 0.9)$	$\omega \sim U(0.1, 0.9)$
Aggregate detection	$\sigma \sim U(0.1, 0.33)$	$\sigma \sim U(0.1, 0.33)$
Segment detection	$\sigma_L \leqslant \sigma_M \leqslant \sigma_H$	$\sigma_L \leqslant \sigma_M \leqslant \sigma_H$
Segment deterrence	$\omega_L \geqslant \omega_M \leqslant \omega_H$	$\omega_L \geqslant \omega_M \leqslant \omega_H$
Population proportions	$p_L = p_M = p_H = 0.333$	$p_L = 0.238, p_M = 0.415, p_H = 0.347$
Population asymmetry	$Ln(h) \sim N(\mu, s^2), s \in \{0, 1, 1.5\}$	$h_L = 0.002, h_M = 0.144, h_H = 0.854$
Sampling differentials		$\lambda_L = \frac{\sigma_L(1-\omega_L)}{\sigma(1-\omega)} = 0.5266$
		$\lambda_M = \frac{\sigma_M(1-\omega_M)}{\sigma(1-\omega)} = 1.5462$
		$\lambda_H = \frac{\sigma_H(1-\omega_H)}{\sigma(1-\omega)} = 0.6707$

The model is evaluated with 10,000 numerical simulations, with summary results presented in two formats in Table 2, These formats reflect the two big questions posed in our introduction. To address the question, "how much potential harm is out there?", we employ a multiplier format, in which detected harm is normalised as $H^S = 1$, and the magnitudes for aggregate, deterred and undetected harms are then expressed as multiples of detected harm. To address the question, "how successful are CAs in removing the harm?", the composition format expresses each type of harm as a percentage of total potential harm.

The distributions of results are summarised parsimoniously with just the medians and the 5th and 95th percentiles for each of the harm types.[18] We refer to the median as the 'typical' value, the lowest 5th percentile as the 'lower bound', and the difference between the 5th and 95th percentiles as the 'range'. Most emphasis will be on the lower bound, which we interpret as the conservative estimate - this is in the spirit of much of the policy literatures in which it is conventional for CAs to employ conservative estimates of the direct impact (for example, CMA, 2016).

Consider first the top panel of Table 2 (for $s = 0$). Typical values suggest

---

[18]Note that these percentiles refer to the distributions taken separately for each harm type and are in that sense independent of each other. For example, in Table 2(i) in 5% of all experiments total harm ≤ 8.54, and in 5% of experiments deterred harm ≤ 2.22, but the two do not perfectly intersect: the set of lowest deterred harms are not necessarily from the same runs as the lowest sets of total harms. Therefore, each column does not provide an exact decomposition.

Table 2: Simulation results with lognormal distribution

	50%	5%	95%
(i) s=0: homogeneous cartels			
Multipliers			
Total	9.51	8.54	13.12
Deterred	4.78	2.22	9.53
Undetected	3.66	2.52	5.71
Detected	1.00	1.00	1.00
Contributions to total harm (%)			
Deterred	50.66	25.02	72.77
Undetected	38.71	19.54	63.68
Detected	10.51	7.61	11.70
(ii) s=1: medium asymmetry			
Multipliers			
Total	9.10	6.57	15.54
Deterred	5.03	1.78	11.84
Undetected	3.00	2.30	4.80
Detected	1.00	1.00	1.00
Contributions to total harm (%)			
Deterred	55.59	24.44	77.17
Undetected	33.19	16.12	61.65
Detected	10.99	6.43	15.21
(iii) s=1.5: high asymmetry			
Multipliers			
Total	9.27	6.17	18.84
Deterred	5.40	1.70	15.30
Undetected	2.78	2.21	4.59
Detected	1.00	1.00	1.00
Contributions to total harm (%)			
Deterred	58.71	24.74	81.22
Undetected	30.27	13.28	60.77
Detected	10.79	5.30	16.20

that about 50% of total harm is deterred, 40% undetected and 10% detected. But retreating to the more cautious 5% lower bound, we can only conclude that deterred and undetected harms are both at least twice the size of detected harm. An important lesson from this panel is that, even without any cartel asymmetry (i.e. $s = 0$), the range of possible outcomes is wide, for example, total harm could be anywhere between 8.5 and 13 times as large as detected harm. Similarly, deterred harm could account for as little as one quarter or as much as three quarters of total harm. This variability, even in this base case, captures the impact of uncertainty about the magnitudes of the aggregate parameters - especially ω - which is all that varies between simulations in this first part of the experiment. The second and third panels to the table show the sensitivity of results to introducing, and then increasing asymmetry into the population ($s = 1$ and 1.5 respectively). While typical values are fairly stable, the lower bounds become more cautious, and the ranges for total and deterred harms (but not undetected harm) increase with more asymmetric populations.[19]

As already explained these general experiments are largely preparatory, but some cautious conclusions can be drawn by reading across the three panels.

Result 1 *(i) Detected harm is only a small proportion of total harm (at most unlikely to exceed one sixth, and possibly as little as 5%). (ii) Deterred harm typically accounts for a much bigger proportion, roughly 50%, but it could be as little as a quarter or as much as 80%. (iii) Undetected harm is usually, but not always, smaller than deterred harm. (iv) There are two sources of variability in these experiments: (a) uncertainty about the values of ω and σ and (b) asymmetry amongst cartels. The results suggest that they have roughly similar impacts.*

IV.2 Evidence based

In the second set of experiments, we continue to model ω and σ as above, and Assumptions 1 and 2 still guide the calibrations of how detection and deterrence

---

[19]For example, the relative ranges, defined by the difference between the 5% and 95% centiles divided by the median, increases from 0.48, for $s = 0$, to $0.99 (s = 1)$, and to 1.37 for $s = 1.5$.

vary across the segments. But beyond this we now explore how much more precision can be bought by introducing some more 'real world' evidence, using the comparison in Figure 2 between legal and illegal cartels.[20] This obviates the need to simulate the population distribution and provides more structure on the relative magnitudes of the sampling differentials.

IV.2.1 Additional structure

First, assume that the legal distribution represents a random sample drawn from the population distribution (what would be observed absent policy), then the three segments L, M, and H, can be identified as the three parts of the distribution delineated by its two points of intersection with the illegal distribution (see Figure 2). These occur at overcharges of 6.4% and 40.0%. In turn, this identifies the proportions of the population and the population harm accounted for by each segment: thus from the legal distribution, in the lower tail $p_L = 23.8\%$, but $h_L = 0.2\%$, and $p_H = 34.7\%$ and $h_H = 85.4\%$ in the upper tail.

Second, comparing the masses in the three segments between the legal and illegal distributions provides direct estimates of the sampling differentials. To see how, note that the low harm segment of the illegal distribution accounts for a proportion, $p_L \frac{\sigma_L(1-\omega_L)}{\sigma(1-\omega)}$, of all observed illegal cartels, while the proportion of the legal sample in the low harm segment is p_L.[21] It follows that the ratio of the illegal to legal proportions in the low segment is $\frac{\sigma_L(1-\omega_L)}{\sigma(1-\omega)}$, that is λ_L/λ. For the other segments λ_M and λ_H are derived identically. These are reported in the 'Evidence based' column of Table 1, and as can be seen, there is relatively more mass in the middle segment of the illegal distribution (1.5462 times the aggregate) but relatively less in the two tails, especially the low harm segment (0.5266).

---

[20]See Bos et al. (2017) for details on this database, and our decomposition into legal and illegal cartels.

[21]The total number of detected illegal cartels is $N\sigma(1-\omega)$, of which $p_L N\sigma_L(1-\omega_L)$ are in the low harm segment.

IV.2.2 Three direct implications

These estimates do not allow us to identify the individual segment detection and deterrence parameters - these remain underidentified and simulation is still necessary. However, they do provide three important new results even before we turn to simulation.

Result 2 *Total potential population harm is 25% greater than would be predicted by the simple multiplier estimate.*

Proof. *Substituting the above estimates of λ_i, p_i and h_i into 7 yields:*

$$H = \frac{H^S/(1-\omega)\sigma}{1+(1.020)\times(0.236)-(0.876)\times(0.507)} = 1.255 H^S/(1-\omega)\sigma. \quad (13)$$

∎

The intuition is that higher harm cartels contribute disproportionately to asymmetry, i.e. $(h_H - p_H) = 0.507 > (p_L - h_L) = 0.236$, (i.e. the Lorenz curve is asymmetric) and although relatively more of the high harm cases are observed $(\lambda_L < \lambda_H)$, this is outweighed by the fact that the unobserved high harm cases are disproportionately harmful. In other words, because the most harmful cases are extremely harmful, when they are unobserved, merely grossing up an estimate of sample (detected) harms using only aggregate detection and deterrence rates (i.e. assuming all cartels are equally harmful) would lead to a 25% underestimate of total population harm. This finding is of wider general insight into how sample selection bias should caution us when using results on detected cartels to draw inferences about all potential cartels.

Result 3 *There is a lower bound on the aggregate deterrence rate, $\omega \geqslant 0.257$.*

Proof. The proof follows by deriving minimum values for each of the three segments ω_i and then aggregating. From Assumption 2 (deterrence is lowest in the middle segment), and Assumption 1 (detection is no lower in high harm than medium harm cases, $\sigma_H \geqslant \sigma_M$), we can back out lower values of ω_M and ω_H. A

necessary condition for ω to be minimised is $\omega_M = 0$; and ω_H is minimised when σ_H is minimised, and since $\frac{\sigma_M(1-\omega_M)}{\sigma(1-\omega)} = 1.546$ and $\frac{\sigma_H(1-\omega_H)}{\sigma(1-\omega)} = 0.671$, simple manipulation yields $\omega_H = 1-0.671/1.546 = 0.566$. In turn, these magnitudes imply: $\frac{\sigma_M}{\sigma} = \frac{\sigma_H}{\sigma} = 1.546(1-\omega)$.To recover minimum ω_L, note that $\frac{\sigma_L(1-\omega_L)}{\sigma(1-\omega)} = 0.527$, so $\frac{(1-\omega_L)}{(1-\omega)} = 0.527/(\sigma_L/\sigma)$. The magnitude of (σ_L/σ) can be backed out from the σ aggregation constraint using the above values for σ_M/σ and σ_H/σ. Substituting into the aggregation constraint for ω yields a pair of simultaneous equations yielding a quadratic for $(1 - \omega_L)$ in terms of $(1 - \omega)$. The solution has two roots: $(\omega = 0.353, \omega_L = 0.660)$ and $(\omega = 0.257, \omega_L = 0.256)$, where the latter is the minimum. ∎

This leads immediately to:

Result 4 *Deterred harm as a share of total harm can never be less than 48.4%*
Proof. *Evaluating deterred harm at these minimum ω_i above gives:* $H^{DR} = \omega_L h_L + \omega_M h_M + \omega_H h_H = 0.256 \times 0.002 + 0 \times 0.144 + 0.566 \times 0.854 = 0.484$ ∎

The intuition for these results is not immediately obvious, but follows from the relative sampling rates in the middle and high harm segments (1.5462 and 0.6707). Combined with the assumption that detection is no lower in the high harm than the medium harm segment, those place a non-trivial lower limit on the deterrence rate for high harm cartels. Since it is these cartels which account for the large majority of potential population harm (85%), even moderate deterrence in this segment will mean that overall a large proportion of total potential harm is deterred.

These are important results for the objectives of this paper. As we have argued, there is a dearth of information in the existing literature on what might be plausible values for the magnitude of aggregate deterrence, and this is a major cause of the wide range of possible outcomes identified in the previous general simulations. We have now been able to rule out very low values for ω (< 0.2574), which provides a now much stronger result of the proportion of harm, which is determined - nearly 50% at the minimum.

IV.2.3 Simulations

We next return to numerical simulation, but now incorporating the new evidence as described in the 'Evidence based' column of Table 1. Adding the estimates of the three (λ_i/λ), and recalling the two aggregation constraints for ω and σ, we now have 5 equations from which to simulate the eight unknown parameters (the aggregate and segment ω_i and σ_i). So simulation is still necessary, but now only for three unknowns - say the magnitudes of ω and σ in aggregate, and for one of the six segment parameters, say σ_L. We can then identify the remaining 5 parameters, subject to the two Assumptions: $\omega_L \geqslant \omega_M \leqslant \omega_H$, and $\sigma_L \leqslant \sigma_M \leqslant \sigma_H$. Subject to these changes, the second round of experiments have been conducted as described for the first round, and are reported in Table 3.

Table 3: Evidence-based experiments

	50%	5%	95%
Multipliers			
Total	14.27	13.47	20.32
Deterred	9.50	7.62	17.06
Undetected	3.32	2.18	5.87
Detected	1.00	1.00	1.00
Contributions to total harm (%)			
Deterred	68.79	52.89	84.16
Undetected	23.99	10.95	40.30
Detected	7.01	4.91	7.42
Success	76.00	59.68	89.03

In a nutshell, the main change is increased magnitudes for deterred harm, both its multiplier and contribution to total harm, and both when measured at the typical (median) or lower bound levels. The reason is of course that we are now able to exclude the possibility of very low ω, as anticipated by Results 3 and 4 above.

More specifically, and concentrating on the lower bounds we have the following result.

Result 5 *Deterred harm is at least seven times greater than detected harm, and accounts for at least half of all potential harm. The magnitudes for undetected harm are much smaller: it may be little more than twice the size of detected harm, and account for only about 10% of total harm. In aggregate, total potential population harm is at least 13 times the magnitude of detected harm.*

Turning, with less emphasis, to the typical (median) values, the deterrence multiplier is over 9 with deterrence accounting for two thirds of total potential harm, while the undetection multiplier in roughly 3, with non-detection accounting for about 23% of total potential harm.

IV.3   Policy Implications

Tentatively, these simulations can provide some insights for evaluating the comparative success of different CAs. Figure 3a depicts the composition of potential harm when the cases are re-ordered by the magnitude of the CA's 'success' rate (the proportion of deterred and detected cases). So, for instance, the median success rate is 76.0% of which 68.79% is deterred and 7.01% is detected harm; this corresponds to the experiment in which aggregate deterrence is $\omega = 0.55$ and aggregate detection is $\sigma = 0.20$.

If we now interpret cases at the lower end of the horizontal axis as 'weakly' performing CAs, and the upper end as 'strong', then Figure 3a shows that the weak CA has much lower deterrence and much higher non-detection than the strong CA. The major difference between the two Authorities is in the percentage of harm deterred. For detected harm, at first sight there appears to be a paradox: the weak CA actually detects a greater percentage of potential harm than the strong CA, in spite of having a much lower detection rate. However, there is no paradox – under a poor CA regime far less harm is deterred and so many more cases actually occur, and this effect outweighs the lower rate of detection.

While this interpretation is offered more as a thought experiment than a literal description of the differences between any real-world CAs, it does drive home two

Figure 3: Depicting CA success

(a) Correlated ω and σ (b) Uncorrelated ω and σ

Note: CA success rate is measured by the sum of detection and deterrence and is increasing from 0 to 1 on the horizontal axis. Correlated ω and σ indicates that a CA strong on detection is also strong on deterrence. Uncorrelated indicates not such tendency.

important implications of this paper. First, deterrence is by far the most important potential impact of anti-cartel policy, and second, estimates of the gains from detected cartels may be a very misleading indicator of the efficacy of cartel policy.

Finally, we revisit the assumption underpinning all the results so far, that ω and σ are perfectly correlated. In policy terms, there are at least three reasons why such a positive correlation should exist. First, firms are more likely to be deterred if confronted with a CA which has a strong detection rate. Second, at any point in time, there is likely to be fewer cartels surviving from previous periods if the CA has been more successful in detection. Third, insofar as deterrence is determined by not only by strong performance in detection but also other activities such as compliance and advocacy, ω and σ are likely to be positively related if a strong CA tends to be better across the range of activities.

Nevertheless, we have re-run the experiments assuming a zero correlation between the two parameters (Figure 3b).[22] As can be seen, all results are robust, both qualitatively and quantitatively.

---

[22]This captures the possibility that leniency discounts now weaken the deterrence effects of detection.

V Likely robustness of results

Our headline result - that at least 50% (and probably much more) of all potential harm is deterred - is the first of its kind, and we hope that it will help serve as a useful yardstick for future work on this topic. Nevertheless, we readily acknowledge that some parts of our analysis are inevitably speculative, and the purpose of this section is to scrutinize them more closely.

In essence, our results are driven, theoretically, by two insights derived from previous literatures: (i) the chances of detection increase with overcharge, and (ii) it is the lowest and highest overcharge cartels which are most likely to be deterred. Empirically, we have employed the distribution of overcharges observed historically for legal distributions and the key assumption is that this approximates the underlying population distribution - what would occur absent cartel law and policy. We now return to these key assumptions and also discuss the unavoidable uncertainty about the magnitudes of the two key aggregate parameters, ω and σ.

The aggregate probabilities, ω and σ

It is assumed throughout that σ lies within the range of estimates reported in the previous literature (0.1, 0.33). While this range is relatively narrow, the same is not true for ω, where the literature offers very few pointers, and where we allow for all but the most extreme values between zero and unity. We showed, in Result 3 that this can be narrowed slightly to never lie below 0.25. Also, in Section 3.1. we showed that the midpoint of this range is consistent with the estimates derived from previous surveys of informed opinions. Nevertheless, our limited knowledge about the precise magnitudes of these parameters accounts for a major source of uncertainty in our results. Faced with this uncertainty, our approach has been to conservatively emphasise the bottom 5% lower bounds to our estimates. These lower bounds are associated with only very low magnitudes of σ and ω apply. Referring back to the experiment reported in Table 3, the 5% lower bound magnitudes in our numerical simulations are only ω=0.353 and σ=0.110.[23]

---

[23]In other experiments, not shown here, we substitute the Beta distribution (2,2) for the uniform in specifying ω. This has a symmetric bell-shape, over the full range (0,1), albeit with much less

Assumptions 1 and 2

While we believe that Assumptions 1 and 2 (A1 and A2) above are fair representations of the previous literature, the results of further simulations in which they are relaxed are now shown in Table 4. Panel 1 reports experiments in which only A1 is imposed, and panel 2 where only A2 is imposed. This shows that A1 is the dominant assumption, in that that it is sufficient for A2 to hold too. In other words, given the estimates of the sampling proportions (the λ_i) derived from the comparison between legal and illegal cartels, if detection increases weakly monotonically with cartel harm, then the U shaped pattern for ω must also hold. The implication, confirmed in Table 4 panel 1, is that these new results, for experiments in which only A1 is assumed are in fact identical to those when both Assumptions 1 and 2 are imposed, as already shown in Table 3. On the other hand, when A1 is relaxed, and A2 only is imposed (panel 2), results, although qualitatively unchanged, are generally less precise. Thus, the multipliers are still of the same broad magnitudes; and deterred harm continues to dominate undetected in the ratio 3:1 at the median. But the "confidence intervals" are now wider, and the lower bounds estimates of deterred now fall to one quarter as opposed to one half previously.

Close inspection of the numerical simulations reveals that the key to this result is that, with A1 relaxed, the probability of detection in the high harm cartels is now often lower than in middle harm and sometimes even than in lower harm cartels. While we do not read the existing literature as offering much support for the proposition that there is a pervasive tendency for the highest harm cartels to be the most likely to go undetected, this result depicts the potential impact of this alternative assumption on our results - deterred harm continues to dominate over most of the probability distribution, albeit with the conservative estimate now reduced to only about one quarter of all potential harm being deterred.

Reliability of the comparison between legal and illegal cartels

---

mass in the tails. Results are qualitatively very similar to those we have reported here for the truncated uniform distribution.

Table 4: Sensitivity of results to A1 and A2

	50%	5%	95%
A1 only imposed			
Multipliers			
Total	14.27	13.47	20.32
Deterred	9.50	7.62	17.06
Undetected	3.32	2.18	5.87
Detected	1.00	1.00	1.00
Contributions to total harm (%)			
Deterred	68.79	52.88	84.16
Undetected	23.99	10.95	40.30
Detected	7.01	4.91	7.42
Success	76.00	59.68	89.03
A2 only imposed			
Multipliers			
Total	12.70	11.22	20.59
Deterred	8.53	2.74	17.33
Undetected	2.98	2.17	7.56
Detected	1.00	1.00	1.00
Contributions to total harm (%)			
Deterred	68.03	24.18	84.06
Undetected	23.83	11.06	66.95
Detected	7.87	4.85	8.91
Success	76.14	32.96	88.94

Empirically, we have illustrated how our framework can be operationalised by calibrating parameters (the λ_i shown in Table 1) using information derived from Connor's well-known database on cartel overcharge. This is an obvious choice, as a meta-analysis of hundreds of source studies across countries and over history, it is by far the most extensive existing resource. It also has the ideal feature for our purposes in that it includes a large subset of legal cartels, which occurred under regimes and/or time periods when they were not illegal. We have used these legal cartels to mimic the distribution of cartel harm (proxied by overcharge), which would be observed absent cartel policy and therefore without deterrence. Nevertheless, this source is not without criticism in the wider literature (for example, Boyer and Kotchoni (2015)) largely concerning the rigour of some of the individual studies from which Connor compiled the database. Moreover, and directly relevant for our use of the legal cartel distribution, there are some confounding factors which complicate any clean comparison between legal and illegal cartels - most notably in time periods and regional location. However, Bos et al. (2017) - from which our estimates derive - have controlled for those of these factors which are observable using propensity score methods. This also shows that the results are robust to excluding those estimates originating from sources, which, by their nature, were less exposed to academic scrutiny at their time of publication. For these reasons, we believe that we have gone as far as possible in this paper, given publicly accessible data at the time of writing, if one wishes to approximate the counterfactual - the distribution of harm which would be observed absent cartel law/policy - with a distribution of observed legal cartels.

Notwithstanding this, there may be alternative approaches to the counterfactual, which sidestep altogether the need for data on legal cartels. Indeed, in the present paper, we have employed one alternative, in the general experiments of Section 4.1, in which a lognormal is used to depict the underlying distribution of potential cartel harm, and exploring sensitivity of results to varying the degree of asymmetry of that distribution. Using as illustrative the medium asymmetry case reported earlier in Table 2, it can be seen that this generated results which are

qualitatively similar, if less absolutely pronounced, to our main evidence based results in Table 3. The multipliers are broadly similar if smaller, and deterred harm typically dominates undetected harm across the distribution, but its share of total potential harm could now be as low as one quarter at the 5% conservative lower bound.

VI   Concluding remarks

This paper raises the ambitious question of how can we quantify the magnitude of the deterrence effect of anti-cartel enforcement? While it is widely acknowledged that deterrence is probably the main effect, this has never been quantified to date. Our approach is to frame the question in two stages: how much potential harm can cartels cause to an economy, and how successful are CAs in rectifying that harm, either by detection or deterrence? The first objective of this paper is methodological - to propose a framework which has the potential to answer these questions. It interprets the cartels detected by a CA as a sample drawn from an otherwise unknown population, which also includes deterred cases and cases which the CA did not detect. This framework identifies three types of information required to quantify the magnitudes of the unknown deterred and undetected harms: (i) the aggregate probabilities of deterrence and detection; (iii) how these probabilities vary with cartel harm; and (iii) the heterogeneity in the population amongst potential cartels. This reminds us that the ubiquitous problem of selection bias is likely to be as important in policy evaluation as in most other areas of empirical Industrial Organisation. There is no reason for assuming that those cartels, which are caught and prohibited by CAs are necessarily representative of other cartels which successfully conceal their existence, or yet other potential cartels, which are deterred by the existence of the law and the CA. On this, the paper underlines a simple but important wisdom: in many aspects of economic and social life, strongly positively skewed distributions prevail - a small proportion of causes generates a very large proportion of effects. This is well understood in general

and is plausibly also likely with cartel-induced harm. If so, even small departures from a random sample can lead to potentially large selection bias (we show, as a by-product of our main results, that ignoring cartel heterogeneity could lead to underestimates of about 25% for total harm, Result 2 in Section 4.)

We calibrate the framework using various results from the existing empirical and theoretical cartel literatures. Theoretical literature provides the rationale for two assumptions: the probability of detection increases weakly monotonically with overcharge, while the probability of deterrence exhibits a U-shaped relationship to overcharge. From the empirical literature we assume that the underlying potential population distribution of overcharge can be approximated by the real world historical distribution of overcharge for observed legal cartels. If so, a comparison of the legal overcharge distribution with the distribution for illegal cartels provides further structure on how detection and deterrence vary with overcharge. Nevertheless, the largest gap in our knowledge relates to the magnitude of the aggregate deterrence probability; here there is little to go on from the previous literature, and for that reason in our simulations we allow this parameter to fall within a very wide range of potential values, and then focus mainly on the lower bounds of our estimates. In that sense, our headline results are deliberately cautious and conservative.

Amongst the results of our experiments, three merit emphasis. First, deterrence is overwhelmingly greater than detected cartel harm - our estimates suggest that, at the very least, 50% of potential harm is deterred. Our typical estimates is much higher, at two thirds. Second, undetected harm is probably much smaller than deterred harm - our upper bound estimate is 40% of potential harm, but the typical estimate is less than 25%. Nevertheless, undetected exceeds detected harm in most of our experiments (see Table 3 for example.) Third, in aggregate, potential harm from the population of cartels which would occur in the absence of cartel law is a large multiple (at least 13) of the harm actually detected by the CA: detected harm is only the tip of an iceberg.

Some of the policy implications are self-evident. Deterrence is a pivotal arm

of cartel policy, and it will be largely this which distinguishes the strong from the weak Authority. On the other hand, harm due to undetected cartels is also likely to be considerable, and it can be very misleading to measure the success of a CA simply by calculating the amount of harm it removes by virtue of cartel busts. Our results imply that a 'weak' CA (relatively ineffectual in deterring and detecting) may actually detect more harm than a 'strong' CA, simply because its inability to deter leaves far more cartels out there to be detected.

It is our hope that the current paper provokes future work in which, for example, alternatives counterfactuals are employed. Within our own research programme, we will work analogously on merger control (for which non-detection is likely to be less important), and more generally the scope for using this sort of framework to set the relative priorities within CAs.

References

Baker, J. B. (2003). The case for antitrust enforcement. *Journal of Economic Perspectives 17*(4), 27–50.

Bigoni, M., S.-O. Fridolfsson, C. L. Coq, and G. Spagnolo (2012). Fines, leniency and rewards in antitrust. *RAND Journal of Economics 43*(2), 368–390.

Block, M. K. and J. S. Feinstein (1986). The spillover effect of antitrust enforcement. *The Review of Economics and Statistics*, 122–131.

Block, M. K., F. C. Nold, and J. G. Sidak (1981). The deterrent effect of antitrust enforcement. *Journal of Political Economy 89*(3), 429–445.

Bolotova, Y. V. (2009). Cartel overcharges: An empirical analysis. *Journal of Economic Behavior & Organization 70*(1), 321–341.

Bos, I., S. W. Davies, J. E. Harrington Jr, and P. L. Ormosi (2017). Does enforcement deter cartels? A tale of two tails. *https://ssrn.com/abstract=2471425*.

Boyer, M. and R. Kotchoni (2015). How much do cartel[s] overcharge? *Review of Industrial Organization 47*(2), 119–153.

Brenner, S. (2009). An empirical study of the European corporate leniency program. *International Journal of Industrial Organization 27*(6), 639–645.

Bryant, P. G. and E. W. Eckard (1991). Price fixing: The probability of getting caught. *The Review of Economics and Statistics 73*(3), 531–536.

Combe, E., C. Monnier, and R. Legal (2008). Cartels: The probability of getting caught in the European Union. *BEER Paper 12*.

Competition and Market Authority (2016). CMA impact assessment 2015-16. Technical report, CMA.

Connor, J. M. (2014). Price-fixing overcharges: Revised 3rd edition. *Available at SSRN: http://ssrn.com/abstract=2400780*.

Damgaard, C. and J. Weiner (2000). Describing inequality in plant size or fecundity. *Ecology 81*(4), 1139–1142.

Davies, S. and O. De (2013). Ringleaders in larger number asymmetric cartels. *The Economic Journal 123*(572), F524–F544.

Deloitte (2007). The deterrent effect of competition enforcement by the OFT. *Office of Fair Trading OFT962*.

Elzinga, K. G. and W. Breit (1973). Antitrust penalties and attitudes toward risk: An economic analysis. *Harvard Law Review 86*, 693.

European Commission DG Competition (2015). 2014 annual activity report. Technical report, DG Competition.

Feinberg, R. M. (1980). Antitrust enforcement and subsequent price behavior. *The Review of Economics and Statistics 62*(4), 609–612.

Fonseca, M. A. and H.-T. Normann (2012). Explicit vs. tacit collusion - the impact of communication in oligopoly experiments. *European Economic Review 56*(8), 1759–1772.

Gastwirth, J. L. (1972). The estimation of the Lorenz curve and Gini index. *Review of Economics and Statistics 54*(3), 306–316.

Geroski, P. (2004). Is competition policy worth it? *Speech on the opening of the Centre for Competition Policy at the University of East Anglia, available at http://www. competitioncommission. org. uk.*

Ginsburg, D. H. and J. D. Wright (2010). Antitrust sanctions. *Competition Policy International 6*(2), 3–39.

Harberger, A. C. (1954). Monopoly and resource allocation. *The American Economic Review 44*(2), 77–87.

Harrington, J. E. and M.-H. Chang (2009). Modeling the birth and death of cartels with an application to evaluating competition policy. *Journal of the European Economic Association 7*(6), 1400–1435.

Harrington, J. E. and Y. Wei (2017). What can the duration of discovered cartels tell us about the duration of all cartels? *The Economic Journal.*

Harrington Jr, J. E. and M.-H. Chang (2015). When can we expect a corporate leniency program to result in fewer cartels? *The Journal of Law and Economics 58*(2), 417–449.

Jensen, S. and L. Sørgard (2016). Fine schedule with heterogenous cartels: Are the wrong cartels deterred? *Journal of Competition Law & Economics 12*(2), 263–277.

Katsoulacos, Y., E. Motchenkova, and D. Ulph (2016). Measuring the effectiveness of anti-cartel interventions: a conceptual framework.

Katsoulacos, Y. and D. Ulph (2013). Antitrust penalties and the implications of empirical evidence on cartel overcharges. *The Economic Journal 123*(572), 558–581.

Kobayashi, B. H. (2001). Antitrust, agency, and amnesty: An economic analysis of the criminal enforcement of the antitrust laws. *George Washington Law Review 715*, 735–739.

Lande, R. H. and J. M. Connor (2012). Cartels as rational business strategy: Crime pays. *Cardozo Law Review 34*(2), 427–490.

Landes, W. M. (1983). Optimal sanctions for antitrust violations. *University of Chicago Law Review 50*, 652.

Levenstein, M. C. and V. Y. Suslow (2006). What determines cartel success? *Journal of Economic Literature 44*(1), 43–95.

London Economics (2011). The impact of competition interventions on compliance and deterrence. *Office of Fair Trading OFT1391*.

Lorenz, M. O. (1905). Methods of measuring the concentration of wealth. *Publications of the American Statistical Association 9*(70), 209–219.

Miller, N. H. (2009). Strategic leniency and cartel enforcement. *American Economic Review 99*(3), 750–768.

Nelson, P. and S. Sun (2002). Consumer savings from merger enforcement: A review of the antitrust agencies' estimates. *Antitrust Law Journal 69*(3), 921–960.

Ormosi, P. L. (2012). Evaluating the impact of competition law enforcement. *OECD Working Papers DAF/COMP/WP2(2012)5*.

Ormosi, P. L. (2014). A tip of the iceberg? The probability of catching cartels. *Journal of Applied Econometrics 29*(4), 549–566.

Werden, G. J. (2008). Assessing the effect of Antitrust enforcement in the United States. *De Economist 156*(4), 433–451.

Werden, G. J., S. D. Hammond, and B. A. Barnett (2012). Deterrence and detection of cartels: Using all the tools and sanctions. *Antitrust Bulletin 56*(2).

Zhou, J. (2011). Evaluating leniency with missing information on undetected cartels: Exploring time-varying policy impacts on cartel duration. *TILEC Discussion Paper 2011*(042).