Learning from Others' Outcomes*

Alexander Wolitzky

MIT

January 9, 2018

Abstract

I develop a simple model of social learning in which players observe others' outcomes but not their actions. A continuum of players arrives continuously over time, and each player chooses once-and-for-all between a safe action (which succeeds with known probability) and a risky action (which succeeds with fixed but unknown probability, depending on the state of the world). The actions also differ in their costs. Before choosing, a player observes the outcomes of K earlier players. There is always an equilibrium in which success is more likely in the good state, and this *alignment* property holds whenever the initial generation of players is not well-informed about the state. In the case of an *outcome-improving innovation* (where the risky action may yield a higher probability of success), players take the correct action as $K \to \infty$. In the case of a *cost-saving innovation* (where the risky action involves saving a cost but accepting a lower probability of success), inefficiency persists as $K \to \infty$ in any aligned equilibrium. Whether inefficiency takes the form of under-adoption or over-adoption also depends on the nature of the innovation. Convergence of the population to equilibrium may be non-monotone.

*For helpful comments, I thank Daron Acemoglu, Abhijit Banerjee, Glenn Ellison, Jeff Ely, Drew Fudenberg, Ben Golub, Harry Pei, Frank Schilbach, and Juuso Toikka; seminar participants at Columbia, MIT, Princeton, PSE, the 2017 Becker-Friedman Institute Economic Theory Conference, and the 2017 Warwick Economic Theory Conference; and the anonymous referees. I thank the NSF for financial support.

1 Introduction

The development and diffusion of new technologies is a fundamental driver of economic growth. But some kinds of technologies seem to be introduced and adopted at a higher rate than others. In particular, innovations that improve observable outcomes—such as crop yields, health outcomes, or manufacturing output—while increasing costs are often more numerous and more successful than innovations that lead to worse outcomes but more than make up for this by reducing costs.

Consider the US healthcare sector, which in recent decades has experienced both rapid technological progress and rapidly rising costs (Newhouse, 1992; Cutler, 2004; Chandra and Skinner, 2012). Nelson et al. (2009) categorize 2,128 cost-effectiveness ratios from 887 medical studies published from 2002 to 2007 according to whether the studied innovations (i) increase or reduce costs, and (ii) increase or reduce quality of care, measured in quality-adjusted life years. They find that the innovations decreased cost and increased quality in 16% of cases, increased cost and decreased quality in 9% of cases, increased both cost and quality in 72% of cases, and decreased both cost and quality in 1.6% of cases. There are thus vastly more innovations that improve quality but also increase costs than innovations that save costs but reduce quality.[1] Furthermore, several innovations in the latter category provide cost savings that greatly outweigh the corresponding reduction in quality, on any reasonable criterion. For example, treating multivessel coronary artery disease with percutaneous coronary intervention ("angioplasty") rather than coronary artery bypass graft surgery ("bypass") is reported to lead to a loss of 13 quality-adjusted life *hours*, while saving \$4,944: a saving of over \$3,000,000 per quality-adjusted life year.

Another leading example—which I will return to throughout the paper—is the overuse of agricultural fertilizer in the developing world. In the context of maize farming in Western Kenya, Duflo, Kremer, and Robinson (2008) document a strong monotonic relationship between the amount of fertilizer used per plant and the resulting yield, with 1 tsp. of fertilizer per plant (the amount recommended by the government, and also the maximum amount they report being used by farmers) producing the highest yield. But fertilizer is not

[1] There are a host of possible explanations for this pattern. I cite this case only as an example of a shortage of cost-saving innovation.

free, and the relationship between amount of fertilizer per plant and *profits* is inverse U-shaped, with a maximum at 1/2 tsp. of fertilizer per plant. Indeed, most farmers (including those who use the recommended 1 tsp. per plant) actually overuse fertilizer to the point where their net returns from fertilizer utilization are strongly negative, and the authors' calculations suggest that reducing fertilizer use from 1 tsp. to 1/2 tsp. per plant would increase farmers' net income by about one quarter. Given the large potential gains, why haven't farmers learned to use less fertilizer?

This paper presents a simple model of social learning, which, among other results, provides a general rationale for the scarcity of cost-saving innovations in fields such as healthcare and agriculture in the developing world (Foster and Rosenzweig, 2010) and manufacturing (Bloom and Van Reenen, 2010; Bloom et al., 2013). I argue that, when individuals learn about the quality of new innovations by observing others' outcomes, it is hard to learn about the effectiveness of a cost-saving innovation, because it is not clear if observing good outcomes is good news or bad news about the innovation. This difficulty in learning about cost-saving innovations prevents efficient adoption of these technologies, which in turn could dis-incentivize research on cost-saving technologies.[2]

In the model, a continuum of players arrive continuously over time, and each player chooses once-and-for-all between a safe action ("status quo") and a risky action ("innovation"). The safe action yields a good outcome ("success") with known probability, while the risky action yields a good outcome with a fixed but unknown probability, which depends on the state of the world. For example, the probability that a maize crop fails to grow in a certain type of soil might be 10% when treated with 1 tsp. of fertilizer (the status quo technology), while it may be 20% or 30% when treated with 1/2 tsp. (the innovation), depending on the (unknown) effectiveness of using less fertilizer on this soil.[3] The risky

[2] This last point may depend on the organization of the R&D sector. Public-spirited governments and non-profits presumably want to fund innovations that will be adopted efficiently, while under some market structures private firms might prefer to fund innovations that will be inefficiently overadopted. The question of how technology adoption interacts with the R&D incentives of private firms is an interesting one but is beyond the scope of the current paper. In any case, public spending is of course a critical component of R&D in many sectors, accounting for over one third of global R&D spending on biomedical technologies (Chakma et al., 2014) and over half of global R&D spending on agriculture (Pray and Fuglie, 2015).

[3] This example is inspired by a study of Duflo, Kremer, Robinson, and Schilbach (see Schilbach (2015) for preliminary results), who try to encourage Kenyan maize farmers to use less fertilizer by disseminating 1/2 tsp. measuring spoons.

action also has a known cost advantage or disadvantage relative to the safe action, which is sufficient for the risky action to be optimal in the good state but not the bad state. For example, the cost-saving of using 1/2 tsp. of fertilizer rather than 1 tsp. is assumed to be sufficient to justify an increase in the rate of crop failures from 10% to 20%, but not to 30%. Before making her choice, a player observes the outcomes ("crop success," "crop failure") of a random sample of K existing players in the population.[4] I study the long-run behavior of the resulting population dynamic.

A first result is that there is always an equilibrium (i.e., a steady state of the population dynamic) in which success is more likely in the good state, so observing success is good news about the state. In the fertilizer example, this says there is always an equilibrium where observing successful harvests is good news about the effectiveness of using 1/2 tsp.. I call such an equilibrium—where success is good news—*aligned*. One of my main results is that there is a compelling reason to focus attention on aligned equilibria. In particular, I show that the equilibrium population dynamic visits only aligned points whenever the initial adoption rates are aligned, and the initial adoption rates will be aligned whenever the initial generation of players is not well-informed about the state. The intuition is that passing from a point where success is good news to a point where success is bad news requires passing through a point where each observation is completely uninformative, and this is impossible because when observations are uninformative the population dynamic drifts back toward the initial point.

These results are important because the efficiency of technology adoption at an aligned equilibrium depends dramatically on the nature of the innovation. There turns out to be a key distinction between the case of an *outcome-improving innovation*—where the risky action may yield a higher probability of success—and that of a *cost-saving innovation*—where the risky action always yields a lower probability of success but saves costs.[5] For example, a new high-yield variety of maize (which is more expensive for farmers to use but may yield a larger harvest with higher probability than traditional maize) is an outcome-improving innovation,

[4] Each outcome is assumed to be binary, which in the fertilizer application corresponds to assuming that each crop either succeeds or completely fails.

[5] This distinction is unrelated to the distinction between good news and bad news learning familiar from the literature on strategic experimentation.

while using less fertilizer is a cost-saving innovation.

If an outcome-improving innovation is adopted efficiently—that is, fully adopted in the good state and fully rejected in the bad state—the resulting adoption rates are aligned. In particular, the success rate in the good state equals the high success rate from using the (effective) innovation, while the success rate in the bad state equals the lower success rate from using the status quo. There is thus no conflict between efficiency and alignment, and, as I show, the innovation is adopted efficiently in equilibrium when players have enough data (i.e., K is large) and the population has enough time to learn.

In contrast, if a cost-saving innovation is adopted efficiently, the resulting adoption rates are *misaligned*: the success rate in the good state equals the lower success rate from using the innovation, and the success rate in the bad state equals the higher success rate from using the status quo. For example, if the innovation of using only 1/2 tsp. of fertilizer is adopted when it is cost-effective and is rejected when it is cost-ineffective (in favor of the high-cost/high-yield status quo of 1 tsp.), then the resulting crop failure rate will be 20% when using 1/2 tsp. is cost-effective and 10% when using 1/2 tsp. is cost-ineffective. This implies that, if 1/2 tsp. is adopted efficiently, observing a successful harvest must be *bad* news about the effectiveness of 1/2 tsp.. But, as I have argued, so long as the initial generation of farmers is not already well-informed, in equilibrium observing a successful harvest is always good news about the effectiveness of 1/2 tsp. This implies that equilibrium adoption of the 1/2 tsp. measure is inefficient, regardless of how much data is available to the farmers and how long they have had to learn.

How can rational agents fail to learn even as they observe more and more outcomes? The only explanation is that each individual observation must become completely uninformative in the $K \to \infty$ limit: that is, the probability of success must become exactly the same in each state. The situation is thus one of confounded learning (McLennan, 1984; Smith and Sørensen, 2000). For example, suppose again that the crop failure rate is 10% with 1 tsp. of fertilizer and either 20% or 30% with 1/2 tsp., depending on the state. Then, as $K \to \infty$, the equilibrium adoption rates of 1/2 tsp. in the bad state (x_0) and the good state (x_1) must satisfy the relationship

$$\underbrace{(1-x_1)(.1)+x_1(.2)}_{\text{failure rate in good state}}=\underbrace{(1-x_0)(.1)+x_0(.3)}_{\text{failure rate in bad state}},$$

or $x_1 = 2x_0$, so that observing a crop failure is completely uninformative about the effectiveness of using 1/2 tsp.. This example also illustrates the inefficiency of aligned equilibria with a cost-saving innovation: efficiency corresponds to adoption rates $(x_0 = 0, x_1 = 1)$, so the requirement that $x_1 = 2x_0$ implies substantial inefficiency.

The main finding of this paper is thus that outcome-improving innovations are adopted efficiently, but cost-saving innovations are adopted inefficiently. I also investigate several extensions of the model and show that this result continues to hold in a range of richer physical and informational environments.

One key feature of the model deserves immediate discussion: I assume learning is *outcome-based*, in that players observe each other's outcomes (e.g., crop yields), but not their actions (e.g., fertilizer utilization) or payoffs (e.g., profits). This contrasts with most existing models of social learning, which assume learning is *action-based*: actions are observed, but not outcomes. Which form of learning is more appropriate depends on the application. For example, Banerjee (1992) famously considers the example of herding on the choice of a restaurant, which certainly seems like a case where action-based learning is more reasonable: one can tell at a glance if tables are full and there is a line out the door, while it is harder to tell if diners are enjoying their meals. In general, action-based learning is likely the more appropriate model for settings where actions correspond to consumption choices and outcomes to subjective utility realizations. On the other hand, in settings where actions correspond to choosing inputs in a production process, outcome-based learning is often more natural.

Many economic applications of social learning fall into this category. Consider the canonical example of the choice of agricultural or health technologies in the developing world. In this setting, it seems natural to assume that farmers can see their neighbors' crop yields more easily than they can observe what seed varieties they planted or how much fertilizer they used; or that parents can see what diseases their neighbors' children contract more clearly than they can observe what preventative medications they took.[6] Another example is given by process innovations in manufacturing, where firms can presumably observe their competitors' output levels more accurately than their production techniques. For instance, a firm might not know if a competitor has adopted a certain lean manufacturing technique, such as just-in-time delivery of inputs by suppliers, while being able to observe how frequently the competitor stocks out of its inventory.[7]

1.1 Related Literature

This paper draws on ideas and modeling techniques from several branches of the literature on social learning and experimentation.

The paper contributes to the literature on social learning and herding following Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992), and Smith and Sørensen (2000). See Chamley (2004) for a survey. The key difference from most of this literature is that in the current paper outcomes are observable and actions are not.

On a technical level, the closest paper in this literature is Banerjee and Fudenberg (2004). Banerjee and Fudenberg assume that actions are perfectly observed, but the models are otherwise closely related. Their paper introduced the approach of studying deterministic social learning dynamics in a continuum population with small sample ("word-of-mouth") learning.[8] An important difference is that under perfect observability of actions the so-called "improvement principle" applies, which states that average welfare in the population is non-decreasing over time; this follows because each player can simply copy an earlier action. The improvement principle leads to asymptotic agreement on actions and often efficiency. The assumption of perfect monitoring thus implies a rather optimistic view of social learning. In contrast, in my model the improvement principle does not hold (as will become clear), and

[6]See Foster and Rosenzweig (2010) for a survey of the large literature on social learning and technology adoption in development. It is interesting to note that under-adoption of cost-saving innovations appears to be a particularly severe problem in the development context. Besides the fertilizer example, the over-prescription of antibiotics in addition to oral rehydration therapy for diarrhea is another well-known and important case; see Das et al. (2016) for some recent evidence.

[7]Bloom and Van Reenen (2010) document striking differences in the quality of management practices across firms and countries. In the context of the Indian textile industry, Bloom et al. (2013) attribute firms' failure to adopt superior production techniques to informational barriers.

[8]Smith and Sørensen (2014) consider a similar model with a discrete population.

equilibrium features heterogeneous actions and persistent inefficiency.

A few papers consider the possibility that players might observe both actions and outcomes in social learning models. Many of these papers focus on documenting the possibility of confounded learning: outcomes becoming uninformative about the state.[9] Confounded learning arises asymptotically as $K \to \infty$ in the cost-saving case of my model. However, my main contribution is characterizing how the extent and form of equilibrium inefficiency depend on the nature of the innovation, not pointing out that confounding arises per se. Another major point of contrast with this literature is that canonical papers such as McLennan (1984), Piketty (1995), and Smith and Sørensen (2000) emphasize that confounded learning is a robust equilibrium outcome in their models, while for any finite K confounded learning *cannot* occur in equilibrium in my model: instead, the fact that the equilibrium population dynamic cannot reach states where confounding would occur is a key factor in determining the equilibrium trajectory. Moreover, the possibility of confounded learning arises in my model because the adoption rate and the success rate conditional on adoption may not be separately identified, which differs from the existing literature that assumes that actions are observable. The underlying analysis is also quite different: McLennan and Smith and Sørensen consider stochastic models and use Markov-martingale arguments to show that the equilibrium trajectory converges with positive probability to a point where learning is confounded, while my model is deterministic (conditional on the state) and, again, confounding cannot arise in equilibrium for finite K.

Banerjee (1993) and Acemoglu and Wolitzky (2014) consider sequential social learning models with imperfect monitoring and limited memory. It is possible to view these models as discrete-time analogues of (a special case of) the current model when the entire population turns over each period. As in my paper, a key issue in these works is that players must keep track of the probability with which each action is played in each state of the world, and must use this information to correctly interpret the observed outcomes.[10] The main analysis and

[9] McLennan (1984) and Kiefer (1989) present classic early examples of confounded learning. Bala and Goyal (1995) and Smith and Sørensen (2000) consider confounded learning in sequential social learning models. Piketty (1995) studies confounded learning in a continuum population social learning model in a political economy application. Easley and Kiefer (1988) and Aghion, Bolton, Harris, and Jullien (1992) provide characterizations of confounded learning in single-agent experimentation problems. See Chamley (2004, Chapter 8) for a textbook treatment.

[10] The same issue arises in search and matching models with incomplete information, and indeed the

results in these papers are however completely different, as they focus on learning dynamics and the possibility of cycles rather than long-run efficiency properties.

Social learning also features in the recent literature on collective experimentation (Bolton and Harris, 1999; Keller, Rady, and Cripps, 2005; see Hörner and Skrzypacz (2016) for a survey.) This literature is less closely related in other respects, because the fact that players in my model act only once eliminates all dynamic strategic considerations. A couple papers do however combine elements of herding and collective experimentation. Murto and Välimäki (2011) and Wagner (2017) consider stopping games with observational learning, focusing on inefficiencies due to excess delay. More closely related, Frick and Ishii (2016) study technology adoption in a model with a continuum of agents. The main modeling differences are that in their model there is no uncertainty about the adoption rate and there is no word-of-mouth learning; instead, their (long-lived) players wait to receive a public and perfectly informative signal of the state, which arrives at a higher rate when more players adopt the innovation. In terms of results, the emphasis in their paper is on the optimal timing of adoption and the shape of the resulting technology adoption curve, while there is no timing decision in my model and most of my results concern long-run efficiency. Kremer, Mansour, and Perry (2014) and Che and Hörner (2017) take a mechanism design approach to incentivizing technology adoption in related models.

2 Model

This section describes the environment, the distinction between outcome-improving and cost-saving innovations, and the solution concept.

2.1 Environment

There are two states, $\theta \in \{0, 1\}$, two actions, $a \in \{0, 1\}$, and two outcomes, $y \in \{0, 1\}$. The prior probability of $\theta = 1$ ("the state is good," "the innovation is effective") is $p \in (0, 1)$. There is a continuum of players, arriving continuously over time, and each player chooses

presence of small sample learning makes my model a kind of search model. For a striking recent contribution to the literature on information aggregation and search, see Lauermann and Wolinsky (2016).

an action once. A player who chooses action $a = 1$ ("risky action," "innovation," "adopt") when the state is θ gets outcome $y = 1$ ("success") with probability π_θ and gets outcome $y = 0$ ("failure") with probability $1 - \pi_\theta$, with $\pi_0 < \pi_1$. A player who chooses action $a = 0$ ("safe action," "status quo") gets outcome $y = 1$ with probability χ and gets outcome $y = 0$ with probability $1 - \chi$, regardless of the state. A player's payoff is $y - ca$, where $c \in \mathbb{R}$ is the cost of taking action 1 (relative to the cost of taking action 0, which is normalized to 0). Note that c can be positive or negative.

Before making her choice, each player observes the outcomes (and *not* the actions or payoffs) of a random sample of K earlier players. The resulting population dynamic and corresponding equilibrium concept are defined formally below.

To rule out trivial cases, throughout the paper I impose assumptions that will imply that neither always choosing the safe action (in both states) nor always choosing the risky action is an equilibrium.

First, I assume that the risky action is optimal in the good state, while the safe action is optimal in the bad state:

$$\pi_1 - c > \chi > \pi_0 - c.$$

This implies the existence of a cutoff belief $p^* \in (0, 1)$ at which a player is indifferent between the two actions, given by $p^*\pi_1 + (1 - p^*)\,\pi_0 - c = \chi$.

The assumption that implies that always choosing the safe action is not an equilibrium is

$$p > p^*. \tag{1}$$

That is, the risky action is optimal under the prior.

The assumption that implies that always choosing the risky action is not an equilibrium is

$$\left(\frac{1 - \pi_0}{1 - \pi_1}\right)^K > \frac{p}{1 - p}\frac{1 - p^*}{p^*}. \tag{2}$$

As will become clear, this says that, if a player believes that everyone else adopts the innovation in both states and all K outcomes she observes are failures ($y = 0$), this is sufficiently bad news that she prefers not to adopt herself.

2.2 Outcome-Improving vs. Cost-Saving Innovations

The model admits quite different interpretations depending on the ordering of χ and π_1.

When $\chi < \pi_1$, the risky action may yield the good outcome with higher probability than the safe action. I call this the *outcome-improving innovation* case. Note that $\chi < \pi_1$ and assumptions (1) and (2) are consistent with c being positive or negative, so that some innovations with $c < 0$ are classified as outcome-improving. The $\chi < \pi_1$ case captures the following situations:

- The safe action is planting a traditional crop variety, and the risky action is planting a new variety that could give higher or lower average yields, depending on the underlying soil suitability.

- The safe action is using a standard medical technology, and the risky action is using a more expensive new technology that promises improved health outcomes, where it is uncertain if the degree of improvement will be sufficient to justify the expense.

I refer to the extreme case where $\chi = 0$ as the *pure outcome-improving innovation* case. An example would be if the risky action corresponds to planting a completely new crop for the first time: depending on the soil conditions, the new crop may be more or less likely to grow if planted, but it will never grow without being planted.

When instead $\chi > \pi_1$, the risky action always yields the good outcome with lower probability than the safe action. Note that the assumption $\pi_1 - c > \chi$ implies $c < 0$ in this case. The uncertainty facing the players is thus whether the reduced probability of receiving the good outcome is substantial enough to outweigh the lower cost of the risky action. Examples captured by this case include:

- The safe action is using the standard amount of fertilizer, and the risky action is using less fertilizer.

- The safe action is planting a traditional crop variety, and the risky action is planting an unfamiliar variety that requires less labor to cultivate.[11]

[11] Bustos, Caprettini, and Ponticelli (2016) argue that genetically engineered, herbicide-resistant

- The safe action is using traditional inventory-management techniques, and the risky action is using just-in-time delivery of inputs.
- The safe action is getting a standard vaccine, and the risky action is forgoing the vaccine.

Since the $\chi > \pi_1$ case involves assuming the risk of a worse outcome to save a cost, I call this the case of a *cost-saving innovation*. Finally, the case where $\chi = 1$ (e.g., the standard vaccine is perfectly effective) is that of a *pure cost-saving innovation*.

2.3 Adoption Rates and Inferences

As described above, each player observes a random sample of earlier players' outcomes before taking her action. (The behavior of the initial generation of players is discussed later.) A player must use her sample to draw an inference about the state, θ. To do so, she must take into account the *equilibrium adoption rate* of the innovation in each state of the world. Due to the continuum population, these adoption rates are modeled as deterministic functions of the state and calendar time: given an equilibrium, the resulting adoption rate in state θ at time t is denoted $X_\theta(t)$. The interpretation is that $X_\theta(t)$ is the equilibrium fraction of players existing in the population at time t who take action 1 when the state is θ. As equilibrium objects, the conditional adoption rates are "known" to a player when she enters the population: a player who enters the population at time t knows that the adoption rate is either $X_0(t)$ or $X_1(t)$, depending on the state.[12]

If the adoption rate in state θ at some time t equals x_θ, then the fraction of existing players at time t who obtained the good outcome $y = 1$ when the state is θ equals

$$\sigma_\theta(x_\theta) := x_\theta \pi_\theta + (1 - x_\theta)\chi.$$

I call $\sigma_\theta(x_\theta)$ the *success rate* in state θ given adoption rate x_θ. When a new player enters

soybeans—which obviate the need for labor-intensive weeding—constitute a cost-saving innovation in this sense. This contrasts with the many high-yield crop varieties introduced during the Green Revolution, which are typically viewed as output-improving innovations.

[12] The model thus implicitly assumes that players know calendar time. Unknown calendar time is discussed in Section 6.

the population at time t, if the state is θ then each of her observations is a success with independent probability $\sigma_\theta(x_\theta)$. Thus, by Bayes' rule, a new player's assessment of the probability that $\theta = 1$ after observing k successes, given equilibrium adoption rates x_0 and x_1, equals

$$p(k; x_0, x_1) := \left[1 + \frac{1-p}{p} \frac{\sigma_0(x_0)^k (1-\sigma_0(x_0))^{K-k}}{\sigma_1(x_1)^k (1-\sigma_1(x_1))^{K-k}}\right]^{-1}. \tag{3}$$

A simple but important observation is that $p(k; x_0, x_1)$ is increasing in k at the pair (x_0, x_1) if and only if the success rate is higher in state 1: $\sigma_1 \geq \sigma_0$, or equivalently

$$x_1(\pi_1 - \chi) \geq x_0(\pi_0 - \chi). \tag{4}$$

There are thus three possible cases (recall that p^* is the cutoff belief at which players are indifferent between the two actions):

1. (4) holds with strict inequality, there is at most one integer k^* satisfying $p(k^*; x_0, x_1) = p^*$, and $p(k; x_0, x_1) \geq p^*$ if and only if $k \geq k^*$.

2. (4) fails, there is at most one integer k^* satisfying $p(k^*; x_0, x_1) = p^*$, and $p(k; x_0, x_1) \geq p^*$ if and only if $k \leq k^*$.

3. (4) holds with equality, and $p(k; x_0, x_1) = p > p^*$ for all k.

I call a pair of adoption rates (x_0, x_1) that satisfies (4) *aligned* and call adoption rates that satisfy the opposite of (4) *misaligned*. An aligned point (x_0, x_1) is thus one where observing more successes is better news about the state in the monotone likelihood ratio sense (Milgrom, 1981), so that the informational content and utility content of the outcome y line up. For example, in the fertilizer application, a pair of adoption rates is aligned if observing crop failures is bad news about the effectiveness of using less fertilizer. Note that, if $\chi \leq \pi_1$, then (x_0, x_1) is aligned whenever $x_0 \leq x_1$; and if $\chi \geq \pi_1$, then (x_0, x_1) is aligned whenever $x_0 \geq x_1$. In particular, points (x_0, x_1) on the 45° line are always aligned. Note also that the efficient point $(x_0 = 0, x_1 = 1)$—corresponding to complete rejection of the innovation in the bad state and complete adoption in the good state—is aligned if and only if $\chi \leq \pi_1$. Thus, the efficient point is aligned in the outcome-improving innovation case and misaligned in the cost-saving innovation case. Let $A \subseteq [0,1]^2$ denote the set of all aligned points (x_0, x_1). Note that A is compact and convex.

2.4 Equilibrium Population Dynamics

I can now define the equilibrium population dynamic. An equilibrium population dynamic captures the process whereby new players enter the population at rate 1 and best-respond to random samples of existing players' outcomes; it is essentially a perfect Bayesian equilibrium of this continuum-player, continuous-time game.[13] More precisely, an equilibrium population dynamic specifies, at each point in time, (i) the equilibrium adoption rate in each state, (ii) a number of observed successes k^* marking the cutoff above or below which new players adopt (depending on whether or not (4) holds), and (iii) the probability s with which new players adopt when they observe exactly k^* successes, if this observation leaves them indifferent.

For the formal definition, let $\phi_\theta(k; x_\theta)$ denote the probability that a new player observes k successes when the state is θ and the adoption rate is x_θ:

$$\phi_\theta(k; x_\theta) := \binom{K}{k} \sigma_\theta(x_\theta)^k (1 - \sigma_\theta(x_\theta))^{K-k}.$$

Given that players follow cutoff strategies, equation (5) below says that, in each state, the derivative of the adoption rate equals the adoption rate among new players minus the adoption rate among existing players.

Definition 1 *An* equilibrium path *(or* equilibrium population dynamic*) is a list of measurable functions of time*

$$(X_0 : \mathbb{R}_+ \to [0,1], X_1 : \mathbb{R}_+ \to [0,1], k^* : \mathbb{R}_+ \to \{0, \ldots, K\}, s : \mathbb{R}_+ \to [0,1])$$

such that

[13] The inflow of new players is modeled as in Banerjee and Fudenberg (2004). There are two mathematically equivalent interpretations. The first is that new players enter at arithmetic rate 1 and existing players exit at arithmetic rate 1, keeping the overall population size constant. The second is that new players enter at exponential rate 1 and no players exit, so the overall population size increases over time. Under either interpretation, $X_\theta(t)$ is the fraction of the existing population taking action 1 in state θ at time t.

1. *Trajectories respect individual optimization: for $\theta = 0, 1$, X_θ is absolutely continuous, with derivative almost everywhere given by*

$$\dot{X}_\theta(t) = \begin{cases} \phi_\theta(k^*(t); X_\theta(t))\, s(t) + \sum_{k=k^*(t)+1}^{K} \phi_\theta(k; X_\theta(t)) - X_\theta(t) \\ \text{if (4) holds at } (X_0(t), X_1(t)); \\ \phi_\theta(k^*(t); X_\theta(t))\, s(t) + \sum_{k=0}^{k^*(t)-1} \phi_\theta(k; X_\theta(t)) - X_\theta(t) \\ \text{if (4) fails at } (X_0(t), X_1(t)) \end{cases}. \tag{5}$$

2. *Cutoffs are consistent with Bayes' rule:*

   $p(k^*(t) - 1; X_0(t), X_1(t)) \leq p^* \leq p(k^*(t); X_0(t), X_1(t))$ *if (4) holds at* $(X_0(t), X_1(t))$*;*

   $p(k^*(t); X_0(t), X_1(t)) \geq p^* \geq p(k^*(t) + 1; X_0(t), X_1(t))$ *if (4) fails at* $(X_0(t), X_1(t))$*.*

3. *Decisions are optimal at the cutoff:*

   $$s(t) \begin{cases} = 1 & \text{if } p(k^*(t); X_0(t), X_1(t)) > p^* \\ \in [0,1] & \text{if } p(k^*(t); X_0(t), X_1(t)) = p^* \end{cases}.$$

Proposition 1 *For any point $\hat{x} \in [0,1]^2$, there exists an equilibrium path (X_0, X_1, k^*, s) with $(X_0(0), X_1(0)) = \hat{x}$.*

Proof. Define the correspondence $F : [0,1]^2 \rightrightarrows [0,1]^2$ by letting $F(x_0, x_1)$ be the set of pairs $(x_0', x_1') \in [0,1]^2$ for which there exist $k^* \in \{0, \ldots, K\}$ and $s \in [0,1]$ such that either

$$\begin{aligned} x_\theta' &= \phi_\theta(k^*; x_\theta)\, s + \sum_{k=k^*+1}^{K} \phi_\theta(k; x_\theta) \ \text{ for } \theta = 0, 1, & (6) \\ p(k^* - 1; x_0, x_1) &\leq p^* \leq p(k^*; x_0, x_1), \text{ and} & \\ & s \begin{cases} = 1 & \text{if } p(k^*; x_0, x_1) > p^* \\ \in [0,1] & \text{if } p(k^*; x_0, x_1) = p^* \end{cases}, & (7) \end{aligned}$$

or

$$\begin{aligned} x_\theta' &= \phi_\theta(k^*; x_\theta)\, s + \sum_{k=0}^{k^*-1} \phi_\theta(k; x_\theta) \ \text{ for } \theta = 0, 1, & (8) \\ p(k^*; x_0, x_1) &\geq p^* \geq p(k^* + 1; x_0, x_1), \text{ and (7).} \end{aligned}$$

Note that, if a pair of absolutely continuous functions $(X_0 : \mathbb{R}_+ \to [0,1], X_1 : \mathbb{R}_+ \to [0,1])$ satisfies

$$\left(\dot{X}_0(t), \dot{X}_1(t)\right) + (X_0(t), X_1(t)) \in F(X_0(t), X_1(t)) \tag{9}$$

almost everywhere, then (X_0, X_1) is an equilibrium path (together with the corresponding values of k^* and s at points where (X_0, X_1) is differentiable, with the values of k^* and s at points of non-differentiability of (X_0, X_1) selected arbitrarily). I claim that F is non-empty, compact- and convex-valued, and upper hemi-continuous. The proof is routine and is deferred to Appendix B (available online). Note also that any solution to (9) satisfies $\dot{X}_\theta(t) \in [-X_\theta(t), 1 - X_\theta(t)]$ almost everywhere, and hence cannot escape the compact set $[0,1]$. Under these conditions, existence of a solution to the differential inclusion (9) for an arbitrary initial point is standard. See, e.g., Aubin and Cellina (1984), Theorem 2.1.4. ■

I do not have a result that the equilibrium path from an arbitrary initial point is always unique.[14] If there are multiple equilibrium paths, the solution concept implicitly assumes all players coordinate on the same one.

Finally, some results will concern steady states of the population dynamic.

Definition 2 *An* equilibrium *(or* steady state*) is a constant equilibrium path.*

I refer to an equilibrium in which (4) holds as an *aligned equilibrium* and refer to an equilibrium in which the opposite of (4) holds as an *misaligned equilibrium*. It is immediate that some equilibrium exists. This follows because equilibria correspond to fixed points of the correspondence F introduced in the proof of Proposition 1, and F satisfies the conditions of Kakutani's fixed point theorem.

3 Invariance and Inefficiency of Aligned Points

My first main result is that any equilibrium path with an aligned initial point visits only aligned points: that is, the set of aligned points is forward invariant. As I will argue, the condition that the initial point is aligned is quite mild. Moreover, with a cost-saving

[14] Standard uniqueness results for differential equations/inclusions require Lipschitz continuity conditions that may fail here (e.g., Aubin and Cellina (1984), Chapter 2.4). In particular, the correspondence F in the proof of Proposition 1 is discontinuous when $p(k^*; x_0, x_1) = p^*$.

innovation, welfare at any aligned point is bounded away from the first-best, regardless of the sample size K. Thus, as long as the initial point is aligned, cost-saving innovations are adopted inefficiently no matter how much time and data are available to the players. In contrast, Section 4 shows that outcome-improving innovations are adopted efficiently when t and K are large.

Theorem 1 *If* (X_0, X_1, k^*, s) *is an equilibrium path and* $(X_0(0), X_1(0))$ *is aligned, then* $(X_0(t), X_1(t))$ *is aligned for all* $t \in \mathbb{R}_+$. *In addition, an aligned equilibrium exists.*

To see the intuition, let L denote the boundary between the aligned and misaligned regions of the unit square: that is, L is the line with equation $x_1(\pi_1 - \chi) - x_0(\pi_0 - \chi) = 0$. In the aligned region, observing a success is good news about θ; in the misaligned region, observing a success is bad news; and close to L, observing a success is almost completely uninformative. Thus, for any fixed sample size K, players' samples contain little information about the state when the population dynamic is sufficiently close to L. In particular, (1) implies that new players always adopt when the population dynamic is close to L. This behavior drives the population dynamic toward the aligned point $(x_0, x_1) = (1, 1)$, and hence drives the dynamic into the interior of the aligned set A. The population dynamic can thus never exit A.[15] See Figures 1 and 2 for illustrations in the cost-saving and outcome-improving cases, respectively—the question marks in the figures indicate that I do not characterize dynamics away from L for a fixed sample size K.[16]

Proof. Assume $X_0(0)(\pi_0 - \chi) \leq X_1(0)(\pi_1 - \chi)$. As X_0 and X_1 are continuous, if there exists a time t' with $X_0(t')(\pi_0 - \chi) > X_1(t')(\pi_1 - \chi)$, then by the intermediate value theorem and the definition of the derivative there must exist another time t where $X_0(t)(\pi_0 - \chi) = X_1(t)(\pi_1 - \chi)$ but it is not the case that X_0 and X_1 are differentiable at t with $\dot{X}_0(t)(\pi_0 - \chi) < \dot{X}_1(t)(\pi_1 - \chi)$ (in particular, the time $\sup\{t < t' : X_0(t)(\pi_0 - \chi) = X_1(t)(\pi_1 - \chi)\}$ must have this property). But, if $X_0(t)(\pi_0 - \chi) = X_1(t)(\pi_1 - \chi)$ then $p(k; X_0(\tau), X_1(\tau)) \approx p > p^*$ for all $k \in \{0, \ldots, K\}$ and all τ in a neighborhood of t, and hence X_θ is differentiable

[15] This argument does not depend on the fact that the point toward which the population dynamic drifts when the success rate is the same in both states is the extreme point $(1, 1)$: all that matters is that this point is aligned. Theorem 1 therefore generalizes to a range of environments without this feature. See Section 6.

[16] Figure 2 is drawn for the case $\chi < \pi_0$. When $\chi \in (\pi_0, \pi_1)$, the line L does not intersect the unit square.

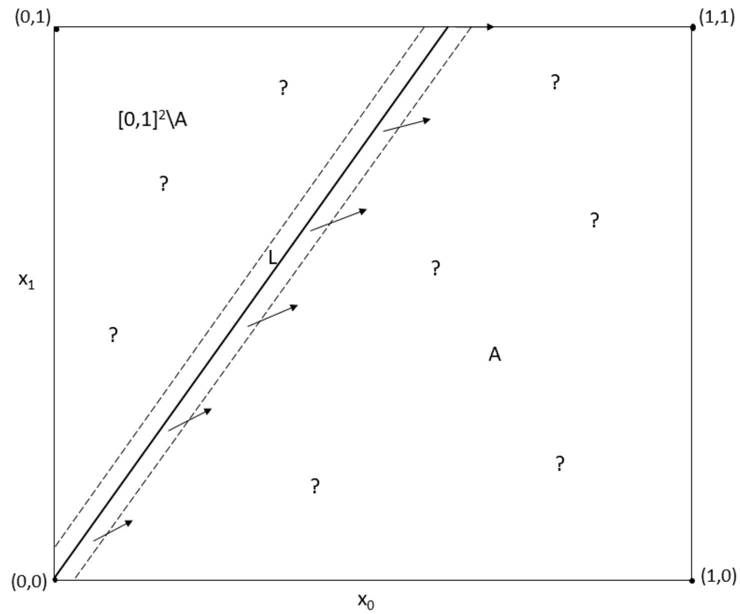

Figure 1: Phase diagram for the equilibrium population dynamic, cost-saving case ($\chi > \pi_1$), arbitrary K.

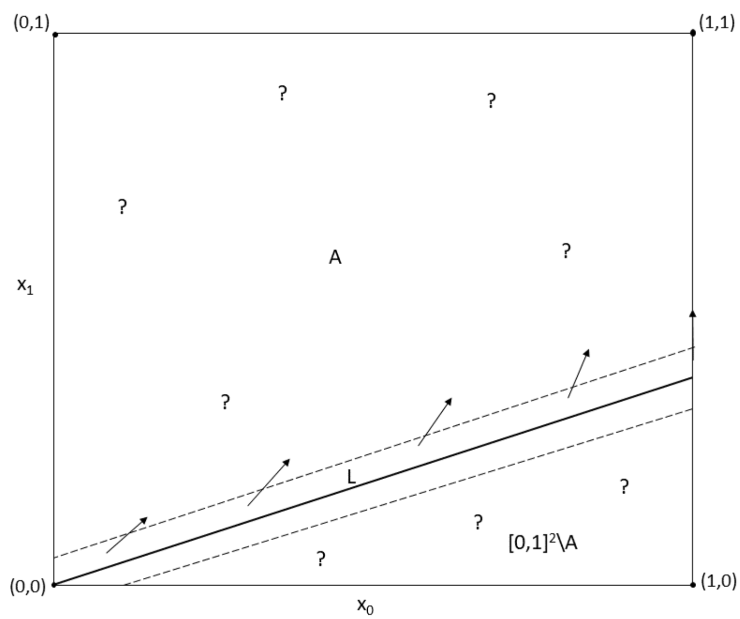

Figure 2: Phase diagram for the equilibrium population dynamic, outcome-improving case with $\chi < \pi_0$, arbitrary K.

at t with $\dot{X}_\theta(t) = 1 - X_\theta(t)$ for $\theta = 0, 1$.[17] Therefore,

$$\begin{aligned}
\dot{X}_0(t)(\pi_0 - \chi) &= \pi_0 - \chi - X_0(t)(\pi_0 - \chi) \\
&= \pi_0 - \chi - X_1(t)(\pi_1 - \chi) \\
&< \pi_1 - \chi - X_1(t)(\pi_1 - \chi) = \dot{X}_1(t)(\pi_1 - \chi).
\end{aligned}$$

It follows that there can be no such time t'. Thus, A is forward invariant.

As A is forward invariant, compact, and convex, and F is non-empty, compact- and convex-valued, and upper hemi-continuous, existence of a fixed point of F in A follows from standard results on differential inclusions. See, e.g., Aubin and Cellina (1984), Corollary 2.2.3. ■

Alignment of the initial point is a mild requirement. If the initial generation of players makes its decisions on the basis of the prior alone, the initial point will be the aligned point $(X_0(0), X_1(0)) = (1, 1)$. If the initial generation also receives some exogenous signals of the state before making its decisions, the initial point will be aligned as long as these signals are unlikely to overturn the prior: see Section 6 for details. Theorem 1 thus justifies focusing on aligned points.

Being able to focus on aligned point is important because, with a cost-saving innovation, welfare at every aligned point is bounded away from the first-best. The intuition is simply that, with a cost-saving innovation, it is efficient for failure to be more likely in the good state: at the first-best, players use the low-cost, high-failure innovation in the good state and use the high-cost, low-failure status quo in the bad state. But, by definition, at an aligned point failure is less likely in the good state. So no aligned point can be close to efficient.

Note that *expected welfare* at the point (x_0, x_1) equals

$$p\left[x_1(\pi_1 - c) + (1 - x_1)\chi\right] + (1 - p)\left[x_0(\pi_0 - c) + (1 - x_0)\chi\right]. \tag{10}$$

[17] Proof: If $p(k; X_0(\tau), X_1(\tau)) > p^*$ for all $k \in \{0, \ldots, K\}$ and all τ in a neighborhood of t, then $\dot{X}_\theta(\tau) = 1 - X_\theta(\tau)$ for all τ at which X_θ is differentiable in a neighborhood of t. But then

$$\lim_{\varepsilon \to 0} \frac{1}{\varepsilon}(X_\theta(t+\varepsilon) - X_\theta(t)) = \lim_{\varepsilon \to 0} \frac{1}{\varepsilon} \int_t^{t+\varepsilon} \dot{X}_\theta(\tau)\, d\tau = \lim_{\varepsilon \to 0} \frac{1}{\varepsilon} \int_t^{t+\varepsilon} (1 - X_\theta(\tau))\, d\tau = 1 - X_\theta(t),$$

where the first equality uses absolute continuity of X and the last uses continuity.

In particular, *first-best* expected welfare equals $p(\pi_1 - c) + (1-p)\chi$.

Proposition 2 *If $\chi > \pi_1$ then expected welfare at every aligned point is uniformly bounded away from efficiency.*

More specifically, if $\chi > \pi_1$ then expected welfare at every aligned point falls short of the first-best by at least

$$(1-p)\left(\frac{\chi - \pi_1}{\chi - \pi_0}\right)(\chi - \pi_0 + c) > 0. \tag{11}$$

Proof. First-best expected welfare exceeds expected welfare at the point (x_0, x_1) by the amount

$$p(1-x_1)(\pi_1 - c - \chi) + (1-p)x_0(\chi - \pi_0 + c).$$

If $\chi > \pi_1$ then at any aligned point $x_0 \geq \frac{\chi-\pi_1}{\chi-\pi_0}x_1$. Hence, this excess is at least

$$p(1-x_1)(\pi_1 - c - \chi) + (1-p)x_1\left(\frac{\chi - \pi_1}{\chi - \pi_0}\right)(\chi - \pi_0 + c).$$

This expression is decreasing in x_1, as $p(\chi - \pi_0)(\pi_1 - c - \chi) > (1-p)(\chi - \pi_1)(\chi - \pi_0 + c)$. (To see this, rearrange this inequality to $p(\pi_1 - \pi_0)(-c) > (\chi - \pi_1)(\chi - \pi_0 + c)$, and note that $-c > \chi - \pi_1$ and $p(\pi_1 - \pi_0) > \chi - \pi_0 + c$, by (1).) Therefore, inefficiency is minimized at $x_1 = 1$, which gives (11). ■

To interpret (11), note that the proof shows that the minimal inefficiency at an aligned point is achieved at the point $(x_0, x_1) = \left(\frac{\chi-\pi_1}{\chi-\pi_0}, 1\right)$. At this point, the correct action is always taken in state 1, and the wrong action is taken in state 0 with probability $\frac{\chi-\pi_1}{\chi-\pi_0}$. As the loss from taking the wrong action in state 0 is $\chi - \pi_0 + c$, it follows that the minimal inefficiency at any aligned point is $(1-p)\left(\frac{\chi-\pi_1}{\chi-\pi_0}\right)(\chi - \pi_0 + c)$.

To get a sense of the magnitude of this loss, note that a completely uninformed player who follows the prior and always takes action 1 suffers a loss of $(1-p)(\chi - \pi_0 + c)$. Thus, the ratio of the minimal loss at an aligned point to the loss of an uninformed player equals $\frac{\chi-\pi_1}{\chi-\pi_0}$. Note that this loss ratio is increasing in χ and π_0 but decreasing in π_1. In this sense, social learning is potentially more powerful when the status quo is less effective and outcomes under the innovation are less noisy.

4 Large Samples: Learning vs. Confounding

My second main result is that outcome-improving innovations are adopted efficiently when samples are large enough and the population has had enough time to learn. Combined with the earlier inefficiency results, this implies a dramatic difference between the outcome-improving and cost-saving cases when samples are large.

More specifically, in the outcome-improving case with any initial point, or in the cost-saving case with a misaligned initial point, the population dynamic converges to the efficient point $(0,1)$ as $K \to \infty$ and $t \to \infty$. In the cost-saving case with an aligned initial point, the population dynamic converges to the line L as $K \to \infty$ and $t \to \infty$.

Let $\|\cdot\|$ and $d(\cdot,\cdot)$ denote Euclidean norm and distance.

Theorem 2 *Fix an initial point $\hat{x} \in [0,1]^2$, and assume that either (i) $\chi < \pi_1$ or (ii) $\chi > \pi_1$ and $\hat{x}$ is misaligned. For every $\varepsilon > 0$, there exist $\bar{K} > 0$ and $T > 0$ such that, for every $K > \bar{K}$, every $t > T$, and every equilibrium path (X_0, X_1, k^*, s) with $(X_0(0), X_1(0)) = \hat{x}$, one has $\|(X_0(t), X_1(t)) - (0,1)\| < \varepsilon$.*

Conversely, fix an initial point $\hat{x} \in [0,1]^2$, and assume that $\chi > \pi_1$ and $\hat{x}$ is aligned. For every $\varepsilon > 0$, there exist $\bar{K} > 0$ and $T > 0$ such that, for every $K > \bar{K}$, every $t > T$, and every equilibrium path (X_0, X_1, k^*, s) with $(X_0(0), X_1(0)) = \hat{x}$, one has $d((X_0(t), X_1(t)), L) < \varepsilon$.

Theorem 2 gives an immediate corollary regarding steady states. In the outcome-improving case, any sequence of steady states indexed by K converges to $(0,1)$ as $K \to \infty$. In the cost-saving case, any sequence of aligned steady states indexed by K converges to L, and any sequence of misaligned steady states bounded away from L converges to $(0,1)$.[18]

Corollary 1 *Assume $\chi < \pi_1$. For any sequence of equilibria $\left(x_0^K, x_1^K, k^{*,K}, s^K\right)$ indexed by K, $\lim_{K\to\infty}\left(x_0^K, x_1^K\right) = (0,1)$.*

Assume $\chi > \pi_1$. For any sequence of aligned equilibria $\left(x_0^K, x_1^K, k^{*,K}, s^K\right)$ indexed by K, $\lim_{K\to\infty} d\left(\left(x_0^K, x_1^K\right), L\right) = 0$. For any $\varepsilon > 0$ and any sequence of misaligned equilibria $\left(x_0^K, x_1^K, k^{*,K}, s^K\right)$ indexed by K such that $d\left(\left(x_0^K, x_1^K\right), L\right) > \varepsilon$ for all K, $\lim_{K\to\infty}\left(x_0^K, x_1^K\right) = (0,1)$.

[18] Note that Theorem 2 allows for the possibility of a sequence of misaligned steady states converging to L as $K \to \infty$.

The intuition for Theorem 2 is that, away from the line L, players learn the state when samples are large enough by the law of large numbers, so the population dynamic drifts toward the efficient point $(0, 1)$. If the initial point is aligned in the outcome-improving case or misaligned in the cost-saving case, it lies on the same side of L as $(0, 1)$, so the population dynamic converges to $(0, 1)$. If the initial point is misaligned in the outcome-improving case, the population dynamic drifts toward $(0, 1)$ until it becomes close to L, and it then crosses L and proceeds to converge to $(0, 1)$. Finally, if the initial point is aligned in the cost-saving case, the population dynamic drifts toward $(0, 1)$ until it becomes close to L, but it cannot cross L (as shown by Theorem 1). The population dynamic therefore remains close to L forever.

The large-K population dynamic is illustrated in Figures 3 and 4. In these figures, the population dynamic drifts toward $(0, 1)$ in the areas outside the gray lines (i.e., far from L), and the population dynamic drifts toward $(1, 1)$ in the area between the dashed lines (i.e., close to L). In the areas between a gray line and a dashed line, the drift of the population dynamic is ambiguous. As $K \to \infty$, both the gray lines and the dashed lines converge to L.

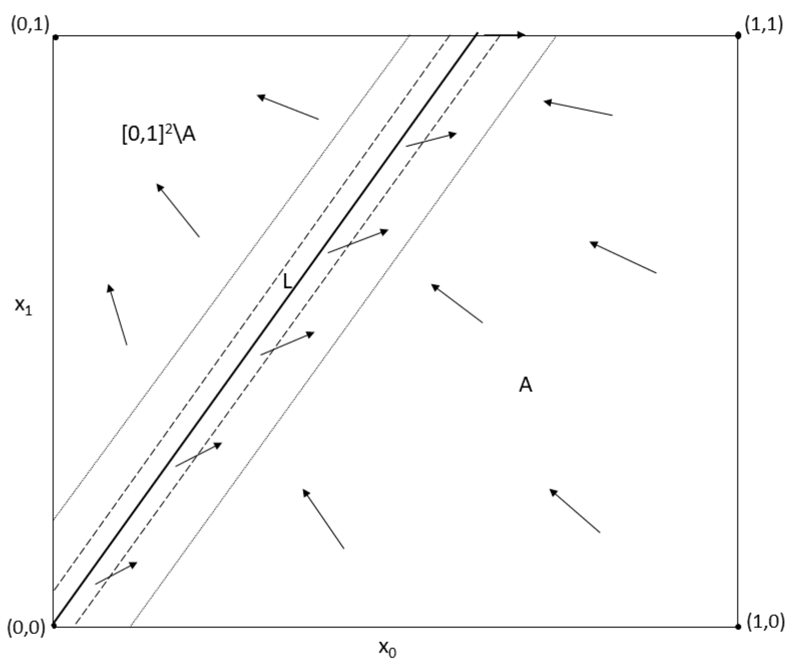

Figure 3: Phase diagram for the equilibrium population dynamic, cost-saving case ($\chi > \pi_1$), large K.

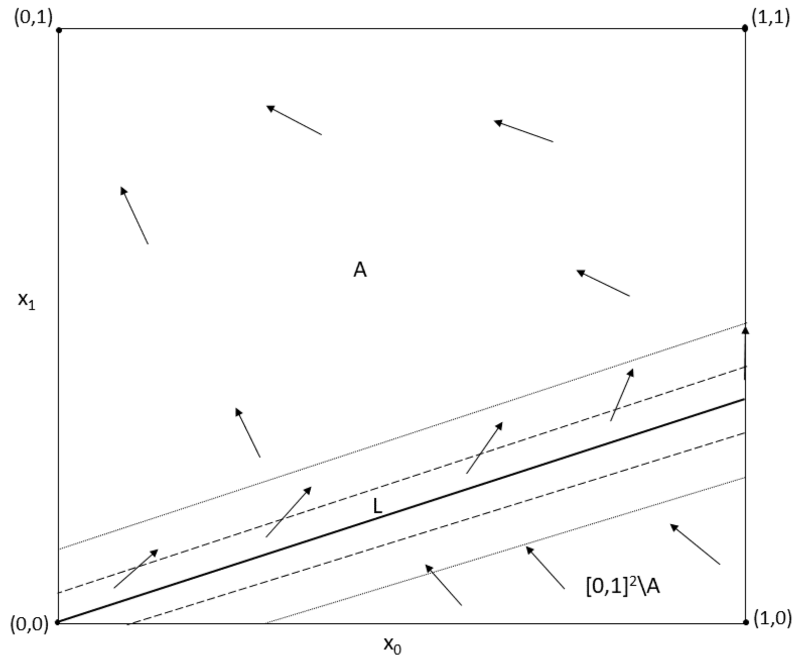

Figure 4: Phase diagram for the equilibrium population dynamic, outcome-improving case with $\chi < \pi_0$, large K.

If one restricts attention to aligned initial points, the intuition for Theorem 2 becomes even simpler. In the outcome-improving case, along any equilibrium path the success rate in the good state is always substantially greater than the success rate in the bad state, so players learn the state. In the cost-saving case, the success rates in the two states become close as the population dynamic approaches L, so learning is confounded.

Theorem 2 yields the main conclusion of this paper: as $K \to \infty$ and $t \to \infty$, outcome-improving innovations are adopted efficiently, but cost-saving innovations are not. The key economic intuition is that it is hard to learn about cost-saving innovations, because it is not clear if observing good outcomes is good news or bad news.

The proof of Theorem 2 is deferred to Appendix A.

5 Small Samples

My most striking findings are the large-K results described above, but the model also yields some interesting results when K is small. In particular, if one fixes K but considers the limits where $\pi_0 \to 0$ and/or $\pi_1 \to 1$, there are again clear differences between the outcome-improving and cost-saving cases. This section briefly describes some notable differences (focusing for simplicity on steady states), and also makes some observations about population dynamics in the special case $K = 1$. See the working paper version (Wolitzky, 2017) for further discussion and results.

5.1 One Observation: Under- vs. Over-Adoption

When $K = 1$, there is a unique steady state, which is aligned, and the steady state does not tend toward efficiency as $\pi_0 \to 0$ and/or $\pi_1 \to 1$.[19] More interestingly, the nature of the inefficiency is different in the pure outcome-improving innovation case ($\chi = 0$) and in all other cases ($\chi > 0$). In the pure outcome-improving innovation case, in the unique equilibrium the innovation is fully rejected in the bad state, but it is not fully adopted in the good state. The innovation is thus *under-adopted.* Outside of this case, the innovation is fully adopted in the good state but is not fully rejected in the bad state: that is, it is *over-adopted.*

Proposition 3 *Assume each player observes only one other player's outcome ($K = 1$). There is a unique equilibrium. Moreover,*

1. *In the pure outcome-improving innovation case, for all $\varepsilon > 0$ there exists $\delta > 0$ such that, if $1 > \pi_1 > 1 - \delta$, then in equilibrium $x_0 < \varepsilon$ and $x_1 < \frac{p-\hat{p}}{p(1-\hat{p})} + \varepsilon$, where $\hat{p} := \frac{c-\pi_0}{1-\pi_0}$. The innovation is thus under-adopted.*

2. *Outside of the pure outcome-improving innovation case, for all $\varepsilon > 0$ there exists $\delta > 0$ such that, if $\pi_1 > 1 - \delta$, then in equilibrium $x_0 = \frac{\chi}{1-\pi_0+\chi}$ and $x_1 > 1 - \varepsilon$. The innovation is thus over-adopted.*

Proof. The proofs of this and all subsequent results in the paper are deferred to Appendix B. ■

To illustrate Proposition 3 in the context of an example, consider an agricultural community that relies on planting traditional maize, and compare the resulting adoption patterns

[19] For the proof, see Proposition 11 in Appendix B.

when the community faces the introduction of a high-yield variety of maize and when it faces the introduction of a completely new crop such as pineapple.[20] In the high-yield maize case, the bad outcome corresponds to a small maize harvest and the good outcome corresponds to a large maize harvest; in the pineapple case, the bad outcome is maize but no pineapple, while the good outcome is maize and pineapple. Suppose that in both cases the new crop is almost certainly successful if the underlying soil conditions in the community are favorable.

The critical distinction between the cases where the new crop is high-yield maize and where it is pineapple is that the status quo crop (traditional maize) might produce the good outcome (a large maize harvest) in the former case, but can never produce the good outcome (maize and pineapple) in the latter. Proposition 3 says that, if each farmer observes only one neighbor's harvest before planting her own crop, the pattern of technology adoption will be completely different in the two cases. In the high-yield maize case, all farmers will (profitably) adopt the new crop when the soil conditions are favorable, but some farmers will (unprofitably) do the same even when the soil is unfavorable. In the pineapple case, no farmers will mistakenly adopt the pineapple when the soil is unfavorable, but some farmers will fail to adopt when the soil is favorable. High-yield maize is thus over-adopted, while pineapple is under-adopted.

More generally, the key prediction of Proposition 3 is that qualitatively new technologies—those that can produce results that cannot be confused with results coming from existing technologies—are under-adopted, while new technologies that are only more likely to produce good results are over-adopted.

5.2 Multiple Observations: Efficiency vs. Persistent Inefficiency

The inefficiency in the $K = 1$ case documented in Proposition 3 vanishes for all $K > 1$ in the pure outcome-improving case, while in the pure cost-saving case inefficiency diminishes only slowly as K increases.[21] Thus, as in the $K \to \infty$ case, equilibrium adoption is dramatically

[20] The example of maize and pineapple is inspired by the influential study of Conley and Udry (2010), though not all aspects of their environment fit my model. Most importantly, they assume that farmers observe their neighbors' planting decisions and fertilizer utilization in addition to their crop yields.

[21] Note that the conditions of Proposition 4 are more restrictive than those of Proposition 3, in that Proposition 4 requires $\pi_0 \to 0$ and $\pi_1 \to 1$ while Proposition 3 requires only $\pi_1 \to 1$.

more efficient with an outcome-improving innovation.

Proposition 4 *Assume $K > 1$.*

1. *In the pure outcome-improving innovation case, for all $\varepsilon > 0$ there exists $\delta > 0$ such that, if $\pi_0 < \delta$ and $\pi_1 > 1-\delta$, then in any equilibrium $x_0 < \varepsilon$ and $x_1 > 1-\varepsilon$. Adoption is thus approximately efficient.*

2. *In the pure cost-saving innovation case, for all $\varepsilon > 0$ there exists $\delta > 0$ such that, if $\pi_0 < \delta$ and $\pi_1 > 1 - \delta$, then in any aligned equilibrium $\left| x_0 - (1 - x_0)^K \right| < \varepsilon$ and $x_1 > 1 - \varepsilon$. Over-adoption thus persists for every finite K.*

   In contrast, for all $\varepsilon > 0$ there exist $\delta > 0$ and $\bar{K} > 0$ such that, if $\pi_0 < \delta$, $\pi_1 > 1 - \delta$, and $K > \bar{K}$, then in any aligned equilibrium $x_0 < \varepsilon$ and $x_1 > 1 - \varepsilon$. Adoption is thus approximately efficient when K is large.

There is of course no contradiction between the last part of Proposition 4 and Proposition 2: in the cost-saving case, for any values of π_0 and π_1 efficiency in any aligned equilibrium is uniformly bounded away from efficiency for all K, but the degree of inefficiency vanishes as $\pi_0 \to 0$ and $\pi_1 \to 1$. Thus, in the pure outcome-improving case, equilibrium is efficient in the $K \to \infty$ limit *as well as* in the $\pi_0 \to 0/\pi_1 \to 1$ limit, while in the pure cost-saving case equilibrium is efficient only if *both* K is large and π_0 is small/π_1 is large.

In the pure outcome-improving case, the stark difference between the $K = 1$ and $K > 1$ cases may be explained by considering the equation for x_1 that results when $x_0 = 0$ and players adopt if and only if they observe at least one success:

$$x_1 = 1 - (1 - x_1\pi_1)^K .$$

When $K = 1$, for any $\pi_1 < 1$ the unique solution to this equation is $x_1 = 0$. In contrast, for any $K > 1$, when π_1 is sufficiently large this equation also admits a positive solution, and this solution converges to 1 as $\pi_1 \to 1$. Thus, when $K = 1$, equilibrium requires players to mix after observing failure (which imposes an upper bound on x_1), while when $K > 1$ it is possible for x_1 to be close to 1 even in a pure strategy equilibrium.

In the pure cost-saving case, when $\pi_0 = 0$ the adoption rate in the bad state satisfies the equation $x_0 = (1 - x_0)^K$. The solution to this equation converges to 0 rather slowly as $K \to \infty$. For example, if $K = 5$ then $x \approx 0.25$, and if $K = 50$ then $x \approx 0.06$. Inefficiency is therefore substantially greater in the pure cost-saving case than in the pure outcome-improving case, even when K is relatively large.

A final observation concerns the comparative statics of expected welfare with respect to K. While Proposition 4 illustrates settings where larger samples lead to higher equilibrium welfare (as one would expect), this comparative static does not hold in general. Indeed, Example 2 in Appendix B shows that welfare does not always unambiguously increase when players observe larger samples even if one restricts attention to stable aligned equilibria.

5.3 Convergence to Equilibrium in the $K = 1$ Case

A question left unanswered by the analysis so far is whether the equilibrium population dynamic always converges to a steady state. The existence of multiple stable aligned steady states—as illustrated in Example 2 in Appendix B—suggests that this may be a difficult question to resolve, and I have not been able to provide a complete answer. However, convergence to the unique steady state does always occur in the $K = 1$ case. Interestingly, this convergence may not be monotonic.

The convergence result for the $K = 1$ case is as follows.

Proposition 5 *When $K = 1$, the unique equilibrium (x_0^*, x_1^*) is globally attracting: for any equilibrium path (X_0, X_1, k^*, s), $\lim_{t\to\infty} (X_0(t), X_1(t)) = (x_0^*, x_1^*)$.*

A simple example illustrates that convergence can be non-monotone.

Proposition 6 *When $K = 1$, $\chi = 0$, and $\pi_0 = 0$, there is a unique equilibrium path starting from initial condition $(X_0(0), X_1(0)) = (1, 1)$, and the population dynamics are as follows:*

1. *$X_1(t)$ decreases at rate $1 - \pi_1$ until reaching its steady-state value x_1^* at some finite time T, and then remains constant forever.*

2. *$X_0(t)$ decreases at rate 1 up to time T, and then converges monotonically to its steady-state value x_0^*. The overall path $X_0(t)$ is thus non-monotone if and only if $X_0(T) < x_0^*$. Finally, there exists $\delta > 0$ such that if $\pi_1 > 1 - \delta$ then $X_0(T) < x_0^*$.*

Figure 5 illustrates the equilibrium path in the case where $X_0(T) < x_0^*$. The intuition for why the path is non-monotone is as follows: First, $X_1(t)$ must decrease over time until reaching its steady-state value x_1^*, as if $X_1(t) > x_1^*$ then observing failure is so informative that a player who observes failure will never adopt the innovation. Once $X_1(t)$ reaches x_1^*, observing failure becomes just informative enough that a player who observes failure is indifferent between the innovation and the status quo. At this point, players who observe failure mix, adopting the innovation with the probability that keeps the adoption rate in the good state constant.[22] Finally, if the adoption rate in the bad state is sufficiently low by the time $X_1(t)$ reaches x_1^* (i.e., if $X_0(T) < x_0^*$), then this mixing probability is high enough that the adoption rate in the bad state starts to increase, before eventually approaching its steady-state level.

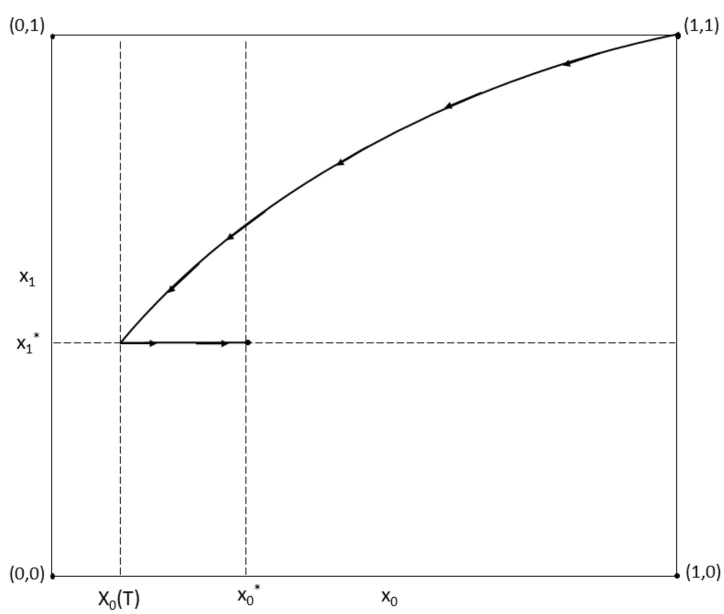

Figure 5: Population dynamics with $K = 1$, $\chi = 0$, $\pi_0 = 0$, initial condition $(X_0(0), X_1(0)) = (1, 1)$, and $X_0(T) < x_0^*$.

[22] Thus, the kink in the population dynamic at time T corresponds to a jump in k^* from 0 to 1.

In addition to featuring non-monotone population dynamics, this example also has the interesting property that the *improvement principle* of Banerjee and Fudenberg does not hold. The improvement principle states that average welfare in the population is non-decreasing over time. The improvement principle always holds with action-based learning, as new players can guarantee themselves the average level of welfare by simply copying a random action. This argument does not apply with outcome-based learning, and indeed the improvement principle fails in the setting of Proposition 6 whenever the population dynamic is non-monotone: this follows because, after time T, average welfare in state 1 is constant while average welfare in state 0 is strictly decreasing.

6 Extensions

This section extends the model in two directions. Section 6.1 considers more general physical environments, which allow for heterogeneous players and multiple states of the world. Section 6.2 considers more general information structures, which allow for additional signals of the state (or other variables measurable with respect to the state, such as the adoption rate) and uncertainty regarding calendar time. The emphasis in all cases is on assessing the robustness of the main results, Theorems 1 and 2. The working paper version (Wolitzky, 2017) contains additional extensions.

6.1 More General Physical Environments

6.1.1 Heterogeneous Players

The baseline model assumes all players are identical. It would be more realistic to assume that players differ in exogenous characteristics that affect their attitude toward the innovation. For example, farmers have information about features of their own soil, and it may be known that, if a new crop variety is effective, it will be particularly effective for certain types of soil. Whether enriching the model on this dimension affects the main results turns out to hinge on whether players' characteristics are observable—for example, if a farmer knows the soil characteristics of the neighbors whose harvests she observes. In particular, unobservable heterogeneity does not affect the results, while observable heterogeneity generically leads to efficient adoption even in the cost-saving case.

Formally, suppose there are $|Q|$ types of players. Each player knows her own type. The constant population share of type q players is α_q. A player's type affects her probability of success when using the innovation: a type q player who uses the innovation in state θ succeeds with probability $\pi_{\theta,q}$. Assume that (1) and (2) are satisfied for each type q.[23]

Suppose first that each player continues to observe a random sample of K other players' outcomes, without observing their types. The population dynamic must now keep track of adoption rates for each type of player. Thus, (i) An equilibrium path consists of a list of functions

$$\left(X_{0,q} : \mathbb{R}_+ \to [0,1]\,, X_{1,q} : \mathbb{R}_+ \to [0,1]\,, k_q^* : \mathbb{R}_+ \to \{0,\ldots,K\}\,, s_q : \mathbb{R}_+ \to [0,1]\right)_{q\in Q}.$$

(ii) A vector of adoption rates $(x_{0,q}, x_{1,q})_{q\in Q}$ is aligned if and only if

$$\sum_{q\in Q} \alpha_q x_{1,q} (\pi_{1,q} - \chi) \geq \sum_{q\in Q} \alpha_q x_{0,q} (\pi_{0,q} - \chi)\,. \tag{12}$$

(iii) The probability that a new player observes k successes when the state is θ and the adoption rate among type-q players is $x_{\theta,q}$ equals

$$\hat{\phi}_\theta\left(k; (x_q)_{q\in Q}\right) = \left(\begin{array}{c} K \\ k \end{array}\right) \left[\chi + \sum_{q\in Q} \alpha_q x_{\theta,q} (\pi_{\theta,q} - \chi)\right]^k \left[1 - \chi - \sum_{q\in Q} \alpha_q x_{\theta,q} (\pi_{\theta,q} - \chi)\right]^{K-k}.$$

(iv) The population trajectory is given by

$$\dot{X}_{\theta,q}(t) = \begin{cases} \hat{\phi}_\theta\left(k_q^*(t); (X_{\theta,q}(t))_{q\in Q}\right) s_q(t) + \sum_{k=k_q^*(t)+1}^{K} \hat{\phi}_\theta\left(k; (X_{\theta,q}(t))_{q\in Q}\right) - X_{\theta,q}(t) \\ \text{if (12) holds at } (X_{0,q}(t), X_{1,q}(t))_{q\in Q}\,; \\ \hat{\phi}_\theta\left(k_q^*(t); (X_{\theta,q}(t))_{q\in Q}\right) s_q(t) + \sum_{k=0}^{k_q^*(t)-1} \hat{\phi}_\theta\left(k; (X_{\theta,q}(t))_{q\in Q}\right) - X_{\theta,q}(t) \\ \text{if (12) fails at } (X_{0,q}(t), X_{1,q}(t))_{q\in Q} \end{cases}.$$

[23] One can also allow types with $\pi_{0,q} > \pi_{1,q}$, in which case the inequalities in (2) must be reversed for these types.

Although the vector of adoption rates now lies in $\mathbb{R}^{2|Q|}$, Theorems 1 and 2 generalize immediately. For Theorem 1, the argument remains that, whenever $(x_{0,q}, x_{1,q})_{q\in Q}$ is close to the boundary of the set of aligned points (now defined by the hyperplane with equation $\sum_{q\in Q} \alpha_q x_{1,q} (\pi_{1,q} - \chi) - \sum_{q\in Q} \alpha_q x_{0,q} (\pi_{0,q} - \chi) = 0$), $p\left(k; (x_{0,q}, x_{1,q})_{q\in Q}\right)$ is close to p for all $k \in \{0, \ldots, K\}$. Thus, so long as (1) is satisfied for each type q, the population dynamic drifts toward the aligned point $(x_{0,q} = 1, x_{1,q} = 1)_{q\in Q}$ whenever it is close enough to the boundary of the aligned set. The dynamic can therefore never exit the aligned set.

However, matters are quite different when each player instead observes a random sample of K other players' outcomes *and types*. In this case, as long as the ratio $(\pi_{1,q} - \chi) / (\pi_{0,q} - \chi)$ is not the same for all types q, when K is sufficiently large the equilibrium path with initial point $(x_{0,q} = 1, x_{1,q} = 1)_{q\in Q}$ converges to the efficient point $(x_{0,q} = 0, x_{1,q} = 1)_{q\in Q}$ in both the outcome-improving and cost-saving cases. The intuition is that, when $x_{1,q} (\pi_{1,q} - \chi)$ hits $x_{0,q} (\pi_{0,q} - \chi)$ for some type q, new players can still infer the state with high probability solely on the basis of the outcomes of players with types other than q. So long as $(\pi_{1,q} - \chi) / (\pi_{0,q} - \chi)$ is not the same for all q, the outcomes of players of some type will always be informative of the state at each point on the line connecting the initial point $(x_{0,q} = 1, x_{1,q} = 1)_{q\in Q}$ to the efficient point $(x_{0,q} = 0, x_{1,q} = 1)_{q\in Q}$.

6.1.2 Multiple States

The issue of whether and how the main results extend with more than two states of the world is somewhat subtle. I therefore provide a more detailed analysis for this case. I find that, (i) unlike in the two-state case, efficient adoption of a cost-saving innovation starting from an aligned point is sometimes possible with multiple states, but (ii) for a large set of parameters, this does not occur, and (iii) efficient adoption of a cost-saving innovation necessarily involves "complex" Bayesian updating, where the information content of a sample is not monotonic in the number of observed successes.

To start with an example, consider the cost-saving case, and suppose there is a third, "very good" state, $\theta = 2$, with $\pi_0 < \pi_1 < \pi_2 < \chi$. Denote the prior probability of state θ by p_θ, and continue to assume that the risky action is optimal under the prior: $p_0\pi_0 + p_1\pi_1 + p_2\pi_2 - c > \chi$.

Can there exist an equilibrium path leading from the initial point $(x_0, x_1, x_2) = (1,1,1)$ to the efficient point $(0,1,1)$? The answer depends on whether the risky action is optimal under the prior, conditional on the state lying in the set $\{0,1\}$ or $\{0,2\}$: that is, on whether $\frac{p_0}{p_0+p_1}\pi_0+\frac{p_1}{p_0+p_1}\pi_1-c$ and $\frac{p_0}{p_0+p_2}\pi_0+\frac{p_2}{p_0+p_2}\pi_2-c$ are greater or less than χ. To see why, note that any equilibrium path leading from $(1,1,1)$ to $(0,1,1)$ must pass through a point (x_0, x_1, x_2) with $x_0(\pi_0-\chi) = x_1(\pi_1-\chi)$ and a point (x_0', x_1', x_2') with $x_0'(\pi_0-\chi) = x_2'(\pi_2-\chi)$. In the first case (say), after every sample the ratio of probability weights assigned to states 0 and 1 under the posterior will be the same as the ratio under the prior. So if the risky action is optimal under the prior (conditional on the state lying in $\{0,1\}$), then the risky action will be optimal after observing any sample at the point (x_0, x_1, x_2), which implies that the equilibrium path can never reach a point with $x_0(\pi_0-\chi) > x_1(\pi_1-\chi)$. On the other hand, if $\chi > \max\left\{\frac{p_0}{p_0+p_1}\pi_0+\frac{p_1}{p_0+p_1}\pi_1, \frac{p_0}{p_0+p_2}\pi_0+\frac{p_2}{p_0+p_2}\pi_2\right\}-c$, then an equilibrium path from $(1,1,1)$ to $(0,1,1)$ may exist for sufficiently large K.

To address this possibility, I present two results that provide conditions guaranteeing that the main results from the two-state model generalize to the case of multiple states.

Suppose there are $n+1$ states, $\Theta = \{0,\ldots,n\}$, with corresponding conditional success probabilities $\pi_0 < \ldots < \pi_n$ and prior probabilities $(p_0,\ldots,p_n)$. Assume $\pi_0 - c < \chi < \pi_n - c$ and let $\theta^* \in \{0,\ldots,n-1\}$ satisfy $\pi_{\theta^*} - c < \chi < \pi_{\theta^*+1} - c$, so that θ^* is the best state at which the status quo remains optimal. Call states $\theta > \theta^*$ *innovation-optimal*, and call states $\theta \leq \theta^*$ *status quo-optimal*. Say that an *asymptotically efficient path exists* if there exists a sequence of equilibrium paths $\left(X_0^K,\ldots,X_n^K\right)$ indexed by K such that $\left(X_0^K(0),\ldots,X_n^K(0)\right) = (1,\ldots,1)$ and

$$\lim_{K\to\infty}\lim_{t\to\infty}\left(X_0^K(t),\ldots,X_n^K(t)\right) = \left(\underbrace{0,\ldots 0}_{\theta^*+1\text{ times}}, \underbrace{1,\ldots 1}_{n-\theta^*\text{ times}}\right).$$

Theorems 1 and 2 imply that, in the two-state case ($n=1$), an asymptotically efficient path fails to exists when $\chi > \pi_1$ and action 1 is optimal under the prior. The following proposition generalizes this result.

Proposition 7 *Let* $a = \sum_{\theta=0}^{\theta^*} p_\theta$ *be the prior probability that the status quo is optimal.*

Suppose there exists a set of innovation-optimal states $\Theta^* \subseteq \{\theta^* + 1, \ldots, n\}$ *with* $\sum_{\theta \in \Theta^*} p_\theta =$ b *such that (i)* $\chi > \pi_{\max \Theta^*}$ *and (ii) the innovation is optimal when* $\theta = 0$ *with probability* $a/(a+b)$ *and* $\theta = \min \Theta^*$ *with probability* $b/(a+b)$*:*

$$\frac{a}{a+b}\pi_0 + \frac{b}{a+b}\pi_{\min \Theta^*} - c > \chi. \tag{13}$$

Then there does not exist an asymptotically efficient path.

For example, suppose that (i) there is only one status quo-optimal state, state 0, and (ii) there is an innovation-optimal state $\hat{\theta}$ with $\chi > \pi_{\hat{\theta}}$ such that the innovation is optimal conditional on the event $\theta \in \left\{0, \hat{\theta}\right\}$. Then Proposition 7 implies that an asymptotically efficient path does not exist.

It seems difficult to substantially weaken the sufficient condition for the non-existence of an asymptotically efficient path in Proposition 7. However, note that the above example of an asymptotically efficient path with $n = 2$ has the property that, at some points in time, the success rate is non-monotone in the state. Equilibrium paths with this feature are arguably implausible, because the inferences players must draw based on their samples are quite complex: for instance, in the example there are times where players take the risky action if they observe a small or large number of successes and take the safe action if they observe an intermediate number of successes.[24] To capture this idea, I say that an equilibrium path is *simple* if at each point in time observing success is always unambiguously good news or bad news: for each t, $X_\theta(t)(\pi_\theta - \chi)$ is either increasing or decreasing in θ.[25]

Proposition 8 *If* $\chi > \pi_\theta$ *for some innovation-optimal state* θ*, then there does not exist a simple asymptotically efficient path.*

With only two states, every equilibrium path is simple. Hence, when there are only two states, Proposition 8 says precisely that there is no asymptotically efficient path in the cost-saving innovation case. Intuitively, with two states, observing success cannot switch from

[24] In particular, this is how players behave at times where the success rate in state 0 has already crossed the succes rate in state 1, but has not yet crossed the success rate in state 2.

[25] Note that the success rate is increasing (resp., decreasing) in θ if and only if observing more (resp., fewer) successes is better news in the monotone likelihood ratio sense.

being good news to being bad news without passing through a point where it is completely uninformative. However, with more than two states, success can potentially switch from being good news to being bad news by passing through a region where it is "mixed news," in that it shifts the likelihood ratio of pairs of states in a non-monotone manner. As the requirement of simplicity rules out this mixed news case, results from the two-state case generalize under this restriction.

6.2 More General Information Structures

6.2.1 Additional Signals of the State

Theorems 1 and 2 are robust to letting players observe additional exogenous signals of the state, so long as these signals are not so informative that the adoption rates that would result from exogenous information alone are misaligned. For example, such signals could result from unmodeled experimentation on the part of each player prior to making her decision. The same results apply a fortiori if players observe signals of other variables that are measurable with respect to the state, such as the aggregate adoption rate or aggregate social welfare.

Formally, suppose that, before a player acts, she observes an additional real-valued signal $\omega \in \mathbb{R}$ that may depend on both the state θ and the player's observed sample of K outcomes, which I denote by ψ. (Thus, $\psi \in \{0,1\}^K$.) Assume ω has an atomless distribution with continuous density $f_\theta(\omega|\psi)$, independent across players. In addition, fixing a sample ψ, assume without loss of generality that $f_1(\omega|\psi)/f_0(\omega|\psi)$ is increasing in ω, so higher signals are better news about the state. Define a critical signal $\omega^*(\psi)$ by

$$\frac{f_0(\omega^*(\psi)|\psi)}{f_1(\omega^*(\psi)|\psi)} = \frac{p}{1-p}\frac{1-p^*}{p^*},$$

or, if no such signal exists, $\omega^*(\psi) = -\infty$. With the above ordering on ω, let $F_\theta(\omega|\psi) = \int_{\omega'<\omega} f_\theta(\omega'|\psi)\,d\omega'$. The required assumption on the informativeness of ω is as follows:

Assumption 1 For every sample ψ, the point $(1 - F_0(\omega^*(\psi)|\psi), 1 - F_1(\omega^*(\psi)|\psi))$ is aligned.

Note that $F_0(\omega^*(\psi)|\psi) \geq F_1(\omega^*(\psi)|\psi)$. Therefore, Assumption 1 is non-trivial only in the cost-saving case ($\chi > \pi_1$), in which case it is equivalent to the inequality

$$\frac{1 - F_1\left(\omega^*\left(\psi\right)|\psi\right)}{1 - F_0\left(\omega^*\left(\psi\right)|\psi\right)} \leq \frac{\chi - \pi_0}{\chi - \pi_1}.$$

When the signal is completely uninformative, $F_0\left(\omega^*\left(\psi\right)|\psi\right) = F_1\left(\omega^*\left(\psi\right)|\psi\right) = 0$, so this inequality holds. Assumption 1 therefore says that the signal cannot be too informative.

It is straightforward to check that the main results hold under Assumption 1. In particular, the assumption implies that, if the population dynamic is close to L (so other players' outcomes are uninformative), then the exogenous signals drive the population dynamic toward the aligned point $\left(1 - F_0\left(\omega^*\left(\psi\right)|\psi\right), 1 - F_1\left(\omega^*\left(\psi\right)|\psi\right)\right)$. The proof of Theorem 1 is thus easily adapted to show that the population dynamic can never exit the aligned region A. The proof of Theorem 2 also remains valid with trivial modifications.

The results are also robust to introducing "noise players," who use the innovation with some exogenously given probability, not necessarily independent of the state. As long as the shares of noise players using the innovation in the two states constitute an aligned point, equilibrium paths cannot exit the aligned region A. For example, if the noise players do use the innovation with the same probability in both states (so the point corresponding to their behavior lies on the 45° line in (x_0, x_1)-space) then this result applies for *any* proportion of noise players in the population. Note also that players who are simply unaware of the innovation—and thus always use the status quo—are a special case of such noise players.

6.2.2 Unknown Calendar Time

The analysis so far assumes that, when a player enters the game, she knows the adoption rate of the innovation conditional on the state. For the definition of a steady state, this assumption is innocuous. However, as adoption rates change over time along a dynamic equilibrium path, the definition of an equilibrium path implicitly assumes that a player knows the time at which she enters the game. There are several ways of relaxing this assumption, but a particularly simple and natural one is to assume that players always draw inferences as if their samples were drawn from a fixed steady state distribution. An interpretation is that players have an improper uniform prior over the time at which they enter the game and expect a steady state to eventually be reached.

Definition 3 *Fix a steady state* (x_0^*, x_1^*, k^*, s^*). *An* unknown calendar time path relative to (x_0^*, x_1^*, k^*, s^*) *is a pair of differentiable functions* $(X_0 : \mathbb{R}_+ \to [0,1], X_1 : \mathbb{R}_+ \to [0,1])$ *such that, for* $\theta = 0, 1$,

$$\dot{X}_\theta(t) = \begin{cases} \phi_\theta(k^*; X_\theta(t))\, s^* + \sum_{k=k^*+1}^{K} \phi_\theta(k; X_\theta(t)) - X_\theta(t) \\ \text{if (4) holds at } (x_0^*, x_1^*); \\ \phi_\theta(k^*; X_\theta(t))\, s^* + \sum_{k=0}^{k^*-1} \phi_\theta(k; X_\theta(t)) - X_\theta(t) \\ \text{if (4) fails at } (x_0^*, x_1^*) \end{cases}.$$

My results are mostly robust to considering this alternative equilibrium concept. In particular, an unknown calendar time path relative to an misaligned equilibrium can never converge to that equilibrium from an aligned initial point. The intuition is that the success rate at the initial point is higher in the good state, and if players use the steady-state inference rule to mistakenly conclude that success is bad news about the state, they switch to the safe action at a higher rate in the good state, which further improves the success rate in the good state.

Proposition 9 *If* (x_0^*, x_1^*, k^*, s^*) *is an misaligned equilibrium,* (X_0, X_1) *is an unknown calendar time path relative to* (x_0^*, x_1^*, k^*, s^*), *and* $(X_0(0), X_1(0))$ *is aligned, then* $(X_0(t), X_1(t))$ *is aligned for all* $t \in \mathbb{R}_+$. *In particular,* $(X_0(t), X_1(t))$ *is bounded away from* (x_0^*, x_1^*) *for all* $t \in \mathbb{R}_+$.

Conversely, a large class of aligned steady states can be reached by an unknown calendar time path starting from a range of aligned initial points, including the point $(x_0, x_1) = (1, 1)$ that results from optimal choice under the prior. Say that an aligned steady state (x_0^*, x_1^*, k^*, s^*) is *stable from above* if, for $\theta = 0, 1$, $\phi_\theta(k^*; x_\theta)\, s^* + \sum_{k=k^*+1}^{K} \phi_\theta(k; x_\theta) < x_\theta$ for all $x_\theta > x_\theta^*$. When $K = 1$, the unique equilibrium is stable from above; in Example 2 in Appendix B, the larger stable aligned equilibrium is stable from above.

Proposition 10 *If* (x_0^*, x_1^*, k^*, s^*) *is an aligned equilibrium that is stable from above,* (X_0, X_1) *is an unknown calendar time path relative to* (x_0^*, x_1^*, k^*, s^*), *and* $(X_0(0), X_1(0))$ *is aligned and satisfies* $X_\theta(0) \geq x_\theta^*$ *for* $\theta = 0, 1$, *then* $\lim_{t\to\infty}(X_0(t), X_1(t)) = (x_0^*, x_1^*)$.

7 Conclusion

This paper has developed a simple model of social learning where learning is *outcome-based*: players observe each other's outcomes, but not their actions. Since outcomes are noisy, the resulting picture of social learning features both inefficiency and persistent diversity of actions. I have focused on the question of how the nature and extent of inefficiency depend on features of the innovation about which the group must learn. The most striking finding is that, while outcome-improving innovations are adopted efficiently when players' samples are large and the population has had enough time to learn, the adoption of cost-saving innovations entails substantial inefficiency whenever the initial generation of players is not already well-informed. While I have not modeled the production of innovations, this difference in adoption patterns can be expected to bias the innovation process against cost-saving innovations, which is consistent with the observed lack of cost-saving innovations in many fields.

Let me conclude by pointing out two implications of the model, one positive and one normative.

First, while the model is admittedly simple and stylized, it does make some clear empirical predictions that would be interesting to test. The key prediction is that the correlation between an environment's underlying suitability to an innovation and the realized adoption rate will be greater for outcome-improving innovations than for cost-saving innovations. For example, in the agricultural context, the prediction is that the correlation between soil suitability and adoption will be greater for outcome-improving innovations (like high-yield crop varieties) than for cost-saving innovations (like labor-saving varieties, such as the genetically engineered soybeans studied by Bustos, Caprettini, and Ponticelli (2016)). While I am not aware of existing data that could directly be brought to bear here, this seems like an interesting direction for empirical research.

Second, consider the situation of a benevolent outsider, like an NGO or a professional organization, that wants to design an intervention to improve the efficiency of technology adoption. In the problematic cost-saving case, publicly releasing information (such as test results) about the technology's effectiveness may not help: the public release of information at the beginning of the game serves only to change the prior p, and the result that the long-run adoption of cost-saving innovations is inefficient is largely independent of the prior. But other interventions could be more effective. For instance, gathering and releasing data on the adoption rate of the innovation is a promising way to restore efficiency: even when long-run adoption is inefficient, the adoption rate will still differ depending on the innovation's effectiveness, so revealing the adoption rate will reveal the state. Another promising approach is providing an information technology that lets each individual conduct her own test of the innovation before deciding whether to adopt it. If these tests are sufficiently informative, the initial adoption rates will be misaligned (i.e., Assumption 1 will be violated), and social learning can then lead the adoption rates to converge to efficiency.

8 Appendix A: Proof of Theorem 2

The proof relies on four lemmas. Lemma 1 derives the asymptotic ($K \to \infty$) formula for the cutoff fraction of observed successes above or below which new players adopt, and shows that convergence to this limit is of order $1/K$. Lemma 2 uses this to show that the adoption rate among new players converges to efficiency as $K \to \infty$, uniformly over points x bounded away from L. This is the key step in the proof. Lemmas 3 and 4 then establish some additional properties of equilibrium paths that are useful in the outcome-improving case when the initial point is misaligned—this case raises some new issues, as now the equilibrium path must cross the line L to reach the efficient point.

Throughout the proof, given a point $x = (x_0, x_1) \in [0,1]^2$, I write $\sigma_0(x) = \sigma_0(x_0)$ and $\sigma_1(x) = \sigma_1(x_1)$. I also sometimes reduce notation by writing x and $\dot{x}$ for $X(t)$ and $\dot{X}(t)$, including the time argument only when necessary.

Lemma 1 *For every K, every equilibrium path (X_0, X_1, k^*, s), and every time t with $X(t) = x$ aligned and $(\sigma_0(x), \sigma_1(x)) \in (0,1)^2$,*

$$\frac{k^*(t)}{K} \in \left[\sigma^*(x) - \frac{1}{K}\zeta(x), \sigma^*(x) - \frac{1}{K}\zeta(x) + \frac{1}{K}\right],$$

where

$$\sigma^*(x) = \frac{\log \frac{1-\sigma_0(x)}{1-\sigma_1(x)}}{\log\left(\frac{1-\sigma_0(x)}{\sigma_0(x)} \frac{\sigma_1(x)}{1-\sigma_1(x)}\right)} \in (\sigma_0(x), \sigma_1(x))$$

and

$$\zeta(x) = \frac{\log\left(\frac{p}{1-p}\frac{1-p^*}{p^*}\right)}{\log\left(\frac{1-\sigma_0(x)}{\sigma_0(x)} \frac{\sigma_1(x)}{1-\sigma_1(x)}\right)}.$$

Similarly, for every K, every equilibrium path (X_0, X_1, k^*, s), and every time t with $X(t) = x$ misaligned and $(\sigma_0(x), \sigma_1(x)) \in (0,1)^2$,

$$\frac{k^*(t)}{K} \in \left[\sigma^*(x) - \frac{1}{K}\zeta(x) - \frac{1}{K}, \sigma^*(x) - \frac{1}{K}\zeta(x)\right],$$

where now $\sigma^*(x) \in (\sigma_1(x), \sigma_0(x))$.

Proof. Consider the aligned case. (The misaligned case is symmetric.) By the definition of an equilibrium path, $p(k^*(t)-1;x_0,x_1) \leq p^* \leq p(k^*(t);x_0,x_1)$. Hence, by Bayes' rule,

$$\frac{\sigma_0(x)^{k^*(t)-1}(1-\sigma_0(x))^{K-k^*(t)+1}}{\sigma_1(x)^{k^*(t)-1}(1-\sigma_1(x))^{K-k^*(t)+1}} \geq \frac{p}{1-p}\frac{1-p^*}{p^*} \geq \frac{\sigma_0(x)^{k^*(t)}(1-\sigma_0(x))^{K-k^*(t)}}{\sigma_1(x)^{k^*(t)}(1-\sigma_1(x))^{K-k^*(t)}}.$$

Letting $\sigma_K^* = \frac{k^*(t)}{K}$, rewrite this as

$$\left[\frac{\sigma_0(x)^{\sigma_K^*-\frac{1}{K}}(1-\sigma_0(x))^{1-\sigma_K^*+\frac{1}{K}}}{\sigma_1(x)^{\sigma_K^*-\frac{1}{K}}(1-\sigma_1(x))^{1-\sigma_K^*+\frac{1}{K}}}\right]^K \geq \frac{p}{1-p}\frac{1-p^*}{p^*} \geq \left[\frac{\sigma_0(x)^{\sigma_K^*}(1-\sigma_0(x))^{1-\sigma_K^*}}{\sigma_1(x)^{\sigma_K^*}(1-\sigma_1(x))^{1-\sigma_K^*}}\right]^K.$$

Taking logs and grouping terms yields

$$\begin{aligned} K\left(\sigma_K^*-\frac{1}{K}\right)\log\left(\frac{\sigma_0(x)}{1-\sigma_0(x)}\frac{1-\sigma_1(x)}{\sigma_1(x)}\right) &\geq \log\left(\frac{p}{1-p}\frac{1-p^*}{p^*}\right)+K\log\frac{1-\sigma_1(x)}{1-\sigma_0(x)} \\ &\geq K\sigma_K^*\log\left(\frac{\sigma_0(x)}{1-\sigma_0(x)}\frac{1-\sigma_1(x)}{\sigma_1(x)}\right). \end{aligned}$$

Dividing by $K\log\left(\frac{\sigma_0(x)}{1-\sigma_0(x)}\frac{1-\sigma_1(x)}{\sigma_1(x)}\right)$ (a negative number) yields the desired result. ■

Lemma 2 *For every $\varepsilon,\eta>0$, there exists $\bar{K}>0$ such that, for every $K>\bar{K}$, every equilibrium path (X_0,X_1,k^*,s), and every time t such that $d(X(t),L)>\eta$,*

$$\left\|\dot{X}(t)+X(t)-(0,1)\right\|<\varepsilon.$$

Proof. Note that $d(x,L)>\eta$ implies $|\sigma_0(x)-\sigma_1(x)|>\|(\chi-\pi_0,\chi-\pi_1)\|\,\eta$. I prove the lemma for $(\sigma_0(x),\sigma_1(x))\in(0,1)^2$. The argument for the case where at least one success rate lies in $\{0,1\}$ is similar but simpler.

By Lemma 1, for every K, every equilibrium path, and every t with $X(t)=x$ aligned and $(\sigma_0(x),\sigma_1(x))\in(0,1)^2$, one has

$$\begin{aligned} \dot{x}_\theta + x_\theta &\leq \Pr(k\geq k^*(t)\,|\theta) \\ &= \Pr\left(\frac{k}{K}\geq\frac{k^*(t)}{K}|\theta\right) \\ &\leq \Pr\left(\sigma\geq\sigma^*(x)-\frac{1}{K}\zeta(x)\,|\theta\right), \end{aligned}$$

where $\sigma := \frac{k}{K}$. Note that

$$\begin{aligned}
\Pr\left(\sigma \geq \sigma^*(x) - \frac{1}{K}\zeta(x)\,|\theta = 0\right) &= \Pr\left(\sigma - \sigma_0(x) \geq \sigma^*(x) - \frac{1}{K}\zeta(x) - \sigma_0(x)\,|\theta = 0\right) \\
&\leq \frac{\mathrm{Var}(\sigma|\theta = 0)}{\left(\sigma^*(x) - \frac{1}{K}\zeta(x) - \sigma_0(x)\right)^2} \\
&= \frac{\sigma_0(x)(1 - \sigma_0(x))}{K\left(\sigma^*(x) - \frac{1}{K}\zeta(x) - \sigma_0(x)\right)^2},
\end{aligned}$$

where the second line uses Chebyshev's inequality and the third follows because outcomes are conditionally independent Bernoulli trials. Let $(L + B_\eta)^c := \left\{x \in [0,1]^2 : d(x, L) > \eta\right\}$. Let

$$C\left(\eta, \bar{K}\right) = \sup_{x \in (L+B_\eta)^c, K > \bar{K}} \frac{\sigma_0(x)(1 - \sigma_0(x))}{\left(\sigma^*(x) - \frac{1}{K}\zeta(x) - \sigma_0(x)\right)^2}.$$

Thus, $\Pr\left(\sigma \geq \sigma^*(x) - \frac{1}{K}\zeta(x)\,|\theta = 0\right) \leq C\left(\eta, \bar{K}\right)/K$ for all $x \in (L + B_\eta)^c$ and $K > \bar{K}$.

I claim that, for all $\eta > 0$, there exists $\bar{K}$ such that $C\left(\eta, \bar{K}\right) < \infty$. To see this, recall the formulas for σ^* and ζ, and note that σ^* in increasing in σ_1 and ζ is decreasing in σ_1 (suppressing the dependence of the σ variables and ζ on x). Hence,

$$\sup_{x \in (L+B_\eta)^c} \frac{\sigma_0(x)(1 - \sigma_0(x))}{\left(\sigma^*(x) - \frac{1}{K}\zeta(x) - \sigma_0(x)\right)^2}$$

is upper-bounded by taking $\sigma_1 = \sigma_0 + \|(\chi - \pi_0, \chi - \pi_1)\|\,\eta$ (recalling the restriction to aligned points, where $\sigma_1 \geq \sigma_0$). Let $\hat{\eta} = \|(\chi - \pi_0, \chi - \pi_1)\|\,\eta$. After substituting out for σ^*, ζ, and σ_1, and again suppressing the dependence on x, the resulting upper bound equals

$$\sup_{\sigma_0 \in (0, 1-\hat{\eta})} \frac{\sigma_0(1 - \sigma_0)}{\left(\frac{\log \frac{1-\sigma_0}{1-\sigma_0-\hat{\eta}}}{\log\left(\frac{1-\sigma_0}{\sigma_0}\frac{\sigma_0+\hat{\eta}}{1-\sigma_0-\hat{\eta}}\right)} - \frac{\frac{1}{K}\log\left(\frac{p}{1-p}\frac{1-p^*}{p^*}\right)}{\log\left(\frac{1-\sigma_0}{\sigma_0}\frac{\sigma_0+\hat{\eta}}{1-\sigma_0-\hat{\eta}}\right)} - \sigma_0\right)^2}.$$

When $K = \infty$, some simple calculus shows that this expression attains its maximum in the open interval $(0, 1 - \hat{\eta})$. Therefore, there exists $\bar{K}$ sufficiently large such that

$$\max_{\sigma_0 \in (0, 1-\hat{\eta})} \frac{\sigma_0(1 - \sigma_0)}{\left(\frac{\log \frac{1-\sigma_0}{1-\sigma_0-\hat{\eta}}}{\log\left(\frac{1-\sigma_0}{\sigma_0}\frac{\sigma_0+\hat{\eta}}{1-\sigma_0-\hat{\eta}}\right)} - \frac{\frac{1}{K}\log\left(\frac{p}{1-p}\frac{1-p^*}{p^*}\right)}{\log\left(\frac{1-\sigma_0}{\sigma_0}\frac{\sigma_0+\hat{\eta}}{1-\sigma_0-\hat{\eta}}\right)} - \sigma_0\right)^2}$$

exists for all $K > \bar{K}$. As the maximand of this expression is decreasing in K for each σ_0, this implies $C\left(\eta, \bar{K}\right) < \infty$.

I have shown that, for all $\eta > 0$, there exist $\bar{K}$ and a finite number $C\left(\eta, \bar{K}\right)$ such that $\Pr\left(\sigma \geq \sigma^*\left(x\right) - \frac{1}{K}\zeta\left(x\right)|\theta = 0\right) \leq C\left(\eta, \bar{K}\right)/K$ for all $x \in \left(L + B_\eta\right)^c$ and $K > \bar{K}$. Hence, for all $\varepsilon, \eta > 0$, there exists $\bar{K}$ such that $\Pr\left(\sigma \geq \sigma^*\left(x\right) - \frac{1}{K}\zeta\left(x\right)|\theta = 0\right) < \varepsilon$ for all $x \in \left(L + B_\eta\right)^c$ and $K > \bar{K}$. Therefore, $\dot{x}_0 + x_0 < \varepsilon$ whenever $x \in \left(L + B_\eta\right)^c$ and $K > \bar{K}$. A symmetric argument implies that $\dot{x}_1 + x_1 \to 1$ uniformly over $x \in \left(L + B_\eta\right)^c$. This yields the conclusion of the lemma in the case where x is aligned. The misaligned case is symmetric. ■

Lemma 3 *For every K, every t, and every equilibrium path $\left(X_0, X_1, k^*, s\right)$,*

$$\dot{X}_1\left(t\right) + X_1\left(t\right) \geq \max\left\{\dot{X}_0\left(t\right) + X_0\left(t\right), \frac{p - p^*}{1 - p^*}\right\}.$$[26]

Proof. At an aligned point x, $\phi_0\left(k; x_0\right)/\phi_1\left(k; x_1\right)$ is decreasing in k, and therefore (as likelihood ratio dominance implies first-order stochastic dominance)

$$\sum_{k=k_0}^{K}\phi_0\left(k; x_0\right) \leq \sum_{k=k_0}^{K}\phi_1\left(k; x_1\right) \text{ for all } k_0 \in \left\{0, \dots, K\right\}.$$

In particular, this inequality holds for $k_0 = k^*$ and $k_0 = k^* + 1$. Hence,

$$\dot{x}_0 + x_0 = \phi_0\left(k^*; x_0\right)s + \sum_{k=k^*+1}^{K}\phi_0\left(k; x_0\right) \leq \phi_1\left(k^*; x_1\right)s + \sum_{k=k^*+1}^{K}\phi_1\left(k; x_1\right) = \dot{x}_1 + x_1.$$

Symmetrically, at an misaligned point, $\phi_0\left(k; x_0\right)/\phi_1\left(k; x_1\right)$ is increasing in k and new players adopt if $k < k^*$, so again $\dot{x}_0 + x_0 \leq \dot{x}_1 + x_1$.

The second part of the proposition follows because the Bayesian information structure that minimizes the probability that a player's posterior is strictly above p^* conditional on $\theta = 1$ sends her posterior to p^* and 1 with probabilities $\frac{1-p}{1-p^*}$ and $\frac{p-p^*}{1-p^*}$, respectively. ■

[26] This implies that, at every steady state, $x_1 \geq \max\left\{x_0, \frac{p-p^*}{1-p^*}\right\}$. The latter fact is not used in the proof of Theorem 2, but it is used in some of the proofs in Appendix B.

Lemma 4 *There exist $\bar{K} > 0$ and $T > 0$ such that, for every $K > \bar{K}$, every $t > T$, and every equilibrium path (X_0, X_1, k^*, s),*

$$X_1(t) \geq \max \left\{ X_0(t), \frac{1}{2} \frac{p - p^*}{1 - p^*} \right\}.$$

Proof. Recall that $\dot{x}_1 + x_1 \geq \frac{p-p^*}{1-p^*}$, by Lemma 3. In particular, if $x_1 \leq \frac{1}{2}\frac{p-p^*}{1-p^*}$ then $\dot{x}_1 \geq \frac{1}{2}\frac{p-p^*}{1-p^*}$. Hence, $x_1(t) \geq \frac{1}{2}\frac{p-p^*}{1-p^*}$ for all $t \geq 1$.

Next, note that the region $\left\{x \in [0,1]^2 : x_1 \geq x_0\right\}$ is forward invariant. This follows because, by Lemma 3, $\dot{x}_1 + x_1 \geq \dot{x}_0 + x_0$, so $x_1 = x_0$ implies $\dot{x}_1 \geq \dot{x}_0$.

It remains to show that x_0 can exceed x_1 for only a finite length of time T. This is proved by deriving a lower bound on $\dot{x}_1 - \dot{x}_0$ that applies whenever $x_0 > x_1$.

First, for all $\eta > 0$, there exists $\bar{K}$ such that, if $K > \bar{K}$, $x_0 > x_1$, and $d(x, L) \geq \eta$, then $\dot{x}_1 - \dot{x}_0 > \frac{1}{2}$. To see this, note that taking $\varepsilon = \frac{1}{2}$ in Lemma 2 implies $\dot{x}_1 + x_1 > \dot{x}_0 + x_0 + \frac{1}{2}$. When $x_0 > x_1$, this gives $\dot{x}_1 - \dot{x}_0 > \frac{1}{2}$.

Second, for all $\eta > 0$, if $x_0 > x_1$, $d(x, L) \leq \eta$, and $\pi_0 > \chi$, then $\dot{x}_1 - \dot{x}_0 \geq \frac{\pi_1 - \pi_0}{\pi_0 - \chi} x_1 - \frac{\|(\pi_0 - \chi, \pi_1 - \chi)\|}{\pi_0 - \chi} \eta$. To see this, note that

$$d(x, L) = \frac{|(\pi_1 - \chi) x_1 - (\pi_0 - \chi) x_0|}{\|(\pi_0 - \chi, \pi_1 - \chi)\|} = \frac{|(\pi_1 - \pi_0) x_1 - (\pi_0 - \chi)(x_0 - x_1)|}{\|(\pi_0 - \chi, \pi_1 - \chi)\|},$$

so $d(x, L) \leq \eta$ and $\pi_0 > \chi$ imply $x_0 - x_1 \geq \frac{\pi_1 - \pi_0}{\pi_0 - \chi} x_1 - \frac{\|(\pi_0 - \chi, \pi_1 - \chi)\|}{\pi_0 - \chi} \eta$. Finally, Lemma 3 implies $\dot{x}_1 - \dot{x}_0 \geq x_0 - x_1$.

One can now derive a formula for T. If $\chi \geq \pi_1$, then $x_0 > x_1$ implies $d(x, L) \geq \frac{\pi_1 - \pi_0}{\|(\pi_0 - \chi, \pi_1 - \chi)\|} x_1$, so $x_1(t) \geq \frac{1}{2}\frac{p-p^*}{1-p^*}$ for all $t \geq 1$ implies $d(x(t), L) \geq \frac{1}{2}\frac{p-p^*}{1-p^*}\frac{\pi_1 - \pi_0}{\|(\pi_0 - \chi, \pi_1 - \chi)\|}$ for all $t \geq 1$. If $\pi_1 > \chi \geq \pi_0$, then $d(x, L) \geq \frac{\pi_1 - \chi}{\|(\pi_0 - \chi, \pi_1 - \chi)\|} x_1$, so $x_1(t) \geq \frac{1}{2}\frac{p-p^*}{1-p^*}$ for all $t \geq 1$ implies $d(x(t), L) \geq \frac{1}{2}\frac{p-p^*}{1-p^*}\frac{\pi_1 - \chi}{\|(\pi_0 - \chi, \pi_1 - \chi)\|}$ for all $t \geq 1$. In either case, there exists $\bar{K}$ such that, if $K > \bar{K}$, $x_0(t) > x_1(t)$, and $t \geq 1$, then $\dot{x}_1(t) - \dot{x}_0(t) > \frac{1}{2}$. Hence, $x_1(t) - x_0(t)$ must reach 0 in time at most $1 + 1/\left(\frac{1}{2}\right) = 3$.

If instead $\chi < \pi_0$, let $\eta = \frac{1}{4}\frac{p-p^*}{1-p^*}\frac{\pi_1 - \pi_0}{\|(\pi_0 - \chi, \pi_1 - \chi)\|}$. Then $x_1(t) \geq \frac{1}{2}\frac{p-p^*}{1-p^*}$ implies $\frac{\pi_1 - \pi_0}{\pi_0 - \chi} x_1(t) - \frac{\|(\pi_0 - \chi, \pi_1 - \chi)\|}{\pi_0 - \chi} \eta \geq \frac{1}{4}\frac{p-p^*}{1-p^*}\frac{\pi_1 - \pi_0}{\pi_0 - \chi}$, so $\dot{x}_1(t) - \dot{x}_0(t) \geq \frac{1}{4}\frac{p-p^*}{1-p^*}\frac{\pi_1 - \pi_0}{\pi_0 - \chi}$ for all $t \geq 1$. (Note that $\frac{1}{4}\frac{p-p^*}{1-p^*}\frac{\pi_1 - \pi_0}{\pi_0 - \chi} < \frac{1}{2}$, so this bound applies whether or not $x(t)$ is close to L.) Hence, $x_1(t) - x_0(t)$

must reach 0 in time at most $1+4\frac{1-p^*}{p-p^*}\frac{\pi_0-\chi}{\pi_1-\pi_0}$. ■

Proof of Theorem 2. Fix $\varepsilon>0$. By Lemma 4, assume without loss of generality that $x_1(t)\geq\max\left\{x_0(t),\frac{1}{2}\frac{p-p^*}{1-p^*}\right\}$ for all $t\geq 0$. (Otherwise, replace t by $t+T$ throughout.)

Assume $\chi<\pi_1$. Note that, for every x such that $x_1\geq\max\left\{x_0,\frac{1}{2}\frac{p-p^*}{1-p^*}\right\}$,

$$\begin{aligned} x_1(\pi_1-\chi)-x_0(\pi_0-\chi) &\geq x_1(\pi_1-\chi)-x_0\max\{\pi_0-\chi,0\} \\ &\geq x_1\min\{\pi_1-\pi_0,\pi_1-\chi\} \\ &\geq \frac{1}{2}\frac{p-p^*}{1-p^*}\min\{\pi_1-\pi_0,\pi_1-\chi\}>0. \end{aligned}$$

Hence, for every such x,

$$d(x,L)\geq\frac{\frac{1}{2}\frac{p-p^*}{1-p^*}\min\{\pi_1-\pi_0,\pi_1-\chi\}}{\|(\pi_1-\chi,\pi_0-\chi)\|}=:\eta.$$

By Lemma 2, there exists $\bar{K}$ such that, for every $K>\bar{K}$ and every x such that $d(x,L)\geq\eta$, $\|\dot{x}+x-(0,1)\|<\frac{\varepsilon}{3}$. Hence, for every such x,

$$\begin{aligned} \frac{d}{dt}\|x(t)-(0,1)\| &= \frac{x_0\dot{x}_0-(1-x_1)\dot{x}_1}{\|x(t)-(0,1)\|} \\ &\leq -\frac{x_0^2+(1-x_1)^2}{\|x(t)-(0,1)\|}+\frac{x_0+(1-x_1)}{\|x(t)-(0,1)\|}\frac{\varepsilon}{3} \\ &\leq -\|x(t)-(0,1)\|+\frac{2\varepsilon}{3}. \end{aligned}$$

Recall that $x_1(t)\geq\max\left\{x_0(t),\frac{1}{2}\frac{p-p^*}{1-p^*}\right\}$ for all t, and therefore $d(x(t),L)\geq\eta$ for all t. So, if $\|x(t)-(0,1)\|>\varepsilon$, then $\frac{d}{dt}\|x(t)-(0,1)\|\leq-\frac{\varepsilon}{3}$. Let $T=\|\hat{x}-(0,1)\|\frac{3}{\varepsilon}$. As $x(t)$ is absolutely continuous, we have $\|x(t)-(0,1)\|=\|\hat{x}-(0,1)\|+\int_{s=0}^{t}\left(\frac{d}{ds}\|x(s)-(0,1)\|\right)ds$. In particular, $\|x(t)-(0,1)\|\leq\|\hat{x}-(0,1)\|-t\frac{\varepsilon}{3}$ if $\|x(s)-(0,1)\|>\varepsilon$ for all $s<t$, and if $\|x(t)-(0,1)\|\leq\varepsilon$ then $\|x(s)-(0,1)\|\leq\varepsilon$ for all $s\geq t$. Hence, $\|x(t)-(0,1)\|\leq\varepsilon$ for all $t>T$.

If instead $\chi>\pi_1$, assume $d(\hat{x},L)>0$, and let $\eta:=d(\hat{x},L)$. By Lemma 2, there exists $\bar{K}$ such that, for every $K>\bar{K}$ and every x such that $d(x,L)\geq\eta$, $\|\dot{x}+x-(0,1)\|<\frac{\varepsilon}{3}\frac{\chi-\pi_1}{\chi-\pi_0}$.

Hence, for every t,

$$\frac{d}{dt}\left\|x\left(t\right)-\left(0,1\right)\right\|\leq-\left\|x\left(t\right)-\left(0,1\right)\right\|+\frac{2\varepsilon}{3}.$$

So, if $\left\|x\left(t\right)-\left(0,1\right)\right\|>\varepsilon$ and $d\left(x\left(t\right),L\right)\geq\eta$, then $\frac{d}{dt}\left\|x\left(t\right)-\left(0,1\right)\right\|\leq-\frac{\varepsilon}{3}$.

In addition, if $x\left(t\right)$ is misaligned then

$$\begin{aligned}
\frac{d}{dt}d\left(x\left(t\right),L\right) &= \frac{\left(\chi-\pi_1\right)\dot{x}_1-\left(\chi-\pi_0\right)\dot{x}_0}{\left\|\left(\chi-\pi_1,\chi-\pi_0\right)\right\|}\\
&\geq \frac{\left(\chi-\pi_1\right)\left(1-x_1-\frac{\varepsilon}{3}\frac{\chi-\pi_1}{\chi-\pi_0}\right)+\left(\chi-\pi_0\right)\left(x_0-\frac{\varepsilon}{3}\frac{\chi-\pi_1}{\chi-\pi_0}\right)}{\left\|\left(\chi-\pi_1,\chi-\pi_0\right)\right\|}\\
&\geq \frac{\left(\chi-\pi_1\right)\left(\varepsilon-\frac{\varepsilon}{3}\frac{\chi-\pi_1}{\chi-\pi_0}\right)+\left(\chi-\pi_0\right)\left(-\frac{\varepsilon}{3}\frac{\chi-\pi_1}{\chi-\pi_0}\right)}{\left\|\left(\chi-\pi_1,\chi-\pi_0\right)\right\|}\\
&\geq \frac{\chi-\pi_1}{\left\|\left(\chi-\pi_1,\chi-\pi_0\right)\right\|}\frac{\varepsilon}{3}.
\end{aligned}$$

Hence, if $\hat{x}$ is misaligned then $x\left(t\right)$ is misaligned and satisfies $d\left(x\left(t\right),L\right)\geq\eta$ for all t. So, letting

$$T=\left\|\hat{x}-\left(0,1\right)\right\|\frac{\left\|\left(\chi-\pi_1,\chi-\pi_0\right)\right\|}{\chi-\pi_1}\frac{3}{\varepsilon},$$

it follows that $\left\|x\left(t\right)-\left(0,1\right)\right\|\leq\varepsilon$ for all $t>T$, by the same argument as in the $\chi<\pi_1$ case.

Conversely, if $x\left(t\right)$ is aligned then

$$\begin{aligned}
\frac{d}{dt}d\left(x\left(t\right),L\right) &= -\frac{\left(\chi-\pi_1\right)\dot{x}_1-\left(\chi-\pi_0\right)\dot{x}_0}{\left\|\left(\chi-\pi_1,\chi-\pi_0\right)\right\|}\\
&\leq -\frac{\chi-\pi_1}{\left\|\left(\chi-\pi_1,\chi-\pi_0\right)\right\|}\frac{\varepsilon}{3}.
\end{aligned}$$

Recall from Theorem 1 that if $\hat{x}$ is aligned then $x\left(t\right)$ is aligned for all t. Hence, $\left\|x\left(t\right)-L\right\|\leq\varepsilon$ for all $t>T$.

Finally, as the conclusion of the theorem holds for all aligned initial points with $d\left(\hat{x},L\right)>0$, it also holds for initial points with $d\left(\hat{x},L\right)=0$. ■

References

[1] Acemoglu, Daron, and Alexander Wolitzky (2014), "Cycles of Conflict: An Economic Model," *American Economic Review*, 104: 1350-1367.

[2] Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien (1991), "Optimal Learning by Experimentation," *Review of Economic Studies*, 58: 621-654.

[3] Aubin, J.-P., and A. Cellina (1984), *Differential Inclusions: Set-Valued Maps and Viability Theory*, Springer: Berlin.

[4] Bala, Venkatesh, and Sanjeev Goyal (1995), "A Theory of Learning with Heterogeneous Agents," *International Economic Review*, 36: 303-323.

[5] Banerjee, Abhijit V. (1992), "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107: 797-817.

[6] Banerjee, Abhijit V. (1993), "The Economics of Rumours," *Review of Economic Studies*, 60: 309-327.

[7] Banerjee, Abhijit, and Drew Fudenberg (2004), "Word-of-Mouth Learning," *Games and Economic Behavior*, 46: 1-22.

[8] Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100: 992-1026.

[9] Bloom, Nicholas, Ben Eifert, Aprajit Mahajan, David McKenzie, and John Roberts (2013), "Does Management Matter? Evidence from India," *Quarterly Journal of Economics*, 128: 1-51.

[10] Bloom, Nicholas, and John Van Reenen (2010), "Why Do Management Practices Differ across Firms and Countries?" *Journal of Economic Perspectives*, 24: 203-224.

[11] Bolton, Patrick, and Christopher Harris (1999), "Strategic Experimentation," *Econometrica*, 67: 349-374.

[12] Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli (2016), "Agricultural Productivity and Structural Transformation: Evidence from Brazil," *American Economic Review*, 106: 1320-1365.

[13] Chakma, Justin, Gordon H. Sun, Jeffrey D. Steinberg, Stephen M. Sammut, and Reshma Jagsi (2014), "Asia's Ascent—Global Trends in Biomedical R&D Expenditures," *New England Journal of Medicine*, 370: 3-6.

[14] Chamley, Christophe (2004), *Rational Herds: Economic Models of Social Learning*, Cambridge University Press: Cambridge.

[15] Chandra, Amitabh, and Jonathan Skinner (2012), "Technology Growth and Expenditure Growth in Health Care," *Journal of Economic Literature*, 50: 645-680.

[16] Che, Yeon-Koo, and Johannes Hörner (2017), "Optimal Design for Social Learning," *Quarterly Journal of Economics*, forthcoming.

[17] Conley, Timothy G., and Christopher R. Udry (2010), "Learning About a New Technology: Pineapple in Ghana," *American Economic Review*, 100: 35-69.

[18] Cutler, David M. (2005), *Your Money or Your Life: Strong Medicine for America's Health Care System*, Oxford University Press: Oxford.

[19] Das, Jishnu, Alaka Holla, Aakash Mohpol, and Karthik Muralidharan, (2016) "Quality and Accountability in Healthcare Delivery: Audit-Study Evidence from Primary Care in India," *American Economic Review*, 106: 3765-3799.

[20] Duflo, Esther, Michael Kremer, and Jonathan Robinson (2008), "How High are Rates of Return to Fertilizer? Evidence from Field Experiments in Kenya," *American Economic Review: Papers & Proceedings*, 98: 482-488.

[21] Easley, David, and Nicholas M. Kiefer (1988), "Controlling a Stochastic Process with Unknown Parameters," *Econometrica*, 56: 1045-1064.

[22] Foster, Andrew D., and Mark R. Rosenzweig (2010), "Microeconomics of Technology Adoption," *Annual Review of Economics*, 2: 395-424.

[23] Frick, Mira and Yuhta Ishii (2016), "Innovation Adoption by Forward-Looking Social Learners," *working paper.*

[24] Hörner, Johannes and Andrzej Skrzypacz (2016), "Learning, Experimentation and Information Design," *working paper.*

[25] Keller, Godfrey, Sven Rady, and Martin Cripps (2005), "Strategic Experimentation with Exponential Bandits," *Econometrica*, 73: 39-68.

[26] Kiefer, Nicholas M. (1989), "A Value Function Arising in the Economics of Information," *Journal of Economic Dynamics and Control*, 13: 201-223.

[27] Kremer, Ilan, Yishay Mansour, and Motty Perry (2014), "Implementing the "Wisdom of the Crowd"," *Journal of Political Economy*, 122: 988-1012.

[28] Lauermann, Stephan, and Asher Wolinsky (2016), "Search with Adverse Selection," *Econometrica*, 84: 243-315.

[29] McLennan, Andrew (1984), "Price Dispersion and Incomplete Learning in the Long Run," *Journal of Economic Dynamics and Control*, 7: 331-347.

[30] Milgrom, Paul R. (1981), "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12: 380-391.

[31] Murto, Pauli, and Juuso Välimäki (2011), "Learning and Information Aggregation in an Exit Game," *Review of Economic Studies*, 78: 1426-1461.

[32] Nelson, Aaron L., Joshua T. Choen, Dan Greenberg, and David M. Kent, (2009), "Much Cheaper, Almost as Good: Decrementally Cost-Effective Medical Innovation," *Annals of Internal Medicine*, 151: 662-667.

[33] Newhouse, Joseph P. (1992), "Medical Care Costs: How Much Welfare Loss?" *Journal of Economic Perspectives*, 6: 3-21.

[34] Piketty, Thomas (1995), "Social Mobility and Redistributive Politics," *Quarterly Journal of Economics*, 110: 551-584.

[35] Pray, Carl E., and Keith O. Fuglie, (2015), "Agricultural Research by the Private Sector," *Annual Review of Resource Economics*, 7: 399-424.

[36] Schilbach, Frank (2015), *Essays in Development and Behavioral Economics*, PhD Thesis, Harvard University.

[37] Smith, Lones, and Peter Sørensen (2000), "Pathological Outcomes of Observational Learning," *Econometrica*, 68: 371-398.

[38] Smith, Lones, and Peter Sørensen (2014), "Rational Social Learning with Random Sampling," *working paper.*

[39] Wagner, Peter (2017), "Who Goes First? Strategic Delay under Information Asymmetry," *Theoretical Economics*, forthcoming.

[40] Wolitzky, Alexander (2017), "Learning from Others' Outcomes," *working paper.*

9 Appendix B: Omitted Proofs and Examples (For Online Publication)

9.1 Properties of F

Non-empty: If (x_0, x_1) satisfies (4) (resp., the opposite of (4)) then $p(k; x_0, x_1)$ is increasing (resp., decreasing) in k. In either case, $p(k; x_0, x_1)$ is monotone in k, so there exists $k^* \in \{0, \ldots, K\}$ such that either $p(k^* - 1; x_0, x_1) \leq p^* \leq p(k^*; x_0, x_1)$ or $p(k^*; x_0, x_1) \geq p^* \geq p(k^* + 1; x_0, x_1)$, where in both cases the first (last) inequality is vacuous if $k^* = 0$ ($k^* = K$). With this value of k^*, let $s = 1$ if $p(k^*; x_0, x_1) \geq p^*$ and let $s = 0$ if $p(k^*; x_0, x_1) < p^*$. Next, with these values of k^* and s, let (x_0', x_1') be computed as in (6) or (8).

Compact-valued: Boundedness is trivial. For closedness, fix (x_0, x_1) and a sequence $(x_0'^n, x_1'^n) \to (x_0', x_1')$, with $(x_0'^n, x_1'^n) \in F(x_0, x_1)$, and let $(k^{*,n}, s^n)$ be arbitrarily chosen corresponding values of k^* and s. Taking a convergent subsequence $(k^{*,n}, s^n) \to (k^{*,\infty}, s^\infty)$, continuity of ϕ_θ implies that, with $k^* = k^{*,\infty}$ and $s = s^\infty$, (x_0', x_1') satisfies the conditions for inclusion in $F(x_0, x_1)$.

Convex-valued: Recall that there is at most one value of $k^* \in \{0, \ldots, K\}$ such that $p(k^*; x_0, x_1) = p^*$. So, if there are distinct elements of $F(x_0, x_1)$, (x_0', x_1') and (x_0'', x_1''), it must be that (x_0', x_1') and (x_0'', x_1'') are computed as in (6) or (8) with distinct values $s', s'' \in [0, 1]$. But then, for all $\beta \in [0, 1]$, letting $s = \beta s' + (1 - \beta) s''$, it follows that $(\beta x_0' + (1 - \beta) x_0'', \beta x_1' + (1 - \beta) x_1'') \in F(x_0, x_1)$.

Upper hemi-continuous: Fix sequences $(x_0^n, x_1^n) \to (x_0, x_1)$ and $(x_0'^n, x_1'^n) \to (x_0', x_1')$, with $(x_0'^n, x_1'^n) \in F(x_0^n, x_1^n)$, and let $(k^{*,n}, s^n)$ be arbitrarily chosen corresponding values of k^* and s. Taking a convergent subsequence $(k^{*,n}, s^n) \to (k^{*,\infty}, s^\infty)$, continuity of ϕ_θ implies that, with $k^* = k^{*,\infty}$ and $s = s^\infty$, (x_0', x_1') satisfies the conditions for inclusion in $F(x_0, x_1)$. ■

9.2 Equilibrium Uniqueness when $K = 1$

Proposition 11 *When $K = 1$, there is a unique equilibrium, and it is aligned. In this equilibrium, players adopt with probability 1 after observing a success and adopt with probability less than 1 after observing a failure.*

Proof. Fix an equilibrium, and suppose players adopt with probability s_1 after observing a success and adopt with probability s_0 after observing a failure. Then, for $\theta = 0, 1$,

$$x_\theta = [\chi + x_\theta (\pi_\theta - \chi)] s_1 + [1 - \chi - x_\theta (\pi_\theta - \chi)] s_0,$$

or

$$x_\theta = \frac{s_0 + \chi (s_1 - s_0)}{1 - (\pi_\theta - \chi)(s_1 - s_0)}. \tag{14}$$

Suppose toward a contradiction that $s_0 = 1$. As $s_0 = s_1 = 1$ would lead to $x_0 = x_1 = 1$, which is not an equilibrium by (2), this implies that $s_0 > s_1$. But $s_0 > s_1$ implies that $x_0 > x_1$, which contradicts Lemma 3. Hence, $s_0 < 1$.

Now, $s_0 < 1$ implies that $p(0; x_0, x_1) \leq p^*$. As $p > p^*$ and p is a convex combination of $p(0; x_0, x_1)$ and $p(1; x_0, x_1)$ (by the law of total probability), this implies that $p(1; x_0, x_1) > p^*$. Hence, $s_1 = 1$.

Next, using (14) and $s_1 = 1$,

$$p(0; x_0, x_1) = \left[1 + \frac{1 - \pi_0}{1 - \pi_1} \frac{1 - (\pi_1 - \chi)(1 - s_0)}{1 - (\pi_0 - \chi)(1 - s_0)} \frac{1 - p}{p}\right]^{-1}.$$

Therefore, $(s_0 = 0, s_1 = 1)$ corresponds to an equilibrium if and only if

$$\frac{1 - \pi_0}{1 - \pi_1} \frac{1 - \pi_1 + \chi}{1 - \pi_0 + \chi} \geq \frac{p}{1 - p} \frac{1 - p^*}{p^*}. \tag{15}$$

On the other hand, $(s_0 = s, s_1 = 1)$ with $s > 0$ corresponds to an equilibrium if and only if

$$\frac{1 - \pi_0}{1 - \pi_1} \frac{1 - (\pi_1 - \chi)(1 - s)}{1 - (\pi_0 - \chi)(1 - s)} = \frac{p}{1 - p} \frac{1 - p^*}{p^*}. \tag{16}$$

The left-hand side of (16) is increasing in s, and by (2) it exceeds the right-hand side when $s = 1$. Hence, by the intermediate value theorem, either there is a unique equilibrium given by $(s_0 = 0, s_1 = 1)$ and (14), or there exists a unique value $s > 0$ such that the unique equilibrium is given by $(s_0 = s, s_1 = 1)$, (16), and (14). ■

9.3 Examples of Misaligned Equilibria

Example 1: An Unstable Misaligned Equilibrium

Let $K = 2$, $\chi = 1$, $\pi_0 = 0$, $\pi_1 = \frac{1}{3}$, $p = \frac{1}{2}$, and $c = -\frac{8}{9}$. I claim that the misaligned point $\left(x_0 = 0, x_1 = \frac{3}{4}\right)$, together with the strategy of adopting if and only if at least one observation is a failure, is an equilibrium. (This point is misaligned because the success rate is 1 in state 0 and $1 - x_1(1 - \pi_1) = \frac{1}{2}$ in state 1.)

This follows because $p^* = \frac{\chi + c - \pi_0}{\pi_1 - \pi_0} = \frac{1}{3}$, while the posterior probability that $\theta = 1$ after observing at least one failure is $1 > p^*$, and the posterior probability that $\theta = 1$ after observing zero failures is

$$\left[1+\frac{1-p}{p}\frac{1}{\left(\frac{1}{2}\right)^2}\right]^{-1}=\frac{1}{5}<p^*.$$

The stated strategy is therefore optimal. In addition, (x_0, x_1) is an stationary point because the probability of observing at least one failure is 0 in state 0 and $1-\left(\frac{1}{2}\right)^2=\frac{3}{4}$ in state 1.

The equilibrium is however unstable, as the probability of observing at least one failure in state 0 when fraction x_0 adopts equals $1-(1-x_0)^2$, which is greater than x_0 for all $x_0 \in (0,1)$.

Example 2: A Stable Misaligned Equilibrium (and Two Stable Aligned Equilibria)

Let $K=3$, $\chi=\frac{9}{10}$, $\pi_0=0$, $\pi_1=\frac{1}{10}$, $p=\frac{1}{2}$, and $c=-\frac{1701}{2000}$.[27] Under the strategy of adopting if and only if at least two observations are failures, the equation for x_θ to be a stationary point is

$$x_\theta=\left(1-\chi+x_\theta\left(\chi-\pi_\theta\right)\right)^3+3\left(1-\chi+x_\theta\left(\chi-\pi_\theta\right)\right)^2\left(\chi-x_\theta\left(\chi-\pi_\theta\right)\right).$$

Consider the point (x_0, x_1) given by taking the smallest solution to this cubic equation for $\theta=0$ and the largest solution for $\theta=1$: $(x_0,x_1)\approx(.07407,.9419)$. This point is misaligned because the success rate is $\chi-x_1\left(\chi-\pi_1\right)\approx 0.1465$ in state 1 and $\chi-x_0\left(\chi-\pi_0\right)\approx 0.8333$ in state 0. It is straightforward to check that this point is stable: for $\theta=0,1$, the above cubic equation has three roots, of which the middle one is unstable. Finally, to see that the proposed strategy is optimal, note that $p^*=\frac{\chi+c-\pi_0}{\pi_1-\pi_0}=0.495$, while the posterior probability that $\theta=1$ after observing two failures is

$$\left[1+\frac{1-p}{p}\frac{\left(1-\chi+x_0\left(\chi-\pi_0\right)\right)^2\left(\chi-x_0\left(\chi-\pi_0\right)\right)}{\left(1-\chi+x_1\left(\chi-\pi_1\right)\right)^2\left(\chi-x_1\left(\chi-\pi_1\right)\right)}\right]^{-1}\approx 0.8217>p^*,$$

while the posterior probability that $\theta=1$ after observing one failure is

$$\left[1+\frac{1-p}{p}\frac{\left(1-\chi+x_0\left(\chi-\pi_0\right)\right)\left(\chi-x_0\left(\chi-\pi_0\right)\right)^2}{\left(1-\chi+x_1\left(\chi-\pi_1\right)\right)\left(\chi-x_1\left(\chi-\pi_1\right)\right)^2}\right]^{-1}\approx 0.1366<p^*.$$

In fact, it is not hard to see that a stable misaligned equilibrium cannot exist when $K=2$, so $K=3$ is the minimum sample size for which a stable misaligned equilibrium can exist.

[27] The explanation for this oddly precise choice of c is that, if $c=-\frac{17}{20}$, the analysis of the example would be exactly the same except that one would have $p^*=p$, which violates (1). As the only role of c in the model is to determine p^*, it suffices to let $c=-\frac{17}{20}-\varepsilon$ for any sufficiently small $\varepsilon>0$.

The reason is that, when $K = 2$, the fraction of players observing at least k failures in state θ is at most quadratic in x_θ, so there is a unique stable misaligned stationary point (x_0, x_1). But, since failure is more likely for a given fraction of adopters in state 0, this unique stable point always has $x_0 > x_1$, so by Lemma 3 it cannot be an equilibrium.

This same example also admits two stable aligned equilibria. Thus, there can be multiple stable aligned equilibria, and they can coexist with a stable misaligned equilibrium.

Specifically, I claim that a point $(x_0', x_1') \approx (0.4681, 0.5)$, together with the strategy of adopting if and only if at least two successes are observed, is an equilibrium; and that so is a point $(x_0'', x_1'') \approx (0.6625, 0.7061)$, together with the strategy of adopting if and only if at least one success is observed. The intuition for this multiplicity is that, when the "bar" for adopting is raised from one observed success to two, this reduces the steady-state adoption rate, which makes failure less likely in both states (as $\chi > \pi_0, \pi_1$), and thus makes failure more informative. This in turn justifies the greater number of observed successes required for adoption.

For the formal construction, note that, under the strategy of adopting if and only if at least two successes are observed, the equation for x_θ to be a stationary point is

$$x_\theta = (\chi - x_\theta (\chi - \pi_\theta))^3 + 3 (\chi - x_\theta (\chi - \pi_\theta))^2 (1 - \chi + x_\theta (\chi - \pi_\theta)).$$

Let (x_0', x_1') be the unique solutions to this equation for $\theta = 0, 1$, given by $(x_0', x_1') \approx (0.4681, 0.5)$. Then the posterior probability that $\theta = 1$ after observing two successes is

$$\left[1 + \frac{1-p}{p} \frac{(1 - \chi + x_0' (\chi - \pi_0)) (\chi - x_0' (\chi - \pi_0))^2}{(1 - \chi + x_1' (\chi - \pi_1)) (\chi - x_1' (\chi - \pi_1))^2}\right]^{-1} \approx 0.5113 > p^*,$$

while the posterior probability that $\theta = 1$ after observing one success is

$$\left[1 + \frac{1-p}{p} \frac{(1 - \chi + x_0' (\chi - \pi_0))^2 (\chi - x_0' (\chi - \pi_0))}{(1 - \chi + x_1' (\chi - \pi_1))^2 (\chi - x_1' (\chi - \pi_1))}\right]^{-1} \approx 0.4900 < p^*.$$

So this is an equilibrium. It is also easily seen to be stable, as the curve $(\chi - x_\theta (\chi - \pi_\theta))^3 + 3 (\chi - x_\theta (\chi - \pi_\theta))^2 (1 - \chi + x_\theta (\chi - \pi_\theta))$ crosses x_θ from above, for $\theta = 0, 1$.

Similarly, under the strategy of adopting if and only if at least one successes is observed, the equation for x_θ to be a stationary point is given by

$$x_\theta = 1 - (1 - \chi + x_\theta (\chi - \pi_\theta))^3.$$

Let (x_0'', x_1'') be the unique solutions, given by $(x_0'', x_1'') \approx (0.6625, 0.7061)$. Then the posterior probability that $\theta = 1$ after observing one successes is

$$\left[1 + \frac{1-p}{p} \frac{\left(1-\chi+x_0''\left(\chi-\pi_0\right)\right)^2\left(\chi-x_0''\left(\chi-\pi_0\right)\right)}{\left(1-\chi+x_1''\left(\chi-\pi_1\right)\right)^2\left(\chi-x_1''\left(\chi-\pi_1\right)\right)}\right]^{-1} \approx 0.5015 > p^*,$$

while the posterior probability that $\theta = 1$ after observing zero successes is

$$\left[1 + \frac{1-p}{p} \frac{\left(1-\chi+x_0''\left(\chi-\pi_0\right)\right)^3}{\left(1-\chi+x_1''\left(\chi-\pi_1\right)\right)^3}\right]^{-1} \approx 0.4655 < p^*.$$

So this is also an equilibrium, and it is also easily seen to be stable.

Finally if one considers this example with $K = 2$ rather than $K = 3$, one finds that there is a unique stable aligned equilibrium $(x_0, x_1) \approx (0.5955, 0.6327)$ (corresponding to the strategy of adopting if and only if at least one success if observed), and welfare in this steady state lies in between that in the two stable aligned steady states that arise when $K = 3$. This shows that welfare does not always unambiguously increase when players observe larger samples, even within the class of stable aligned equilibria.

9.4 Proof of Proposition 3

1. If $\chi = 0$, then (15) is violated (as $\frac{p}{1-p}\frac{1-p^*}{p^*} > 1$), so the unique equilibrium is given by (16) and (14). Solving for x_0, x_1, and s gives

$$\begin{aligned}
x_0 &= \frac{\left(p-p^*\right)\left(1-\pi_1\right)}{p^*\left(1-p\right)\left(\pi_1-\pi_0\right)}, \\
x_1 &= \frac{\left(p-p^*\right)\left(1-\pi_0\right)}{p\left(1-p^*\right)\left(\pi_1-\pi_0\right)}, \text{ and} \\
s &= \frac{\left(p-p^*\right)\left(1-\pi_0\right)\left(1-\pi_1\right)}{p^*\left(1-p\right)\left(1-\pi_0\right)\pi_1 - p\left(1-p^*\right)\pi_0\left(1-\pi_1\right)}.
\end{aligned}$$

Noting that $p^* \to \hat{p}$ as $\pi_1 \to 1$, it follows that $x_0 \to 0$ and $x_1 \to \frac{p-\hat{p}}{p(1-\hat{p})}$ as $\pi_1 \to 1$.

2. If $\chi > 0$, then (15) holds when π_1 is close enough to 1. In this case, (14) gives $x_\theta = \frac{\chi}{1-\pi_\theta+\chi}$ for $\theta = 0, 1$. Hence, $x_0 = \frac{\chi}{1-\pi_0+\chi}$ and $x_1 \to 1$ as $\pi_1 \to 1$. ■

9.5 Proof of Proposition 4

1. Fix a sequence of parameters $(\pi_0^n, \pi_1^n) \to (\pi_0, \pi_1) = (0, 1)$ and fix a corresponding sequence of equilibria $(x_0^n, x_1^n, k^{*,n}, s^n) \to (x_0, x_1, k^*, s)$. Note that (x_0, x_1, k^*, s) must be an equilibrium. Suppose toward a contradiction that $x_1 < 1$. By Lemma 3, $x_1 \geq \frac{p-p^*}{1-p^*}$, so

$x_1\pi_1 \in (0,1)$. On the other hand, $x_0\pi_0 = 0$. Therefore, $p(k; x_0, x_1) > p^*$ for all $k \geq 1$, and hence $k^* = 0$. The steady state equation then implies that

$$x_1 = 1 - (1 - x_1)^K (1 - s) \geq 1 - (1 - x_1)^2 = x_1 (2 - x_1).$$

But this implies that $x_1 = 1$, a contradiction.

To show that $x_0 = 0$, let $\hat{s}$ be the probability with which players adopt after observing K failures in the equilibrium (x_0, x_1, k^*, s). (Thus, $\hat{s} = 0$ if $k^* > 0$, and $\hat{s} = s$ if $k^* = 0$.) Then the steady state equation implies that $x_0 = \hat{s}$. Next, note that $p(0; x_0, 1) = p(0; 1, 1) < p^*$ (by $\pi_0 = 0$ and (2)). Hence, $\hat{s} = 0$.

2. Fix a sequence of parameters $(\pi_0^n, \pi_1^n) \to (\pi_0, \pi_1) = (0, 1)$ and a corresponding sequence of aligned equilibria $(x_0^n, x_1^n, k^{*,n}, s^n) \to (x_0, x_1, k^*, s)$. Note that (x_0, x_1, k^*, s) must be an aligned equilibrium. I claim that $x_0 > 0$. To see this, note that, in any aligned equilibrium $p(K; x_0, x_1) > p > p^*$, and therefore players adopt with probability 1 after observing K successes. Thus, if $x_0 = 0$ and $\chi = 1$, then in state 0 players would observe all successes with probability 1; and therefore x_0 would equal 1, a contradiction.

Next, as $x_0 > 0$, $\pi_0 = 0$, and $\chi + x_1 (\pi_1 - \chi) = 1$, $p(k; x_0, x_1) = 0$ for all $k < K$, so players adopt with probability 0 after observing even a single failure. On the other hand, I have shown that players adopt with probability 1 after observing all successes, so

$$x_\theta = (1 - x_\theta (1 - \pi_\theta))^K \text{ for } \theta = 0, 1.$$

As $\pi_0 = 0$ and $\pi_1 = 1$, this implies that $x_0 = (1 - x_0)^K$ and $x_1 = 1$.

The last part of the proposition follows as the solution to the equation $x_0 = (1 - x_0)^K$ converges to 0 as $K \to \infty$. ■

9.6 Proof of Proposition 5

Given adoption rates (x_0, x_1), the posterior belief that $\theta = 1$ after observing failure equals

$$\left[1 + \frac{1-p}{p} \frac{1 - \chi - x_0 (\pi_0 - \chi)}{1 - \chi - x_1 (\pi_1 - \chi)}\right]^{-1}.$$

This posterior equals p^* if and only if

$$\frac{1 - \chi - x_0 (\pi_0 - \chi)}{1 - \chi - x_1 (\pi_1 - \chi)} = \frac{p}{1-p} \frac{1 - p^*}{p^*}.$$

This equation defines a line $\hat{L}$ in (x_0, x_1) space. Let H be half-space where the posterior exceeds p^* and let H^c be the half-space where the posterior is less than p^*; thus $\hat{L}$ marks the boundary between H and H^c. Recall from the proof of Proposition 11 that there are two possible cases: either the equilibrium is $\left(x_0 = \frac{\chi}{1-\pi_0+\chi}, x_1 = \frac{\chi}{1-\pi_1+\chi}\right)$ and this point lies in the half-space H^c, or the equilibrium lies on the line $\hat{L}$.

At any point $(x_0, x_1) \in H$, it follows that $\dot{x}_\theta = 1 - x_\theta$ for $\theta = 0, 1$, so the vector $(\dot{x}_0, \dot{x}_1)$ points from (x_0, x_1) toward the point $(1, 1)$. By (2), the point $(1, 1)$ lies in the complementary half-space H^c. Hence, if the initial point $(x_0(0), x_1(0))$ lies in H, the distance between $(x_0(t), x_1(t))$ and the line $\hat{L}$ is decreasing in t and reaches 0 in finite time.

Similarly, if $(x_0, x_1) \in H^c$, then $\dot{x}_\theta = \chi - x_\theta(1 - \pi_\theta + \chi)$ for $\theta = 0, 1$. Hence, $(x_0(t), x_1(t))$ converges monotonically toward the point $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$, so long as $(x_0(t), x_1(t))$ remains in H^c. Thus, if the equilibrium is $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$ then the population dynamic converges monotonically to the equilibrium starting from any point in H^c, and otherwise the population dynamic converges monotonically toward the point $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$ until it hits the line $\hat{L}$ (which again occurs in finite time).

Next, if $(x_0(t), x_1(t)) \in \hat{L}$ then $\dot{x}_\theta \geq \chi - x_\theta(1 - \pi_\theta + \chi)$ for $\theta = 0, 1$. Hence, if the equilibrium is $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$, then the population dynamic converges toward this point from any point in $\hat{L}$. Combining the observations made so far, it follows that when the equilibrium is $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right)$, it is globally attracting.

Finally, if $(x_0(t), x_1(t)) \in \hat{L}$ and $(x_0^*, x_1^*) \in \hat{L}$, then the population dynamic remains in $\hat{L}$ forever: this follows because, as I have shown, the gradient $(\dot{x}_0, \dot{x}_1)$ points toward $\hat{L}$ whenever $(x_0, x_1) \notin \hat{L}$. Next, for any point $(x_0, x_1) \in \hat{L}$, there is a unique mixing probability conditional on observing failure, $s((x_0, x_1))$, such that the gradient $(\dot{x}_0, \dot{x}_1)$ is parallel to $\hat{L}$, and in addition the mixing probability $s((x_0, x_1))$ is itself continuous in (x_0, x_1).[28] As the vector $(\dot{x}_0, \dot{x}_1)$ is continuous in (x_0, x_1) and the mixing probability s, it may therefore also be viewed as a continuous function of (x_0, x_1). Furthermore, as any stationary point in $\hat{L}$ is an equilibrium, (x_0^*, x_1^*) is the unique point in $\hat{L}$ such that $(\dot{x}_0, \dot{x}_1) = (0, 0)$. Hence, as $(\dot{x}_0, \dot{x}_1)$ is continuous in (x_0, x_1), it must be that $(\dot{x}_0, \dot{x}_1)$ points toward the steady state, and in addition $(\dot{x}_0(t), \dot{x}_1(t))$ can converge to 0 only if $(x_0(t), x_1(t))$ converges to (x_0^*, x_1^*). Therefore, $(x_0(t), x_1(t))$ must converge to (x_0^*, x_1^*) starting from any initial point in $\hat{L}$. As

[28] To see this, note that if $s = 0$, then $(\dot{x}_0, \dot{x}_1) = (\chi - x_0(1 - \pi_0 + \chi), \chi - x_1(1 - \pi_1 + \chi))$, which points into H (when $\left(\frac{\chi}{1+\chi-\pi_0}, \frac{\chi}{1+\chi-\pi_1}\right) \in H$, or equivalently when the equilibrium is in $\hat{L}$), and if $s = 1$ then $(\dot{x}_0, \dot{x}_1) = (1 - x_0, 1 - x_1)$, which points into H^c. Denote these vectors by $(\dot{x}_0^0, \dot{x}_1^0)$ and $(\dot{x}_0^1, \dot{x}_1^1)$, and let $\dot{x}_\theta^s = (1 - s)\dot{x}_\theta^0 + s\dot{x}_\theta^1$ for $\theta = 0, 1$. By the intermediate value theorem, there is a unique mixing probability $s((x_0, x_1))$ such that $\dot{x}^{s((x_0,x_1))}$ is parallel to $\hat{L}$, and $s((x_0, x_1))$ is continuous in (x_0, x_1) because $(\dot{x}_0^0, \dot{x}_1^0)$ and $(\dot{x}_0^1, \dot{x}_1^1)$ are continuous in (x_0, x_1).

I have shown that $(x_0(t), x_1(t))$ reaches $\hat{L}$ in finite time starting from any initial point in $[0,1]^2$, it follows that (x_0^*, x_1^*) is globally attracting. ■

9.7 Proof of Proposition 6

As $(x_0(0), x_1(0))$ is aligned, Theorem 1 implies that $(x_0(t), x_1(t))$ is aligned for all t. Hence, players adopt with probability 1 after observing a success. On the other hand, a player's posterior after observing a failure at time t is given by

$$p(0; x_0(t), x_1(t)) = \left[1 + \frac{1-p}{p}\frac{1}{1 - x_1(t)\pi_1}\right]^{-1}.$$

This posterior is less than p^* at time 0 by (2), and it remains less than p^* until $x_1(t)$ reaches the value

$$x_1^* = \frac{1}{\pi_1}\left(1 - \frac{1-p}{p}\frac{p^*}{1-p^*}\right) < 1.$$

(Note that this equation defines the line $\hat{L}$ introduced in the proof of Proposition 5.) Letting T be the first time when $x_1(t)$ reaches x_1^*, it follows that $\dot{x}_\theta(t) = -x_\theta(t)(1-\pi_\theta)$ for all $t < T$ and $\theta = 0, 1$. Combined with the initial condition $(x_0(0), x_1(0)) = (1, 1)$, this gives $x_\theta(t) = \exp(-(1-\pi_\theta)t)$ for $\theta = 0, 1$.

Next, as shown in the proof of Proposition 5, $(x_0(t), x_1(t))$ remains on the line $\hat{L}$ for all $t > T$: that is, $x_1(t) = x_1^*$ for all $t > T$. It follows that $s(t) = s$ for all $t > T$, where s is given by $x_1^* = x_1^*\pi_1 + (1 - x_1^*\pi_1)s$, or

$$s = \frac{1-\pi_1}{\pi_1}\left(\frac{p}{1-p}\frac{1-p^*}{p^*} - 1\right).$$

In addition, for $t > T$, $\dot{x}_0(t) = s - x_0(t)$, so $x_0(t)$ converges monotonically to its steady-state value of s.

Finally, the time T satisfies

$$T = \frac{1}{1-\pi_1}\left[\log \pi_1 - \log\left(1 - \frac{1-p}{p}\frac{p^*}{1-p^*}\right)\right].$$

Hence,

$$x_0(T) = \exp(-T) = \left(\frac{1}{\pi_1}\left(1 - \frac{1-p}{p}\frac{p^*}{1-p^*}\right)\right)^{\frac{1}{1-\pi_1}}.$$

In particular, $x_0(T) < s$ if and only if

$$1 - \frac{1-p}{p}\frac{p^*}{1-p^*} < \pi_1 \left(\frac{1-\pi_1}{\pi_1}\right)^{1-\pi_1} \left(\frac{p}{1-p}\frac{1-p^*}{p^*} - 1\right)^{1-\pi_1}.$$

The right-hand side of this inequality goes to 1 as $\pi_1 \to 1$, so $x_0(T) < s$ whenever π_1 is close enough to 1. ■

9.8 Proof of Proposition 7

Fix $\varepsilon \in (0, (\chi - \pi_{\max \Theta^*}) / (1 + \chi - \pi_{\max \Theta^*}))$. Suppose an asymptotically efficient path exists. Then there exists $\bar{K} > 0$ such that if $K > \bar{K}$ then $\left(X_0^K(0), \ldots, X_n^K(0)\right) = (1, \ldots, 1)$ and $\lim_{t\to 0} X_\theta^K(t) < \varepsilon$ (resp., $> 1-\varepsilon$) for all $\theta \le \theta^*$ (resp., $> \theta^*$). For such a K, the success rate at $t = 0$ conditional on the event $\theta \le \theta^*$ equals $(1/a)\sum_{\theta=0}^{\theta^*} p_\theta \pi_\theta$ and the success rate at $t = 0$ conditional on the event $\theta \in \Theta^*$ equals $(1/b)\sum_{\theta\in\Theta^*} p_\theta \pi_\theta$, which is larger. On the other hand, as $t \to \infty$ the success rate conditional on the event $\theta \le \theta^*$ converges to a number greater than $(1-\varepsilon)\chi$, while the success rate conditional on the event $\theta \in \Theta^*$ converges to a number less than $\varepsilon + (1-\varepsilon)\pi_{\max \Theta^*}$, which is smaller. Hence, there must exist a time t^* such that (i) at $t = t^*$, the success rate conditional on the event $\theta \le \theta^*$ equals the success rate conditional on the event $\theta \in \Theta^*$, and (ii) for all $t > t^*$, the success rate conditional on the event $\theta \le \theta^*$ is larger than the success rate conditional on the event $\theta \in \Theta^*$.

Now, at $t = t^*$, after observing any sample a player's relative assessment of the probability of the events $\theta \le \theta^*$ and $\theta \in \Theta^*$ equals the prior probability $a/(a+b)$. Thus, (13) implies that (after observing any sample at $t = t^*$) action 1 is optimal conditional on the event $\theta \in \{1, \ldots, \theta^*\} \cup \Theta^*$. In addition, action 1 is optimal at any state $\theta \notin (\{1, \ldots, \theta^*\} \cup \Theta^*)$. Hence, $\dot{X}_\theta(t^*) = 1 - X_\theta(t^*)$ for all θ. Therefore,

$$\begin{aligned}
\frac{1}{a}\sum_{\theta=0}^{\theta^*} p_\theta \dot{X}_\theta(t^*)(\pi_\theta - \chi) &= \left[\frac{1}{a}\sum_{\theta=0}^{\theta^*} p_\theta(\pi_\theta - \chi)\right] - \left[\frac{1}{a}\sum_{\theta=0}^{\theta^*} p_\theta X_\theta(t^*)(\pi_\theta - \chi)\right] \\
&= \left[\frac{1}{a}\sum_{\theta=0}^{\theta^*} p_\theta(\pi_\theta - \chi)\right] - \left[\frac{1}{b}\sum_{\theta\in\Theta^*} p_\theta X_\theta(t^*)(\pi_\theta - \chi)\right] \\
&< \left[\frac{1}{b}\sum_{\theta\in\Theta^*} p_\theta(\pi_\theta - \chi)\right] - \left[\frac{1}{b}\sum_{\theta\in\Theta^*} p_\theta X_\theta(t^*)(\pi_\theta - \chi)\right] \\
&= \frac{1}{b}\sum_{\theta\in\Theta^*} p_\theta \dot{X}_\theta(t^*)(\pi_\theta - \chi).
\end{aligned}$$

But this implies that, just after time t^*, the success rate conditional on the event $\theta \le \theta^*$ is smaller than the success rate conditional on the event $\theta \in \Theta^*$, a contradiction. ■

9.9 Proof of Proposition 8

Suppose $\pi_\theta < \chi$ for some innovation-optimal state θ, and suppose a simple asymptotically efficient path exists. As in the proof of Proposition 7, for large enough K, at $t = 0$ the success rate in each innovation-optimal state is greater than the success rate in each status-quo optimal state, and the situation is reversed for large enough t. Hence, there must exist a time t^* at which the success rates in a status-quo optimal state and an innovation-optimal state cross for the first time: that is, a time t^* such that (i) $X_\theta(t^*)(\pi_\theta - \chi) \leq X_{\theta'}(t^*)(\pi_{\theta'} - \chi)$ for all $\theta \leq \theta^* < \theta'$, and (ii) there exists $\varepsilon > 0$ and $\theta \leq \theta^* < \theta'$ such that $X_\theta(t)(\pi_\theta - \chi) > X_{\theta'}(t)(\pi_{\theta'} - \chi)$ for all $t \in (t^*, t^* + \varepsilon)$.

The proof is completed by considering separately the case where $X_\theta(t^*)(\pi_\theta - \chi) = X_{\theta'}(t^*)(\pi_{\theta'} - \chi)$ for all θ, θ' and the case where $X_\theta(t^*)(\pi_\theta - \chi) < X_{\theta'}(t^*)(\pi_{\theta'} - \chi)$ for some θ, θ', and deriving a contradiction in each.

In the first case, the success rate is equal in all states at time t^*, and hence $\dot{X}_\theta(t^*) = 1 - X_\theta(t^*)$ for all θ. But, as in the proofs of Theorem 1 and Proposition 7, this implies that there cannot be a pair of states $\theta < \theta'$ with $X_\theta(t^*)(\pi_\theta - \chi) = X_{\theta'}(t^*)(\pi_{\theta'} - \chi)$ and $X_\theta(t)(\pi_\theta - \chi) > X_{\theta'}(t)(\pi_{\theta'} - \chi)$ for all $t \in (t^*, t^* + \varepsilon)$, a contradiction.

In the second case, there are three states with either (i) $\theta_0 < \theta \leq \theta^* < \theta'$ and

$$X_{\theta_0}(t^*)(\pi_{\theta_0} - \chi) < X_\theta(t^*)(\pi_\theta - \chi) = X_{\theta'}(t^*)(\pi_{\theta'} - \chi)$$

or (ii) $\theta \leq \theta^* < \theta' < \theta_0$ and

$$X_\theta(t^*)(\pi_\theta - \chi) = X_{\theta'}(t^*)(\pi_{\theta'} - \chi) < X_{\theta_0}(t^*)(\pi_{\theta_0} - \chi).$$

Consider the first case (the second is symmetric). Then, as $X_\theta(t)$ is continuous for all θ, for sufficiently small $\varepsilon > 0$,

$$X_{\theta_0}(t^* + \varepsilon)(\pi_{\theta_0} - \chi) < X_{\theta'}(t^* + \varepsilon)(\pi_{\theta'} - \chi) < X_\theta(t^* + \varepsilon)(\pi_\theta - \chi).$$

But then the path is not simple. ■

9.10 Proof of Proposition 9

By assumption, $X_0(0)(\pi_0 - \chi) \leq X_1(0)(\pi_1 - \chi)$. As X_0 and X_1 are continuous, if there exists a time t' with $X_0(t')(\pi_0 - \chi) > X_1(t')(\pi_1 - \chi)$, then there must exist another time

t where $X_0(t)(\pi_0 - \chi) = X_1(t)(\pi_1 - \chi)$ but it is not the case that $\dot{X}_0(t)(\pi_0 - \chi) < \dot{X}_1(t)(\pi_1 - \chi)$. By Theorem 2, a misaligned equilibrium cannot exist in the outcome-improving case, so $\pi_1 < \chi$. Hence, $X_0(t)(\pi_0 - \chi) = X_1(t)(\pi_1 - \chi)$ implies $X_0(t) < X_1(t)$. But, by definition of $\dot{X}_\theta(t)$, if $X_0(t)(\pi_0 - \chi) = X_1(t)(\pi_1 - \chi)$ and $X_0(t) < X_1(t)$, then $\dot{X}_0(t) > \dot{X}_1(t)$, and hence $\dot{X}_0(t)(\pi_0 - \chi) < \dot{X}_1(t)(\pi_1 - \chi)$. So there can be no such time t'. ■

9.11 Proof of Proposition 10

It follows immediately from the definition of $\dot{X}_\theta(t)$ and stability from above that $\dot{X}_\theta(t)$ is bounded below 0 for all t such that $X_\theta(t)$ is bounded above x_θ^*. It is also straightforward to argue by contradiction that $X_\theta(t)$ can never cross x_θ^*, completing the proof. ■