LOYOLA eCOMMONS

Computer Science: Faculty Publications and Other Works

Faculty Publications

7-1-2012

Collaboration and Health Care Diagnostics: an Agent Based Model Simulation

Sebastian Linde
Purdue University, linde@purdue.edu

George Thiruvathukal
Loyola University Chicago, gkt@cs.luc.edu

Recommended Citation

Collaboration and Health Care Diagnostics: an Agent Based Model Simulation

Sebastian B. Linde*and George K. Thiruvathukal†

Date: 10.23.13 (Draft #2.3)

Abstract

This paper presents a simple ABM of health care diagnostics using the NetLogo architecture. We simulate patient and doctor populations where patients differ by their health condition, and doctors differ by their area of expertise. We assume that an accurate diagnosis results from a patient-doctor match if and only if the doctor's area of expertise overlaps with the patient's health condition. The model allows for doctors to collaborate and in so doing share their knowledge with each other. Collaboration of this kind increases the odds of a successful diagnosis following from any patient-doctor match. In the analysis, we compare the patient population outcomes with and without physician collaboration, and further explore the parameter space of the model. A number of feasible model extensions are also enumerated.

1 Introduction

Health care diagnostics is the art of matching patient symptoms with a documented condition. However, such a match is by no means direct or error free, and can in many cases be the outcome of a long journey for the patient; encompassing several consultations with many different professionals before arriving (if at all) at the correct diagnosis. Diagnostic, and medical, errors are in fact unsettlingly common. Roughly 1 in 10 diagnosis provided are wrong(Graber et al., 2012; Wachter, 2010), and an estimated 44,000 to 98,000 US hospital deaths (annually) result as a direct consequence of misdiagnosis(Kohn et al., 2000). Meanwhile, medication errors are estimated to harm more than 1.5 million people every year and the extra medical cost of treating drug-related injuries occurring in hospitals alone amounts to over $3.5 billion a year(Aspden et al., 2007). Diagnostic errors are the basis for 40% of ambulatory malpractice claims that costs approximately $300,000 per claim on average(Singh and

---

*Sebastian B. Linde is a doctoral student in economics at Purdue University (email: linde@purdue.edu)

†George K. Thiruvathukal is a professor of computer science at Loyola University Chicago (email: gkt@luc.edu)

Graber, 2010). While there are several factors that contribute toward making the diagnostic process inefficient (see: (Singh and Graber, 2010)), this paper provides a model for looking at the issues brought about by the presence of incomplete information–both on behalf of the patient, who is often unaware of the details of his condition, and the health care professional, who possesses expertise on only a limited number of conditions. In this setting, information (expertise) is taken to be spread over a large population of health care professionals, and as such there are direct benefits from collaboration as this provides them with access to each other's expertise.

This paper is outlined as follows: in Section 2. we draw out the of the model, explain the health updating rules, and the matching process used in the ABM. Section 3. then explains the implementation of the model in the NetLogo environment; whereas Section 4. presents some simulation results under both the complete, and incomplete, information settings. Next, Section 5. outlines some feasible extensions of the model, while Section 6. concludes.

1.1 NetLogo and ABM

Before addressing the model specifications in detail, we briefly introduce the modeling platform used, NetLogo, and also comment on why ABM is appropriate for dealing with the problem at hand as opposed to other (GE) economic modeling.

NetLogo[1](Wilensky, 1999; Tisue and Wilensky, 2004) is a programmable modeling environment for simulating interactive phenomena–be it natural or social. It is particularly well suited for modeling complex systems developing over time, and its ability to give instructions to hundreds (or even thousands) of agents enables it to explore connections between micro-level behavior and the macro-level consequences that result as a consequence of such interactions. Another convenient feature of the platform is that it allows for a thorough study of the model's parameter space by means of experimentation. NetLogo has a built in software tool called "BehaviorSpace" that allows the user to perform recursive experiments with the model, while systematically varying the model's parameter settings. Recording such data allows us to construct distributions of the simulation outcomes, and also see how different parameter combinations effect the behavior of the model. The study of interaction between multiple heterogenous agents, and experimentation of this kind, is generally not feasible with the mainstream economic models.

Mainstream economic models, as put by Gallegati (Gallegati, 2008), are based on the classical physics assumptions of: reductionism, determinism and mechanism. Under this approach there is no difference between micro and macro since the whole is nothing but a summation of its components. In order to arrive at tractable solutions, these General Equilibrium and Representative Agent models often need to assume that all agents are homogenous in nature

---

[1]NetLogo was developed by Uri Wilensky in 1999 and is still in continuous development. See: http://ccl.northwestern.edu/netlogo/ [accessed 30 May 2012]

and behavior, and thus predictable. In contrast to this, the ABM methodology embraces heterogeneity of agents and caters for the study of aggregate dynamics and empirical regularities that are not know *a priori*. Since the problem of health care diagnostics concern the complex interplay of heterogenous health care professionals as well as heterogenous agents (as will soon be explained in greater detail) the choice of an ABM approach seems natural and called for. Furthermore, as was mentioned, many standard models fail to capture evolutionary behavior over time, while an ABM is capable of capturing this in great detail, giving insight into the unfolding of our final outcomes.[2]

2 A Simple Public Health Model

2.1 Model Setup

Agents come in two types (or breeds) in this model: patients and doctors. The patient population is composed of agents $i \in \{1, 2, ..., N\} = \Omega$, and the health care professional (which we call doctor) population is made up of agents $j \in \{1, 2..., M\} = \Theta$. Patients initial health level at time $t = 0$ is given by the random variable $h_{0i} \sim N(\mu_h, \sigma_h^2)$ which is normally distributed for all i. We further take that patient utility is a function of lifetime (inter-temporal) health, and hence given by:

$$u_i = u_i(h_{0i}, h_{1i}, ..., h_{Ti}) = u_i(H_i) \qquad (2.1)$$

Where H_i is the lifecycle health of agent i. Next, we say that all agents with $h_{ki} \in (0, \mu_h - \sigma_h)$ are sick with a condition $c_i \sim U(a, b)$, where the conditions are taken to be uniformly distributed. When endowed with a condition, we assume that the agent attempts to find professional help.

Doctors, have an area of expertise given by the random variable $e_j \sim U(a, b)$ that he/she is able to provide a diagnosis for. The scope of expertise of any one doctor can be expanded by doctors engaging in collaboration with one-another. For instance, if doctors $1, 2, 3..., m$ are taken to collaborate, (denoting this as: $A = \{1, 2, 3, ..., m\} \subseteq \Theta$), then we assume that they all have access to the collective expertise $E = \bigcup_{j \in A} \{e_j\}$ which they can readily make use of in their patient consultations. [3]

This gives us the basic setup of the patient doctor population, with their respective characteristics. Next, we look at the health updating rules, how to derive the lifetime–individual and aggregate–health, and explore the matching algorithms by which sick patients are matched with a doctor.

---

[2]While one can conceive a large scale longitudinal study as an alternative to an ABM approach to the problem at hand, such a study would be preventively complicated to manage, and would require several years of dedicated investment of both time and financial resources. As such, the ABM simulation approach seems to lend itself well to this problem, at this time.

[3]NB: E is the union of all individual expertise since some may be overlapping, e.g. if $E = \{e_k\} \cup \{e_l\}$ where $k \neq l$ but $e_k = e_l$, then $E = \{e_i\} = \{e_j\}$ and there are no evident benefits from cooperation.

2.1.1 Health Decay When Healthy and Sick

Starting with the patient population, if a patient is healthy, then his level of health decays at a rate γ per period t of time, while the decay is δ in the case that the patient is sick. We assume that $\gamma \leq \delta$ holds true. Hence, in the case that the patient is healthy, we use the following health updating rule:

$$h_i(t+1) = h_i(t) - \gamma \tag{2.2}$$

and in the case of where the patient is sick (and not in a consultation–explained below), the updating rule is:

$$h_i(t+1) = h_i(t) - \delta \tag{2.3}$$

If patient is sick, he/she will employ a search algorithm in order to try and find a doctor to help him/her receive a diagnosis. Before explaining the search process employed, lets first attend to the mechanism of a consultation, and to how we derive the lifetime health.

2.1.2 Consultation Updating Rule: with (and without) Cooperation

Supposing that patient i is in a consultation with doctor j at time t, then the health of the patient at time $t+1$ is given by:

$$h_i(t+1) = \begin{cases} h_i(t) - \delta + \alpha x & if \ c_i = e_j \\ h_i(t) - \delta & if \ c_i \neq e_j \end{cases} \tag{2.4}$$

Here, $h_i(t)$ denotes patient i's health at time t, and x is a random variable $x \sim U(p,q)$ with a uniform distribution denoting the benefit that the diagnosis has on agent i's health, with $\alpha \in [0,1]$ being a parameter indicating the level of accuracy of the diagnosis. If, on the other hand, doctor j collaborates with other doctors then we rewrite the above updating rule as:

$$h_i(t+1) = \begin{cases} h_i(t) - \delta + \alpha x & if \ c_i \in E_j \\ h_i(t) - \delta & if \ c_i \notin E_j \end{cases} \tag{2.5}$$

Different from before, patient i is here diagnosed so long as at least one of the doctors (in the team) has expertise on his condition, i.e. $c_i \in E$.

2.1.3 Individual Lifetime Health and PopulationAggregate Health

Let the total lifecycle health of patient i be given by:

$$H_i = H_{1i} + H_{2i} + H_{3i} \tag{2.6}$$

Here, H_{1i} denotes the aggregate health accumulated during the period η that the agent is healthy. It is given by:

$$H_{1i} = h_0 + (h_0 - \gamma) + (h_0 - 2\gamma) + \cdots + (h_0 - \eta\gamma) = h_0(\eta + 1) - \left(\frac{\eta(\eta+1)}{2}\right)\gamma.$$

Let κ denote an exogenously given threshold value of health beyond which the patient becomes sick. That is, if $h_i \leq \kappa$ then i is sick. Thus, setting $h_0(\eta + 1) - \left(\frac{\eta(\eta+1)}{2}\right)\gamma = \kappa$ and solving the quadratic we get η, and since the initial condition h_0 and the decay parameter γ together with κ are known, so is η. As such, we can solve for H_{1i} for any i. Next, H_{2i} denotes the aggregate health while sick. This can be solved similar to the previous case, however, here we set:

$$H_{2i} = h_0(\zeta + 1) - \left(\frac{\zeta(\zeta+1)}{2}\right)\delta = 0$$

to solve for $\zeta = \frac{2h_0}{\gamma}$, and therethrough we get H_{2i}. The final part of the total lifecycle health for a given patient is given by:

$$H_{3i} = P(diagnosis) * E(Z)$$

where $P(diagnosis)$ is the probability that i will be diagnosed over his lifetime and hence depends on the search (matching) algorithm employed. $E(Z)$ is the expected health benefit from obtaining a diagnosis, where $Z \sim U(a, b)$. Note, below we solve H_{3i} for each of the two search algorithms.

As such, we have that $H_i = H_{1i} + H_{2i} + H_{3i}$ has its first two parts, H_{1i} and H_{2i}, fully determined by the initial condition of health, h_0, while H_{3i} is the more interesting part that depends on the underlying search algorithm used, as well as the patients expected health benefit from a successful diagnosis. From this result, we simply write the aggregate lifecycle health in society as:

$$H = \sum_{i=1}^{n} H_i \tag{2.7}$$

It is worth pointing out that while we say that aggregate health is simply the sum of all the individual health levels, the patient health is not homogenous. Instead, we have a large group of unique patients with heterogenous health conditions.

With the health updating rules and aggregate health defined, we now look at the search algorithms and define the specific form of H_{3i} for each of these.

2.2 Search (Matching) Algorithms

2.2.1 Incomplete Information (Random Search):

In this setting sick agents move across the (simulation) space at random. If one of the patients happens to be within a specified radius r away from one of the doctors, i.e. if the distance is $d(i, j) = |i - j| \leq r$, and the doctor is

available, then patient i and doctor j engage in a consultation. In so doing, either (2.4) or (2.5) above is invoked as the updating rule depending on whether the doctors cooperate or not. If successful in receiving a diagnosis the patient stops, otherwise he continuous until he either receives a diagnosis, exhausts all doctors, or dies. In summary, this model can be thought to reflect incomplete information on behalf of the patients when it comes to their ability to directly locate a doctor, and where there is a significant amount of preliminary search conducted by the patient before there is a meeting.

Given this algorithm, we have that the probability of being diagnosed at any given period t (while i is searching) is given by: $P(diagnosis|t) = \frac{m}{p} * \frac{|E_i|}{m} = \frac{|E_i|}{p}$, where m is number of doctors, p number of states (different types of patches) in the illness space, and $|E_i|$ is the cardinality of $E_i \subset E$ where E_i contains all the doctors with relevant expertise for patient i. Solving the lifetime benefit from health care, H_{3i}, we get:

$$H_{3i} = \left(\frac{|E_i|}{p} + \left(\frac{p - |E_i|}{p} \right) \frac{|E_i|}{p} + \cdots + \left(\frac{p - |E_i|}{p} \right)^\zeta \frac{|E_i|}{p} \right) E(Z) =$$

$$= \left(\sum_{k=0}^{\zeta} \left(\frac{p - |E_i|}{p} \right)^k \frac{|E_i|}{p} \right) E(Z) = \left(\frac{|E_i|}{p} \cdot \frac{1 - \left(\frac{p - |E_i|}{p} \right)^k}{1 - \left(\frac{p - |E_i|}{p} \right)} \right) E(Z)$$

which, as we can see, simplifies since it is a simple geometric series.

2.2.2 Partially Perfect Information (Non-random Search):

Although the patients still have incomplete information regarding their own condition, and the expertise of the doctors, they have perfect information on where each of the doctors is located. As such, the patients employ a more structured approach toward trying to get a diagnosis from the doctor population. To start, patient i chooses doctor $j \in \Theta$ randomly out of the doctor population. In the case that doctor j is available they have a consultation, again invoking the relevant updating rule of (...) or (...). If the patient is unsuccessful in receiving a diagnosis from j then he picks a new doctor $k \in \Theta \setminus \{j\}$ and repeats the process until he either receives a diagnosis, exhausts the population of doctors, or dies.

The expected benefit from health care for patient i is in this case calculated as:

$$H_{3i} = \left(\frac{|E_i|}{m} + \left(\frac{m - |E_i|}{m} \right) \frac{|E_i|}{m - 1} + \cdots + \left(\frac{m - |E_i|}{m} \right)^\zeta \frac{|E_i|}{m - \zeta} \right) E(Z) =$$

$$= \left(\sum_{k=0}^{\zeta} \left(\frac{m - |E_i|}{m} \right)^k \frac{|E_i|}{m - k} \right) E(Z)$$

6

where the terms are as previously defined, and the summation is taken as follows: if $m > \zeta$ we stop at the time t when $h_i = 0$ (the agent dies); and conversely, if $m < \zeta$ we stop when $m = 0$ (all doctors have been exhausted).

Since both of these matching algorithms are of interest we run our ABM simulation experiments on both separately. The next section explains the specification of the outlined model in the NetLogo environment before presenting the results of our simulations.

3 NetLogo Model Implementation

Having outlined the basic logic of our model, we now turn to look at its implementation, and specification, in the NetLogo environment.[4]
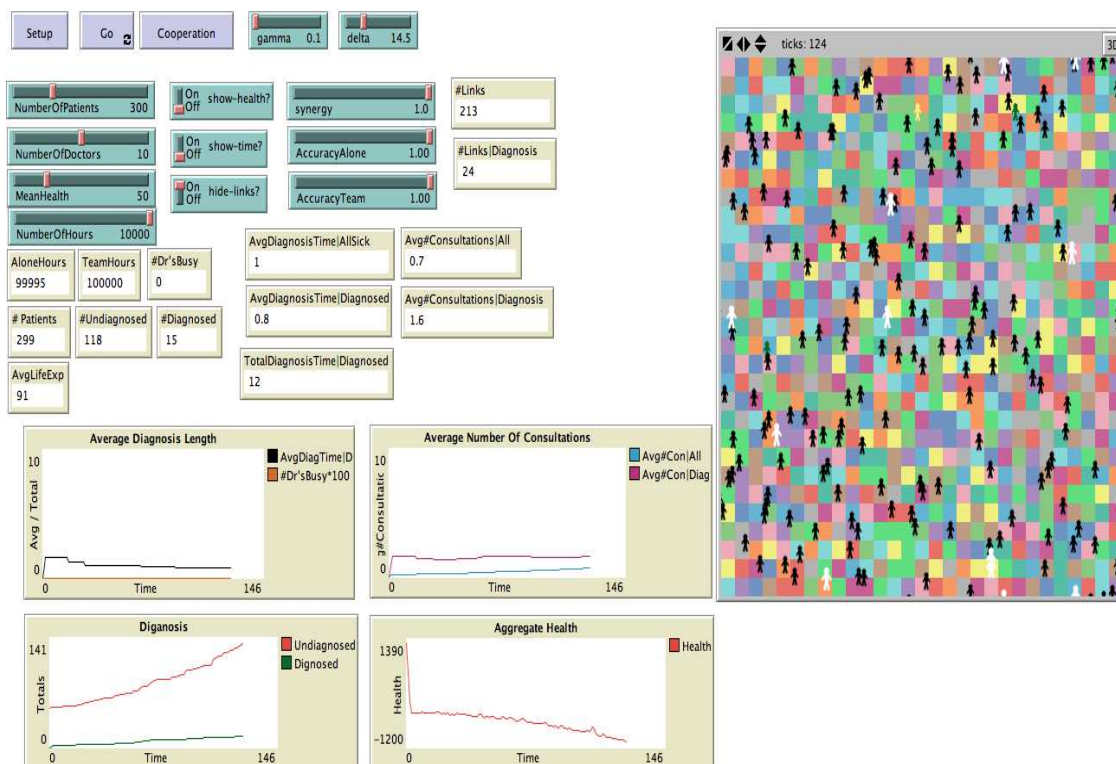
3.1 Interface Overview

Figure 3.1: Netlogo Interface

---

[4]The model source code can be found at:...

Figure 3.1 shows the NetLogo interface. On the LHS we see the buttons/switches that allow us to set the initial parameter values of the model. Varying these allows us to fully explore the parameter space of the model. While the program is running, one can study the evolving behavior using the plots (seen in the lower left corner), and further view the interaction of the agents in the simulation world (the large square on the RHS).

In the agent simulation window, the white agents are doctors, black agents are healthy individuals, while colored agents are sick (with their color indicating the particular condition). Meanwhile, we also note that the space is made up of colorful patches–these indicate the expertise space of the doctors. For example, if doctor j is standing on a red patch, then this means that he has expertise on condition "red", and as such he can diagnose patients with this illness. If doctors collaborate, and say, doctor j stands on a red patch, and doctor k stands on a green patch, then they are both able to diagnose conditions "red" and "green". This follows from our assumption that if doctors collaborate this means that they have access to each other's expertise.

Below, we go into a more through treatment of the Setup and Recursive procedures of the ABM model.

3.2 Setup Procedures

The Setup button in the upper RHS of Figure 3.1 initiates the patient and doctor procedures. The patient setup procedure, creates the number of patients that has been specified (using one of the sliders below the Setup button), randomly allocates them to a patch on the space, and assigns each an initial health endowment from a normal distribution of health (who's mean is set by one of the sliders). If a given patient is allocated a health level that is less than one standard deviation from the mean health he is considered sick, and as such, assigned a condition from a uniform distribution of illnesses. The type of condition is indicated by the patient's color, however, the patient is assumed unaware of the specifics of his own condition.

Doctors are instantiated in a similar fashion. Their number can be specified by one of the sliders, and they are all randomly assigned to a patch who's color indicates that particular doctors area of expertise. By pressing the Cooperate button (to the right of the Setup button) we indicate that we want the doctors to cooperate in their respective efforts to diagnose the patients. As seen in Figure 3.2, the presence of a cooperation regime is visually displayed by white links that indicate a stream of two way information (expertise) flow between the health care professionals.

Last, we have that setup of the patches. These are randomly allocated a color (same as the conditions) from a uniform distribution of colors (conditions).

3.3 Recursive Procedures

Once the setup of the model has been implemented, the recursive procedures dictate the rules by which the system evolves over time. These rules are initiated

Figure 3.2: Presence (left) and Absence (right) of Cooperation Between Doctors

by pressing the Go button (just right of the Setup button). The nature of the evolving process for the agents depends on the search algorithm we choose for the model, and as such we here take a look at both of these.

Under the Incomplete Information setting the patient procedure work to have the sick patients randomly move around the space (as seen in Figure 3.2). When the patient lands on a patch of a doctor they engage in a consultation. The consultation is successful in providing the patient with a diagnosis in the case that the doctors area of expertise is the same as that of the patient (in the case of no cooperation between doctors), or when it matches that of at least one of the doctors expertise (in the case of cooperation). If the patient receives a diagnosis of his condition we set the patients color to green, and set his health level to increase by an amount drawn from a uniform distribution. The idea here is that the benefit of a diagnosis may have a varied effect on different people, since in some cases a diagnosis may help the agent to take measures that helps improve his/her health, but this is not necessarily true in all cases.

Under complete information, the matching process is far less arbitrary. Here, the agents randomly chooses a new available doctor in each period until: a successful diagnosis is provided, all doctors are exhausted, or the patient dies. The particular why in which this actually works is as follows: each period an agent randomly draws a doctor, if he is available, i.e. not in a consultation with somebody else, then the patient forms a "grey" link with that doctor (as seen in Figure 3.3). The next period this repeats, however, the agent will not form a link with any doctor that he has previously linked with.

9

Figure 3.3: Complete Information Doctor and Patient Matching

3.4 Analysis and Behavior Space

The NetLogo model allows for continuous tracking and analysis of the evolving system over time. This is a nice feature, and as can be seen in Figure 3.4, it allows for the study of how, for example, the number of diagnosed patients and aggregate health changes over time in our simulated society. It further allows for the comparison of such inter-temporal change across different setups, for example, in the presence, and absence, of cooperation between doctors. Again, this can seen in Figure 3.4 by comparing the top level output plots (for the setting of cooperation), with those below (for no cooperation).

An additional feature of NetLogo is the possibility of fully exploring the parameter space of the model. The Behavior Space feature allows one to run, and record, the model hundreds of times and as such analyze the effect of different parameter changes upon the model behavior. In the next section, titled Results, we document the main findings of our model, and further explore the robustness of these results to changes in the parameter setup.
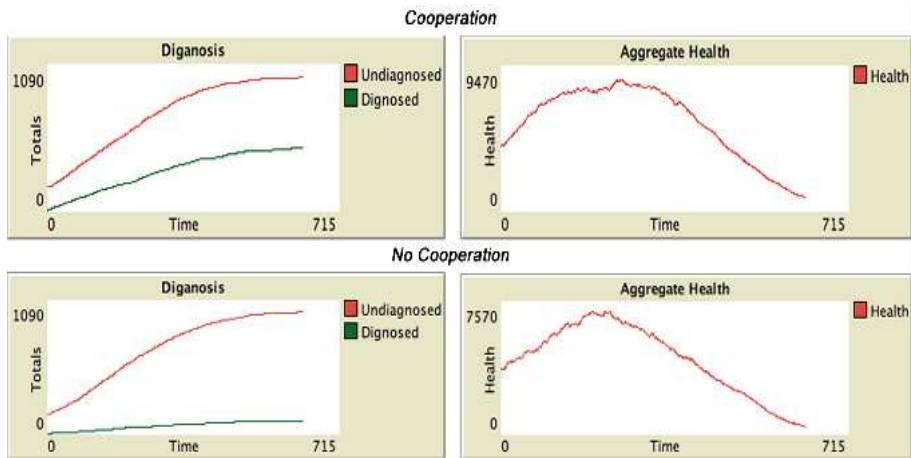
Figure 3.4: Diagnosis and Aggregate Health Outcomes (Incomplete Information)

4 Results

Here we look at the model under incomplete and complete information. Each of these settings are first analyzed separately before being compared. The analysis consists of comparing the overall diagnostic efficiency of the system in the presence and absence of cooperation. In so doing, we look at: the number of patients successfully diagnosed under each setting, the level of aggregate health, the average time it takes for a patient to be successfully diagnosed, and the average number of consultations required for a successful diagnosing to result.

4.1 Incomplete Information Setting

The basic model setup used as a frame of reference here in the results is:

Parameter:	Level:
Number Of Patients	1000
Number Of Doctors	10
Mean Of Health Dist.	50
Condition/Expert Space	14
Accuracy Alone	1
Accuracy Team	1
Synergy	1
Gamma	0.1
Delta	0.2

Table 1: Initial Settings

4.1.1 Number of Diagnosed Patients

The model was run 20 times, with and without cooperation, in order to remove any spurious results brought about by the random process that underpin the model. The results show a statistically significant (p-value $= 2*e^{-16}$) difference between the number of patients successfully diagnosed when there is, and is no, collaboration among doctors.


Figure 4.1: Number of Patients Diagnosed Under Cooperation

As illustrated in Figure 3.1, we see that the number of successful diagnosis under cooperation is almost four times that achieved when doctors do not work in teams. In order to investigate the robustness of this finding we looked at the sensitivity of the result to changes in the mean health level, μ_h, and the level of decay of health, δ, when the patient is sick. Increasing the mean health level has the effect that agents have a longer time to search for a doctor with the proper expertise, hence it improves the chances of the patient being diagnosed. this is seen in the left image of Figure 3.2.

From the figure we see that the statistically significant difference between the number of patients diagnosed under the two settings of cooperation remains. In fact, as patients are given more time to search in this incomplete information setting, the probability that they be diagnosed in a collaborative setting grows greater (as illustrated by the fact that the slope of the red line is steeper than that of the blue in the left graph of Figure 3.2).

Increasing the level of health decay while sick, δ, (see RHS graph in Figure 3.2) has the effect of decreasing the number of diagnosed patients under both regimes of cooperation. This is because the level of health decay has an inverse relationship with the amount of time that the patient has in-order to find a doctor, let alone the one with the right expertise. Nevertheless, a significantly greater number of patients are still being diagnosed in the setting of
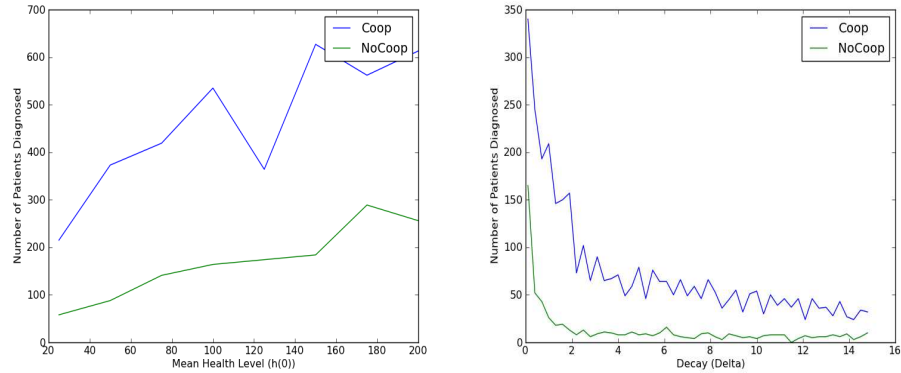
Figure 4.2: Sensitivity to Number Diagnosed Due to Changes in Mean Health and Decay

collaboration.

4.1.2 Level of Aggregate Health

Since cooperation, as just seen, implies a larger portion of diagnosed patients this also translates into a larger aggregate level of health in society at large. The scale and significance of the effect that collaboration between health care professionals may have here is, however, dependent on several factors previously described in Section 2.

From the experiments run in the ABM we see the following:

- The size of the health benefit, $E(Z)$, that a diagnosis can provide is positively correlated with overall health.

- The size of the condition/expertise space is inversely related to the level of health in society. That is, as the number of possible conditions grow, *ceteris paribus*, this decreases the chances of a patient being diagnosed since the probability of any doctor being an expert on that condition is reduced.

- More doctors translate into a higher probability of diagnosis, and thus into a larger health level overall.

- If initial condition of health, h_0, is increased, it provides an overall health benefit that is more than $1 : 1$ to the increase in h_0.

4.1.3 Time to Diagnosis

We find that although the population of diagnosed patients under cooperation is significantly larger (p-value $= 4.09 * e^{-15}$) than that under no cooperation,

Figure 4.3: Time Taken for diagnosed patients to get their diagnosis

the speed with which the patients are diagnosed when doctors collaborate is still significantly faster. This difference appears stable, but the overall timing of diagnosing changes as follows when we change the parameters of the model:

- increasing the number of doctors is inversely related to the timing of diagnosis.

- Increasing the size of the illness space increases the average time of diagnosis.

- Likewise, increasing the time allowed for search also increases the time to diagnosis since a lot of patients that previously died without receiving a diagnosis are now diagnosed.

- The time taken is not dependent on the number of patients in this simple setup.

Seeking to reduce the time to diagnosis is important for several reasons: firstly, it may help increase the life expectancy of the patients by means of early intervention. Second, it may prevent the spread of disease that would otherwise affect the health of other individuals in society too, and thereby have a significantly adverse effect on overall aggregate health. And Third, longer time is often reflective of an inefficient health care system with a lot of deadweight costs and inefficient use of health care resources. We can look a this cost in more detail by studying the number of doctor consultations each of the diagnosed patients need to go through in order to get their diagnosis.
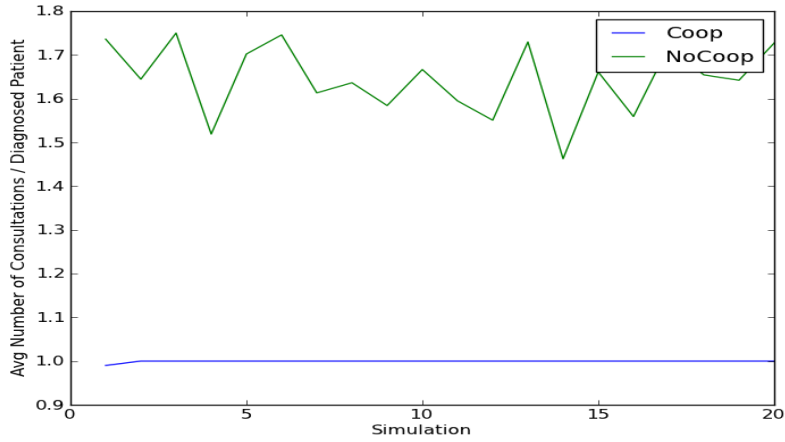
Figure 4.4: Number of Consultations for Diagnosed Patients to get their Diagnosis

4.1.4 Number of Consultations to Diagnosis

We use the term "deadweight loss" to signify expenditure that would have been avoided by the consumer had he/she better information available. For example, in the case that a patient i with condition c_i is matched with a doctor j that has expertise e_j, where $c_i \neq e_j$, then both parties would be better of from avoiding this encounter–the patient will not get his moneys worth since he does not get diagnosed, and the doctor could be allocating his time towards other patient whom he can actually help diagnose. Figure 3.4 above shows that the average number of consultations that a patient needs to go through (and hence pay for) when doctors work in a teams, as opposed to when they work alone, is significant (p-value = $2*e^{-16}$) . This difference grows even larger when we allow agents more time to search in this algorithm, because here agents that get matched at most only get to see 2 doctors, meaning that if they are to accurately get diagnosed they need to be lucky in ending up with the right doctor from the start.

Nevertheless, we see that significant efficiency improvements are established under the cooperative, making it attractive not only from a public health perspective, but furthermore from an economic standpoint too.

4.2 Partially Perfect Information Setting

The basic model setup used as a frame of reference here is the same as before, the only difference is in the search algorithm employed.

15

4.2.1 Number of Diagnosed Patients

Figure 4.5: Number of Patients Diagnosed

Unlike in the incomplete information setting, here, the search algorithm employed by the patients is so effective that the total number of patients is not significantly different given our initial setup (p-value = 0.18) . However, this result is no robust in that it is highly sensitive to the level of decay chosen for when the patient is sick. That is, a high enough δ, which implies that in a large enough population of doctors the patient must find the right doctor fast, and cannot rely on visiting every doctor before getting diagnosed.

From Figure 3.6, we see that the drop in the number of patients diagnosed when there is no cooperation drops significantly, while the overall level of diagnosis of cooperating doctors remains stable and robust to this change.

4.2.2 Level of Aggregate Health

Same analysis as for the previous algorithm is true here. Significant difference in the max level of health (p-value = 0.0185).

4.2.3 Time to Diagnosis

As illustrated in Figure 3.7, we see that there is a significant difference (p-value = $2*e^{-16}$) between the time taken to successfully diagnose a patient in the presence (and absence) of collaboration between doctors. Furthermore, comparing the result under complete information, to that of incomplete information, we see that the average time to diagnosis for patients that are diagnosed is now greater that it was under the incomplete information setup. The reason for this is that under complete information we employ a more efficient search algorithm. As
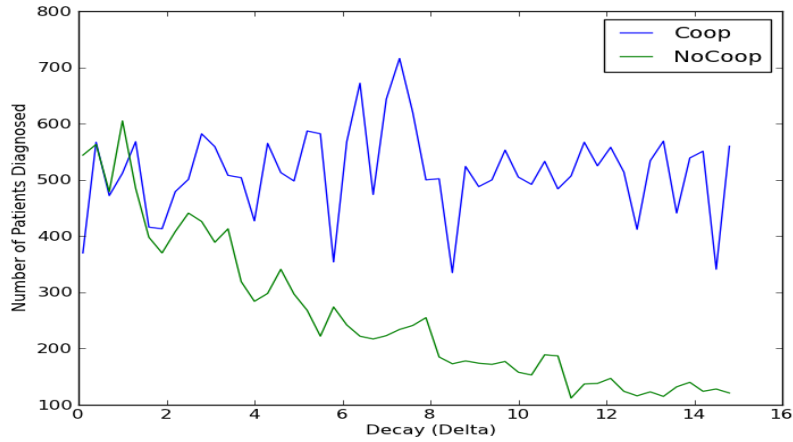
16

Figure 4.6: Number of Patients Diagnosed as δ is changed

such, we have that agents now get to meet more doctors in the time that they are alive, allowing agents that would have previously died a better chance at being diagnosed.

4.2.4 Number of Consultations to Diagnosis

Due to the nature of the search algorithm the number of consultations is similar to the time to diagnosis. This is also evident from the similarity of the Figures 3.7 and 3.8. And yet again, we note a significant difference (p-value = $2*e^{-16}$) in favor of collaboration between doctors. In fact, we see that cooperation brings with it a more efficient matching process, with fewer consultations, and thereby lower costs for the patients. Likewise, the efficiency of the matching process is also positive news for public health provision in that we have a better use of medical resources under a setting of cooperation.

4.3 Summary

While the models analyze the impact of cooperation under two different matching algorithms the results of both indicate that cooperation is to be preferred from a health and economic point of view in both cases. More specifically, in model 1 we saw a significant difference on every account of: number of diagnosis, overall health, time to diagnosis, and number of consultations. While Model 2 did not demonstrate a significant difference between the two settings in terms of the number of people being diagnosed, the difference becomes significant if we allow for a larger health decay for when the patient is sick. But even though the difference in number of diagnosed patients may not have been radically different, the speed with which they were diagnosed, and the number of consultations
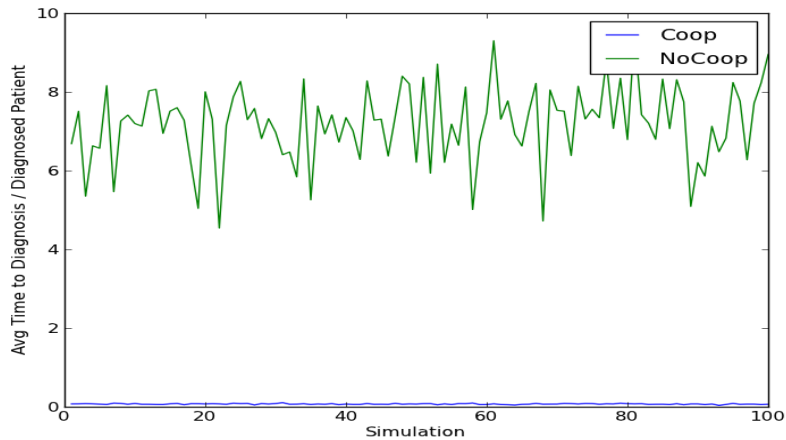
17

Figure 4.7: Time to Diagnosis for Diagnosed Patients

required for a diagnosis, was significant.

These results seem to suggest that cooperation between health care professionals on a micro-level can have positive macro-level effects both in terms of overall public health, as well as on the economic expenditure on health care.

5 Possible Extensions

The following list provides some interesting extensions for future work:

- **Make the model multi-generational.** The current model only reviews one generation, so allowing for new generations to come into being may be of interest, and would allow for an overlapping generations type of study to be made. With that said, the additional benefit that such a study could give for the purpose at hand seemed minimal, and was as such not pursued.

- **Allow patients to cooperate and share information.** Here we assume that patients do not leverage their social networks in order to find a suitable doctor that can evaluate their symptoms. However, in the real world collective intelligence is often at play, and could be a factor helping to speed up the matching process when doctors do not fully cooperate. This seems to me a fruitful avenue of future work as it could yield insight into the relationship between social structures and their implications for favorable economic and health outcomes. In a simple case, one can assume that each of the agents has information about a limited number of doctors, together with their area of expertise, and that he/she can share this with

18

Figure 4.8: Number of Consultations to diagnosis for diagnosed patients

immediate neighbors, and further with the neighbors of neighbors with some level of information decay accounted for.

- **Adding health care budget.** Every visit to the doctor costs patients money, either directly or indirectly via insurance. Including such a budget could limit the number of doctors any given patient may sample in order to get a diagnosis. While the model indirectly captures economic cost by the way of counting the number of consultations that a patient engages in, it does not constraint the number of such encounters that the agent can engage in–something that a budget would do. This basic extension (or modification) can easily be brought about by E.g. adding cost K_j for each consultation with doctor j. This would then be added to the utility function so that the agent weight expected benefits from a possible diagnosis against the costs of a doctors appointment.

- **Appointment/queuing system.** The model does not block out more than one patient seeing a doctor at a given time, so including an appointment system could help with this shortcoming.

- **Time budget for Doctors.** This introduces a constraint on the number of patients that any given doctor can see. We have included this into the code of the ABM, but excluded it out of the analysis for the purpose of keeping things as simple as possible. Nevertheless, working with a time budget may be of interest, especially if we consider that it may take more/less time for doctors to see patients in a team, verses in a one-on-one, setting. Then this modification could provide better insight into the cost benefit tradeoffs (if any) that come from health care professionals working in teams.

19

- **Accuracy of Diagnosis.** Aided by real world data on what the potential diagnosis accuracies are of a doctor working in a team as opposed to alone would be of great interest, and would help make the analysis more accurate. The literature on malpractice in the medical profession seems to here suggest that teams would have a higher rate of accuracy–something that is brought about by the increased level of peer review when working in a team environment.

6   Concluding Remarks

Micro-level behavior can have significant Macro-level effects. In the model presented here we looked at the effect that cooperations between health care professionals on a micro-level can do for the overall level of health in our model, on the number of patients successfully diagnosed, and the efficiency with which they are diagnosed. In doing this, we looked at two models–one of complete, another of incomplete, information. The results under each setting, although different, seems to unanimously suggest that collaboration between health professionals is indeed something of importance. As we saw, an efficient system results in a lower level of deadweight expenditure, which translates into less healthcare spending for patients and less unnecessary use of scarce medical resources.

While there are numerous possible extensions to the model (as outlined in Section 4.) that can be investigated with future work, there is also a need for more research into measuring the real world efficiencies of medical teams. Such research could provide data with which the current model could be compared, to further see if the predictions it makes are reasonable.

Appendix

Online Resources:

To explore the NetLogo model interactively visit: http://ccl.northwestern.edu/netlogo/ to first download NetLogo.

Next, visit our Bitbucket repository at: https://bitbucket.org/slinde/collaboration-and-health-care-diagnostics-an-agent-based-model/src (here you can download the Netlogo model that we developed and explored in this paper).

References

Aspden, P., Wolcott, J., Bootman, J., and Cronenwett, L. (2007). *Preventing Medication Errors: Quality Chasm Series.* The National Academies Press.

Gallegati, M. (2008). Agent based models in economics and complexity.

Graber, M. L., Wachter, R. M., and Cassel, C. K. (2012). Bringing diagnosis

into the quality and safety equations. *JAMA: The Journal of the American Medical Association*, 308(12):1211–1212.

Hayek, F. (1945). The use of knowledge in society. *The American Economic Review*, 35(4):519–530.

Kohn, L. T., Corrigan, J., and Donaldson, M. S. (2000). *To err is human: building a safer health system*, volume 6. Joseph Henry Press.

Newman-Toker, D. E. and Pronovost, P. J. (2009). Diagnostic errors—the next frontier for patient safety. *JAMA: the journal of the American Medical Association*, 301(10):1060–1062.

Singh, H. and Graber, M. (2010). Reducing diagnostic error through medical home-based primary care reform. *JAMA*, 304(4):463–464.

Tisue, S. and Wilensky, U. (2004). Netlogo: A simple environment for modeling complexity. In *International Conference on Complex Systems*, pages 16–21.

Wachter, R. M. (2010). Why diagnostic errors don't get any respect—and what can be done about them. *Health Affairs*, 29(9):1605–1610.

Wilensky, U. (1999). *NetLogo User Manual*. Northwestern University.