Risking reputations: when to endorse new drugs?

Marion Dumas, Jameson Watts

June 20, 2013

1 Introduction

In this short project, we decided to focus on the strategies doctors might adopt to manage their reputation when faced with the choice of publicly adopting a new prescription drug. The intuition is that more public figures face a tradeoff when deciding whether to take a public stance on an issue when it is still unclear which stance is going to prove the most popular in the long run. An early runner of a stance (or judgement) that eventually becomes popular (or correct) is often ex-post judged positively for having been a forerunner, able to judge things correctly before anyone else. However, forerunners whose ideas don't immediately stand up to the test of experience are often shunned in the process of collecting this experience. We are interested here in exploring this tension. Notably, we do not assume that doctors' adoption decision influence others' decisions. Instead, they adopt waiting time strategies to manage their reputation. Thus, our model is not one of product diffusion. In addition, doctors' adoption decision does not influence people's perception of how well the drug functions (they are not shaping collective opinion about the quality of the product). In this model, we assume that this information is independently accumulated through experience with the drug.

2 Model

Setting and assumptions There are N doctors and a sequence of M drugs. A drug is characterized by the mean satisfaction (μ) derived from it by the patient pool and the variability σ in this satisfaction among patients. When the drug is released, doctors do not know μ, but they know how controversial the drug is (they know whether σ is rather high or low). Doctors observe other doctors' drug adoption behavior, as well as their success with patients to whom they've prescribed the drug. Because each doctor only observes a small number of doctors relative to the full doctor population, the pattern of who observes whom defines a network. This observational network determines the size of each doctor's audience and we assume that doctors care about the size of their audience, which we take to be equivalent to reputation. We model a process in which a doctor who adopts early on a drug that eventually turns out to be highly valuable can hope to see his audience increase, but in the process can get his reputation damaged if his patients are observed to be disappointed with the product by those who observe him. In modeling this process, we ask whether agents will choose different waiting times for low and high risk drugs in order to manage their reputation, and whether the choice of waiting time will be affected by a doctor's reputation. In each drug release episode, there are T periods of play, at each one of which doctors either start prescribing a drug or instead remain undecided and observe other doctors. The strategies of interest are doctors' waiting time until adoption $P_w \in (1, T)$. In each period of play, doctors that have adopted the drug prescribe it to m patients and their success with the drug is assumed to be the sum of the satisfactions of these treated patients. When the drug variability is high, doctors thus face a higher risk. When doctors are unsuccessful with a drug, they lose some of their audience. However, at the end of a drug experimentation episode, doctors gain some audience thanks to the propagation of information about their success, as a function of both how good the drug proved to be and how early they chose to adopt it.

Algorithm

1. Initialization: doctors are initially connected according to an Erdos-Renyi graph of density n. Each doctor has a "low risk drug" and a "high risk drug" strategy (i.e. waiting times P_w, denoted H and L hereafter). Both strategies are initially randomly distributed.

2. For each of M drug experimentation episodes:

   (a) A new drug comes on the market, characterized by the distribution of satisfactions s_i patients will derive from it, assumed to be $N(\mu, \sigma)$.μ is drawn from $U(0, max_\mu)$ and unknown to the doctors, while σ is drawn from $U(1, max_\sigma)$. Doctors perceive the drug to be a high risk drug if $\sigma > \tau_\sigma$ and low risk otherwise.

   (b) In each period, from 1 to T, each doctor either adopts the drug or doesn't.

   - If he does, he prescribes it to m patients and augments his success score S_d by $\sum_{i \in patients} s_i$.
   - If he doesn't, he observes the doctors who have adopted it (and thus are in his observation network) and evaluate the likely value of the drug on the basis of the successes of these doctors, i.e. by observing $S_N = \sum_{d \in observed\ that\ have\ adopted} S_d$. If $S_N < 0$, he removes from his observation network the doctors that have adopted the drug, each with probability r.

   (c) At the end of the drug episode, μ, the true average value of the drug is revealed and doctors re-evaluate reputations by adding enough doctors to their observational network to maintain the same number n of out-degrees. To do so, each doctor looks to their 2-neighbhorhood (the doctors observed by those they themselves observe). If this pool of doctor is too small, they search additional doctors at random. They then pick doctors within that pool according to the following weighing function:

$$w_i = \frac{1}{P_{w,i}} \mu \tag{1}$$

   . Thus, the weight captures both the early-mover advantage and the value of the drug and doctors link up with doctors with high weights at a higher frequency than those with low weights.

   (d) Finally, doctors adapt their strategy ($P_{w,i} = H$ or $P_{w,i} = L$ depending on the drug's risk type) by moving their own strategy one period closer to that of the doctor in their network who had the greatest increase in reputation[1].

[1] originally, we let agents average their strategy with the most successful in their net-

3 Model runs

3.1 Change in strategies

Graphical representation Figure 1 shows the changes in strategies adopted for high and low risk drugs by doctors in low and high reputation states, respectively, varying the value of the drugs (the maximum value in the uniform distribution of μ) and the probability r of being "shunned" in case that the drug initially leads to patient dissatisfaction. Other parameters were set as indicated in table 1. The key observations are:

- Low risk drugs do not generate very much change in strategies and differentiation in behavior between low and high reputation doctors.
- On the contrary, high risk drugs do generate change and differentiation in behavior between low and high reputation doctors, at least initially.
- In the case of high risk drugs, the early moving strategy is initially associated with high reputation individuals.
- Over time, low reputation individuals learn the early moving strategy.
- The rate of learning is high when the drugs are drawn from a distribution that has a higher average value and when the probability r of initially being shunned by doctors is higher.

work, but the strategies converged too fast for agents to learn from changes in their reputation.

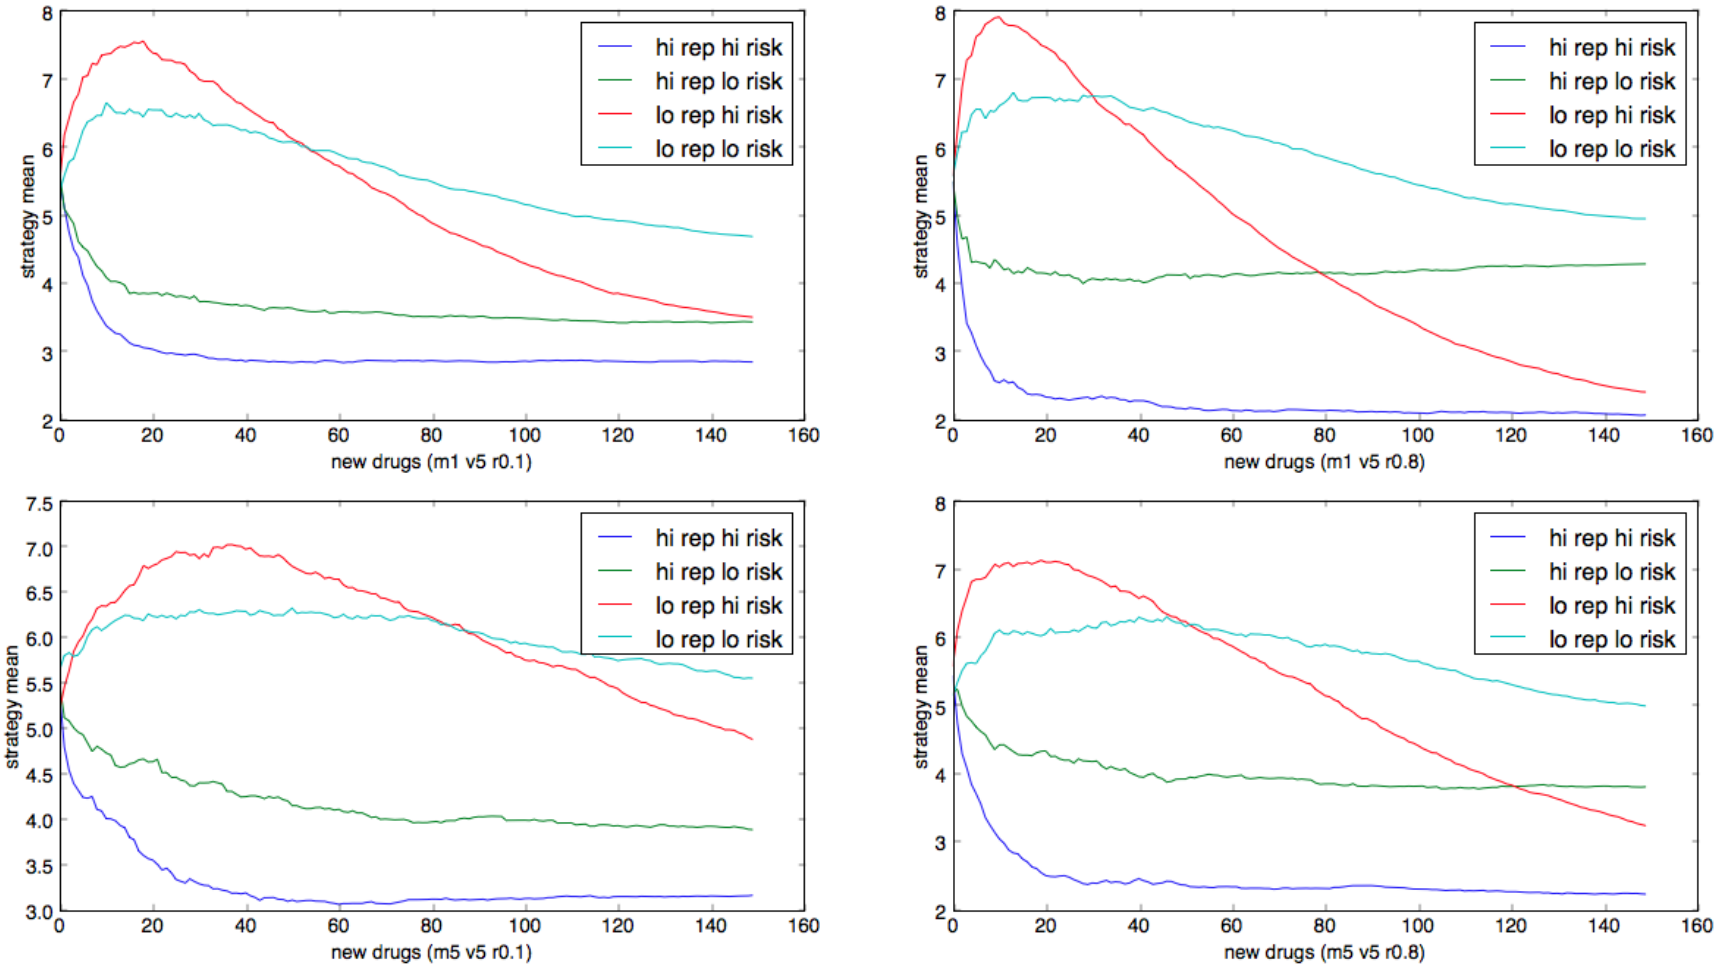

Figure 1: Dynamics of strategies, means over 25 runs.

parameter	value
M (drug episodes)	150
N (doctors)	100
n (network density)	0.2
T (periods of experimentation)	10
m (# of prescriptions per doctor and period)	5
σ_{max} (maximum variability of drug)	5

Table 1: Parameter values

Interpretation of the process What seems to be driving the dynamics we observe is that higher risk drugs lead to a lot of changes in the observational network, which in turn brings opportunities for early movers to expand their reputation at the end of a drug episode. This is because shunning has the effect of leading people to re-evaluate reputations once the true value of the drug is revealed and to cluster onto early movers. It is also in part due to our assumption that drugs all turn out to be valuable. Thus, people who initially have an early moving strategy gain in reputation early on and keep that strategy. Those with late moving strategies do not lose as much reputation (since they are not as exposed to the early variability of the drug), thus those with medium moving strategies may initially learn from them. However, because high risk drugs lead to a lot of turnover in doctors' follower networks, the early moving reputation of high reputation doctors is learned relatively fast by lower reputation doctors. In the graphs, it seems that high and low reputation strategies on high risk drug converge. It would be interesting to run the simulations longer because as they converge, we should expect some differentiation to happen again.
As a crude first confirmation that the turnover in reputation drives the process, Figure 2 shows the relatively obvious (but reassuring) result that when $r = 0$, strategies hardly change.

3.2 Changes in network configuration

As a result of these dynamics, the network radically changes. Figure 3 shows the distribution of incoming links (which is the size of a doctor's audience, or his reputation in our interpretation) in the first period, contrasted with that in the last period and also shows the growing range of reputations (by showing the increasing gap between the 25th and 75th percentile). We see that the

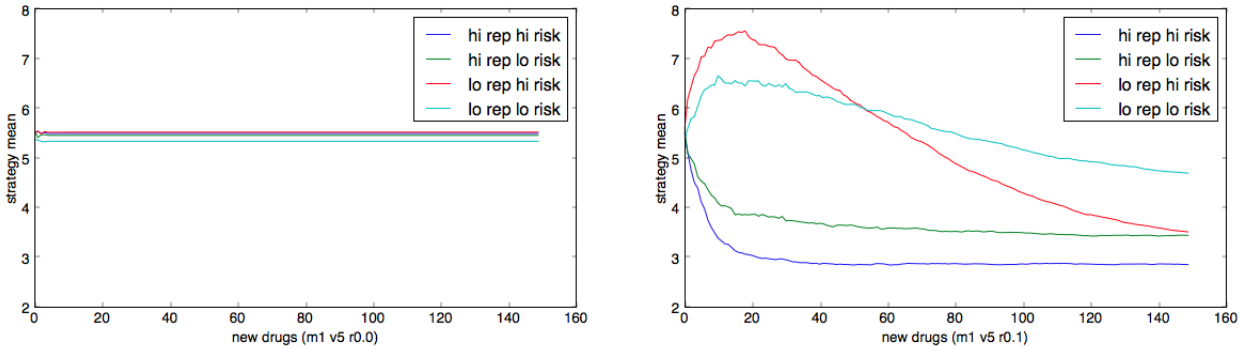

Figure 2: Mean over 25 runs when r=0, contrasted to $r > 0$.

frequency distribution of incoming links goes from that of a standard Erdos-Renyi graph (Poisson distributed) to extremely skewed. This is consistent with the idea that those with early moving strategies initially rapidly gain in reputation and that this process is self-sustaining.

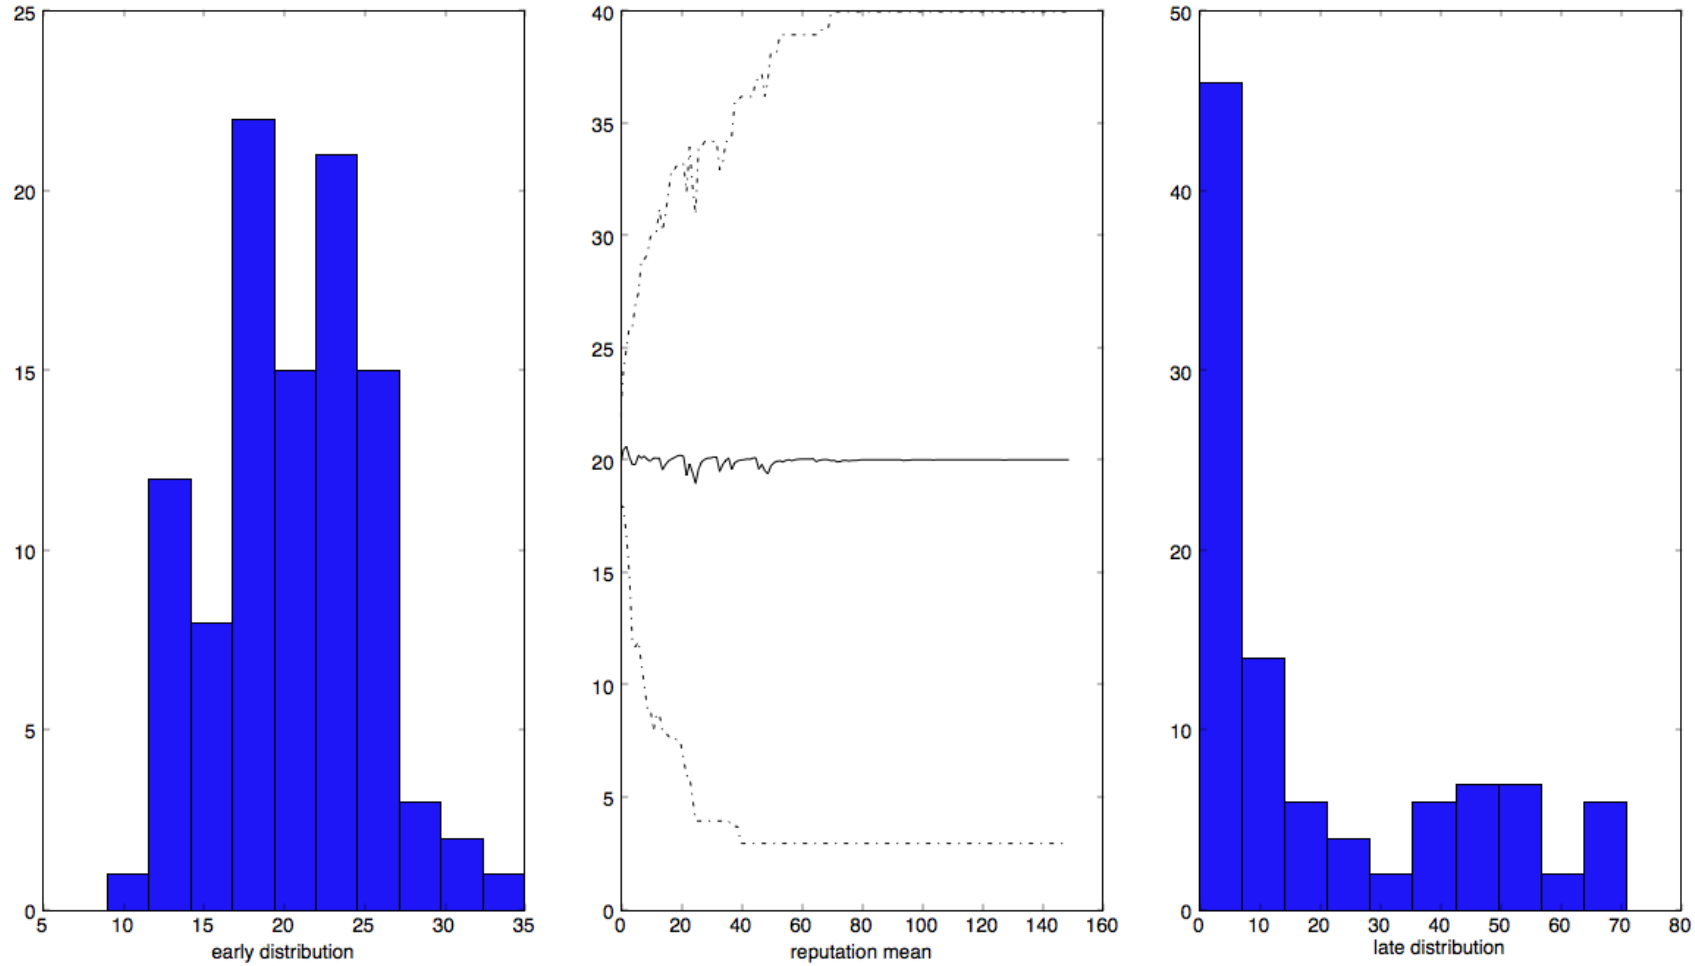

Figure 3: Change in in-degree distributions from first period to last period, and change over time of the spread.

4 Conclusion

Extensions The analysis of the process should be extended by verifying that our explanation of the underlying process is correct and on its basis predicting what should happen if we were to let the process run longer as well as examining that the individual changes in reputation and strategies are consistent with our hypothesized process. Another way of testing our understanding of the model is by predicting outcomes for a scenario in which drugs can also turn out to have negative value. In this context, we expect that high reputation individuals will become more conservative than observed here.
The model can be extended by allowing for more experimentation in the learning process, so that strategies don't stabilize too easily. Another is to allow the agents to learn the cutoff τ for differentiating high and low risk drugs. Finally, it would be useful to understand how much this process is driven by constraints we have or could have put on the network.

Generalizability The process we modeled here could apply to any situation where reputation is influenced by how one takes a public stance on issues that are perceived as initially uncertain but are then resolved and where it pays to be have been on the "right" side of how the issue is eventually resolved early on (thus being later seen as a smart forerunner). This could thus apply to public reactions to product crises on public comment platforms (e.g. Twitter). Of course it would generalize to risky products other than drugs, where experts are judged for the quality of their judgement vis-a-vis the product. This could perhaps apply to stances politicians take before a bill is voted on if they care about being on the winning side (as this may increase political capital for future coalition building efforts). The part of our model that is perhaps not obviously generalizable is the idea that reputations are more rapidly spoiled than they are constituted (as it stands, in our model, the deletion of observational ties happens in every period of the drug experimentation episode, which is more often and on the basis of less informational feedback than the creation of new observational ties, which happens at the end of each drug experimentation episode).