

ENDOGENOUS EXPERIMENTATION IN ORGANIZATIONS*

Germán Gieczewski[†] Svetlana Kosterina[‡]

June 2020

Abstract

We study policy experimentation in organizations with endogenous membership. An organization initiates a policy experiment and then decides when to stop it based on its results. As information arrives, agents update their beliefs, and become more pessimistic whenever they observe bad outcomes. At the same time, the organization's membership adjusts endogenously: unsuccessful experiments drive out conservative members, leaving the organization with a radical median voter. We show that there are conditions under which the latter effect dominates. As a result, policy experiments, once begun, continue for too long. In fact, the organization may experiment forever in the face of mounting negative evidence. This result provides a rationale for obstinate behavior by organizations, and contrasts with models of collective experimentation with fixed membership, in which under-experimentation is the typical outcome.

JEL codes: D71, D83

Keywords: experimentation, dynamics, median voter, endogenous population

Word count: 11607

*We would like to thank Daron Acemoglu, Avidit Acharya, Alessandro Bonatti, Glenn Ellison, Robert Gibbons, Faruk Gul, Matias Iaryczower, Navin Kartik, Wolfgang Pesendorfer, Nicola Persico, Andrea Prat, Caroline Thomas, Juuso Toikka, the audiences at the 2018 American Political Science Association Conference (Boston), the 2018 Columbia Political Economy Conference, the 2019 Canadian Economic Theory Conference (Montreal) and the 2020 Utah Winter Organizational and Political Economics Conference, as well as seminar participants at MIT, Princeton, Indiana University Kelley School of Business, University of Calgary and UCEMA for helpful comments. All remaining errors are our own.

[†]Department of Politics, Princeton University. Address: 311 Fisher Hall, Princeton NJ 08540. Phone number: 6172330642. E-mail: germang@princeton.edu.

[‡]Department of Politics, Princeton University. Address: 300 Fisher Hall, Princeton NJ 08540. Phone number: 9179756491. Email: svetlana.kosterina@princeton.edu.

1 Introduction

Organizations frequently face opportunities to experiment with promising but untested policies. Two conclusions follow from our understanding of experimentation to date. First, experimentation should respond to information. That is, when a policy experiment performs badly, agents should become more pessimistic about it, and if enough negative information accumulates, they should abandon it. Second, when experimentation is collective, the temptation to free-ride and fears that information will be misused by other agents lower incentives to experiment.¹ Thus organizations should experiment too little. However, history is littered with examples of organizations that have stubbornly persisted with unsuccessful policies to the bitter end. This presents a puzzle for our understanding of decision-making in organizations.

Consider, for example, the case of Theranos, a Silicon Valley start-up founded by Elizabeth Holmes in 2003. Theranos sought to produce a portable machine capable of running hundreds of medical tests on a single drop of blood. If successful, Theranos would have revolutionized medicine, but its vision was exceedingly difficult to realize. Over the course of ten years, the firm invested over a hundred million dollars into trying to attain Holmes’s vision, while devoting little effort to developing a more incremental improvement over existing technologies as a fall-back plan. Theranos eventually launched in 2013 with a mixture of inaccurate and fraudulent tests, and the ensuing scandal irreversibly damaged the company.

Up to the point where Theranos began to engage in outright fraud, a pattern repeated itself. The company would bring in high-profile hires and create enthusiasm with its promises, but once inside the organization, employees and board members would gradually become disillusioned by the lack of progress.² As a result, many left the company,³ with those who were more pessimistic about Theranos’s prospects being more likely to leave than those who saw Holmes as a visionary. While the board came close to removing Holmes as CEO early on, she managed to retain control for many years after, because too many of the people who had lost faith in her leadership

¹See, for example, Keller, Rady and Cripps (2005) and Strulovici (2010).

²For instance, Theranos’s lead scientist, Ian Gibbons, told his wife that nothing at Theranos was working, years after joining the company. See Carreyrou (2018).

³For example, board member Avie Tevanian, while mulling over a decision to buy more shares of the company at a low price, was asked by a friend: “Given everything you now know about this company, do you really want to own more of it?” See Carreyrou (2018).

had left the organization before they could have a majority.

The takeaway from this example is that Theranos experimented for too long in spite of not succeeding precisely because the pessimistic members of the organization kept leaving. Motivated by this and similar examples, we propose an explanation for obstinate behavior by organizations that rests on three key premises. First, agents disagree about the merits of different policies, that is, they have heterogeneous prior beliefs. Second, the membership of organizations is fluid: agents are free to enter and leave in response to information. Third, the organization’s policies are responsive to the opinions of its members.

We operationalize these assumptions in the following model. An organization chooses between a safe policy and a risky policy in each period. The safe policy yields a flow payoff known by everyone, while the risky policy yields an uncertain flow payoff, which may be higher or lower than that of the safe policy depending on its type. There is a continuum of agents. In every period, each agent decides whether to work for the organization or independently. The payoff of working for the organization depends on its policy. The outside option yields a guaranteed flow payoff.

All agents want to maximize their returns but hold heterogeneous prior beliefs about the quality of the risky policy. As long as agents work for the organization, they remain voting members of the organization and vote on the policy it pursues. We assume that the median voter—that is, the member with the median prior belief—chooses the organization’s policy. Whenever the risky policy is used, the results are publicly observed.

We show that experimentation in organizations is inefficient in two ways that are novel to the literature. First, there is over-experimentation from the point of view of all agents. Over-experimentation takes a particularly stark form: the organization experiments forever regardless of the outcome. Second, the organization may experiment more if the risky policy is bad than if it is good. In other words, the organization’s policy may respond to information in a perverse way.

Our main result provides a simple necessary and sufficient condition under which perpetual experimentation is the unique equilibrium outcome. The condition requires that, at each history, the pivotal agent prefers perpetual experimentation to no ex-

perimentation. We then analyze the comparative statics. We show that perpetual experimentation is more likely when the outside option is more attractive, the organization’s safe policy is less attractive, agents are more patient, and the distribution of prior beliefs contains more optimists in the MLRP sense.

Two forces affect the amount of experimentation in our model. On the one hand, the median member of an organization is reluctant to experiment today if she anticipates losing control of the organization tomorrow as a result. On the other hand, if no successes are observed, as time passes, only the most optimistic members remain in the organization, and these are precisely the members who want to continue experimenting the most. The first force makes under-experimentation more likely, while the second pushes the organization to over-experiment. Ex-ante, it is not obvious which force will dominate. We show that the second force often dominates. This underpins the main result of our paper.

We next show that our result of perpetual experimentation is robust. Our baseline model features perfectly informative good news. We show that perpetual experimentation also obtains under bad news and imperfectly informative good news. In addition, a novel result arises when news are imperfectly informative: for appropriately chosen parameter values, there is an equilibrium in which the organization stops experimenting with a strictly positive probability *only if enough successes are observed*. Hence, counterintuitively, the organization is more likely to experiment forever if the technology is bad than if it is good.⁴ The implication is that self-selection of agents into organizations may not only induce excessive experimentation overall—it may also cause organizations to actively radicalize in the face of failure. Conversely, success may make organizations more conservative and prone to backing away from the very strategies that brought them success.

Finally, we show that our result is robust to general voting rules; settings in which the members’ flow payoffs, or the learning rate, depend on the size of the organization; an alternative specification in which agents have common values; and a specification in which the size of the organization is fixed rather than variable and the organization hires agents based on ability.

The rest of the paper proceeds as follows. Section 2 discusses the applications of

⁴Approximately, for this to happen, it is sufficient for the distribution of prior beliefs to be single-dipped enough over some interval.

the model. Section 3 reviews the related literature. Section 4 introduces the baseline model. Section 5 analyzes the set of the equilibria in the baseline model. Section 6 considers other learning processes. Section 7 develops other extensions of the model, such as general voting rules and settings where the members' flow payoffs depend on the size of the organization.

2 Applications

Our model has a variety of applications besides the one discussed in the Introduction. In this Section, we discuss how our assumptions map to several settings such as cooperatives, non-profits, activist organizations, firms and political parties.

We first consider experimentation in a cooperative. Agents are individual producers who own factors of production. In a dairy cooperative, for example, each member owns a cow. The agent can manufacture and sell his own dairy products independently or he can join the cooperative. If he joins, his milk will be processed at the cooperative's plants, which benefit from economies of scale. The cooperative can choose from a range of dairy production policies, some of which are riskier than others. For instance, it can limit itself to selling fresh milk and yogurt, or it can develop a new line of premium cheeses that may or may not become profitable. Dairy farmers have different beliefs about the market viability of the latter strategy. Should this strategy be used, only the more optimistic farmers will choose to join or remain in the cooperative. Moreover, cooperatives typically allow their members to elect directors.

In the case of activist organizations, agents are citizens seeking to change the government's policy or the behavior of multinational corporations. Agents with environmental concerns can act independently by writing to their elected representatives, or they can join an organization, such as Greenpeace, that has access to strategies not available to a citizen acting alone, such as lobbying, demonstrations, or direct action—for instance, confronting whaling ships. While all members of the organization want to bring about a policy change, their beliefs as to the best means of achieving this goal differ. Some support safe strategies, such as lobbying, while others prefer riskier ones, such as direct action.

An organization that employs direct action will drive away its moderate members, increasingly so if its attempts are unsuccessful. The resulting self-selection can sustain a base of support for extremist strategies. Our model can thus explain the behavior of fringe environmental groups, such as Extinction Rebellion, Animal Liberation Front and Earth Liberation Front that engage in ecoterrorism and economic sabotage in spite of the apparent ineffectiveness of their approach.⁵ The same logic applies to other forms of activism, as well as to charitable organizations choosing between different poverty alleviation tactics, for example, cash transfers as opposed to microcredit.⁶

Our model is also relevant to the functioning of political parties. Here agents are potential politicians or active party members, and the party faces a choice between a widely understood mainstream platform—for example, social democracy—and an extremist one which may be vindicated or else fade into irrelevance. A communist platform that claims the collapse of capitalism is imminent is an example of the latter. Again, the selection of extremists into extremist parties, which intensifies when such parties are unsuccessful, explains their rigidity in the face of setbacks. For example, the French Communist Party declined from a base of electoral support of roughly 20% in the postwar period to less than 3% in the late 2010s.⁷ Despite this dramatic decline, partly caused by the demise of the Soviet Union, they have preserved the main tenets of their platform, such as the claim that the capitalist system is on the verge of collapse.

3 Related Literature

This paper is related to the literature on strategic experimentation with multiple agents (Keller et al. 2005, Keller and Rady 2010, Keller and Rady 2015, Strulovici

⁵See, for example, <https://www.theguardian.com/commentisfree/2019/apr/19/extinction-rebellion-climate-change-protests-london>.

⁶Note that in these examples agents should be modeled as having common values, since agents benefit from a change in public policy regardless of how much their actions contributed to it. Although we write our main model for the case of private values, we show in Section 7 that our main results survive in the common values setting. Another way to accommodate common values in our model is to endow agents with expressive payoffs, whereby agents benefit not just from a policy change but also from having participated in the efforts that brought it about.

⁷See, for example, Bell (2003), Bréchon (2011) and Damiani and De Luca (2016).

2010), as well as the literature on dynamic decision-making in clubs (Acemoglu et al. (2008, 2012, 2015), Roberts 2015, Bai and Lagunoff 2011, Gieczewski 2019). The idea that members of a declining organization may react by leaving or pushing for policy changes has been discussed qualitatively by Hirschman (1970).

In Keller, Rady and Cripps (2005), multiple agents with common priors control two-armed bandits of the same type which may have breakthroughs at different times. In their model, there is under-experimentation due to free-riding. In contrast, we study an organization making a single collective decision in each period about whether to experiment. While the organization experiments, members can exit and collect an outside option, but at the cost of their voting rights. The selection of optimists into the organization causes excessive experimentation.⁸

In Strulovici (2010) a community of agents decides by voting whether to collectively experiment with a risky technology. Agents have common priors, but experimentation gradually reveals some of them to be *winners* and others to be *losers* from the risky technology. In equilibrium, there is too little experimentation because agents fear being trapped into using a technology that turns out to be bad for them.

A similar motive to under-experiment is present in our model. Indeed, consider an agent who would prefer to experiment today but not tomorrow. If she anticipates that learning will result in an extreme optimist coming to power tomorrow, then she may choose not to experiment today, lest she be forced to over-experiment or switch to her inefficient outside option. However, there are two important differences between our model and Strulovici's. First, in our model, agents can switch to an outside option at the cost of their voting rights. This novel assumption is what allows for the organization to be captured by extremists. Second, in Strulovici's model, learning exacerbates the conflict between agents, while in our model learning helps agents converge to a common belief.

The literature on decision-making in clubs studies dynamic policy-making in settings where current policy choices determine current flow payoffs and future control of the club. Some papers in this literature focus on discrete policy spaces (Acemoglu et al. 2008, 2012, 2015, Roberts 2015), as we do, while others consider continuous

⁸While there is free-riding in the sense that outsiders benefit from the option value of experimentation, it is not socially costly because we assume the learning rate to be independent of the size of the organization. See Section 7 for an extension with an endogenous learning rate.

policy spaces (Bai and Lagunoff 2011, Gieczewski 2019). In all these papers, tensions arise due to conflicting preferences; in contrast, in our model, they are the result of heterogeneous beliefs. Furthermore, we allow these beliefs to change endogenously and stochastically due to experimentation, whereas the existing work considers deterministic models.⁹ Like the present paper, Gieczewski (2019) studies a setting in which agents are free to join or leave an organization, and only exert control over the policy if they are members. The rest of the clubs literature, on the other hand, considers settings in which the policy choice influences the set of the decision-makers directly. Our main results can only be obtained in a model with learning and heterogeneous beliefs, and are thus unique to our paper. In particular, in the equilibrium with perpetual experimentation that we characterize, the organization persists in using a policy that, in the limit, is preferred by almost nobody. The observation that such an equilibrium can exist is novel to the literature.

4 The Model

Time $t \in [0, \infty)$ is continuous. There is an organization that has access to a *risky* policy and a *safe* policy. The risky policy is either *good* or *bad* and its type is persistent. We use the notation $\theta = G, B$ for each respective scenario.

The world is populated by a continuum of agents, represented by a continuous density f over $[0, 1]$. The position of an agent in the interval $[0, 1]$ indicates her beliefs: an agent $x \in [0, 1]$ has a prior belief that the risky policy is good with probability x . All agents discount the future at rate γ .

At every instant, each agent chooses whether to be a member of the organization. Agents can enter and leave the organization at no cost.¹⁰ Agents who choose not to be members work independently and obtain a guaranteed *autarkic* flow payoff a . The flow payoffs of members depend on the organization's policy.

Whenever the organization uses the safe policy ($\pi_t = 0$), all members receive a guaranteed flow payoff s . When the risky policy is used ($\pi_t = 1$), its payoffs depend

⁹The only exception is Acemoglu et al. (2015), but they only consider *exogenous* stochastic shocks.

¹⁰Our main results survive if we assume that agents cannot reenter after exiting.

on the state of the world. If the risky policy is good, it succeeds according to a Poisson process with rate r . If the risky policy is bad, it never succeeds. Each time the risky policy succeeds, all members receive a lump-sum unit payoff. At all other times, the members receive zero while the risky policy is used.

We assume that $0 < a < s < r$. That is, the organization's safe policy is preferable to working independently.¹¹ Moreover, the risky policy would be the best choice were it known to be good, but the bad risky policy is the worst of all options.

When the risky policy is used, its successes are observed by everyone. By Bayes' rule, the posterior belief of an agent with prior x who has seen k successes after experimenting for a length of time t is

$$\frac{x}{x + (1 - x)L(k, t)}$$

where $L(k, t) = \mathbb{1}_{k=0}e^{rt}$. (Note that all posteriors jump to 1 after a success). Since $L(k, t)$ serves as a sufficient statistic for the agents' beliefs, suppressing the dependence on k and t , we take $L = L(k, t)$ to be the state variable in our model and define $p(L, x)$ as the posterior of an agent with prior x when the state variable is L . We use L_t to denote the value of L at time t .

We focus on Markov Perfect Equilibria, that is, equilibria in which the strategies condition only on the information about the risky policy revealed so far and on the incumbent policy. We let $m(L, \pi)$ denote the equilibrium median member given the state variable L and the incumbent policy π .

The structure of the game is as follows. At each instant $t > 0$, *policy* and *membership* decisions are made. That is, first the median member of the organization, $m(L_t, \pi_t)$, chooses a policy to be used by the organization in the immediate future.¹² After this, all agents are allowed to enter or leave the organization.

¹¹Our model features one organization with access to a risky technology. We can, however, allow for the existence of other organizations that only have access to a safe technology. a can be interpreted as the (maximal) productivity of these alternative organizations. The assumption $a < s$ means that the organization with access to the risky technology also enjoys a competitive advantage in the use of the safe technology. If $a \geq s$, then our main results still go through but become less interesting as there is no longer an opportunity cost to having the organization experiment.

¹²For the median to be well-defined, the set of members must be Lebesgue-measurable. This will not be an issue, since in equilibrium the set of members is always an interval.

To simplify the presentation of the model, we will make the following assumptions. First, we assume that the organization is using the risky policy in the beginning of the game, that is, $\pi_0 = 1$.¹³ Second, we assume that a switch to the safe policy is irreversible.¹⁴

Formally, a pure Markov strategy profile is given by a policy function $\alpha(L)$ and a membership function $\beta(x, L, \pi)$, where $\alpha(L) = 1$ if the organization continues experimenting in state L and $\alpha(L) = 0$ if it stops, and $\beta(x, L, \pi) = 1$ if x is a member in state (L, π) .¹⁵

Definition 1. An equilibrium satisfies the following:

- (i) Agents choose whether to be members based on their flow payoffs: x is a member at time t if and only if $s + \pi_t(p(L_t, x)r - s) > a$.
- (ii) The organization continues experimenting at time t if and only if $m(L_t, 1)$'s payoff from the equilibrium continuation is greater than the payoff from switching to the safe policy, $\frac{s}{\gamma}$.

The reason that agents make membership decisions based on their flow payoffs is that there is a continuum of agents, so an agent obtains no value from her ability to vote. Part (ii) of the definition embeds an important additional assumption about the timing of policy and membership decisions: it implies that, for the organization to stop experimenting, a majority of those who chose to be members *under experimentation* must be in favor of stopping. In other words, we rule out equilibria in which a large number of agents who dislike the current policy join the organization and immediately change its policy.

The equilibrium we define can be obtained as a limit of the equilibria of a discrete-time game in which membership and policy decisions are made at times

¹³Starting with the safe policy at $t = 0$ is equivalent to starting with the risky policy, unless the population median finds the continuation in the latter scenario inferior to the payoff from never experimenting, in which case the safe policy is used forever.

¹⁴We show in the Appendix that this assumption is without loss of generality: in a more general model in which the organization can change policies any number of times, switches to the safe policy are permanent in every equilibrium. The reason is that switching to the safe policy brings in more pessimistic members, and hence yields control to an even more pessimistic median than the one who chose to stop experimenting.

¹⁵We can also define non-Markov strategies by making α and β functions of the history h rather than the payoff-relevant state (L, π) , but this will not be necessary for our analysis.

$t \in \{0, \epsilon, 2\epsilon, \dots\}$ with $\epsilon > 0$ small. In this game, at each time t in $\{0, \epsilon, 2\epsilon, \dots\}$, first the incumbent median chooses a policy π_t for time $[t, t + \epsilon)$, and then all agents choose whether to be members. The agents who choose to be members at time t – and hence accrue the flow payoffs generated by policy π_t – are the incumbent members at time $t + \epsilon$. The median of this set of members then chooses $\pi_{t+\epsilon}$. In this game, equilibria in which all agents coordinate to join the organization at a particular time and vote to stop experimentation do not exist because each agent can profitably deviate by waiting one more period to join.

An additional tie-breaking rule is required to eliminate undesirable equilibria of the following variety. For example, $\alpha(L) = 0$ for all L – that is, all pivotal agents immediately switch to the safe policy – is always compatible with equilibrium: any agent who deviates towards experimenting will see her decision immediately overturned, so Condition (ii) holds vacuously. To rule out such equilibria, we will require optimal behavior even when the agent’s policy choice only affects the path of play for an infinitesimal amount of time.¹⁶

- (iii) If $m(L_t, 1)$ expects experimentation to stop immediately regardless of her action but strictly prefers experimentation to continue for any small enough length of time $\epsilon > 0$ rather than stop, then $\alpha(L_t) = 1$.

5 Equilibria in the Baseline Model

In this section we characterize the equilibria of the model described above. The presentation of the results is structured as follows. We first explain who the members of the organization are depending on what has happened in the game so far. We then reduce the problem of equilibrium characterization to finding an optimal stopping time. Next, we state our first main result, which asserts that the equilibrium in which the organization experiments forever is unique whenever it exists (Proposition 1). We present comparative statics in Proposition 2. Finally, in Proposition 3 we characterize the set of equilibria in the alternative case where experimentation cannot go on forever.

¹⁶Condition (iii) is in the spirit of weak dominance: an agent who prefers experimentation should experiment if she expects her successors to tremble and continue experimenting with some positive probability.

We start with two useful observations. First, because the bad risky policy never succeeds, the posterior belief of every agent with a positive prior jumps to 1 if a success is observed. Since $r > s$, a , if a success is ever observed, the risky policy is always used thereafter, all agents enter the organization and remain members forever.

Second, whenever the risky policy is being used, the set of members is the set of agents for whom $p(L, x)r \geq a$. It is clear that, for any $L > 0$ (that is, if no successes have been observed), $p(L, x)$ is increasing in x . That is, agents who are more optimistic at the outset remain more optimistic after observing additional information. Hence the set of members is an interval of the form $[y_t, 1]$, where y_t is defined by $p(L_t, y_t) = \frac{a}{r}$. Note that y_t is uniquely determined as a function of L_t , and $m(L_t, 1)$ can be calculated as the median of F restricted to $[y_t, 1]$.

These observations, together with our assumption that switching to the safe policy is irreversible, imply that any equilibrium path must have the following structure. The risky policy is used until some time $T \in [0, \infty]$. If it succeeds by then, it is used forever. Otherwise the organization switches to the safe policy at time T .¹⁷ While no successes are observed, agents become more pessimistic over time and the organization becomes smaller. As soon as a success occurs or the organization switches to the safe policy, all agents join and remain members of the organization forever, and no further learning occurs.

Proposition 1 states our first main result. Before stating it, we introduce several definitions. We let $V(x)$ denote the continuation utility of an agent with belief x , provided that she expects experimentation to continue forever. This is the payoff that agent x will get from staying in the organization until some time $t(x)$ when her posterior reaches $\frac{a}{r}$ (if a success is not observed by then), then exiting the organization and working on her own, rejoining the organization if and only if a success is observed. We let m_t denote the median voter at time t provided that the organization has experimented unsuccessfully up to time t , (that is, $m_t = m(e^{rt}, 1)$), and we let $p_t(m_t)$ denote m_t 's posterior in this case (that is, $p_t(m_t) = p(e^{rt}, m_t)$). Note that $V(x)$, m_t and $p_t(m_t)$ are all uniquely defined and exogenous functions of the primitives; explicit formulas for them are given in Claims 9.1, 9.2 and 9.3 in the Appendix.

It is clear that there exists an equilibrium in which the organization experiments

¹⁷If $T = \infty$, the risky policy is used forever.

forever if and only if $V(p_t(m_t)) \geq \frac{s}{\gamma}$ for all t . Indeed, for such an equilibrium to exist, it must be that for all t the pivotal agent at time t , m_t , weakly prefers to continue experimenting. By construction, $V(p_t(m_t))$ and $\frac{s}{\gamma}$ are her payoffs from continuing and experimentation, respectively, under the expectation that future pivotal agents will never stop experimenting.

We can show that the set of parameters for which the condition $V(p_t(m_t)) \geq \frac{s}{\gamma}$ is satisfied for all t is non-empty. In particular, the condition is satisfied if the payoff a from the outside option is not too low compared to the payoff s from the safe policy.¹⁸ Proposition 1 shows that the equilibrium with perpetual experimentation is unique, whenever it exists.

Proposition 1. *Whenever there is an equilibrium in which the organization experiments forever, this is the unique equilibrium.*

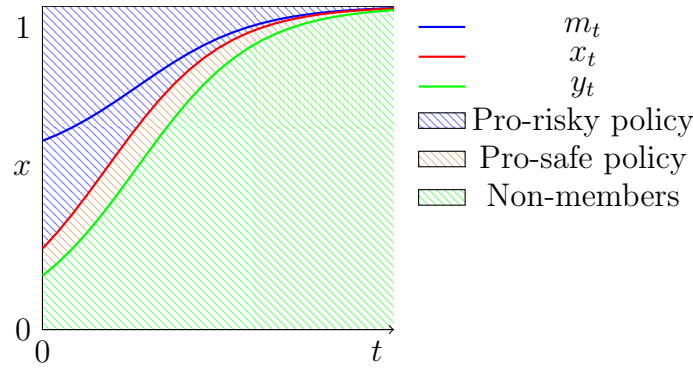


Figure 1: Median voter, indifferent voter, and marginal member on the equilibrium path

The intuition for the equilibrium dynamics is illustrated in Figure 1. As the organization experiments unsuccessfully on the equilibrium path, all agents become more pessimistic. That is, $p_t(x)$ is decreasing in t for fixed x . Letting x_t denote the agent indifferent about continuing experimentation at time t , so that $V(p_t(x_t)) = \frac{s}{\gamma}$, this implies that x_t must be increasing in t . Thus there is a shrinking mass of agents in favor of the risky policy (the agents shaded in blue in Figure 1) and a growing mass of agents against it (the agents shaded in red and green). For high t , almost all agents agree that experimentation should be stopped.

However, growing pessimism induces members to leave. Hence the marginal

¹⁸An explicit threshold such that the condition holds if a exceeds this threshold is derived in Corollary 9.1 in the Appendix.

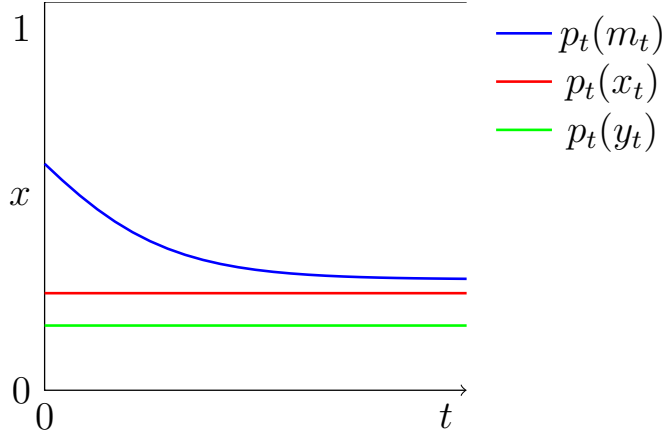


Figure 2: Posterior beliefs on the equilibrium path

member becomes more extreme, and so does the median member. If $m_t \geq x_t$ for all t , that is, if the prior of the median is always higher than the prior of the indifferent agent, then the agents in favor of the risky policy always retain a majority within the organization, due to most of their opposition forfeiting their voting rights.

Figure 2 shows the same result in the space of posterior beliefs. The accumulation of negative information puts downward pressure on $p_t(m_t)$ as t grows but selection forces prevent it from converging to zero. Instead, $p_t(m_t)$ converges to a belief strictly between 0 and 1, which is above the critical value $p_t(x_t)$ in this example. Hence the median voter always remains optimistic enough to continue experimenting.

To establish whether this equilibrium entails over-experimentation, we need a definition of over-experimentation in a setting with heterogeneous priors. We will use the following notion. Consider an alternative model in which an agent with initial belief x controls the policy at all times. It is well-known that whenever $0 < x < 1$, the agent would experiment until some finite time depending on x . We say that an equilibrium of our model features over-experimentation from x 's point of view if experimentation continues longer than that. By this definition, when the condition $V(p_t(m_t)) \geq \frac{s}{\gamma}$ for all t is satisfied, there is over-experimentation from the point of view of *all* agents except those with prior belief exactly equal to 1.

The level of experimentation in equilibrium is determined by the interaction of two opposing forces, in addition to the usual incentives present in the canonical single-agent bandit problem. When the pivotal agent decides whether to stop experimenting

at time t , she takes into account the difference in the expected flow payoffs generated by the safe policy and the risky one, as well as the option value of experimenting further. However, because the identity of the median voter changes over time, the pivotal agent knows that if she chooses to continue experimenting, the organization will stop at a time chosen by some other agent, which she likely considers suboptimal. This force encourages her to stop experimentation while the decision is still in her hands, leading to under-experimentation. It is similar to the force behind the under-experimentation result in Strulovici (2010) in that, in both cases, agents prefer a sub-optimal amount of experimentation because they expect a loss of control over future decisions if they allow experimentation to continue. It is also closely related to the concerns about slippery slopes faced by agents in the clubs literature (see, for example, Bai and Lagunoff (2011) and Acemoglu et. al. (2015)).

The second force stems from the endogeneity of the median voter's position in the distribution. As discussed above, the more pessimistic a fixed observer becomes about the risky policy, the more extreme the median voter is. This effect is so strong that, as time passes, the posterior belief of the median after observing no successes does not converge to zero, and the median voter may choose to continue experimenting when no successes have been observed for an arbitrarily long time.

Next, we explain why the equilibrium with perpetual experimentation is unique. The key here is that if an agent prefers to experiment forever rather than not at all, then she also prefers to experiment for any finite amount of time T rather than not at all. Thus if the median conjectured that the organization will experiment for some finite time T instead of forever, the median still would not want to stop experimentation.

The central result that we use in the proof here is that the value function $W_T(x)$ of an agent with prior x who expects experimentation to continue for time T is single-peaked in T . That is, there is a time T^* such that if the agent had control over the policy at all times, the agent would experiment for time T^* , and the farther away the actual length of experimentation is from T^* , the less happy the agent is. Since $W_0(x) = \frac{s}{\gamma}$ and $W_T(x)$ is increasing in T before the peak, the agent prefers to experiment for time T before the peak rather than not at all. Moreover, since $\lim_{T \rightarrow \infty} W_T(x) = V(x)$ and $V(x) > \frac{s}{\gamma}$ by our hypothesis, the fact that $W_T(x)$ is decreasing in T after the peak implies that the agent also prefers to experiment for

time T after the peak rather than not at all. This establishes our result.

We can check whether there is perpetual experimentation in equilibrium by calculating $\inf_t V(p_t(m_t))$. For certain families of belief densities, either an exact expression or a lower bound for $\inf_t V(p_t(m_t))$ can be given.¹⁹ For instance, if the density f is non-decreasing, then the following holds:

$$\gamma \inf_{t \geq 0} V(p_t(m_t)) = \gamma V\left(\frac{2a}{r+a}\right) = \frac{2ra}{r+a} + \left(\frac{1}{2}\right)^{\frac{\gamma}{r}} \frac{a(r-a)}{r+a} \frac{r}{\gamma+r}$$

The value function of the pivotal agent has a simple interpretation. Under a non-decreasing density f , as the organization experiments unsuccessfully, the posterior belief of the median converges to $\frac{2a}{r+a}$ from above. The first term on the right-hand side then represents the agent's expected flow payoff from experimentation: it is the product of the probability $\frac{2a}{r+a}$ that the asymptotic median assigns to the risky policy being good, and the expected flow payoff r from the good risky policy. The second term is the option value derived from the agent's ability to leave the organization when she becomes pessimistic enough, and to return if there is a success.

Our next major result concerns the comparative statics of our model. We show that perpetual experimentation is more likely when the payoffs from the risky policy and the outside option are high, the payoff from the organization's safe policy is low, the agents are patient, and there are many optimists.

Proposition 2. *If there is an equilibrium with perpetual experimentation under parameters (r, s, a, γ, f) , then the same holds for any set of parameters (r, s', a', γ', f') such that $s' \leq s$, $a' \geq a$, $\gamma' \leq \gamma$ and f' MLRP-dominates f .²⁰*

The intuition behind this result is as follows. An increase in the payoff s from the safe policy makes the safe policy more attractive and has no effect on the expected payoff from perpetual experimentation. An increase in patience is equivalent to an increase in the rate of learning from experimentation, and a higher learning rate allows agents to make better entry-exit decisions.

An increase in the number of optimists leaves the value function and the marginal

¹⁹See Proposition 9 in the Appendix for more details.

²⁰We say that g MLRP-dominates f if $x \mapsto \frac{g(x)}{f(x)}$ is non-decreasing for $x \in [0, 1]$.

member unchanged but results in a more optimistic median who is more likely to support experimentation.²¹ An increase in a has two effects that favor experimentation. First, it increases the payoff from perpetual experimentation because the agent expects to quit the organization and collect the outside option payoff with some probability. Second, it induces agents to quit, which leaves the organization with a more radical median voter.

Finally, an increase in r has three effects. On the one hand, it makes experimentation more attractive, both by directly increasing the payoff from the good risky policy and by increasing the learning rate. At the same time, it induces more agents to stay in the organization, which makes the median more pessimistic. Hence the overall effect is ambiguous. However, for a well-behaved family of densities, the first two effects dominate, so that perpetual experimentation is also more likely for high r .²²

If there does not exist an equilibrium with perpetual experimentation, there may be multiple equilibria featuring different levels of experimentation, supported by different off-path behavior.²³ To state the results, we let \hat{T} denote the time such that the initial median is indifferent between switching to the safe policy and continuing to experiment if she anticipates that, should she continue, the organization will stop experimentation at time \hat{T} .

Remark. *Any equilibrium stopping time must lie in $[0, \hat{T}]$.*

Note that experimentation must stop by time \hat{T} , as otherwise the initial median would switch to the safe policy immediately. Proposition 3 speaks to the extent of experimentation when there does not exist an equilibrium with perpetual experimentation.

Proposition 3.

²¹For example, if f is uniform, m_t is the midpoint between y_t and 1, while if f is increasing, then m_t is closer to 1 than y_t .

²²Specifically, for each $\omega \geq 0$, let $f_\omega(x)$ denote a density in the power-law family given by $f_\omega(x) = (\omega + 1)(1 - x)^\omega$. If $f = f_\omega$ for some ω , then $\inf_t V(p_t(m_t))$ is increasing in r . See Proposition 9 and the proof of Proposition 2 in the Appendix for details.

²³More generally, a pure strategy equilibrium can be described by a sequence $t_0 < t_1 < t_2 < \dots$ of stopping times as follows. For any $t \in (t_{n-1}, t_n]$, if the risky policy was used in the period $[0, t]$ and no successes were observed, the organization continues using it until time t_n . If the risky policy has not succeeded by t_n , the organization switches to the safe policy at t_n .

1. *There are parameters under which²⁴ all $t \in [0, \hat{T}]$ are equilibrium stopping times.*
2. *If for all $t \in [0, \bar{t}]$, m_t prefers to experiment forever rather than not at all, then in any equilibrium the organization experiments up to at least \bar{t} .*

It can be shown that, because W_T is single-peaked in T , the initial median's ideal stopping time lies between 0 and \hat{T} . Then the first part of the Proposition implies that both over- and under-experimentation are possible depending on which equilibrium is played. Under-experimentation obtains if an early median voter expects that, should she continue experimenting, the next stopping time will be too far in the future.

The second part of the Proposition obtains for the following reason. If all pivotal agents up to some time \bar{t} prefer experimenting forever to not at all, then, because W_T is single-peaked, these agents will never stop experimenting. Therefore, the equilibrium stopping time must be at least \bar{t} . This means that our result of perpetual experimentation survives in an approximate form even when, for instance, the support of the prior distribution is truncated away from 1.²⁵

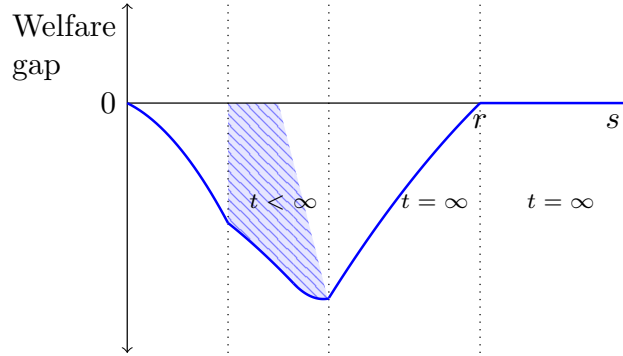


Figure 3: Welfare gap between the equilibrium and the socially optimal stopping time

Figure 3 illustrates the welfare effect of varying the quality of the outside option. The blue curve is the difference between the initial median's equilibrium utility and

²⁴The condition on the parameters roughly amounts to requiring that $p_t(m_t)$ does not decrease too steeply in t . For example, $p_t(m_t)$ being constant in t would be sufficient. See Proposition 10 in the Appendix for the exact condition and further details.

²⁵Formally, suppose that for some density f there is perpetual experimentation, and let f_y be f truncated at y , that is, $f_y(x) \propto f(x) \mathbb{1}_{x \leq y}$. Denote by \bar{t}_y the minimal equilibrium stopping time as a function of y . Then $\bar{t}_y \rightarrow \infty$ as $y \rightarrow 1$.

her utility from experimentation at her most preferred time. The shaded blue region represents the range of welfare outcomes that obtain when multiple equilibria exist. The reason this region occupies only part of the space between the two dotted lines is that, as a increases, some equilibria disappear, so the range of obtainable welfare outcomes shrinks.

In this example, the welfare gap in the worst equilibrium, compared to the initial median's optimal stopping time, is U-shaped. The reason is as follows. When $a \geq s$, the organization experiments forever. Here perpetual experimentation is optimal from the point of view of all agents. This is because, since the organization's safe policy is no better than the outside option, agents who lose faith in the risky policy prefer to switch to the outside option rather than use the organization's safe policy.

For $a < s$ relatively close to s , the organization still experiments forever, but this is now inefficient from the point of view of all agents, with the size of the welfare loss increasing as the gap between s and a grows. The inefficiency comes from the fact that pessimistic agents are denied access to the organization's safe policy.

For even lower values of a , perpetual experimentation is impossible because too few agents leave the organization. Here a range of outcomes is possible. These outcomes may include the initial median's preferred stopping time. Finally, when a is very close to 0, we can show that the equilibrium is again unique. In this equilibrium, few agents ever leave, but the median still becomes more optimistic over time as bad news arrive. This means that, generally, the decision is still made by an agent more optimistic than the initial median. Moreover, the gap between the initial median and the median who stops is growing in the number of people who leave, which itself is increasing in a .

6 Other Learning Processes

The baseline model presented above has two salient features. First, experimentation has a low probability of generating a success, which increases agents' posterior beliefs substantially, and a high probability of generating no successes, which lowers their posteriors slightly. In other words, the baseline model is a model of good news. Second, because the risky policy can only succeed when it is good, good news are

perfectly informative.

In this Section, we relax these assumptions, presenting variants of the model which allow for imperfectly informative good news and for bad news. In the first case, we show that our finding of over-experimentation is robust to imperfectly informative news. We also show that the organization may respond perversely to information, becoming more reluctant to experiment after a success. In the case of perfectly informative bad news, in contrast, there is typically under-experimentation.

6.1 A Model of Bad News

In this section we consider the same model as in Section 4, except that the risky policy now generates different flow payoffs. In particular, if the risky policy is good, it generates a guaranteed flow payoff r . If it is bad, it generates a guaranteed flow payoff r but also experiences failures, which arrive according to a Poisson process with rate r . Each failure lowers the payoffs of all members by 1. Thus, as in the baseline model, the expected flow payoff from using the risky policy is r when it is good and 0 when it is bad. The learning process, however, is different.

The dynamics of organizations under bad news differ substantially from those in the baseline model. As is usual in models of bad news, as long as no failures are observed, all agents become more optimistic about the risky technology, so the organization expands over time instead of shrinking. This gradual expansion continues either forever or until some time T unless a failure occurs, in which case the organization switches to the safe technology and all agents previously outside the organization become members. Interestingly, the switch to the safe technology must happen upon observing a failure but may happen even if no failures are observed.

As before, m_t is the median member at time t provided that the risky policy has been used up to time t and no failures have been observed. $p_t(m_t)$ is the median's posterior belief at time t , and $V(p_t(m_t))$ is her continuation value when experimentation is expected to continue forever unless there is a failure. Let us use \underline{t} to denote the earliest time when an agent with $V(p_t(m_t)) < \frac{\varepsilon}{\gamma}$ is pivotal. Proposition 4 provides a characterization of the equilibria in this variant of the model.

Proposition 4.

1. If $V(p_t(m_t)) > \frac{s}{\gamma}$ for all t , then there is a unique equilibrium. In it, the organization experiments forever.
2. If $V(p_t(m_t)) < \frac{s}{\gamma}$ for some t , then in any equilibrium the organization stops experimenting at a finite time $T < \underline{t}$.²⁶

Proposition 4 shows that perpetual experimentation is the unique equilibrium outcome if all pivotal agents prefer it to the safe policy. If, however, some pivotal agents are pessimistic enough to stop experimentation, the organization will always switch to the safe policy even *before* any of these pessimists become pivotal. Note that, even when perpetual experimentation arises in the bad news setting, it does not constitute over-experimentation, as it is possible only when all agents agree that perpetual experimentation is optimal.

To understand these results, it is instructive to consider the associated single-agent bandit problem. In a model of bad news, the agent switches to the safe policy permanently upon observing a failure, and becomes more optimistic over time if she pursues the risky policy and observes no failures. The more optimistic the agent becomes, the more she wants to continue using the risky policy. Hence the agent will either want to experiment forever or not at all.

Two implications follow from this observation. First, pessimistic agents with $V(p_t(m_t)) < \frac{s}{\gamma}$ always switch to the safe policy when they are pivotal: by assumption, they prefer no experimentation to perpetual experimentation, and thus also to any other continuation. Second, optimistic agents with $V(p_t(m_t)) > \frac{s}{\gamma}$ have stronger incentives to experiment if they expect experimentation to continue in the future: only then can they collect the option value of learning about the policy. For them, current and future experimentation are strategic complements.

These implications lead to the organization experimentation strictly before \underline{t} , as follows. Optimistic agents with $V(p_t(m_t)) > \frac{s}{\gamma}$ are willing to experiment if they expect perpetual experimentation in the continuation. However, agents who are pivotal shortly before \underline{t} know that any experimentation they attempt will be short-lived. Thus they will choose to stop experimenting even if they are optimists. In turn, their expected behavior may induce even earlier pivotal agents to switch to the safe policy as well.

²⁶This is true so long as $\underline{t} > 0$. If $\underline{t} = 0$, then $T = 0$.

To summarize, in a bad news setting, over-experimentation is never possible from the point of view of any pivotal agent, while under-experimentation is possible, and always obtains when experimentation does not continue forever. These results stand in stark contrast to those of our previous models. The results depend on a special feature of the perfectly informative bad news learning process: bad news create common knowledge that the risky policy is bad. Because of this, there is no room for organizational capture by optimists who would disagree with the majority.

6.2 A Model of Imperfectly Informative (Good) News

We first treat the case of imperfectly informative news, which allows for much richer dynamics than the baseline model: agents' beliefs, rather than decreasing monotonically or else jumping permanently to 1, can change in both directions as successes and failures arrive. For brevity, we consider the case of good news, but similar results can be obtained for imperfectly informative bad news.

The model is the same as in Section 4 except for the payoffs generated by the risky policy. If the risky policy is good, it generates successes according to a Poisson process with rate r . If it is bad, it generates successes according to a Poisson process with rate c . We now assume that $r > s > a > c > 0$.

As before, the effect of past information on the agents' beliefs can be aggregated into a one-dimensional sufficient statistic. Suppose the risky policy has been used for a length of time t and k successes have occurred during that time. Define

$$L(k, t) = \left(\frac{c}{r}\right)^k e^{(r-c)t}$$

Then the posterior of an agent with prior x at time t after observing the organization use the risky policy for a length of time t and achieve k successes is

$$\frac{x}{x + (1 - x)L(k, t)}$$

We again suppress the dependence of $L(k, t)$ on k and t and use L to denote our sufficient statistic. Note that high L indicates bad news about the risky policy.

We let $V_x(L)$ denote the value function of an agent with prior x given that the state is L and the organization experiments forever in the continuation. In addition, we denote x 's ex ante utility in the same situation by $V(x) = V_x(1)$. The next Proposition shows that, as in Section 5, experimentation can continue forever regardless of how badly the risky policy performs.

Proposition 5. *If $V_{m(L)}(L) > \frac{s}{\gamma}$ for all L , then there is a unique equilibrium. In this equilibrium, the organization experiments forever.*

We can show that there exist parameters such that the inequality $V_{m(L)}(L) > \frac{s}{\gamma}$ is satisfied. For instance, if f is non-decreasing, then the posterior of the median converges to $\frac{2(a-c)}{(r-c)+(a-c)}$, so $\inf_t V(p_t(m_t)) = V\left(\frac{2(a-c)}{(r-c)+(a-c)}\right)$. Then it is enough to verify that $V\left(\frac{2(a-c)}{(r-c)+(a-c)}\right) \geq \frac{s}{\gamma}$.²⁷

The following result illustrates the novel outcomes that can arise under imperfectly informative news.

Proposition 6. *There exist parameters such that there is an equilibrium in which the organization experiments more when the risky policy is bad than when it is good.*

The intuition for the result in Proposition 6 is as follows. We first show that, for an appropriately chosen density f , an equilibrium of the following form exists: whenever $L = L^*$, the organization stops experimenting with probability ϵ , and at all other times the organization continues experimenting for sure. For this to work, f must be such that the median is most pessimistic when $L = L^*$.²⁸ Moreover, s must be such that $V_{m(L^*)}(L^*) = \frac{s}{\gamma}$, so that the median is indifferent about experimentation at L^* , while other agents prefer to continue experimenting when they are pivotal.

The striking feature of this equilibrium is that only happens for an intermediate value of L . In particular, if L^* is smaller than the initial L , the only way experimentation will stop is *if it succeeds* enough times for L to decrease all the way to L^* , which is more likely to happen when the risky policy is good.

²⁷However, in this case it is not possible to give an exact expression for V , owing to the complicated behavior of L over time.

²⁸This occurs, for instance, if f is very high in a small neighborhood of $y(L^*)$. Then, when $L > L^*$, all the pessimists to the left of $y(L^*)$ leave, so that $m(L)$ is more optimistic, while when $L < L^*$, pessimists become members, yielding a lower $m(L)$. This construction is not possible for some densities of prior beliefs. In particular, if f is uniform or follows a power law distribution, then $p(L, m(L))$ is decreasing in L (Lemma 5).

7 Other Extensions

7.1 General Voting Rules

We assume throughout the paper that the median member of the organization is pivotal. This assumption is not essential to our analysis: our results can be extended to other voting rules under which the agent at the q -th percentile is pivotal.

It is instructive to consider how the results change as we vary q . Letting q_t represent the pivotal agent at time t , it is clear that q_t and $p_t(q_t)$ are increasing in q for all t . To illustrate further, assume that f is uniform. Then, as $t \rightarrow \infty$, the posterior belief of the pivotal agent converges to $\frac{a}{qa+(1-q)r}$ rather than $\frac{2a}{r+a}$. It follows that more stringent supermajority requirements are functionally equivalent to more optimistic leadership of the organization, and make it easier to sustain an equilibrium with excessive experimentation.

7.2 Size-Dependent Payoffs

In some settings the payoffs that an organization generates may depend on its size. In this section we discuss how different operationalizations of this assumption affect our results. We show that our main result is robust to this extension, and discuss how different kinds of size-dependent payoffs may exacerbate or prevent over-experimentation.

We consider three types of size-dependent payoffs. For the first two, we suppose that when the set of members of the organization has measure μ , the safe policy yields a flow payoff $sg(\mu)$, the good risky policy yields instantaneous payoffs of size $g(\mu)$ generated at rate r , and the bad risky policy yields zero payoffs. We assume that $g(1) = 1$, so that r , s and 0 are the expected flow payoffs from the good risky policy, the safe policy and the bad risky policy respectively when all agents are in the organization. For the first type of payoffs we consider, $g(\mu)$ is increasing in μ , so there are economies of scale. For the second type, $g(\mu)$ is decreasing in μ , so there is a congestion effect.

In general, the effect of size-dependent payoffs on the level of experimentation

is ambiguous because of two countervailing effects. On the one hand, when there is a congestion effect, as the organization shrinks, higher flow payoffs increase the benefits from experimentation, which makes experimentation more attractive.²⁹ We call this the *payoff effect*. On the other hand, because increasing flow payoffs provide incentives for agents to stay in the organization, the organization shrinks at a lower speed, which causes the median voter in control of the organization to be more pessimistic about the risky policy. We call this the *control effect*. When there are economies of scale, these effects are reversed.

When there are economies of scale, the set of members may not be uniquely determined as a function of the state at time t . This is because the more members there are, the higher payoffs are, so the membership stage may have multiple equilibria. We will assume, for simplicity, that the set of members is uniquely determined.³⁰ It is sufficient to assume that g does not increase too fast.

The following Proposition presents our first result.

Proposition 7. *Suppose that $f = f_\omega$.³¹ Let $\bar{g} = \lim_{\mu \rightarrow 0} g(\mu)$, and let $V_{g,t}(p_t(m_t))$ denote the utility of the pivotal agent at time t if she expects experimentation to continue forever. If*

$$\lambda^{\frac{\gamma}{r}} a \frac{r}{\gamma + r} + \frac{a}{\lambda} \frac{\gamma}{\gamma + r} > r \frac{r}{\gamma + r} + a \frac{\gamma}{\gamma + r}$$

then $\lim_{t \rightarrow \infty} V_{g,t}(p_t(m_t))$ is strictly increasing in \bar{g} for all $\bar{g} \in [a, \infty)$. In this case, perpetual experimentation obtains for a greater set of parameter values with a congestion effect and for a smaller set of parameter values with economies of scale, relative to the baseline model.

Conversely, if the reverse inequality holds strictly, then $\lim_{t \rightarrow \infty} V_{g,t}(p_t(m_t))$ is strictly decreasing in \bar{g} for all $\bar{g} \in [a, \infty)$.

The intuition for the Proposition is as follows. By the same argument as in the baseline model, Proposition 1 holds, and a sufficient condition to obtain experi-

²⁹While the safe policy could also yield high payoffs when the organization is small, all agents will enter as soon as the safe policy is implemented, so these high payoffs can never be captured.

³⁰Formally, we require that the equation $\frac{y_t}{y_t + (1 - y_t)e^{rt}} = \frac{a}{g(1 - F(y_t))b}$ has a unique fixed point for all $t \geq 0$.

³¹Recall that $f_\omega(x)$ is a density with support $[0, 1]$ such that $f_\omega(x) = (\omega + 1)(1 - x)^\omega$ for $x \in [0, 1]$.

mentation forever is that $V_{g,t}(p_t(m_t)) \geq \frac{s}{\gamma}$ for all t . While it is difficult to calculate $V_{g,t}(p_t(m_t))$ explicitly for all t , calculating its limit as $t \rightarrow \infty$ is tractable and often allows us to determine whether the needed condition holds for all t . We show that the limit depends only on \bar{g} rather than the entire function g . Moreover, it is a hyperbola in \bar{g} , so it is either increasing or decreasing in \bar{g} everywhere. In the first case, size-dependent payoffs affect the equilibrium mainly through the *payoff effect*, so experimentation is more attractive with a congestion effect and less so with economies of scale. In the second case, the *control effect* dominates, and the comparative statics are reversed. These statements are precise as $t \rightarrow \infty$ (that is, conditional on the risky policy having been used for a long time). We can show that when congestion effects make experimentation more likely in the limit, they do so for all t .³²

The inequality in the Proposition determines which case we are in. Because $r > \lambda^{\frac{\gamma}{r}}a$ and $\frac{a}{\lambda} > a$, if r is large enough relative to γ , then over-experimentation is easier to obtain with economies of scale than in the baseline model, and easier to obtain in the baseline model than with a congestion effect. The opposite happens if γ is large relative to r . The reason is that, under economies of scale, the pivotal decision-maker is very optimistic about the risky policy but expects to receive a low payoff from the first success. If $\frac{r}{\gamma}$ is large, so that successes arrive at a high rate or the agent is very patient, the first success is expected to be one of many, while if $\frac{r}{\gamma}$ is small, further successes are expected to be heavily discounted. Conversely, with a congestion effect, for large t the pivotal decision-maker is almost certain that the risky policy is bad but believes that, with a low probability, it will net a very large payoff before she leaves.

The third way to operationalize size-dependent payoffs that we consider deals with changes to the learning rate rather than to flow payoffs. Here we suppose that when the organization is of size μ , the good risky policy generates successes at a rate $r\mu$. Each success pays a total of 1, which is split evenly among members, so that each member gets $\frac{1}{\mu}$. All other payoffs are the same as in the baseline model. An example that fits this setting is a group of researchers trying to find a breakthrough. If there are fewer researchers, breakthroughs are just as valuable but happen less often. When f is uniform, and using V to denote the continuation utility under perpetual

³²It can be shown that when congestion effects make experimentation less likely in the limit, they may not do so for all t .

experimentation, we have

$$\gamma \inf V_t(p_t(m_t)) = \gamma \lim_{t \rightarrow \infty} V_t\left(\frac{2a}{r+a}\right) = \frac{2ra}{r+a}$$

In other words, the asymptotic median's expected payoff from experimentation comes only from the flow payoff of the risky policy; the option value of experimentation vanishes as the learning rate converges to zero. It follows that perpetual experimentation is less likely to obtain here than in the baseline model, but is still the unique outcome if $\frac{2ra}{r+a} > \frac{s}{\gamma}$. Note that, in this case, as agents who join the organization increase the learning rate, they confer a positive externality on outsiders which is not internalized. Hence, there is free-riding as in Keller et al. (2005). It is simultaneously possible that too few agents partake in experimentation—given that the organization's policy is risky—and that the risky policy is used for too long.

7.3 Organizations of Fixed Size

For the sake of simplicity and clarity, our main model assumes that the organization allows agents to enter and exit freely and adjusts in size to accommodate them. While free exit is a reasonable assumption in all of our applications, the assumptions of free entry and flexible size are often violated: in the short run, organizations may need to maintain a stable size to sustain their operations. In this Section, we discuss a variant of our model incorporating these concerns.

Assume now that agents differ in two dimensions: their prior belief $x \in [0, 1]$ and their ability $\chi \in [\underline{\chi}, \bar{\chi}]$. Suppose that the density of agents at each pair (x, χ) is of the form $f(x)h(\chi)$, where f is a probability density function and h is a degenerate density such that, for any $\chi > \underline{\chi}$, $\int_{\underline{\chi}}^{\chi} h(\tilde{\chi})d\tilde{\chi} = \infty$ but $\int_{\chi}^{\bar{\chi}} h(\tilde{\chi})d\tilde{\chi} < \infty$. In other words, prior beliefs and ability are independently distributed, and for each belief x there is a deep pool of agents, if candidates of low enough ability are considered.

Assume that the organization must maintain a fixed size μ ; that it observes only ability, and not beliefs, from its hiring process; and that it benefits from hiring high-ability agents. Suppose that agents are compensated equally for their ability inside or outside the organization (that is, their propensity to be members is independent of

ability). Then, in equilibrium, at time t only candidates with prior belief $x \geq y_t$ are willing to work at the organization. Here y_t is given by $p_t(y_t) = \frac{a}{r}$, as in the baseline model. The organization hires all candidates of ability at least χ_t , where χ_t is chosen so that $(1 - F(y_t))(1 - H(\chi_t)) = \mu$.

Since x and χ are independently distributed, the median belief within the organization at time t is still $p_t(m_t)$. From this fact we can derive an analog of Proposition 1 and show that perpetual experimentation can also obtain in this case.³³ Moreover, over-experimentation becomes even more likely if prior beliefs and ability are positively correlated, or if the organization is able to observe beliefs to some extent and prefers optimistic agents.

7.4 Common Values

Although our model features agents with private values, our results can be extended to a model with common values, which is more appropriate for some of our applications, such as environmental organizations or civil rights activism.

We discuss this extension in the context of our example of civil rights activism. Instead of each agent x generating some private income flow, she now makes a flow *contribution* to a rate of change in the relevant laws which can be attributed to the activism of agent x . The mapping from membership and policy decisions to the outcomes is the same as in Section 4, but now agents care only about the overall rate of changes in the law, and not about their own contribution.

Formally, we let $U_x(\sigma_y, \sigma)$ denote the private utility of agent x when she plays the equilibrium strategy of agent y and the equilibrium path is dictated by the strategy profile σ . Then in the private values case x 's equilibrium utility is $U_x(\sigma_x, \sigma)$, while in the common values case it is $\int_0^1 U_x(\sigma_y, \sigma) f(y) dy$. Note that, even though all agents share the objective of maximizing the aggregate rate of legal change, their utility functions still differ due to differences in prior beliefs. However, it is still optimal for

³³In fact, perpetual experimentation is easier to obtain in this case. Letting χ_0 be the ability threshold when everyone wants to be a member, that is, when there has been a success or the safe policy is being used, we can show that the incentives to advocate for experimentation are the same as in the baseline model for agents (x, χ) with $\chi \geq \chi_0$. However, agents (x, χ) with $\chi < \chi_0$ have a dominant strategy to advocate for experimentation, because they know that a switch to the safe policy would see them fired immediately.

them to make membership decisions that maximize their flow contributions at each point in time, just as in Section 4.³⁴

Let us conjecture a strategy profile in which the organization experiments forever, and let $\tilde{V}_t(x)$ denote the continuation utility at time t of an agent who has posterior belief x at time t under this strategy profile.³⁵ Then Proposition 1 holds with a similar proof, replacing $V(p_t(m_t))$ in the original proof with $\tilde{V}_t(p_t(m_t))$. Moreover, the following lower bound for the value function holds:

Proposition 8. *For any $x \in [0, 1]$ and any $t \geq 0$,*

$$V(x) \geq \tilde{V}_t(x) \geq \min \left\{ x \frac{r}{\gamma}, x \frac{r}{\gamma} + (1-x) \frac{a}{\gamma} - x \frac{r-a}{\gamma+r} \right\}$$

The lower bound on $\tilde{V}_t(x)$ is the minimum of two expressions. The first one is the expected payoff from the organization experimenting forever and all agents remaining members forever. The second one is the payoff from the organization experimenting forever and (almost) no one remaining a member.

Proposition 8 can be used to obtain a sufficient condition for perpetual experimentation. For instance, when the density of the prior beliefs f is uniform, experimentation continues forever as long as

$$\min \left\{ \frac{2ra}{r+a}, \frac{2ra}{r+a} + \frac{(r-a)a}{r+a} \left(1 - 2 \frac{\gamma}{\gamma+r} \right) \right\} \geq s$$

In other words, over-experimentation can still occur in equilibrium for reasonable parameter values – in particular, if a is close enough to s .

The common values setting differs from the baseline model in two important ways. First, the fact that $\tilde{V}_t(x) \leq V(x)$ means that agents' payoffs from experimentation are always weakly lower in the common values case than in the private values case. As a result, over-experimentation occurs for a smaller set of parameter values in

³⁴Agents are now indifferent about their membership decisions: the membership status of a set of agents of measure zero has no impact on anyone's payoffs. However, it is natural to assume that each agent joins when doing so would be optimal if her behavior had a positive weight in her own utility function. For instance, this is the case in a model with a finite number of agents.

³⁵ t matters in this case, in contrast to the model in Section 4, because the membership strategies of other agents, which depend on t rather than x , enter the agent's utility function.

the common values case. The reason is that, under common values, an agent considers the entry and exit decisions of other agents suboptimal, and her payoff is affected by these decisions as long as experimentation continues. In contrast, in the private values case, agents' payoffs depend only on their own entry and exit decisions, which are chosen optimally given their beliefs.

Second, in the private values case, experimentation can continue forever even if agents are impatient, as long as the density f does not decrease too quickly near 1 and other parameters are chosen appropriately (for example, a is close to s). This occurs because the pivotal agent is optimistic enough that the expected flow payoff from experimentation is higher than s , even without taking the option value into account. In contrast, in the common values case, the expected flow payoff from experimentation goes to a as $t \rightarrow \infty$ if there are no successes, no matter how optimistic the agent is. Indeed, here agents care about the contributions of all players, and they understand that for large t most players will become outsiders and generate a , regardless of the quality of the policy. Thus perpetual experimentation is only possible if agents are patient enough.

References

- Acemoglu, Daron, Georgy Egorov, and Kostantin Sonin**, “Coalition Formation in Non-Democracies,” *Review of Economic Studies*, 2008, *75* (4), 987–1009.
- , – , **and** – , “Dynamics and Stability of Constitutions, Coalitions and Clubs,” *American Economic Review*, 2012, *102* (4), 1446–1476.
- , – , **and** – , “Political Economy in a Changing World,” *Journal of Political Economy*, 2015, *123* (5), 1038–1086.
- Bai, Jinhui and Roger Lagunoff**, “On the Faustian Dynamics of Policy and Political Power,” *Review of Economic Studies*, 2011, *78* (1), 17–48.
- Bell, David S**, “The French Communist Party: from revolution to reform,” in Jocelyn Evans, ed., *The French party system*, Manchester University Press, 2003, chapter 2.
- Bréchon, Pierre**, *Les partis politiques français*, La Documentation française, 2011.
- Carreyrou, John**, *Bad Blood: Secrets and Lies in a Silicon Valley Startup*, Penguin Random House, 2018.
- Damiani, Marco and Marino De Luca**, “From the Communist Party to the Front de gauche. The French radical left from 1989 to 2014,” *Communist and Post-Communist Studies*, 2016, *49* (4), 313–321.
- Durrett, Rick**, *Probability: Theory and Examples*, 4 ed., Cambridge University Press, 2010.
- Gieczewski, Germán**, “Policy Persistence and Drift in Organizations,” 2019.
- Hirschman, Albert O.**, *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*, Cambridge, Massachusetts: Harvard University Press, 1970.
- Keller, Godfrey and Sven Rady**, “Strategic experimentation with Poisson bandits,” *Theoretical Economics*, 2010, *5* (2), 275–311.
- **and** – , “Breakdowns,” *Theoretical Economics*, 2015, *10* (1), 175–202.

– , – , and **Martin Cripps**, “Strategic Experimentation with Exponential Bandits,” *Econometrica*, 2005, *73* (1), 39–68.

Roberts, Kevin, “Dynamic voting in clubs,” *Research in Economics*, 2015, *69* (3), 320–335.

Strulovici, Bruno, “Learning While Voting: Determinants of Collective Experimentation,” *Econometrica*, 2010, *78* (3), 933–971.

A Appendix (For Online Publication)

A.1 Proofs

We begin with three useful Lemmas. Lemma 1 shows that agents with more optimistic priors derive higher payoffs from any expected continuation that involves experimentation. Lemma 2 proves several properties of the agents' value function when the policy path involves switching to the safe policy if there are no successes by some time T . In particular, we prove that the value function is single-peaked in T . Lemma 3 shows that increasing the number of optimists (in the MLRP sense) results in a more optimistic pivotal agent after any history.

Lemma 1. *Let $V_x(L, \pi)$ denote the value function of an agent with prior x when the initial state is (L, π) . Then for all (L, π) , $x \mapsto V_x(L, \pi)$ is strictly and continuously increasing for all agents x that are in the organization while it experiments, at a set of times of positive measure with a positive probability (on the equilibrium path).*

Proof of Lemma 1.

Consider two agents $x' > x$. Let $V_{x'}^x(L, \pi)$ denote the payoff to agent x' from copying the equilibrium strategy of agent x . When x and x' are outside the organization, their flow payoffs are equal to a and do not depend on their priors.

When x' is in the organization, if the organization is using the risky policy, at a continuation where the state variable is \tilde{L} , x' 's expected flow payoff is $p(\tilde{L}, x')b$. Because $x' > x$, we have $p(\tilde{L}, x') > p(\tilde{L}, x)$ and thus $p(\tilde{L}, x')r > p(\tilde{L}, x)b$, so x' 's flow payoff is higher than x 's when x and x' are members.

We then have $V_{x'}^x(L, \pi) > V_x(L, \pi)$ if x is in the organization while it experiments, at a set of times of positive measure with a positive probability. Because $V_{x'}(L, \pi) \geq V_{x'}^x(L, \pi)$, we have $V_{x'}(L, \pi) > V_x(L, \pi)$, as required. The continuity is proved similarly. ■

Corollary 1. *If the organization experiments forever on the equilibrium path, then $V_x(L, \pi) = V_{p(L, x)}(1, \pi)$. Moreover, $V_{p(L, x)}(1, \pi)$ is increasing in $p(L, x)$.*

Let $W_T(x)$ denote the continuation value of an agent with current belief x in

an equilibrium in which the organization stops experimenting after a length of time T .³⁶ Let $T^* = \operatorname{argmax}_T W_T(x)$ denote the optimal amount of time that an agent with prior x would want to experiment for if she was always in control of the organization.

Lemma 2. (i) $T \mapsto W_T(x)$ is differentiable for all $T \in (0, \infty)$ and right-differentiable at $T = 0$.

$$(ii) \quad W_0(x) = \frac{s}{\gamma} \text{ and } \left. \frac{\partial W_T(x)}{\partial T} \right|_{T=0} = \max\{xr, a\} - s + \frac{xr(r-s)}{\gamma}.$$

(iii) $T \mapsto W_T(x)$ is strictly increasing for $T \in [0, T^*]$ and strictly decreasing for $T > T^*$.

(iv) If $V(x) > \frac{s}{\gamma}$, then $W_T(x) > \frac{s}{\gamma}$ for all $T > 0$.

Proof of Lemma 2.

Fix $T_0 \geq 0$ and $\epsilon > 0$. We can write

$$\begin{aligned} W_{T_0+\epsilon}(x) - W_{T_0}(x) &= e^{-\gamma T_0} Q_{T_0}(x) (W_\epsilon(p_{T_0}(x)) - W_0(p_{T_0}(x))) \\ &= e^{-\gamma T_0} Q_{T_0}(x) \left(W_\epsilon(p_{T_0}(x)) - \frac{s}{\gamma} \right) \end{aligned}$$

where $Q_T(x)$ is the probability that there is no success up to time T , based on the prior x , and $p_T(x)$ is the posterior belief of an agent with prior x in this case. (See Lemma 8 for details.) Then

$$\begin{aligned} \left. \frac{\partial W_T(x)}{\partial T} \right|_{T=T_0} &= \lim_{\epsilon \searrow 0} \frac{W_{T_0+\epsilon}(x) - W_{T_0}(x)}{\epsilon} = \\ &= \lim_{\epsilon \searrow 0} e^{-\gamma T_0} Q_{T_0}(x) \frac{W_\epsilon(p_{T_0}(x)) - \frac{s}{\gamma}}{\epsilon} = e^{-\gamma T_0} Q_{T_0}(x) \left. \frac{\partial W_T(p_{T_0}(x))}{\partial T} \right|_{T=0} \end{aligned}$$

The case with $\epsilon < 0$ is analogous. Then it is enough to prove that $T \mapsto W_T(x)$ is right-differentiable at $T = 0$. This can be done by calculating $W_T(x)$ explicitly for small $T > 0$. We do this for $x > \frac{a}{r}$. In this case, for T sufficiently small, x is in the organization because the experiment with the risky policy is going to stop before she

³⁶This is defined for $T \in [0, \infty]$, where $W_\infty(x) = V(x)$.

wants to leave. Using that $1 - Q_T(x) = x(1 - e^{-rT})$,

$$\begin{aligned} W_T(x) &= \int_0^T xre^{-\gamma t} dt + \int_T^\infty e^{-\gamma t} (Q_T(x)s + (1 - Q_T(x))r) dt = \\ &= xr \frac{1 - e^{-\gamma T}}{\gamma} + \frac{e^{-\gamma T}}{\gamma} (s + x(1 - e^{-rT})(r - s)) \end{aligned}$$

This implies that $\left. \frac{\partial W_T(x)}{\partial T} \right|_{T=0} = xr - s + \frac{xr(r-s)}{\gamma}$.

Similarly, it can be shown that if $x \leq \frac{a}{r}$, then $\left. \frac{\partial W_T(x)}{\partial T} \right|_{T=0} = a - s + \frac{xr(r-s)}{\gamma}$. This proves (i) and (ii).

For (iii), note that, by our previous result, $\left. \frac{\partial W_T(x)}{\partial T} \right|_{T=T_0}$ is positive (negative) whenever $\left. \frac{\partial W_T(p_{T_0}(x))}{\partial T} \right|_{T=0}$ is positive (negative). In addition, it follows from our calculations that $y \mapsto \left. \frac{\partial W_T(y)}{\partial T} \right|_{T=0}$ is increasing and $T_0 \mapsto p_{T_0}(x)$ is decreasing. Moreover, for large T_0 , $p_{T_0}(x)$ is close to zero, so $\left. \frac{\partial W_T(p_{T_0}(x))}{\partial T} \right|_{T=0}$ is negative. It follows that $T \mapsto W_T(x)$ is single-peaked. If $\left. \frac{\partial W_T(x)}{\partial T} \right|_{T=0} > 0$, then the peak is the unique T^* satisfying $\left. \frac{\partial W_T(x)}{\partial T} \right|_{T=T^*} = 0$. If $\left. \frac{\partial W_T(x)}{\partial T} \right|_{T=0} \leq 0$, then $T^* = 0$.

Hence if $0 < T \leq T^*$, then $W_T(x) > W_0(x) = \frac{s}{\gamma}$ because in this case $T \mapsto W_T(x)$ is increasing by (iii), and if $T > T^*$, then $W_T(x) \geq \lim_{T \rightarrow \infty} W_T(x) = V(x)$ because in this case $T \mapsto W_T(x)$ is decreasing by (iii). This proves (iv). \blacksquare

Lemma 3. *Let $m(L)$ and $\tilde{m}(L)$ denote the median voter when the state variable is $(L, 1)$ and the density is f and \tilde{f} respectively. Suppose that \tilde{f} MLRP-dominates f . Then $\tilde{m}(L) \geq m(L)$ for all L .*

Proof of lemma 3.

Let $y(L)$ denote the indifferent agent given information L under either density (note that $y(L)$ is given by the condition $p(L, y(L)) = \frac{a}{r}$, which is independent of the density). By definition, we have $\int_{y(L)}^{m(L)} f(x) dx = \int_{m(L)}^1 f(x) dx$. Suppose that $\tilde{m}(L) < m(L)$. This is equivalent to

$$\int_{y(L)}^{m(L)} \frac{\tilde{f}(x)}{f(x)} f(x) dx = \int_{y(L)}^{m(L)} \tilde{f}(x) dx > \int_{m(L)}^1 \tilde{f}(x) dx = \int_{m(L)}^1 \frac{\tilde{f}(x)}{f(x)} f(x) dx$$

Since \tilde{f} MLRP-dominates f , $\frac{\tilde{f}(x)}{f(x)}$ is weakly increasing. Thus

$$\int_{y(L)}^{m(L)} \frac{\tilde{f}(m(L))}{f(m(L))} f(x) dx \geq \int_{y(L)}^{m(L)} \frac{\tilde{f}(x)}{f(x)} f(x) dx > \int_{m(L)}^1 \frac{\tilde{f}(x)}{f(x)} f(x) dx \geq \int_{m(L)}^1 \frac{\tilde{f}(m(L))}{f(m(L))} f(x) dx$$

which is a contradiction. ■

Proof of Proposition 1.

We will prove the following statement that encompasses Proposition 1 and the remark preceding it. If $V(p_t(m_t)) > \frac{s}{\gamma}$ for all t , there is an essentially unique equilibrium. In this equilibrium, the organization experiments forever. If $\inf_{t \geq 0} V(p_t(m_t)) < \frac{s}{\gamma}$, there is no equilibrium in which the organization experiments forever.

We first argue that if $\inf_{t \geq 0} V(p_t(m_t)) \geq \frac{s}{\gamma}$, then experimenting forever is an equilibrium if the organization is experimenting at time $t = 0$.

Consider the following strategy profile: $\alpha(L) = 1$ for all L and $\beta(x, L, \pi)$ is given by Condition (i). We can check that Conditions (i)-(iii) hold, so this is an equilibrium (in the class of equilibria that we restrict attention to).

Next, we argue that if $V(p_t(m_t)) > \frac{s}{\gamma}$ for all $t \geq 0$, then any equilibrium must be of this form. Let σ be an equilibrium. Clearly, the risky policy must always be used after a success.³⁷ By assumption, after a switch from the risky policy to the safe policy, the safe policy is used forever.

Now suppose for the sake of contradiction that σ is *not* the equilibrium with perpetual experimentation given above – that is, suppose that for some L we have $\alpha(L) \neq 1$. By Condition (ii), we must have $\frac{s}{\gamma} \geq V_{m(L,1)}(L)$.

Let T be the (possibly random) time until the policy first switches to 0, starting in state L . Then $V_{m(L,1)}(L) = E[W_T(p(L, m(L, 1)))]$, where $W_T(x)$ is as in Lemma 2 and the expectation is taken over T .

Recall that $V(p(L, m(L, 1))) > \frac{s}{\gamma}$ by assumption. By Lemma 2, this implies that $W_T(p(L, m(L, 1))) > \frac{s}{\gamma}$ for all $T > 0$. If $E[T] > 0$, it follows that $E[W_T(p(L, m(L, 1)))] > \frac{s}{\gamma}$, implying that $V_{m(L,1)}(L) > \frac{s}{\gamma}$, a contradiction. If $E[T] =$

³⁷This follows trivially from Conditions (ii) and (iii).

0, then $V_{m(L,1)}(L) = \frac{s}{\gamma}$ but, by the same argument, $W_\epsilon(p(L, m(L, 1))) > \frac{s}{\gamma}$ for all $\epsilon > 0$, so $\alpha(L) = 1$ by Condition (iii), a contradiction.

Finally, suppose that $\inf_{t \geq 0} V(p_t(m_t)) < \frac{s}{\gamma}$, so that $V(p_{t_0}(m_{t_0})) < \frac{s}{\gamma}$ for some t_0 . Then an equilibrium in which $\alpha(L) = 1$ for all L cannot exist, as by assumption we would have $\alpha(e^{rt_0}) = 0$ by Condition (ii), so experimentation would stop at time t_0 if unsuccessful. \blacksquare

Proposition 9. *The value function V in Section 5 satisfies the following:*

(i) *If f is non-decreasing, then*

$$\gamma \inf_{t \geq 0} V(p_t(m_t)) = \gamma V\left(\frac{2a}{r+a}\right) = \frac{2ra}{r+a} + \left(\frac{1}{2}\right)^{\frac{\gamma}{r}} \frac{a(r-a)}{r+a} \frac{r}{\gamma+r}$$

(ii) *Given $\omega > 0$, let $f_\omega(x)$ denote a density with support $[0, 1]$ such that $f_\omega(x) = (\omega + 1)(1 - x)^\omega$ for $x \in [0, 1]$. Let f be a density with support $[0, 1]$ that MLRP-dominates f_ω . Let $\lambda = \frac{1}{2^{\omega+1}}$. Then*

$$\gamma \inf_{t \geq 0} V(p_t(m_t)) \geq \gamma V\left(\frac{a}{\lambda r + (1-\lambda)a}\right) = \frac{ra}{\lambda r + (1-\lambda)a} + \lambda^{\frac{\gamma+r}{r}} \frac{a(r-a)}{\lambda r + (1-\lambda)a} \frac{r}{\gamma+r}$$

(iii) *Let f be any density with support $[0, 1]$. Then*

$$\gamma \inf_{t \geq 0} V(p_t(m_t)) \geq \gamma V\left(\frac{a}{r}\right) = a + \frac{a(r-a)}{\gamma+r}$$

Proof of Proposition 9.

We prove each inequality in three steps.

First, we show that the median posterior belief is uniformly bounded below for all t , with different bounds depending on the density. When f is uniform, we have $p_t(m_t) \geq \frac{2a}{r+a}$. For any $\omega > 0$, when $f = f_\omega$, we have $p_t(m_t) \geq \frac{a}{\lambda r + (1-\lambda)a}$ for $\lambda = \frac{1}{2^{\omega+1}}$. Finally, if f has full support, we have $p_t(m_t) \geq \frac{a}{r}$. The first two assertions follow from the following Claim:

Claim 9.1. *Suppose that the initial distribution of priors is f_ω for some $\omega \geq 0$. Then*

$$p_t(m_t) = \frac{a + (1 - \lambda)(r - a)e^{-rt}}{\lambda(r - a) + a + (1 - \lambda)(r - a)e^{-rt}}$$

Proof of Claim 9.1. The posterior belief of agent x at time t is given by $p_t(x) = \frac{xe^{-rt}}{xe^{-rt} + 1 - x}$. Using the fact that $p_t(y_t) = \frac{a}{r}$ for the marginal member y_t , we set $\frac{y_te^{-rt}}{y_te^{-rt} + 1 - y_t} = \frac{a}{r}$. Solving for y_t , we obtain

$$y_t = \frac{\frac{a}{r}}{\frac{a}{r} + (1 - \frac{a}{r})e^{-rt}} = \frac{a}{a + (r - a)e^{-rt}}$$

The median m_t must satisfy the condition $2 \int_{m_t}^1 f_\omega(x)dx = \int_{y_t}^1 f_\omega(x)dx$, so that $2(1 - m_t)^{\omega+1} = (1 - y_t)^{\omega+1}$. Hence $1 - m_t = \lambda(1 - y_t)$, which implies that

$$m_t = 1 - \lambda + \lambda y_t = 1 - \lambda + \lambda \frac{a}{a + (r - a)e^{-rt}} = \frac{a + (1 - \lambda)(r - a)e^{-rt}}{a + (r - a)e^{-rt}}$$

Substituting the above expression into the formula for $p_t(x)$, we obtain

$$p_t(m_t) = \frac{a + (1 - \lambda)(r - a)e^{-rt}}{\lambda(r - a) + a + (1 - \lambda)(r - a)e^{-rt}}$$

In particular, if $\omega = 0$, then f is uniform and we have $p_t(m_t) = \frac{2a + (r - a)e^{-rt}}{r + a + (r - a)e^{-rt}}$. ■

Indeed, it then follows that $p_t(m_t) \searrow \frac{2a}{r+a}$ as $t \rightarrow \infty$ when f is uniform, and $p_t(m_t) \searrow \frac{a}{\lambda r + (1 - \lambda)a}$ when $f = f_\omega$. The third claim is implied by the fact that $m_t \geq y_t$ and $p_t(y_t) = \frac{a}{r}$.

Second, we argue that these bounds hold not just for the aforementioned densities but also for any that dominate them in the MLRP sense. This follows from Lemma 3 and the fact that the function $x \mapsto p_t(x)$ is strictly increasing.

Third, we observe that, since $V(x)$ is strictly increasing and continuous in x (by Lemma 1), we have $\inf_{t \geq 0} V(p_t(m_t)) = V(\inf_{t \geq 0} p_t(m_t))$. Hence, to arrive at the bounds in the Proposition, it is enough to evaluate V at the appropriate beliefs.

To calculate $V\left(\frac{a}{\lambda r + (1 - \lambda)a}\right)$, we use the following two Claims:

Claim 9.2. *Let $t(x)$ denote the time it will take for an agent's posterior belief to go*

from x to $\frac{a}{r}$ (provided that no successes are observed during this time), at which time she would leave the organization. Then

$$V(x) = xr \frac{1}{\gamma} + (1-x)e^{-\gamma t(x)} \frac{a}{\gamma} - x(r-a) \frac{e^{-(\gamma+r)t(x)}}{\gamma+r}$$

Proof of Claim 9.2.

Let $P_t = x(1 - e^{-rt})$ denote the probability that an agent with prior belief x assigns to having a success by time t . Then

$$V(x) = x \int_0^{t(x)} r e^{-\gamma \tau} d\tau + \int_{t(x)}^{\infty} (P_{\tau} r + (1 - P_{\tau}) a) e^{-\gamma \tau} d\tau$$

The first term is the payoff from time 0 to time $t(x)$, when the agent stays in the organization. The second term is the payoff after time $t(x)$, when the agent leaves the organization and obtains the flow payoff a thereafter, unless the risky technology has had a success (in which case the agent returns to the organization and receives a guaranteed expected flow payoff r). We have

$$\begin{aligned} V(x) &= x \int_0^{t(x)} r e^{-\gamma \tau} d\tau + \int_{t(x)}^{\infty} a e^{-\gamma \tau} d\tau + \int_{t(x)}^{\infty} P_{\tau} (r - a) e^{-\gamma \tau} d\tau \\ &= xr \frac{1 - e^{-\gamma t(x)}}{\gamma} + e^{-\gamma t(x)} \frac{a}{\gamma} + x(r - a) \left(\frac{e^{-\gamma t(x)}}{\gamma} - \frac{e^{-(\gamma+r)t(x)}}{\gamma+r} \right) \\ &= xr \frac{1}{\gamma} + (1-x)e^{-\gamma t(x)} \frac{a}{\gamma} - x(r-a) \frac{e^{-(\gamma+r)t(x)}}{\gamma+r}. \end{aligned}$$

■

Claim 9.3. Let $t^y(x)$ denote the time it takes for an agent's posterior belief to go from x to y . Then

$$t^y(x) = -\frac{\ln\left(\frac{y}{1-y} \frac{1-x}{x}\right)}{r} \quad t(x) = -\frac{\ln\left(\frac{a(1-x)}{(r-a)x}\right)}{r}$$

If $x = \frac{2a}{r+a}$, then $e^{-rt(x)} = \frac{1}{2}$. If $x = \frac{a}{\lambda r + (1-\lambda)a}$, then $e^{-rt(x)} = \lambda$.

Proof of Claim 9.3.

We solve $p_t(x) = \frac{x e^{-rt}}{x e^{-rt} + 1 - x} = y$ for t . Then we obtain $e^{-rt^y(x)} = \frac{y}{1-y} \frac{1-x}{x}$ or,

equivalently, $t^y(x) = -\frac{\ln\left(\frac{y}{1-y} \frac{1-x}{x}\right)}{r}$.

In particular, $t(x) = t^{\frac{a}{r}}(x) = -\frac{\ln\left(\frac{a(1-x)}{(r-a)x}\right)}{r}$. Substituting $x = \frac{2a}{r+a}$ into $e^{-rt(x)} = \frac{a(1-x)}{(r-a)x}$ and simplifying, we obtain $e^{-rt} = \frac{1}{2}$. Substituting $x = \frac{a}{\lambda r + (1-\lambda)a}$ into $e^{-rt(x)} = \frac{a(1-x)}{(r-a)x}$ and simplifying, we obtain $e^{-rt(x)} = \lambda$. \blacksquare

Thus, taking $x = \frac{a}{\lambda r + (1-\lambda)a}$, we have $e^{-rt(x)} = \lambda$ and $e^{-\gamma t(x)} = \lambda^{\frac{\gamma}{r}}$. Substituting this value of x into the formula for $V(x)$ from Claim 9.2, we obtain

$$\gamma V\left(\frac{a}{\lambda r + (1-\lambda)a}\right) = \frac{ra}{\lambda r + (1-\lambda)a} + \frac{(r-a)a}{\lambda r + (1-\lambda)a} \frac{r}{\gamma + r} \lambda^{\frac{\gamma+r}{r}}$$

In particular, for $\omega = 0$, this becomes

$$\gamma V\left(\frac{2a}{r+a}\right) = \frac{2ra}{r+a} + \left(\frac{1}{2}\right)^{\frac{\gamma}{r}} \frac{(r-a)a}{r+a} \frac{r}{\gamma + r}$$

On the other hand, for $x = \frac{a}{r}$, we have $t(x) = 0$. Substituting this in, we obtain

$$\gamma V\left(\frac{a}{r}\right) = a + \frac{r-a}{r}a - \frac{a}{r}(r-a) \frac{\gamma}{\gamma + r} = a + \frac{(r-a)a}{\gamma + r}$$

An additional argument is required to show that the bound is tight in part (i).

Take f to be any non-decreasing density. Let \tilde{m}_t denote the median at time t under f , and let m_t denote the median at time t under the uniform density. It is sufficient to show that the asymptotic posterior of the median is $\frac{2a}{r+a}$ under f , that is, that $\lim_{t \rightarrow \infty} p_t(\tilde{m}_t) = \lim_{t \rightarrow \infty} p_t(m_t) = \frac{2a}{r+a}$.

Claim 9.4. *Let $m(L)$ and $\tilde{m}(L)$ denote the median voters when the state variable is L under the uniform density and a non-decreasing density f respectively. Suppose that $y(L) \rightarrow 1$ as $L \rightarrow \infty$. Then $\frac{1-\tilde{m}(L)}{1-m(L)} \rightarrow 1$ as $L \rightarrow \infty$.*

Proof of Claim 9.4.

Given a state variable L and the marginal member $y(L)$ corresponding to it, let $f_{0L} = f(y(L))$ and $f_1 = f(1)$. By Lemma 3, we have $m(L) \leq \tilde{m}(L) \leq \hat{m}(L)$ where $\hat{m}(L)$ is the median corresponding to a density \hat{f} such that $\hat{f}(x) = f_{0L}$ for $x \in [y(L), \hat{m}(L)]$ and $\hat{f}(x) = f_1$ for $x \in [\hat{m}(L), 1]$.

By construction, because $\hat{m}(L)$ is the median, we have $f_{0L}(\hat{m}(L) - y(L)) = f_1(1 - \hat{m}(L))$, so $\hat{m}(L) = \frac{f_{0L}y(L) + f_1}{f_{0L} + f_1}$. Thus $1 - \hat{m}(L) = \frac{f_{0L}(1 - y(L))}{f_{0L} + f_1}$ and, because $m(L) = \frac{y(L) + 1}{2}$ so that $1 - m(L) = \frac{1 - y(L)}{2}$, we have $\frac{1 - \hat{m}(L)}{1 - m(L)} = \frac{2f_{0L}}{f_{0L} + f_1}$.

Since f is increasing, using the fact that $f(x) \rightarrow \sup_{y \in [0,1]} f(y)$ as $x \rightarrow 1$, we find that $f(x) \rightarrow f(1)$ as $x \rightarrow 1$. Then, as $t \rightarrow \infty$, we have $y(L) \rightarrow 1$, $f_{0L} = f(y(L)) \rightarrow f_1$ and $\frac{1 - \hat{m}(L)}{1 - m(L)} \rightarrow 1$. \blacksquare

Claim 9.5. *Let x_t, \tilde{x}_t be two time-indexed sequences of agents such that $x_t \leq \tilde{x}_t$ for all t and $x_t \rightarrow 1$ as $t \rightarrow \infty$. If $\frac{1 - x_t}{1 - \tilde{x}_t} \rightarrow 1$, then $\frac{p_t(\tilde{x}_t)}{p_t(x_t)} \rightarrow 1$.*

Proof of Claim 9.5.

Using the formula for the posterior beliefs, we have

$$\frac{p_t(\tilde{x}_t)}{p_t(x_t)} = \frac{\tilde{x}_t}{\tilde{x}_t + (1 - \tilde{x}_t)L_t} \frac{x_t + (1 - x_t)L_t}{x_t} = \frac{\tilde{x}_t x_t + (1 - x_t)L_t}{x_t \tilde{x}_t + (1 - \tilde{x}_t)L_t}.$$

Since $x_t \rightarrow 1$ and $\tilde{x}_t \geq x_t$ for all t , $\tilde{x}_t \rightarrow 1$, whence $\frac{\tilde{x}_t}{x_t} \rightarrow 1$. In addition, since $\frac{1 - x_t}{1 - \tilde{x}_t} \rightarrow 1$, $\frac{(1 - x_t)L_t}{(1 - \tilde{x}_t)L_t} \rightarrow 1$. As a result, for all t ,

$$\min \left\{ \frac{x_t}{\tilde{x}_t}, \frac{(1 - x_t)L_t}{(1 - \tilde{x}_t)L_t} \right\} \leq \frac{x_t + (1 - x_t)L_t}{\tilde{x}_t + (1 - \tilde{x}_t)L_t} \leq \max \left\{ \frac{x_t}{\tilde{x}_t}, \frac{(1 - x_t)L_t}{(1 - \tilde{x}_t)L_t} \right\}$$

so $\frac{x_t + (1 - x_t)L_t}{\tilde{x}_t + (1 - \tilde{x}_t)L_t} \rightarrow 1$, which concludes the proof. \blacksquare

Claim 9.4 shows that $\frac{1 - \tilde{m}_t}{1 - m_t} \rightarrow 1$. Note that we have $\tilde{m}_t \geq m_t$ for all t by Lemma 3 and $m_t \rightarrow 1$ as $t \rightarrow \infty$. Then Claim 9.5 applies. Claim 9.5 applied to the sequences \tilde{m}_t and m_t guarantees that the ratio $\frac{p_t(\tilde{m}_t)}{p_t(m_t)}$ converges to 1. \blacksquare

Corollary 9.1. *There exist parameters such that $\inf V(p_t(m_t)) \geq \frac{s}{\gamma}$.*

Proof of Corollary 9.1. We can, in fact, prove that, holding the other parameters constant, there is $a^* < s$ such that, if $a \in [a^*, s)$, then $\inf_t V(p_t(m_t)) \geq \frac{s}{\gamma}$. This follows from Proposition 9(iii): indeed, we can take a^* to be such that $a^* + \frac{a^*(r - a^*)}{\gamma + r} = s$. \blacksquare

Proof of Proposition 2 and the assertions following it. First note that when $f = f_\omega$, the fact that $\inf_t V(p_t(m_t))$ is increasing in r follows from the formula for

$V\left(\frac{a}{\lambda r + (1-\lambda)a}\right)$ given in the proof of Proposition 9. Lemma 3 implies that an MLRP-increase in f increases $\inf_t V(p_t(m_t))$. As for changes in a , note that an increase in a clearly increases $V(x)$ for each x , and also increases y_t , and hence m_t , for each t . Finally, note that a decrease in γ is equivalent to an increase in the learning rate.³⁸ This leaves y_t and m_t unchanged but increases $V(x)$ for all x , as agents have strictly more information to base their entry and exit decisions on. ■

Next we will characterize equilibria without perpetual experimentation. To do this, define a *stopping function* $\varphi : [0, \infty) \rightarrow [0, \infty]$ as follows. For each $t \geq 0$, $\varphi(t) \geq t$ is such that m_t is indifferent about switching to the safe policy at time t if she expects a continuation where experimentation will stop at time $\varphi(t)$ should she fail to stop at t . If the agent never wants to experiment regardless of the expected continuation, then $\varphi(t) = t$, while if she always does, $\varphi(t) = \infty$.

Proposition 10. *Any pure strategy equilibrium σ in which the organization does not experiment forever is given by a sequence of stopping times $t_0(\sigma) \leq t_1(\sigma) \leq t_2(\sigma) \leq \dots$ such that $t_n(\sigma) = \varphi(t_{n-1}(\sigma))$ for all $n > 0$ and $t_0(\sigma) \leq \varphi(0)$.*

There exists $t \in [0, \varphi(0)]$ for which $(t, \varphi(t), \varphi(\varphi(t)), \dots)$ constitutes an equilibrium. Moreover, if φ is weakly increasing, then $(t, \varphi(t), \varphi(\varphi(t)), \dots)$ constitutes an equilibrium for all $t \in [0, \varphi(0)]$.

Proof of Proposition 10. We first argue that the stopping function φ is well-defined. Let t be the current time and let t^* be the time at which m_t would choose to stop experimenting if she had complete control over the policy. Recall the definition of $W_{T-t}(x)$ from Lemma 2: $W_{T-t}(x)$ is the value function starting at time t of an agent with belief x at time t given a continuation equilibrium path on which the organization experiments until T and then switches to the safe technology. Then, equivalently, $t^* = \operatorname{argmax}_T W_{T-t}(x)$.

There are three cases. If $t^* = t$, then $\varphi(t) = t$. If $t^* > t$, that is, if x wants to experiment for a positive amount of time, and $V(p_t(m_t)) < \frac{s}{\gamma}$, then $W_{T-t}(p_t(m_t))$ is strictly increasing in T for $T \in [t, t^*]$ and strictly decreasing in T for $T > t^*$, as shown in Lemma 2, and there is a unique $\varphi(t) > t^*$ for which $W_{\varphi(t)-t}(p_t(m_t)) = \frac{s}{\gamma}$. Finally,

³⁸In other words, it is equivalent to increasing the success rate of the good risky policy to rq , for $q > 1$, and lowering the payoff per success to $\frac{1}{q}$.

if $t^* > t$ and $V(p_t(m_t)) \geq \frac{s}{\gamma}$, then $\varphi(t) = \infty$.

Next, note that φ is continuous. If $\varphi(t_0) \in (t_0, \infty)$, then for t in a neighborhood of t_0 , $\varphi(t)$ is defined by the condition $W_{\varphi(t)-t}(p_t(m_t)) = \frac{s}{\gamma}$, where $p_t(m_t)$ is continuous in t , and $W_T(x)$ is differentiable in (T, x) at $(T, x) = (\varphi(t), p_t(m_t))$ and strictly decreasing in T ,³⁹ so the continuity of φ follows from the Implicit Function Theorem. The proofs for the cases when $\varphi(t_0) = 0$ or $\varphi(t_0) = \infty$ are similar.

Consider a pure strategy equilibrium σ in which the organization does not experiment forever on the equilibrium path. Let $t_0(\sigma)$ be the time at which experimentation stops on the equilibrium path. Clearly, we have $t_0(\sigma) \leq \varphi(0)$, as otherwise m_0 would switch to the safe policy at time 0. As before, if a success occurs or if the organization switches to the safe policy, everyone joins the organization permanently.

Consider what happens at time $t_0(\sigma)$ if $m_{t_0(\sigma)}$ deviates and continues experimenting. Suppose first that $\varphi(t_0(\sigma)) \in (t_0(\sigma), \infty)$. In a pure strategy equilibrium, there must be a time $t_1(\sigma) \geq t_0(\sigma)$ for which experimentation stops in this continuation, and it must satisfy $t_1(\sigma) = \varphi(t_0(\sigma))$. To see why, suppose that $t_1(\sigma) > \varphi(t_0(\sigma))$. In this case, for $\epsilon > 0$ sufficiently small, $m_{t_0(\sigma)+\epsilon}$ would strictly prefer to stop experimenting, a contradiction. On the other hand, if $t_1(\sigma) < \varphi(t_0(\sigma))$, then $m_{\varphi(t_0(\sigma))}$ would strictly prefer to deviate from the equilibrium path and not stop.

Next, suppose that $\varphi(t_0(\sigma)) = \infty$, that is, $m_{t_0(\sigma)}$ weakly prefers to continue experimenting regardless of the continuation. Then it must be that $t_1(\sigma) = \infty$ and $V(p_{t_0(\sigma)}(m_{t_0(\sigma)})) = \frac{s}{\gamma}$, and in this case we must still have $t_1(\sigma) = \varphi(t_0(\sigma))$.

Now suppose that $\varphi(t_0(\sigma)) = t_0(\sigma)$, that is, $m_{t_0(\sigma)}$ weakly prefers to stop regardless of the continuation. In this case, the implied sequence of points is $(t_0(\sigma), t_0(\sigma), \dots)$. This does not fully describe the equilibrium, as it does not specify what happens conditional on not experimentation by $t_0(\sigma)$, but still provides enough information to characterize the equilibrium path fully, as in any equilibrium experimentation must stop at $t_0(\sigma)$.

Next, we show that if φ is increasing and $t \in [0, \varphi(0)]$, then $(t, \varphi(t), \varphi(\varphi(t)), \dots)$ constitutes an equilibrium. Our construction already shows that $m_{t_n(\sigma)}$ is indifferent between switching to the safe policy at time $t_n(\sigma)$ and continuing to experiment.

³⁹The fact that $W_T(x)$ is differentiable in T at $(T, x) = (\varphi(t), p_t(m_t))$ and strictly decreasing in T is implied by Lemma 2.

To finish the proof, we have to show that for t not in the sequence of the stopping times, m_t weakly prefers to continue experimenting. Fix $t \in (t_n(\sigma), t_{n+1}(\sigma))$. Since $t > t_n(\sigma)$ and φ is increasing, we have $\varphi(t) \geq \varphi(t_n(\sigma)) = t_{n+1}(\sigma)$. Hence the definition of $\varphi(t)$ and the fact that $T \mapsto W_T(x)$ is single-peaked by Lemma 2 imply that $W_{t_{n+1}(\sigma)-t}(p_t(m_t)) \geq \frac{s}{\gamma}$, and Conditions (ii) and (iii) imply that m_t weakly prefers to continue experimenting.

Finally, we show that even if φ is not increasing, this construction yields an equilibrium for at least one value of $t \in [0, \varphi(0)]$. Note that our construction fails if and only if there is $t \in (t_k(\sigma), t_{k+1}(\sigma))$ for which $\varphi(t) < t_{k+1}(\sigma)$. Motivated by this, we say t is *valid* if $\varphi(t) = \inf_{t' \geq t} \varphi(t')$, and say t is *n-valid* if $t, \varphi(t), \dots, \varphi^{(n-1)}(t)$ are all valid. Let $A_0 = [0, \varphi(0)]$ and, for $n \geq 1$, let $A_n = \{t \in [0, \varphi(0)] : t \text{ is } n\text{-valid}\}$.

Suppose that $\varphi(t) > t$ and $\varphi(t) < \infty$ for all t . Clearly, $A_n \supseteq A_{n+1}$ for all n , and the continuity of φ implies that A_n is closed for all n . In addition, A_n must be non-empty for all n by the following argument. Take $t_0 = t$ and define a sequence $\{t_0, t_{-1}, t_{-2}, \dots, t_{-k}\}$ by $t_{-i} = \max\{\varphi^{-1}(t_{-i+1})\}$ for $i \leq -1$, and $t_{-k} \in [0, \varphi(0)]$. By construction, $t_{-k} \in A_0$ is k -valid, and, because $\varphi(t) < \infty$ for all t , if we choose t large enough, we can make k arbitrarily large.⁴⁰ Then $A = \bigcap_0^\infty A_n \neq \emptyset$ by Cantor's intersection theorem, and any sequence $(t, \varphi(t), \dots)$ with $t \in A$ yields an equilibrium. The same argument goes through if $\varphi(t) = \infty$ for some values of t but there are arbitrarily large t for which $\varphi(t) < \infty$.

If $\varphi(t) = t$ for some t , let $\bar{t} = \min\{t \geq 0 : \varphi(t) = t\}$. If there is $\epsilon > 0$ such that $\varphi(t) \geq \varphi(\bar{t})$ for all $t \in (\bar{t} - \epsilon, \bar{t})$, then we can find a finite equilibrium sequence of stopping times by setting $t_0 = \bar{t}$ and using the construction in the previous paragraph. If there is no such ϵ , then the previous argument works.⁴¹ The only difference is that, to show the non-emptiness of A_n , we take $t \rightarrow \bar{t}$ instead of making t arbitrarily large.

If $\varphi(t) > t$ for all t and there is \tilde{t} for which $\varphi(t) = \infty$ for all $t \geq \tilde{t}$, without loss of generality, take \tilde{t} to be minimal (that is, let $\tilde{t} = \min\{t \geq 0 : \varphi(t) = \infty\}$). Then we can find a finite sequence of stopping times compatible with equilibrium by taking

⁴⁰Under the assumption that $\varphi(t) < \infty$ for all t , since φ is continuous, the image of φ^l restricted to the set $[0, \varphi(0)]$ is compact and hence bounded for all l . Thus for any t larger than the supremum of this image, k must be larger than l .

⁴¹If there is $\epsilon > 0$ with the aforementioned property, then $\varphi^{-1}(\bar{t})$ is strictly lower than \bar{t} and reaching $[0, \varphi(0)]$ takes finitely many steps. If there is no such ϵ , then $\varphi^{-1}(\bar{t}) = \bar{t}$ and there exists a sequence converging to \bar{t} .

$t_0 = \tilde{t}$, assuming that m_{t_0} stops at t_0 and using the above construction. ■

Proof of Proposition 3. Part 1 follows from the last claim in Proposition 10: indeed, it holds whenever the parameters are such that τ is increasing (note also that $\hat{T} = \tau(0)$). Part 2 follows from the same proof as Proposition 1. In addition, the Remark preceding Proposition 3 follows from the first claim in Proposition 10. ■

A.2 A Model of Bad News

Lemma 4. *In a model of bad news, the value function of an agent with prior x who is in the organization and expects the organization to continue forever unless a failure is observed is*

$$V(x) = (xr + (1 - x)s)\frac{1}{\gamma} - (1 - x)s\frac{1}{\gamma + r}$$

Proof of lemma 4. Note that an agent receives an expected flow payoff of r only if the technology is good and the organization has not switched to the safe technology upon observing a failure. Because a good technology cannot experience a failure, as long as experimentation continues, an agent with posterior belief x receives an expected flow payoff of r with probability x .

Let $P_t = x + (1 - x)e^{-rt}$ denote the probability that an agent with prior belief x assigns to not having a failure by time t . Then

$$\begin{aligned} V(x) &= \int_0^\infty (xr + (1 - P_\tau)s) e^{-\gamma\tau} d\tau = \int_0^\infty (xr + (1 - x)(1 - e^{-r\tau})s) e^{-\gamma\tau} d\tau \\ &= (xr + (1 - x)s)\frac{1}{\gamma} - (1 - x)s\frac{1}{\gamma + r} \end{aligned}$$

■

Assumption 1. *The parameters r, s, a, γ, f are such that for all $t' > t$, $\frac{\partial}{\partial t} W_{t'-t}(p_t(m_t)) \neq 0$ whenever $W_{t'-t}(p_t(m_t)) = \frac{s}{\gamma}$.*

Assumption 1 guarantees that the agents' value functions are well-behaved: that is, for each t' , the function $t \mapsto W_{t'-t}(p_t(m_t))$ crosses the threshold $\frac{s}{\gamma}$ finitely many times, and is never tangent to it. Under this assumption, Proposition 11 characterizes the equilibrium in the bad news model.

Proposition 11. *Under Assumption 1, there is a unique equilibrium. The equilibrium can be described by a finite, possibly empty set of intervals $I_0 = [t_0, t_1]$, $I_1 = [t_2, t_3]$, \dots , I_n such that $t_0 < t_1 < t_2 < \dots$ as follows: conditional on the risky policy having been used during $[0, t]$ with no failures, the median m_t switches to the safe policy at time t if and only if $t \in I_k$ for some k .*

Proof of Proposition 11. We first argue that there exists T such that for all $t \geq T$, if no failures have been observed during $[0, t]$, then $V(p_t(m_t)) > \frac{s}{\gamma}$ and $p_t(m_t)r > s$. Note that, because in a model of bad news agents do not leave the organization, we have $\liminf_{t \rightarrow \infty} m_t > 0$. Moreover, $\lim_{t \rightarrow \infty} e^{-rt} = 0$. This implies that $\lim_{t \rightarrow \infty} p_t(m_t) = \lim_{t \rightarrow \infty} \frac{m_t}{m_t + e^{-rt}(1-m_t)} = 1$, so $\lim_{t \rightarrow \infty} p_t(m_t)r = r > s$. Provided that no failures have been observed during $[0, t]$, we have $\lim_{t \rightarrow \infty} V(p_t(m_t)) = V(1)$ because V is continuous, and $V(1) = \frac{r}{\gamma} > \frac{s}{\gamma}$. Next, we argue that these agents will always experiment.

Claim 11.1. *If $p_t(m_t)r > s$, then in any equilibrium m_t continues experimenting.*

Proof of claim 11.1. Suppose for the sake of contradiction that this is not the case. Let $t + t_+$ denote the first time after t when the equilibrium prescribes a switch to the safe policy.⁴² Because experimentation will stop after a period of length t_+ , the payoff to m_t from experimenting is

$$\begin{aligned} W_{t_+}(p_t(m_t)) &= \int_0^{t_+} (p_t(m_t)r + (1 - P_\tau)s)e^{-\gamma\tau} d\tau + \int_{t_+}^{\infty} se^{-\gamma\tau} d\tau \\ &\geq \int_0^{t_+} p_t(m_t)re^{-\gamma\tau} d\tau + \int_{t_+}^{\infty} se^{-\gamma\tau} d\tau = p_t(m_t)r \frac{1 - e^{-\gamma t_+}}{\gamma} + s \frac{e^{-\gamma t_+}}{\gamma} \end{aligned}$$

The payoff to experimentation is $\frac{s}{\gamma}$. Then, since $p_t(m_t)r > s$ by assumption, we have $\frac{1 - e^{-\gamma t_+}}{\gamma} p_t(m_t)r + \frac{e^{-\gamma t_+}}{\gamma} s > \frac{s}{\gamma}$, so m_t strictly prefers to continue experimenting, a contradiction. ■

Our results so far are already enough to deal with one important case. If $V(p_t(m_t)) > \frac{s}{\gamma}$ for all t , then the organization experiments forever. The reason is as follows. For $t \geq T$, all pivotal agents m_t continue experimenting by Claim 11.1. Let $\mathcal{T} \subseteq [0, T)$ be the set of times for which the pivotal agent at that time stops

⁴²We write the argument assuming that $t_+ > 0$. If $t_+ = 0$, the proof follows a similar argument leveraging Condition (iii).

experimenting in equilibrium, and assume \mathcal{T} is nonempty. Let $t^* = \sup \mathcal{T}$. If $t^* \in \mathcal{T}$, then m_{t^*} stops experimenting even though $V(p_t(m_t)) > \frac{s}{\gamma}$ and m_{t^*} gets perpetual experimentation by continuing, a contradiction. If $t^* \notin \mathcal{T}$, a similar argument can be made leveraging Condition (iii).

Suppose then that there exists $t \leq T$ such that $V(p_t(m_t)) < \frac{s}{\gamma}$.

Claim 11.2. *Suppose that on the equilibrium path, the organization continues experimenting for time t_+ unless a failure occurs and then switches to the safe policy. Then the value function of an agent with prior x in this equilibrium is given by*

$$(xr + (1-x)s) \frac{1 - e^{-\gamma t_+}}{\gamma} - (1-x)s \frac{1 - e^{-(\gamma+r)t_+}}{\gamma+r} + e^{-\gamma t_+} \frac{s}{\gamma}$$

Proof of claim 11.2. Because experimentation ends after a period of length t_+ , the value function of agent x is given by

$$\begin{aligned} & \int_0^{t_+} (xr + (1 - P_\tau)s) e^{-\gamma \tau} d\tau + \int_{t_+}^{\infty} s e^{-\gamma \tau} d\tau \\ &= (xr + (1-x)s) \frac{1 - e^{-\gamma t_+}}{\gamma} - (1-x)s \frac{1 - e^{-(\gamma+r)t_+}}{\gamma+r} + e^{-\gamma t_+} \frac{s}{\gamma} \end{aligned}$$

■

Claim 11.3. *Suppose that in some equilibrium m_{t_0} stops experimenting. If for all $t \in [\underline{t}, t_0)$ we have $p_t(m_t)r < s$, then for all $t \in [\underline{t}, t_0)$, m_t stops experimenting.*

Proof of claim 11.3. Suppose for the sake of contradiction that this is not the case. Then there exists a non-empty subset $B \subseteq [\underline{t}, t_0)$ such that for all $t \in B$, m_t continues experimenting.

There are two cases. In the first case, B has a non-empty interior. In this case, for all $\epsilon > 0$ small, there must exist $\tau \in [\underline{t}, t_0)$ such that, starting at time τ , experimentation continues up to time $\tau + \epsilon$ and then stops.⁴³

By claim 11.2, the payoff to m_τ from continuing experimentation is $W_\epsilon(p_\tau(m_\tau)) = (p_\tau(m_\tau)r + (1 - p_\tau(m_\tau))s) \frac{1 - e^{-\gamma \epsilon}}{\gamma} - (1 - p_\tau(m_\tau))s \frac{1 - e^{-(\gamma+r)\epsilon}}{\gamma+r} + e^{-\gamma \epsilon} \frac{s}{\gamma}$, which is of the form $\frac{s}{\gamma} + (p_\tau(m_\tau)r - s)\epsilon + \mathcal{O}(\epsilon^2)$. The payoff to experimentation is $\frac{s}{\gamma}$. Then, since

⁴³To find such τ , let \tilde{t} be in the interior of B , and let $\tilde{t} = \inf\{t \geq \tilde{t} : t \notin B\}$. Then $\tau = \tilde{t} - \epsilon$ works for all $\epsilon > 0$ small enough.

$p_\tau(m_\tau)r < s$ by assumption, for ϵ small enough m_τ strictly prefers to stop experimenting, a contradiction.

In the second case, the interior of B is empty. In this case, the proof follows a similar argument leveraging Condition (iii). \blacksquare

Let $t_{2n+1} = \sup \left\{ t : V(p_t(m_t)) < \frac{s}{\gamma} \right\}$ denote the largest time for which the median stops experimenting.

Let $T_1 = \{t : p_t(m_t)r \leq s\}$ and $T_2 = \{t : p_t(m_t)r > s\}$. Our genericity assumption (Assumption 1) implies that T_1 and T_2 are finite collections of intervals. Enumerate the intervals such that $T_1 = \cup_{i=0}^n [\underline{t}_i, \bar{t}_i]$.

Suppose first that $p_t(m_t)r \leq s$ for all $t < t_{2n+1}$. In this case, by claim 11.3, for all $t \leq t_{2n+1}$, m_t stops experimentation. Then we set $n = 0$, $t_0 = 0$ and $I_0 = [t_0, t_1]$.

Suppose next that there exists $t < t_{2n+1}$ such that $p_t(m_t)r > s$. Set $t_{2n} = \sup \{t < t_{2n+1} : p_t(m_t)r > s\}$. Note that, because F admits a continuous density, $t \mapsto p_t(m_t)$ is continuous, which implies that we must have $p_{t_{2n}}(m_{t_{2n}})r - s = 0$. Then claim 11.3 implies that for all $t \in [t_{2n}, t_{2n+1}]$, m_t stops experimentation. Note also that $t_{2n} < t_{2n+1}$ as $s = \gamma V(p_{t_{2n+1}}(m_{t_{2n+1}})) > p_{t_{2n+1}}(m_{t_{2n+1}})r$.

Let us conjecture a continuation equilibrium path on which, starting at t , the organization experiments until t_{2n} . Recall that $W_{t_{2n}-t}(x)$ denotes the value function of an agent with belief x (at time t) given this continuation equilibrium path. We then let $t_{2n-1} = \sup \left\{ t < t_{2n} : W_{t_{2n}-t}(p_t(m_t)) \leq \frac{s}{\gamma} \right\}$.

Note that, because, by construction, for $t \in (t_{2n-1}, t_{2n})$ we have $W_{t_{2n}-t}(p_t(m_t)) > \frac{s}{\gamma}$, the median m_t continues experimentation for all $t \in (t_{2n-1}, t_{2n})$.

Since F admits a continuous density, $t \mapsto W_{t_{2n}-t}(p_t(m_t))$ is continuous, which implies that we must have $t_{2n-1} = \max \left\{ t < t_{2n} : W_{t_{2n}-t}(p_t(m_t)) \leq \frac{s}{\gamma} \right\}$. Note that it is then consistent with equilibrium for the median $m_{t_{2n}}$ to stop experimenting.

Now note that if $W_{t_{2n}-t_{2n-1}}(p_{t_{2n-1}}(m_{t_{2n-1}})) = \frac{s}{\gamma}$, then $p_{t_{2n-1}}(m_{t_{2n-1}})r < s$. By continuity, there exists an interval $[\underline{t}_i, \bar{t}_i]$ in T_1 such that $t_{2n-1} \in [\underline{t}_i, \bar{t}_i]$ (and \underline{t}_i satisfies $\underline{t}_i = \min \{t < t_{2n-1} : p_t(m_t)r \leq s\}$).

Set $t_{2n-2} = \underline{t}_i$. Because $p_t(m_t)r \leq s$ for all $t \in [t_{2n-2}, t_{2n-1}]$, claim 11.1 implies that, for all $t \in [t_{2n-2}, t_{2n-1}]$, m_t stops experimenting.

We then proceed inductively in the above manner, finding the largest t strictly less than t_{2n-2} such that $W_{t_{2n-2}-t}(p_t(m_t)) \leq \frac{s}{\gamma}$. Because T_1 is finite collection of intervals, the induction terminates in a finite number of steps.

The equilibrium is generically unique for the following reason. Under Assumption 1, each t_{2k+1} satisfies not only $W_{t_{2k+2}}(p_{t_{2k+1}}(m_{t_{2k+1}})) = \frac{s}{\gamma}$ but also $\frac{\partial}{\partial t} W_{t_{2k+2}-t}(p_t(m_t))|_{t=t_{2k+1}} > 0$, that is, $W_{t_{2k+2}-t}(p_t(m_t)) < \frac{s}{\gamma}$ for all $t < t_{2k+1}$ close enough to t_{2k+1} . Thus, even if we allow $m_{t_{2k+1}}$ to continue experimenting, all agents in $(t_{2k+1} - \epsilon, t_{2k+1})$ must stop as they strictly prefer to do so. Likewise, each t_{2k} satisfies not only $p_{t_{2k}}(m_{t_{2k}})b - s = 0$ but also $\frac{\partial}{\partial t} p_t(m_t)|_{t=t_{2k}} < 0$, that is, $p_t(m_t)r - s > 0$ for all $t < t_{2k}$ close enough to t_{2k} . Thus, even if we allow $m_{t_{2k}}$ to stop experimenting, all agents in $(t_{2k} - \epsilon, t_{2k})$ must stop as they strictly prefer to do so. ■

Proof of Proposition 4. Part 1 is proved as part of Proposition 11. Part 2 follows from the characterization given in Proposition 11, in particular, from the observation that $t_{2n} < t_{2n+1}$. ■

A.3 A Model of Imperfectly Informative (Good) News

Proposition 12. *If $V_{m(L)}(L) > \frac{s}{\gamma}$ for all L , then there is a unique equilibrium. In it, the organization experiments forever. Moreover, if f is non-decreasing, then $V_{m(L)}(L) \geq V\left(\frac{2(a-c)}{(r-c)+(a-c)}\right) \geq \frac{1}{\gamma} \frac{(r-c)a+(a-c)r}{(r-c)+(a-c)}$, so there exist parameter values such that $V_{m(L)}(L) > \frac{s}{\gamma}$ for all L .*

Proof of Proposition 12. The proof is similar to the proof for the baseline model (Propositions 1 and 9). If $V_{m(L)}(L) > \frac{s}{\gamma}$ for all L , perpetual experimentation is clearly an equilibrium, as each pivotal agent $m(L)$ has a choice between $V_{m(L)}(L)$ and $\frac{s}{\gamma}$, and strictly prefers the former. The equilibrium is unique by the following argument. Suppose for the sake of contradiction that there is another equilibrium in which experimentation stops whenever $L \in \mathcal{L} \neq \emptyset$. Let $W_{\mathcal{L}}(x)$ denote the continuation utility of an agent with current belief x when she expects the organization to stop whenever $L \in \mathcal{L}$.

For L close enough to 0, it can be shown that pivotal agents will prefer to experiment no matter what equilibrium continuation they expect. That is, $W_{\mathcal{L}}(x) \geq$

$\frac{s}{\gamma}$ for all \mathcal{L} and x close enough to 1. In other words, there is $L_0 > 0$ such that $\mathcal{L} \subseteq (L_0, +\infty)$.

Let $L_1 = \inf \mathcal{L}$. It can be shown that, because $m(L_1)$ would rather experiment forever than not at all, she would also prefer to experiment until L hits \mathcal{L} . That is, if $V_{m(L_1)}(L_1) > \frac{s}{\gamma}$ then $W_{\mathcal{L}}(p(L_1, m(L_1))) > \frac{s}{\gamma}$. To see why, suppose that $W_{\mathcal{L}}(p(L_1, m(L_1))) \leq \frac{s}{\gamma}$. Clearly, this implies $W_{\mathcal{L}}(p(L, m(L_1))) < \frac{s}{\gamma}$ for any $L > L_1$. Note that changing the set of states from \mathcal{L} to \emptyset changes the payoff $m(L_1)$ gets from some continuations—namely, continuation starting at states $L > L_1$ —from $\frac{s}{\gamma}$ to objects of the form $W_{\mathcal{L}}(p(L, m(L_1)))$ for $L > L_1$. Hence we must have $V_{m(L_1)}(L_1) < \frac{s}{\gamma}$, a contradiction.

This proves the first statement. Next, we provide an explicit bound on V when f is non-decreasing.

Claim 12.1. *If the organization is experimenting at time t , then an agent with prior belief x is in the organization at time t if and only if $L(k, t) \leq \frac{x(r-a)}{(1-x)(a-c)}$.*

Proof of Claim 12.1. Because agents make their membership decisions based on expected flow payoffs, agent x is a member at time t if and only if $p(L, x)r + (1 - p(L, x))c \geq a$, that is, if $p(L, x) \geq \frac{a-c}{r-c}$. Since $p(L, x) = \frac{x}{x+(1-x)L(k, t)}$, this is equivalent to $L(k, t) \leq \frac{x(r-a)}{(1-x)(a-c)}$. ■

Claim 12.2. *If f is uniform, $\lim_{L \rightarrow \infty} p(L, m(L)) = \frac{2(a-c)}{(r-c)+(a-c)}$.*

Proof of Claim 12.2. By Claim 12.1, the marginal member $y(L)$ satisfies $L = \frac{y(L)(r-a)}{(1-y(L))(a-c)}$, whence $y(L) = \frac{a-c}{a-c+(r-a)\frac{1}{L}}$. If f is uniform, we have $m(L) = \frac{1+y(L)}{2}$, so $m(L) = \frac{1}{2} \frac{2L(a-c)+r-a}{L(a-c)+r-a}$. Recall that $p(L, m(L)) = \frac{1}{1+(\frac{1}{m(L)}-1)L}$. Then $p(L, m(L)) = \frac{2L(a-c)+r-a}{L(2(a-c)+r-a)+r-a}$, so $\lim_{L \rightarrow \infty} p(L, m(L)) = \frac{2(a-c)}{(r-c)+(a-c)}$. ■

Assume that f is uniform. By Lemma 1, $x \mapsto V(x)$ is strictly increasing and, by Lemma 5, $L \mapsto p(L, m(L))$ is decreasing, so $L \mapsto V(p(L, m(L)))$ is decreasing. Thus, for all L , $V(p(L, m(L))) \geq \lim_{L' \rightarrow \infty} V(p(L', m(L'))) \geq \frac{s}{\gamma}$. By Claim 12.2, $\lim_{L' \rightarrow \infty} V(p(L', m(L'))) \geq V\left(\frac{2(a-c)}{(r-c)+(a-c)}\right)$. By Lemma 3, this result extends to all non-decreasing densities f .

Next, we show that $V\left(\frac{2(a-c)}{(r-c)+(a-c)}\right) \geq \frac{1}{\gamma} \frac{(r-c)a+(a-c)r}{(r-c)+(a-c)}$. Note that, in an equilibrium in which the organization experiments forever, the payoff of an agent x is

bounded below by her payoff from staying in the organization forever. This is given by $\frac{r}{\gamma}$ if the risky policy is good and $\frac{c}{\gamma}$ if not. Then $V(x) \geq x\frac{r}{\gamma} + (1-x)\frac{c}{\gamma}$, which implies that $V\left(\frac{2(a-c)}{(r-c)+(a-c)}\right) \geq \frac{1}{\gamma} \frac{(r-c)a+(a-c)r}{(r-c)+(a-c)}$.

Finally, note that there exist parameter values such that $\frac{1}{\gamma} \frac{(r-c)a+(a-c)r}{(r-c)+(a-c)} \geq \frac{s}{\gamma}$ is satisfied. In general, for any values of r , a and c satisfying $r > a > c > 0$, there is $r^*(r, a, c)$ such that the condition holds if $s \leq s^*(r, a, c)$, and, moreover, $r^*(r, a, c) \in (a, r)$. ■

Proof of Proposition 5. Follows from Proposition 12. ■

The following two results (Lemma 5 and Lemma 6) show how to construct prior distributions for which the mapping $L \mapsto p(L, m(L))$ has a strict interior minimum, as needed for Proposition 13. Lemma 6 provides a tractable necessary and sufficient condition, and Lemma 5 gives an explicit construction.

Lemma 5. *If the distribution of priors is power law, then $L \mapsto p(L, m(L))$ is decreasing. Moreover, if $L_0 m'(L_0) < m(L_0)(1 - m(L_0))$, then $L \mapsto p(L, m(L))$ is strictly decreasing at $L = L_0$, and if $L_0 m'(L_0) > m(L_0)(1 - m(L_0))$, then $L \mapsto p(L, m(L))$ is strictly increasing at $L = L_0$.*

Proof of lemma 5. The density of the power law distribution is given by $f_\omega(x) = (1-x)^\omega(\omega+1)$. Hence we have $F(z) = 1 - (1-z)^{\omega+1}$ and the CDF of the distribution with support on $[y, 1]$ is given by $\frac{(1-y)^{\omega+1} - (1-z)^{\omega+1}}{(1-y)^{\omega+1}}$.

Recall that $m(L)$ and $y(L)$ denote the median and the marginal members of the organization respectively when the state variable is L . The above argument implies that the median must satisfy $\frac{(1-y(L))^{\omega+1} - (1-m(L))^{\omega+1}}{(1-y(L))^{\omega+1}} = \frac{1}{2}$. Equivalently, we must have $(1-m(L))^{\omega+1} = \frac{1}{2}(1-y(L))^{\omega+1}$. Then the median must satisfy $1-m(L) = (1-y(L))2^{-\frac{1}{\omega+1}}$, or $m(L) = 1 - \kappa + \kappa y(L)$ for $\kappa = 2^{-\frac{1}{\omega+1}}$.

Note that $p(L, m(L)) = \frac{1}{1 + (\frac{1}{m(L)} - 1)L}$. Then $\frac{\partial}{\partial L} p(L, m(L)) \propto -\frac{\partial}{\partial L} \left(1 + \left(\frac{1}{m(L)} - 1\right)L\right)$ and $\frac{\partial}{\partial L} \left(1 + \left(\frac{1}{m(L)} - 1\right)L\right) = \frac{\partial}{\partial L} \left(\left(\frac{1}{m(L)} - 1\right)L\right) = \frac{1}{m(L)} - 1 - \frac{L}{(m(L))^2} m'(L)$.

This implies that if $L_0 m'(L_0) < m(L_0)(1 - m(L_0))$, then $L \mapsto p(L, m(L))$ is strictly decreasing at $L = L_0$, and if $L_0 m'(L_0) > m(L_0)(1 - m(L_0))$, then $L \mapsto p(L, m(L))$ is strictly increasing at $L = L_0$.

After some algebra, using the fact that $y(L) = \frac{a-c}{a-c+(r-a)\frac{1}{L}}$, we get that if the distribution of priors is power law, then $Lm'(L) < m(L)(1 - m(L))$ is equivalent to $0 < (1 - \kappa)(1 - \zeta)$, where $\zeta = \frac{a-c}{r-c}$. Since κ and ζ are between 0 and 1, this always holds. \blacksquare

Lemma 6. *There exist distributions for which there exist states $L_1 < L_2$ such that L_1 is a unique minimizer of $p(L, m(L))$ and $L \mapsto p(L, m(L))$ is strictly increasing on (L_1, L_2) .*

Proof of lemma 6.

Consider a distribution with a density $f(x) = d$ for $x \in [0, b]$ and $f(x) = d'$ for $x \in [b, 1]$. Note that we must have $db + d'(1 - b) = 1$ so that f integrates to 1. Define $\bar{r} = r - c$, $\bar{a} = a - c$, $y = y(L)$, $m = m(L)$, $z = p(L, m(L))$. Let L_1 be such that $m(L_1) = b$ and let L_2 be such that $y(L_2) = b$. Clearly, $0 < L_1 < L_2$. For $L > L_2$, $m(L)$ and $p(L, m(L))$ are the same as in the uniform case. In particular, $p(L, m(L)) = \frac{2L\bar{a} + \bar{r} - \bar{a}}{L(\bar{a} + \bar{r}) + \bar{r} - \bar{a}}$, which is decreasing in L . Moreover, with the notation we have defined, the formula for $y(L)$ can be written as $y = \frac{L\bar{a}}{L\bar{a} + \bar{r} - \bar{a}}$.

For $L \in (L_1, L_2)$, we have $d(b - y) + d'(m - b) = d'(1 - m)$, that is, $m = \frac{1+b}{2} - \frac{db}{2d'} + \frac{d}{2d'}y$. Equivalently, $m = (1 - \frac{1}{2d'}) + \frac{d}{2d'}y = (1 - \frac{1}{2d'}) + \frac{d}{2d'} \frac{L\bar{a}}{L\bar{a} + \bar{r} - \bar{a}}$. Then

$$\frac{1}{z} - 1 = \frac{L(1 - m)}{m} = L \frac{L \frac{1-d}{2d'} \bar{a} + \frac{1}{2d'} (\bar{r} - \bar{a})}{(1 - \frac{1}{2d'} + \frac{d}{2d'}) L\bar{a} + (1 - \frac{1}{2d'}) (\bar{r} - \bar{a})}$$

For $L < L_1$, we have $d(m - y) = d(b - m) + d'(1 - b)$, that is, $2dm = db + d'(1 - b) + dy = 1 + dy$, so $m = \frac{1}{2d} + \frac{1}{2}y$, and

$$\frac{1}{z} - 1 = \frac{L(1 - m)}{m} = L \frac{L (\frac{1}{2} - \frac{1}{2d}) \bar{a} + (1 - \frac{1}{2d}) (\bar{r} - \bar{a})}{(\frac{1}{2d} + \frac{1}{2}) L\bar{a} + \frac{1}{2d} (\bar{r} - \bar{a})}.$$

Now take $d' = \frac{1}{2}$ and any $d > 1$ (note that choosing roth pins down $b = \frac{1}{2d-1}$). Then we can verify that $L \mapsto \frac{1}{p(L, m(L))} - 1$ is increasing on $(0, L_1)$ and decreasing on (L_1, L_2) . In other words, $L \mapsto p(L, m(L))$ is decreasing on $(0, L_1)$ and (L_2, ∞) but increasing on (L_1, L_2) , so L_1 is a local minimizer for $p(L, m(L))$.

Moreover, we can verify that under some extra conditions L_1 is a global minimizer: note that $\lim_{L \rightarrow \infty} \frac{1}{p(L, m(L))} - 1 = \frac{\bar{r} - \bar{a}}{2\bar{a}}$, while $\frac{1}{p(L_1, m(L_1))} - 1 = \frac{L_1(1-d)\bar{a} + \bar{r} - \bar{a}}{d\bar{a}}$.

Since $m(L_1) = b$, we have

$$\begin{aligned} \frac{1}{p(L_1, m(L_1))} - 1 &= \frac{L_1}{m(L_1)} - L_1 = \frac{L_1}{b} - L_1 = \frac{L_1(1-d)\bar{a} + \bar{r} - \bar{a}}{d\bar{a}} \\ L_1 &= \frac{\bar{r} - \bar{a}}{\bar{a}(\frac{d}{b} - 1)} \\ \frac{1}{p(L_1, m(L_1))} - 1 &= \frac{L_1}{b} - L_1 = \frac{\bar{r} - \bar{a}}{\bar{a}} \frac{\frac{1}{b} - 1}{\frac{d}{b} - 1} = \frac{\bar{r} - \bar{a}}{\bar{a}} \frac{1-b}{d-b} \\ &= \frac{\bar{r} - \bar{a}}{\bar{a}} \frac{2d-2}{2d^2-d-1} = \frac{\bar{r} - \bar{a}}{\bar{a}} \frac{1}{d+\frac{1}{2}} \end{aligned}$$

so L_1 is a global minimizer if we take $d \in (1, \frac{3}{2})$. ■

Proposition 13. *There exist $r, s, a, c, f, \epsilon \in (0, 1]$ and $L^* > 0$ such that an equilibrium of the following form exists: whenever $L = L^*$, the organization stops experimenting with probability ϵ , and whenever $L \neq L^*$, the organization continues experimenting with probability one.*

Proof of Proposition 13. For convenience, we multiply all the value functions in this proof by γ . Let $V_x^\epsilon(L)$ denote the value function of agent x given that the state is $L(k, t) = L$ and the behavior on the equilibrium path is as described in the Proposition. Note that $V_x^0(L)$ is the value function of agent x when the state is L and there is perpetual experimentation.

We claim that we can choose the density f such that there is a unique global minimum of $L \mapsto V_{m(L)}^0(L)$, which we will call L^* , and in addition so that $V_{m(L)}^0(L)$ has a kink at L^* . By Corollary 1, this is equivalent to $p(L, m(L))$ being uniquely minimized at L^* with a kink at L^* . This claim then follows from Lemma 6.⁴⁴

Note that $V_{m(L)}^0(L)$ does not depend on s and the density f constructed in the proof of Lemma 6 does not depend on s , which implies that L^* does not depend on s . Then we can choose s such that

$$V_{m(L^*)}^0(L^*) = s \tag{1}$$

⁴⁴Technically, we also need the condition that $V_{m(L^*)}^0(L^*) < \lim_{L \rightarrow \infty} V_{m(L)}^0(L)$, but this is also satisfied by the construction in Lemma 6.

Then, because L^* is the unique minimizer of $L \mapsto V_{m(L)}^0(L)$, we have $V_{m(L)}^0(L) > s$ for all $L \neq L^*$.

We aim to show that if we change the equilibrium to require that experimentation stops at $L = L^*$ with an appropriately chosen probability $\epsilon > 0$, the constraints $V_{m(L^*)}^\epsilon(L^*) = s$ and $V_{m(L)}^\epsilon(L) \geq s$ for all $L \neq L^*$ still hold.

It is useful at this point to write the value function recursively, as follows. For any strategy profile and any $L, L' \in \mathbb{R}$, define

$$\begin{aligned}\zeta_{x,L'}(L) &= \int_0^\infty \gamma e^{-\gamma t} Pr[\exists t' \in [0, t] : L_{t'} = L' | L_0 = L] dt \\ \tilde{V}_{x,L'}(L) &= \frac{\int_0^\infty \gamma e^{-\gamma t} E[u_x(h^t) \mathbf{1}_{\nexists t' \in [0, t] : L_{t'} = L'} | L_0 = L] dt}{1 - \zeta_{x,L'}(L)}\end{aligned}$$

where $u_x(h^t)$ is agent x 's flow payoff at time t and history h^t and the expectation is taken with respect to the stochastic process induced by the equilibrium strategy and the stochastic process $\left(\tilde{L}_\tau\right)_\tau$.

Intuitively, $\zeta_{x,L'}(L)$ is the weighted discounted probability that the stochastic process $(L_t)_t$ hits the value L' at least once, and the numerator of $\tilde{V}_{x,L'}(L)$ is the expected utility of agent x starting with $L_0 = L$ but setting the continuation value to zero when $(L_t)_t$ hits L' . Then the value function can be written recursively as

$$V_x(L) = (1 - \zeta_{x,L'}(L)) \tilde{V}_{x,L'}(L) + \zeta_{x,L'}(L) V_x(L')$$

Taking $L' = L^*$, this implies that for any $\epsilon \in [0, 1]$, we have

$$V_x^\epsilon(L) = (1 - \zeta_{x,L^*}(L)) \tilde{V}_{x,L^*}(L) + \zeta_{x,L^*}(L) V_x^\epsilon(L^*) \quad (2)$$

where $\zeta_{x,L^*}(L)$ is independent of ϵ , since changing ϵ has no impact on the policy path except when $L = L^*$. Let

$$\begin{aligned}\zeta_{x,L'}(L^+) &= \int_0^\infty \gamma e^{-\gamma t} Pr[\exists t' \in (0, t] : L_{t'} = L' | L_0 = L] dt \\ \tilde{V}_{x,L'}(L^+) &= \frac{\int_0^\infty \gamma e^{-\gamma t} E[u_x(h^t) \mathbf{1}_{\nexists t' \in (0, t] : L_{t'} = L'} | L_0 = l] dt}{1 - \zeta_{x,L'}(L^+)}\end{aligned}$$

Observe that $\tilde{V}_{x,L^*}^\epsilon(L^{*+}) = \lim_{L \searrow L^*} \tilde{V}_{x,L^*}^\epsilon(L)$ and $\zeta_{x,L^*}(L^{*+}) = \lim_{L \searrow L^*} \zeta_{x,L^*}(L)$.

Let $\tilde{W}_x^\epsilon = \tilde{V}_{x,L^*}^\epsilon(L^{*+})$ and $W_x^\epsilon = \lim_{L \searrow L^*} V_x^\epsilon(L)$. W_x^ϵ is the expected continuation value of agent x when $L = L^*$ and the median member, $m(L^*)$, has just decided *not* to stop experimenting. This is closely related to $V_x^\epsilon(L^*)$, which is the expected continuation value where the expectation is taken before $m(L^*)$ has decided whether to stop experimenting or not. Specifically, we have

$$V_x^\epsilon(L^*) = \epsilon s + (1 - \epsilon)W_x^\epsilon = \epsilon s + (1 - \epsilon) \left((1 - \zeta_{x,L^*}(L^{*+})) \tilde{W}_x^\epsilon + \zeta_{x,L^*}(L^{*+}) V_x^\epsilon(L^*) \right)$$

Solving this for $V_x^\epsilon(L^*)$, we obtain

$$\begin{aligned} V_x^\epsilon(L^*) &= \frac{\epsilon s + (1 - \epsilon)(1 - \zeta_{x,L^*}(L^{*+})) \tilde{W}_x^\epsilon}{1 - (1 - \epsilon)\zeta_{x,L^*}(L^{*+})} \\ &= V_x^0(L^*) + \epsilon \frac{s - V_x^0(L^*)}{1 - (1 - \epsilon)\zeta_{x,L^*}(L^{*+})} \end{aligned} \quad (3)$$

where the second equality follows from the fact that $\tilde{W}_x^\epsilon = V_x^0(L^*)$ because \tilde{W}_x^ϵ is the continuation value of the agent conditional on the event that $(L_t)_t$ never hits L^* again, which means that in this case experimentation continues forever. Hence, substituting (3) into (2), we obtain

$$\begin{aligned} V_x^\epsilon(L) &= (1 - \zeta_{x,L^*}(L)) \tilde{V}_{x,L^*}(L) + \zeta_{x,L^*}(L) \left(V_x^0(L^*) + \epsilon \frac{s - V_x^0(L^*)}{1 - (1 - \epsilon)\zeta_{x,L^*}(L^{*+})} \right) \\ &= \zeta_{x,L^*}(L) \epsilon \frac{s - V_x^0(L^*)}{1 - (1 - \epsilon)\zeta_{x,L^*}(L^{*+})} + V_x^0(L) \end{aligned} \quad (4)$$

At the same time, because we have assumed that $V_{m(L)}(L)$ is minimized at L^* with a kink at L^* , there exist $\delta > 0$ and $K > 0$ such that for all $L \in (L^* - \delta, L^* + \delta)$

$$V_{m(L)}^0(L) = V_{p(L,m(L))}^0(1) \geq V_{p(L^*,m(L^*))}^0(1) + K|L - L^*| = s + K|L - L^*| \quad (5)$$

where the first equality follows from Corollary 1, and the last equality follows from (1). On the other hand, for $L \notin (L^* - \delta, L^* + \delta)$ there exists $K' > 0$ such that

$$V_{m(L)}^0(L) = V_{p(L,m(L))}^0(0) \geq V_{p(L^*,m(L^*))}^0(0) + K' = s + K' \quad (6)$$

where the first equality follows from Corollary (1), the inequality follows from the fact that $p(L, m(L)) - p(L^*, m(L^*))$ is bounded away from zero in this case, and the second equality follows from (1).

By (4), $V_{m(L)}^\epsilon(L) \geq s$ is equivalent to $V_{m(L)}^0(L) \geq s - \zeta_{m(L), L^*}(L) \epsilon \frac{s - V_{m(L)}^0(L^*)}{1 - (1 - \epsilon)\zeta_{m(L), L^*}(L^{*+})}$. If $V_{m(L)}^0(L^*) - s \leq 0$, then we are done, so suppose that $V_{m(L)}^0(L^*) - s > 0$.

Suppose that $L \in (L^* - \delta, L^* + \delta)$. Then, by (5), it is sufficient that

$$\begin{aligned} s + K|L - L^*| &\geq s - \zeta_{m(L), L^*}(L) \epsilon \frac{s - V_{m(L)}^0(L^*)}{1 - (1 - \epsilon)\zeta_{m(L), L^*}(L^{*+})} \\ \iff K|L - L^*| &\geq \zeta_{m(L), L^*}(L) \epsilon \frac{V_{m(L)}^0(L^*) - s}{1 - (1 - \epsilon)\zeta_{m(L), L^*}(L^{*+})} \\ \iff K|L - L^*| &\geq \epsilon \frac{V_{m(L)}^0(L^*) - s}{1 - \zeta_{m(L), L^*}(L^{*+})} \iff \epsilon \leq \frac{K|L - L^*|(1 - \zeta_{m(L), L^*}(L^{*+}))}{V_{m(L)}^0(L^*) - V_{m(L^*)}^0(L^*)} \end{aligned}$$

where we have used that $\zeta_{m(L), L^*}(L) \in (0, 1)$.

Suppose next that $L \notin (L^* - \delta, L^* + \delta)$. Then, because $V_{m(L)}^0(L) \geq s + K'$ by (6), it is sufficient that

$$\begin{aligned} s + K' &\geq s - \zeta_{m(L), L^*}(L) \epsilon \frac{s - V_{m(L)}^0(L^*)}{1 - (1 - \epsilon)\zeta_{m(L), L^*}(L^{*+})} \\ \iff K' &\geq \zeta_{m(L), L^*}(L) \epsilon \frac{V_{m(L)}^0(L^*) - s}{1 - (1 - \epsilon)\zeta_{m(L), L^*}(L^{*+})} \\ \iff K' &\geq \epsilon \frac{r - s}{1 - \zeta_{m(L), L^*}(L^{*+})} \iff \epsilon \leq K'(1 - \zeta_{m(L), L^*}(L^{*+})) \frac{1}{r - s} \end{aligned}$$

where we have used that $\zeta_{m(L), L^*}(L) \in (0, 1)$ and $V_{m(L)}^0(L^*) \leq r$.

Let $\epsilon_1 = \inf_{L \in (L^* - \delta, L^* + \delta)} K(1 - \zeta_{m(L), L^*}(L^{*+})) \frac{|L - L^*|}{V_{m(L)}^0(L^*) - V_{m(L^*)}^0(L^*)}$ and $\epsilon_2 = \inf_{L \notin (L^* - \delta, L^* + \delta)} K'(1 - \zeta_{m(L), L^*}(L^{*+})) \frac{1}{r - s}$. To have $\min\{\epsilon_1, \epsilon_2\} > 0$ we need to verify that $\sup_x \zeta_{x, L^*}(L^{*+}) < 1$ and that $\sup_{L \in (L^* - \delta, L^* + \delta)} \left| \frac{V_{m(L)}^0(L^*) - V_{m(L^*)}^0(L^*)}{L - L^*} \right|$ is finite. $\sup_x \zeta_{x, L^*}(L^{*+}) < 1$ is immediate. The fact that $\sup_{L \in (L^* - \delta, L^* + \delta)} \left| \frac{V_{m(L)}^0(L^*) - V_{m(L^*)}^0(L^*)}{L - L^*} \right|$ is finite follows from the fact that $\frac{\partial}{\partial x} V_x^0(L^*)$ and $m'(L)$ are bounded.

Then choosing $\epsilon \in (0, \min\{\epsilon_1, \epsilon_2\})$ delivers the result. ■

Proof of Proposition 6. Take the example constructed in Proposition 13, and assume that $L_0 > L^*$.⁴⁵

Let $P_\theta(L_0)$ be the probability that, conditional on starting at L_0 and the state being $\theta \in \{G, B\}$, the organization stops experimenting at any finite time $t < \infty$. We will show that $P_G(L_0) > P_B(L_0)$ for L_0 large enough. In fact, we will prove a stronger result: we will show that there is $C > 0$ such that $P_G(L_0) \geq C > 0$ for all $L_0 > L^*$, but $\lim_{L_0 \rightarrow \infty} P_B(L_0) = 0$.

Let $Q_\theta(L_0, L^*)$ denote the probability that there exists $t < \infty$ such that $L_t \in ((\frac{c}{r}) L^*, L^*]$ when the state is $\theta \in \{G, B\}$. $Q_\theta(L_0, L^*)$ is the probability that L_t ever crosses over to the left of L^* .

We claim that $Q_G(L_0, L^*) = 1$ for all $L_0 > L^*$ but $\lim_{L_0 \rightarrow \infty} Q_B(L_0, L^*) = 0$.

Let $l(k, t) = \ln L(k, t)$, and note that $l(k, t) = \ln \left(\left(\frac{c}{r} \right)^k e^{(r-c)t} \right) = k \ln \left(\frac{c}{r} \right) + \ln(e^{(r-c)t}) = k(\ln(c) - \ln(r)) + (r - c)t$. Let $l_0 = \ln(L_0)$.

When $\theta = G$, we then have $(l_t)_t = l_0 + (r - c)t - [\ln(r) - \ln(c)] N(t)$, where $(N(t))_t$ is a Poisson process with rate r , that is, $N(t) \sim P(rt)$. This can be written as a random walk: for integer values of t , $l_t - l_0 = \sum_{i=0}^t S_i$, where $S_i = r - c - [\ln(r) - \ln(c)] N_i$, and $N_i \sim P(b)$ are iid. Note that $E[S_i] = r - c - r(\ln(r) - \ln(c)) < 0$.⁴⁶ Then, by the strong law of large numbers, we have $\frac{l_t}{t} \xrightarrow[t \rightarrow \infty]{} E[S_i] < 0$ a.s., whence $l_t \xrightarrow[t \rightarrow \infty]{} -\infty$ a.s., implying the first claim.

On the other hand, when $\theta = B$, we have $(l_t)_t = l_0 + (r - c)t - \ln(r - c)N(t)$, where $(N(t))_t$ is a Poisson process with rate c . This can be written as a random walk with positive drift: $l_t - l_0 = \sum_{i=0}^t S_i$, where $S_i = r - c - [\ln(r) - \ln(c)] N_i$, $N_i \sim P(c)$, and $E[S_i] = r - c - c(\ln(r) - \ln(c)) > 0$. As above, by the strong law of large numbers, we have $l_t \xrightarrow[t \rightarrow \infty]{} \infty$ a.s.

⁴⁵Technically, our definition of L_t requires that $L_0 = 1$, but we can relax this assumption by considering a continuation of the game starting at some $t_0 > 0$, where, by assumption, the number of successes at time t_0 is such that the state variable at t_0 is L_0 . This example can be fit into our original framework by redefining the density of prior beliefs \tilde{f} to be the density of the posteriors held by agents when $L = L_0$ and f is as in Proposition 13. With this relabeling, L_0 would equal 1 and L^* would shift to some value less than 1. We find it is easier to think in terms of shifting L_0 and leaving f unchanged.

⁴⁶Let $\frac{r}{c} = 1 + x$. Then $r - c - r(\ln(r) - \ln(c)) = c(x - (1 + x) \ln(1 + x))$, where $x - (1 + x) \ln(1 + x)$ is negative for all $x > 0$. Similarly, $r - c - c(\ln(r) - \ln(c)) = c(x - \ln(1 + x))$, where $x - \ln(1 + x)$ is positive for all $x > 0$.

Note that $Q_B(L, \frac{c}{r}L) = q$ is independent of L because $(l_t)_t$ follows a random walk. Now suppose for the sake of contradiction that $\limsup_{L \rightarrow \infty} Q_B(L, L^*) > 0$. We claim that this implies $q = 1$. Suppose towards a contradiction that $q < 1$. Fix $J \in \mathbb{N}$. Then, for L_0 large enough that $(\frac{c}{r})^{2J+1} L_0 > L^*$,

$$Q_B(L_0, L^*) \leq \prod_{j=0}^J Q_B\left(\left(\frac{c}{r}\right)^{2j} L_0, \left(\frac{c}{r}\right)^{2j+1} L_0\right) = q^{J+1}$$

This implies that, whenever $\limsup_{L \rightarrow \infty} Q_B(L, L^*) > 0$, we have $q = 1$, as the above equation must hold for arbitrarily large J . Hence $(l_t)_t$ is recurrent, that is, it visits the neighborhood of every $l \in \mathbb{R}$ infinitely often (Durrett 2010: pp. 190–201). However, this contradicts the fact that $\lim_{t \rightarrow \infty} l_t = \infty$ a.s. Therefore, $\limsup_{L \rightarrow \infty} Q_B(L, L^*) = 0$.

This implies that $P_B(L_0) \leq Q_B(L_0, L^*) \rightarrow 0$ as $L_0 \rightarrow \infty$. On the other hand, $P_G(L_0) \geq Q_G(L_0, L^*) \inf_{L \in ((\frac{c}{r})L^*, L^*]} P_G(L) > 0$. The first inequality holds for the following reason. With probability 1, if $L_t = L^*$ for some t , there must be $t' < t$ such that $L_{t'} \in (\frac{c}{r}L^*, L^*)$, which happens with probability $Q_G(L_0, L^*)$. Conditional on this event, the probability of hitting state L^* in the continuation is $P_G(L_{t'})$. Note that $\inf_{L \in ((\frac{c}{r})L^*, L^*]} P_G(L) > 0$ because it is equal to $P_G((\frac{c}{r})L^*)$. ■

A.4 Other Extensions

Proof of Proposition 7. Fix an equilibrium σ in which organization experiments forever and let μ_t be the size of the organization at time t on the equilibrium path. Let $g_t = g(\mu_t)$. The *first success* that happens at time t yields the per-capita payoff of g_t , and all further successes pay 1 (because all agents enter the organization after the first success).

Let $P_t = 1 - e^{-rt}$ denote the probability that there is a success by time t given that the risky technology is good. In the above problem, an agent with belief x who

expects experimentation to continue forever has utility

$$V_{(g_t)_t}(x) = x \int_0^{t^*} e^{-\gamma t} (P_t r + (1 - P_t) g_t r) dt + x \int_{t^*}^{\infty} e^{-\gamma t} (P_t r + (1 - P_t) a) dt + (1 - x) \int_{t^*}^{\infty} e^{-\gamma t} a dt$$

where t^* is the time at which the agent leaves, that is, when her posterior reaches $\frac{a}{g_t r}$. Now consider the case in which $g_t = g$ for all t . Then the above expression equals

$$V_g(x) = x \left(\frac{r}{\gamma} - \frac{r}{\gamma + r} + \frac{gr(1 - e^{-(\gamma+r)t^*})}{\gamma + r} + \frac{e^{-(\gamma+r)t^*} a}{\gamma + r} - \frac{e^{-\gamma t^*} a}{\gamma} \right) + \frac{e^{-\gamma t^*} a}{\gamma}$$

Suppose that $f = f_\omega$, as in Proposition 9. By the same arguments as in that Proposition, if y_t satisfies $p_t(y_t) = \frac{a}{gr}$ for all t , then $p_t(m_t) \searrow \frac{a}{\lambda(gr-a)+a}$ as $t \rightarrow \infty$, and, by Claim 9.3 in the proof of Proposition 9, we have $t^* = -\frac{\ln(\lambda)}{r}$.⁴⁷ Then

$$V_g \left(\frac{a}{\lambda(gr-a)+a} \right) = \frac{a}{\lambda gr + (1-\lambda)a} \left(\frac{r}{\gamma} - \frac{r}{\gamma + r} + \frac{gr(1 - \lambda^{\frac{\gamma+r}{r}})}{\gamma + r} + \frac{\lambda^{\frac{\gamma+r}{r}} a}{\gamma + r} - \frac{\lambda^{\frac{\gamma}{r}} a}{\gamma} \right) + \frac{\lambda^{\frac{\gamma}{r}} a}{\gamma}$$

Since this is a hyperbola in g , it is either increasing in g for all $g > 0$ or decreasing in g for all $g > 0$. In particular, when the congestion effect is maximal, that is, when $g \rightarrow \infty$, we have

$$\lim_{g \rightarrow \infty} \gamma V_g \left(\frac{a}{\lambda(gr-a)+a} \right) = \frac{a}{\lambda} \left(1 - \lambda^{\frac{\gamma+r}{r}} \right) \frac{\gamma}{\gamma + r} + \lambda^{\frac{\gamma}{r}} a = \frac{a}{\lambda} \frac{\gamma}{\gamma + r} + \lambda^{\frac{\gamma}{r}} a \frac{r}{\gamma + r}$$

On the other hand, when the economies of scale are maximal, that is, as $g \rightarrow \frac{a}{r}$,⁴⁸

$$\lim_{g \rightarrow \frac{a}{r}} \gamma V_g \left(\frac{a}{\lambda(gr-a)+a} \right) = \gamma \left(\frac{r}{\gamma} - \frac{r}{\gamma + r} + \frac{a(1 - \lambda^{\frac{\gamma+r}{r}})}{\gamma + r} + \frac{\lambda^{\frac{\gamma+r}{r}} a}{\gamma + r} \right) = r \frac{r}{\gamma + r} + a \frac{\gamma}{\gamma + r}$$

Thus $\frac{a}{\lambda} \frac{\gamma}{\gamma + r} + \lambda^{\frac{\gamma}{r}} a \frac{r}{\gamma + r} > r \frac{r}{\gamma + r} + a \frac{\gamma}{\gamma + r}$ is equivalent to $\lim_{g \rightarrow \infty} V_g \left(\frac{a}{\lambda(gr-a)+a} \right) >$

⁴⁷Note that this is the same t^* as in the baseline model.

⁴⁸If $g < \frac{a}{r}$, we enter a degenerate case in which the organization becomes empty immediately.

$\lim_{g \rightarrow \frac{a}{r}} V_g \left(\frac{a}{\lambda(gr-a)+a} \right)$. Because $V_g \left(\frac{a}{\lambda(gr-a)+a} \right)$ is either increasing in g for all $g > 0$ or decreasing in g for all $g > 0$, this condition implies that $V_g \left(\frac{a}{\lambda(gr-a)+a} \right)$ is increasing in g for all $g > 0$. The argument in the case when the inequality is reversed is similar.

In addition, note that if $V_g \left(\frac{a}{\lambda(gr-a)+a} \right)$ is increasing in g , then we can guarantee that, with a congestion effect,

$$V_{(g_\tau)_{\tau \geq t}}(p_t(m_t)) > V_{g_t}(p_t(m_t)) > V_{g_t} \left(\frac{a}{\lambda(g_t r - a) + a} \right) > V \left(\frac{a}{\lambda(r - a) + a} \right)$$

Here the first inequality follows because $g_\tau \mapsto V_{g_t, \dots, g_\tau, \dots}(x)$ is increasing, under congestion effect $\mu \mapsto g(\mu)$ is decreasing and under perpetual experimentation $t \mapsto \mu_t$ is decreasing, so $t \mapsto g_t = g(\mu_t)$ is increasing. The second inequality follows because $x \mapsto V_{g_t}(x)$ is strictly increasing and $p_t(m_t) \searrow \frac{a}{\lambda(gr-a)+a}$ as $t \rightarrow \infty$. The last inequality follows because $g \mapsto V_g \left(\frac{a}{\lambda(gr-a)+a} \right)$ is increasing, under congestion effect $\mu \mapsto g(\mu)$ is decreasing and in the baseline model we have $g(\mu) = g(1) = 1$ for all μ .

Thus the condition to obtain experimentation forever is slacker with a congestion effect than in the baseline model at every t , not just in the limit. By the same argument, the condition for experimentation forever is tighter for all t under economies of scale.⁴⁹ ■

Proof of Proposition 8.

We have

$$\begin{aligned} \tilde{V}_t(x) = & x \int_t^\infty e^{-\gamma(\tau-t)} [P_\tau r + (1 - P_\tau)(aF(y_\tau) + r(1 - F(y_\tau)))] d\tau \\ & + (1 - x) \int_t^\infty e^{-\gamma(\tau-t)} aF(y_\tau) d\tau, \end{aligned}$$

where $P_\tau = 1 - e^{-r(\tau-t)}$ is the probability that there has been a success by time τ , conditional on the state being good and there being no success up to time t , and $F(y_\tau)$ is the fraction of the population that are outsiders at time τ , conditional on no

⁴⁹If $g \mapsto V_g \left(\frac{a}{\lambda(gr-a)+a} \right)$ is decreasing, it is more difficult to make general statements about what happens away from the limit because in this case the effect of moving away from the limit goes against the result: for instance, under congestion effect the condition becomes tighter in the limit but increasing g_t slackens the condition.

successes. We can rewrite this equation as

$$\begin{aligned}\tilde{V}_t(x) &= x \int_t^\infty e^{-\gamma(\tau-t)} [r - (1 - P_\tau)(r - a)F(y_\tau)] d\tau + (1 - x) \int_t^\infty e^{-\gamma(\tau-t)} aF(y_\tau) d\tau \\ &= x \frac{r}{\gamma} - x \int_t^\infty e^{-(\gamma+r)(\tau-t)} (r - a)F(y_\tau) d\tau + (1 - x) \int_t^\infty e^{-\gamma(\tau-t)} aF(y_\tau) d\tau \quad (7)\end{aligned}$$

Let $F_\tau = F(y_\tau)$ and note that $\tau \mapsto F_\tau$ is weakly increasing.

The upper bound for $\tilde{V}_t(x)$ is now obtained as follows. Note that, given $\tau \geq t$, the derivative of (7) with respect to F_τ is proportional to $-xe^{-r(\tau-t)}(r - a) + (1 - x)a$. It follows that if the agent could choose F_τ everywhere at will to maximize her payoff, she would choose $F_\tau = 1$ for $\tau \geq t(x)$ and $F_\tau = 0$ for $\tau < t(x)$, where $t(x)$ is defined by the condition $xe^{-r(t(x)-t)}(r - a) = (1 - x)a$ (obtained by setting the derivative equal to 0). The result of this choice is $V(x)$, her utility in the private values case, in which she only cares about her own entry and exit decisions and gets to choose them optimally. Because in the common values case the entry and exit decisions of other agents are not optimal from x 's point of view, $\tilde{V}_t(x)$ must be weakly lower than $V(x)$.

As for the lower bound, assume for the sake of argument that F_τ is constant for all y_τ and equal to $\bar{F} \in [0, 1]$. Then the expression in (7) is

$$x \frac{r}{\gamma} - x \frac{(r - a)\bar{F}}{\gamma + r} + (1 - x) \frac{s\bar{F}}{\gamma}$$

which is linear in \bar{F} and is minimized either when $\bar{F} = 0$ or when $\bar{F} = 1$. In the first case, the expression equals $x \frac{r}{\gamma}$. In the second case, it equals $x \frac{r}{\gamma} - x \frac{r-a}{\gamma+r} + (1 - x) \frac{a}{\gamma}$.

To finish the proof, we argue that whenever F_τ is weakly increasing in τ , the expression in 7 is higher than the expression that is obtained when F_τ is replaced by a suitably chosen constant \bar{F} . Hence the lower bound obtained for constant F_τ applies in all cases.

The argument is as follows. Take $\bar{F} = F_{t(x)}$. Then for $\tau > t(x)$, F_τ is weakly greater than \bar{F} and $\tilde{V}_t(x)$ is increasing in the value of F at τ . Conversely, for $\tau < t(x)$, F_τ is weakly lower than \bar{F} and $\tilde{V}_t(x)$ is decreasing in the value of F at τ . Hence the agent's utility is weakly higher under F_τ than under a constant $F_{t(x)}$. \blacksquare

B Model with Unrestricted Policy Changes (For Online Publication)

In this Section we present a more general model which relaxes the assumption that switching to the safe policy is irreversible. Instead, we allow the organization to change its policy π_t any number of times. The main result is that, under a sensible restriction on the behavior of agents under indifference (which is uniquely selected by the discrete-time limit discussed in Section 4), switches to the safe policy are permanent in every equilibrium. Hence our assumption of irreversible policy changes is without loss of generality.

B.1 Definition of Equilibrium

We let π_{t-} and π_{t+} denote the left and right limits of the policy path at time t respectively, whenever the limits are well-defined. We require that π_t , the current policy at time t , is chosen by the decision-maker who is pivotal given the incumbent policy π_{t-} . Similarly, π_{t+} is chosen by the decision-maker who is pivotal given π_t . That is, for the policy to change from π to π' along the path of play, the decision-maker induced by π must be in favor of the change.

We define a membership function β so that $\beta(x, L, \pi) = 1$ if agent x chooses to be a member of the organization given information L and policy π , and $\beta(x, L, \pi) = 0$ otherwise. We define a policy correspondence α so that $\alpha(L, \pi)$ is the set of policies that the median voter, $m(L, \pi)$, is willing to choose.⁵⁰ We emphasize that $\alpha(L, \pi)$ need *not* be the set of policies that the median voter finds optimal in the sense of maximizing her utility given the behavior of the other agents – that is, $\alpha(L, \pi)$ is not an equilibrium notion. Our notion of strategy profile summarizes the above requirements:

Definition 2. A Markov strategy profile is given by a membership function $\beta : [0, 1] \times \mathbb{R}_+ \times \{0, 1\} \rightarrow \{0, 1\}$, a policy correspondence $\alpha : \mathbb{R}_+ \times \{0, 1\} \rightarrow \{\{0\}, \{1\}, \{0, 1\}\}$,

⁵⁰ $\alpha(L, \pi)$ can take the values $\{0\}$, $\{1\}$ and $\{0, 1\}$. Defining $\alpha(L, \pi)$ in this way is convenient because some paths of play cannot be easily described in terms of the instantaneous switching probabilities of individual agents. α should be understood as a choice rule in the decision-theoretic sense.

and a stochastic path of play consisting of information and policy paths $(L_t, \pi_t)_t$ satisfying the following:

- (a) Conditional on the policy type θ , $(L_t, \pi_t)_{t \geq 0}$ is a progressively measurable Markov process with paths that have left and right limits at every $t \geq 0$, satisfying $(L_0, \pi_0) = (1, 1)$.
- (b) Letting $(\tilde{k}_\tau)_\tau$ denote a Poisson process with rate r or 0 if $\theta = G$ or B respectively, letting $(\tilde{L}_\tau)_\tau$ be given by $\tilde{L}_\tau = L(\tilde{k}_\tau, \tau)$, and letting $n(t) = \int_0^t \pi_{t'} dt'$ denote the amount of experimentation up to time t , we have $L_t = \tilde{L}_{n(t)}$.
- (c) $\pi_t \in \alpha(L_t, \pi_{t-})$ for all $t \geq 0$.
- (d) $\pi_{t+} \in \alpha(L_t, \pi_t)$ for all $t \geq 0$.

We define $V_x(L, \pi)$ as the continuation utility of an agent with prior belief x given information L and incumbent policy π . In other words, $V_x(L, \pi)$ is the utility agent x expects to get starting at time t_0 when the state follows the process $(L_t, \pi_t)_{t \geq t_0}$ given that $(L_{t_0}, \pi_{t_0}) = (L, \pi)$.

Definition 3. An equilibrium σ is a strategy profile such that:

- (i) $\beta(x, L, \pi) = 1$ if $s + \pi(p(L, x)r - s) > a$ and $\beta(x, L, \pi) = 0$ otherwise.
- (ii) If $V_{m(L, \pi)}(L, \pi') > V_{m(L, \pi)}(L, 1 - \pi')$, then $\alpha(L, \pi) = \{\pi'\}$.

Part (i) of the definition of equilibrium says that agents make membership decisions that maximize their flow payoffs. Part (ii) says that the pivotal agent chooses her preferred policy based on her expected utility, assuming that the equilibrium strategies are played in the continuation.

As in Section 4, an additional condition is needed to rule out undersirable equilibria that arise when $V_{m(L, \pi)}(L, 1) = V_{m(L, \pi)}(L, 0)$ for the trivial reason that the continuation is independent of $m(L, \pi)$'s actions. To eliminate such equilibria, we will consider short-lived deviations optimal if they would be profitable when extended for a short amount of time. To formalize this, we define $\bar{V}_x(L, \pi, \epsilon)$ as x 's continuation utility under the following assumptions: the state is (L, π) at time t_0 , the policy π is *locked in* for a length of time $\epsilon > 0$ irrespective of the equilibrium path of play, and the equilibrium path of play continues at time $t_0 + \epsilon$. We will impose the requirement

that equilibria satisfy the following:

- (iii) If $V_{m(L,\pi)}(L, 1) = V_{m(L,\pi)}(L, 0)$ but $\bar{V}_{m(L,\pi)}(L, \pi', \epsilon) - \bar{V}_{m(L,\pi)}(L, 1 - \pi', \epsilon) > 0$ for all $\epsilon > 0$ small enough, then $\alpha(L, \pi) = \{\pi'\}$.

B.2 Analysis

The results are structured as follows. Lemmas 7, 8 and 10 are technical statements. Lemma 9 shows that agents strictly prefer the risky policy after a success. Proposition 14 shows that switches to the safe policy are permanent.

Lemma 7. *For any policy path $(\pi_t)_t$ with left and right-limits everywhere, there is another policy path $(\hat{\pi}_t)_t$ such that $\hat{\pi}_0 = \pi_0$, $(\hat{\pi}_t)_t$ is càdlàg for all $t > 0$, and $(\hat{\pi}_t)_t$ is equal to $(\pi_t)_t$ almost everywhere.⁵¹*

Proof of Lemma 7.

Define $\hat{\pi}_0 = \pi_0$ and $\hat{\pi}_t = \pi_{t+}$ for all $t > 0$. Let $\mathcal{T} = \mathbb{R}_{\geq 0} \setminus \{t \geq 0 : \pi_{t-} = \pi_t = \pi_{t+}\}$. Because $(\pi_t)_t$ has left and right-limits everywhere, \mathcal{T} must be countable—otherwise \mathcal{T} would have an accumulation point t_0 , and either the left-limit or right-limit of $(\pi_t)_t$ at t_0 would not be well-defined. Then, since $\hat{\pi}_t = \pi_t$ for all $t \notin \mathcal{T}$, $(\hat{\pi}_t)_t$ and $(\pi_t)_t$ only differ on a countable set. Moreover, it is straightforward to show that, for all $t > 0$, $\hat{\pi}_{t-} = \pi_{t-}$ and $\hat{\pi}_{t+} = \pi_{t+} = \hat{\pi}$, so $(\hat{\pi}_t)_t$ is càdlàg. ■

Corollary 2. *For any strategy profile $(\beta, \alpha, (L_t, \pi_t)_t | (L, \pi, \theta))$ the stochastic process $(L_t, \hat{\pi}_t)_t | (L, \pi, \theta)$ (where $(\hat{\pi}_t)_t$ is as in Lemma 7) has càdlàg paths, satisfies Conditions (a) and (b), and induces a path of play that yields the same payoffs as the strategy.*

Lemma 8 (Recursive Decomposition). *Let $\Theta \subseteq \mathbb{R}^n$ be a closed set, let $(\theta_t)_t$ be a right-continuous progressively measurable Markov process with support contained in Θ , let f be a bounded function, and let*

$$U(\theta_0) = \int_0^\infty e^{-\gamma t} E_{\theta_0}[f(\theta_t)] dt$$

Let Ψ be a closed subset of Θ and define a stochastic process $(\psi_t)_t$ with a co-domain $(\Psi \cup \{\emptyset\})$ as follows: $\psi_t = \theta$ if there exists $t' \leq t$ such that $\theta_{t'} = \theta$ and

⁵¹Hence it is payoff-equivalent to π_t and generates the same learning path $(L_t)_t$.

$\theta_{t''} \notin \Psi$ for all $t'' < t'$. If this is not true for any $\theta \in \Psi$, then $\psi_t = \emptyset$.⁵² Then

$$U(\theta_0) = \int_0^\infty e^{-\gamma t} E_{\theta_0} [f(\theta_t) \mathbf{1}_{\{\psi_t = \emptyset\}}] dt + \int_\Psi U(\theta) d\tilde{P}$$

where \tilde{P} is defined as follows: P_{ψ_t} is the probability measure on $\Psi \cup \emptyset$ induced by ψ_t , and $\tilde{P} = \gamma \int_0^\infty e^{-\gamma t} P_{\psi_t} dt$.

Proof of Lemma 8.

$$\begin{aligned} U(\theta_0) &= \int_0^\infty e^{-\gamma t} E_{\theta_0} [f(\theta_t)] dt = \\ &\int_0^\infty e^{-\gamma t} E_{\theta_0} [f(\theta_t) \mathbf{1}_{\{\psi_t = \emptyset\}}] dt + \int_0^\infty e^{-\gamma t} E_{\theta_0} [f(\theta_t) \mathbf{1}_{\{\psi_t \in \Psi\}}] dt \end{aligned}$$

So it remains to show that

$$\int_\Psi U(\theta) d\tilde{P} = \int_0^\infty e^{-\gamma t} E_{\theta_0} [f(\theta_t) \mathbf{1}_{\{\psi_t \in \Psi\}}] dt$$

Define a random variable z with a co-domain $(\Psi \times [0, \infty)) \cup \{\emptyset\}$ as follows: $\xi = (\theta, t)$ if $\theta_t = \theta \in \Psi$ and $\theta_{t'} \notin \Psi$ for all $t' < t$. If this is not true for any $\theta \in \Psi$ and $t \geq 0$, then $\xi = \emptyset$.⁵³ Let P_ξ be the probability measure on $(\Psi \times [0, \infty)) \cup \{\emptyset\}$ induced by ξ . Let $\theta(\xi)$ and $t(\xi)$ be the random variables equal to the first and second coordinates of ξ , conditional on $\xi \neq \emptyset$. Note that $\psi_t = \theta$ if and only if $\xi = (\theta, t')$ for some $t' \leq t$. Then we can write

$$\begin{aligned} &\int_0^\infty e^{-\gamma t} E_{\theta_0} [f(\theta_t) \mathbf{1}_{\{\psi_t \in \Psi\}}] dt = \int_0^\infty e^{-\gamma t} E_{\theta_0} [f(\theta_t) \mathbf{1}_{\{\xi \in \Psi \times [0, \infty)\}} \mathbf{1}_{\{t \geq t(\xi)\}}] dt = \\ &= \int_{\Psi \times [0, \infty)} \left(\int_{t(\xi)}^\infty e^{-\gamma t} E_{\theta_0} [f(\theta_t) | \xi] dt \right) dP_\xi = \int_{\Psi \times [0, \infty)} e^{-\gamma t(\xi)} U(\theta(\xi)) dP_\xi = \\ &= \int_{\Psi \times [0, \infty)} \left(\int_{t(\xi)}^\infty \gamma e^{-\gamma t} dt \right) U(\theta(\xi)) dP_\xi = \int_0^\infty \int_{\Psi \times [0, \infty)} \gamma e^{-\gamma t} \mathbf{1}_{\{t \geq t(\xi)\}} U(\theta(\xi)) dP_\xi dt = \\ &= \int_0^\infty \gamma e^{-\gamma t} \left(\int_{\Psi \times [0, \infty)} \mathbf{1}_{\{t \geq t(\xi)\}} U(\theta(\xi)) dP_\xi \right) dt = \int_0^\infty \gamma e^{-\gamma t} \left(\int_\Psi U(\theta) dP_{\psi_t} \right) dt = \int_\Psi U(\theta) d\tilde{P} \end{aligned}$$

as desired. ■

⁵²In other words, ψ_t takes the value of the first $\theta \in \Psi$ that $(\theta_t)_t$ hits.

⁵³In other words, ξ takes the value of the first $\theta \in \Psi$ that $(\theta_t)_t$ hits, and the time when it hits.

Lemma 9. *In any equilibrium, $\alpha(0, 1) = \alpha(0, 0) = 1$.*

Proof of Lemma 9.

Observe that $L_{t_0} = 0$ implies $L_t = 0$ for all $t \geq t_0$ no matter what policy path is followed, and hence $p(L_t, x) = 1$ for all t and x . For the rest of the argument, we can then write $V(0, \pi)$ instead of $V_x(0, \pi)$. By Lemma 8, there is $\rho \in [0, 1]$ such that

$$V(0, 0) = \rho \frac{s}{\gamma} + (1 - \rho)V(0, 1) \quad (8)$$

It follows that there exist $\eta \in [0, 1]$ and $\eta' \in [0, 1]$ such that $\eta \geq \eta'^{54}$ and

$$V(0, 0) = \eta \frac{s}{\gamma} + (1 - \eta) \frac{r}{\gamma} \quad V(0, 1) = \eta' \frac{s}{\gamma} + (1 - \eta') \frac{r}{\gamma}$$

η and η' are the discounted fractions of the expected time that the organization spends on the safe policy, when starting in states $(0, 0)$ and $(0, 1)$, respectively.

Observe that if $\eta > \eta'$, then $V(0, 0) < V(0, 1)$. In particular, $V_{m(0, \pi)}(0, 1) > V_{m(0, \pi)}(0, 0)$ for all π , which implies that $\alpha(0, \pi) = 1$ for all π by Condition (ii). If $\eta = \eta'$, then $V(0, 0) = V(0, 1)$. Because $\bar{V}(0, 0, \epsilon) < \bar{V}(0, 1, \epsilon)$ for any $\epsilon > 0$, by Condition (iii), in this case we must also have $\alpha(0, \pi) = 1$ for all π . ■

Lemma 10. *For any state (L, π) , there is a CDF G with support⁵⁵ contained in $[0, \infty]$ such that*

$$V_x(L, \pi) = \int_0^\infty W_T(p(L, x)) dG(T)$$

for all $x \in [0, 1]$, where $W_T(y)$ is as defined in Lemma 2.

Similarly, for any state (L, π) and any $\epsilon > 0$, there is a distribution G_ϵ with support contained in $[0, \infty]$ such that

$$\bar{V}_x(L, \pi, \epsilon) = \int_0^\infty W_T(p(L, x)) dG_\epsilon(T)$$

⁵⁴We have $\eta \geq \eta'$ for the following reason. $V(0, 1) = \eta' \frac{s}{\gamma} + (1 - \eta') \frac{r}{\gamma}$ and (8) imply that $\eta \frac{s}{\gamma} + (1 - \eta) \frac{r}{\gamma} = V(0, 0) = \rho \frac{s}{\gamma} + (1 - \rho)V(0, 1) = (\rho + (1 - \rho)\eta') \frac{s}{\gamma} + (1 - \rho)(1 - \eta') \frac{r}{\gamma}$. Then $\eta = \rho + (1 - \rho)\eta'$, which implies that $\eta \geq \eta'$, as required.

⁵⁵ G is a degenerate CDF that can take the value ∞ with positive probability. Equivalently, G satisfies all the standard conditions for the definition of a CDF, except that $\lim_{T \rightarrow \infty} G(T) \leq 1$ instead of $\lim_{T \rightarrow \infty} G(T) = 1$. This is needed to allow for the case where experimentation continues forever.

for any $x \in [0, 1]$.

Proof of Lemma 10.

We prove the first statement. The proof of the second statement is analogous.

Note that we can, without loss of generality, assume that the distribution over future states (L, π) induced by the continuation starting in state (L, π) satisfies the following: the policy is equal to 1 in the beginning and, if it ever changes from 1 to 0, it never changes back to 1. Indeed, suppose that π switches from 1 to 0 at time t and switches back at a random time $t + \nu$, where ν is distributed according to some CDF H . Let $p = \int_0^\infty e^{-\gamma\nu} dH(\nu)$. Then a continuation path on which the policy only switches to 0 at time t with probability $1 - p$ and never returns to 1 after switching induces the same discounted distribution over future states.

Under the above assumption and given that the policy always remains at 1 after a success by Lemma 9, the path of play can be described as follows: experimentation continues uninterrupted until a success or a permanent stop. Then we can let G be the CDF of the stopping time, conditional on no success being observed. ■

Proposition 14. *In any equilibrium, for any L , if $0 \in \alpha(L, 1)$, then $\alpha(L, 0) = 0$.*

Proof of Proposition 14.

If $L = 0$, then $\alpha(0, \pi) = 1$ for all π by Lemma 9, so the statement is vacuously true. Suppose then that $L > 0$. Suppose for the sake of contradiction that the statement is false.

Observe that for all L there is $\rho_L \in [0, 1]$ independent of x such that

$$V_x(L, 0) = \rho_L \frac{s}{\gamma} + (1 - \rho_L) V_x(L, 1) \quad (9)$$

for all x . This follows from Lemma 8, with the added observation that ρ_L (equivalently, \tilde{P} in the notation of Lemma 8) is independent of x in this case because the stochastic process governing (L, π) is independent of x if $\pi = 0$.⁵⁶ We now consider

⁵⁶If (L_t, π_t) has càdlàg paths, this follows from Lemma 8. If not, then Lemma 8 cannot be applied because the stochastic process in question is not necessarily right-continuous. However, we can use Corollary 2 of Lemma 7 to obtain a payoff-equivalent path of play with càdlàg paths and then apply Lemma 8 to it.

three cases.

Case 1: Suppose that $\rho_L > 0$, and that the expected amount of experimentation under the continuation starting in state $(L, 1)$ is positive. By Lemma 1, $V_x(L, 1)$ is strictly increasing in x . Then equation (9) implies that $V_x(L, 1) - V_x(L, 0)$ is strictly increasing in x . Since $m(L, 1) > m(L, 0)$, we have $V_{m(L,1)}(L, 1) - V_{m(L,1)}(L, 0) > V_{m(L,0)}(L, 1) - V_{m(L,0)}(L, 0)$. Since $1 \in \alpha(L, 0)$ implies that $V_{m(L,0)}(L, 1) - V_{m(L,0)}(L, 0) \geq 0$, we have $V_{m(L,1)}(L, 1) - V_{m(L,1)}(L, 0) > 0$, and thus $\alpha(L, 1) = 1$, a contradiction.

Case 2: Suppose that $\rho_L = 0$. We make two observations. First, $V_x(L, 0) = V_x(L, 1)$ for all x . Second, the expected amount of experimentation under the continuation starting in state $(L, 1)$ is positive. Indeed, $\rho_L = 0$ implies that, conditional on the state at t being $(L_t, \pi_t) = (L, 0)$, we have $\inf\{t' > t : \pi_{t'} = 1\} = t$ a.s. Since Condition (d) requires that π_{t+} exists and the result that $\inf\{t' > t : \pi_{t'} = 1\} = t$ a.s. rules out that $\pi_{t+} = 0$ with a positive probability, it must be that $\pi_{t+} = 1$ a.s. In turn, this implies that $\inf\{t' > t : \pi_{t'} = 0\} > t$ a.s. Then $E[\inf\{t' > t : \pi_{t'} = 0\}] - t > 0$.

By definition, we have

$$\bar{V}_x(L, 0, \epsilon) = \rho_\epsilon \frac{s}{\gamma} + (1 - \rho_\epsilon) V_x(L, 1) \quad (10)$$

for $\rho_\epsilon = 1 - e^{-\gamma\epsilon}$.

In the following argument, for convenience, we subtract $\frac{s}{\gamma}$ from every value function.⁵⁷ Note that, by definition, because we lock policy 1 in for time ϵ , G_ϵ satisfies $1 - G_\epsilon(T) = \min\left\{\frac{1-G(T)}{1-G(\epsilon)}, 1\right\}$. Then for $T \in [0, \epsilon]$, $1 - G_\epsilon(T) = 1$ and for

⁵⁷That is, we let $\check{V}_x(L, \pi) = V_x(L, \pi) - \frac{s}{\gamma}$, $\check{W}_T(x) = W_T(x) - \frac{s}{\gamma}$, $\check{\bar{V}}_x(L, \pi, \epsilon) = \bar{V}_x(L, \pi, \epsilon) - \frac{s}{\gamma}$. For the rest of this proof, we work with the normalized functions \check{V} , \check{W} , $\check{\bar{V}}$, but drop the operator \circ to simplify notation.

$T > \epsilon$, $1 - G_\epsilon(T) = \frac{1-G(T)}{1-G(\epsilon)}$. Hence for $\epsilon > 0$ sufficiently small we have

$$\begin{aligned}\bar{V}_x(L, 1, \epsilon) &= \int_0^\infty W_T(p(L, x)) dG_\epsilon(T) = \\ &= \int_0^\epsilon W_T(p(L, x)) dG_\epsilon(T) + \int_\epsilon^\infty W_T(p(L, x)) dG_\epsilon(T) = \\ &= 0 + \frac{1}{1-G(\epsilon)} \int_\epsilon^\infty W_T(p(L, x)) dG(T) = \\ &= \frac{V_x(L, 1)}{1-G(\epsilon)} - \frac{1}{1-G(\epsilon)} \int_0^\epsilon W_T(p(L, x)) dG(T) = \frac{V_x(L, 1)}{1-G(\epsilon)} + G(\epsilon)\mathcal{O}(\epsilon)\end{aligned}$$

The first part of the third equality follows from the fact that $G_\epsilon(T) = 0$ for all $T \in [0, \epsilon]$ and the second part of the third equality follows from the fact that $G_\epsilon(T) = \frac{G(T)-G(\epsilon)}{1-G(\epsilon)}$ for $T > \epsilon$. The last equality follows from the fact that $\lim_{\epsilon \rightarrow 0} G(\epsilon) = 0$ since $\inf\{t' > t : \pi_{t'} = 1\} = t$ a.s. and from the fact that $\left. \frac{\partial_+ W_T(x)}{\partial T} \right|_{T=0} = \max\{xr, a\} - s + \frac{xr(r-s)}{\gamma}$ by part (ii) of lemma 2.⁵⁸

Suppose for the sake of contradiction that $V_{m(L,0)}(L, 1) < 0$. Note that (10) then implies that $V_{m(L,0)}(L, 1) < \bar{V}_{m(L,0)}(L, 0, \epsilon)$. Therefore, we have $\bar{V}_{m(L,0)}(L, 1, \epsilon) \leq V_{m(L,0)}(L, 1) < \bar{V}_{m(L,0)}(L, 0, \epsilon)$ for all $\epsilon > 0$ sufficiently small and hence $\alpha(L, 0) = 0$, a contradiction. Hence $V_{m(L,0)}(L, 1) \geq 0$. It follows that, because $m(L, 1) > m(L, 0)$ and, by Lemma 1, $x \mapsto V_x(L, 1)$ is strictly increasing, we have $V_{m(L,1)}(L, 1) > 0$. Then $\bar{V}_x(L, 1, \epsilon) = \frac{V_x(L, 1)}{1-G(\epsilon)} + G(\epsilon)\mathcal{O}(\epsilon)$ implies that $\bar{V}_{m(L,1)}(L, 1, \epsilon) \geq V_{m(L,1)}(L, 1)$.⁵⁹ Moreover, because $V_{m(L,1)}(L, 1) > 0$, we have $V_{m(L,1)}(L, 1) > \bar{V}_{m(L,1)}(L, 0, \epsilon)$, as $\bar{V}_{m(L,1)}(L, 0, \epsilon)$ is a convex combination of $V_{m(L,1)}(L, 1)$ and 0.⁶⁰ Then $\bar{V}_{m(L,1)}(L, 1, \epsilon) \geq V_{m(L,1)}(L, 1) > \bar{V}_{m(L,1)}(L, 0, \epsilon)$ for all $\epsilon > 0$ sufficiently small. By Condition (iii), this implies that $\alpha(L, 1) = 1$, a contradiction.

Case 3: Suppose that the expected amount of experimentation starting in state $(L, 1)$ is zero. In this case $V_x(L, 0) = V_x(L, 1) = \frac{s}{\gamma}$ for all x , and $\bar{V}_x(L, 0, \epsilon) = \frac{s}{\gamma}$ for all x and $\epsilon > 0$. Again, we subtract $\frac{s}{\gamma}$ from every value function for simplicity.

⁵⁸In greater detail, $\int_0^\epsilon W_T(p(L, x)) dG(T) \approx \int_0^\epsilon (k_0 T + k_1) dG(T) \leq \int_0^\epsilon (k_0 \epsilon + k_1) dG(T) = (k_0 \epsilon + k_1) \int_0^\epsilon dG(T) = (k_0 \epsilon + k_1)(G(\epsilon) - G(0)) = (k_0 \epsilon + k_1)G(\epsilon) = k_0 \epsilon G(\epsilon) = G(\epsilon)\mathcal{O}(\epsilon)$ where we have used the fact that we subtracted $\frac{s}{\gamma}$ from every value function to get rid of the constant k_1 .

⁵⁹ $\bar{V}_{m(L,1)}(L, 1, \epsilon) \geq V_{m(L,1)}(L, 1)$ is then equivalent to $G(\epsilon)(1 - G(\epsilon))\mathcal{O}(\epsilon) \geq -G(\epsilon)V_{m(L,1)}(L, 1)$, which is satisfied for $V_{m(L,1)}(L, 1) > 0$.

⁶⁰Recall that we have subtracted $\frac{s}{\gamma}$ from every value function.

By definition, for all $\epsilon > 0$ the path starting in state $(L, 1, \epsilon)$ has a positive expected amount of experimentation. Moreover, G_ϵ defined in Lemma 10 is FOSD-decreasing in ϵ (that is, if $\epsilon' < \epsilon$, then $G_{\epsilon'} \geq G_\epsilon$) and hence, taken as a function of ϵ , has a pointwise limit G (that is, $G_\epsilon(T) \xrightarrow{\epsilon \rightarrow 0} G(T)$ for all $T \geq 0$). Then

$$\bar{V}_x(L, 1, \epsilon) \xrightarrow{\epsilon \rightarrow 0} \int_0^\infty W_T(p(L, x)) dG(T)$$

Since $1 \in \alpha(L, 0)$, there exists a sequence $\epsilon_n \searrow 0$ such that $\bar{V}_{m(L,0)}(L, 1, \epsilon_n) \geq 0$ for all n ,⁶¹ whence $\lim_{\epsilon \rightarrow 0} \bar{V}_{m(L,0)}(L, 1, \epsilon) \geq 0$.

There are now two cases. First, if $E_G[T] > 0$, we can use the following argument. $\lim_{\epsilon \rightarrow 0} \bar{V}_{m(L,0)}(L, 1, \epsilon) \geq 0$ implies that $\lim_{\epsilon \rightarrow 0} \bar{V}_{m(L,1)}(L, 1, \epsilon) > 0$ because $m(L, 1) > m(L, 0)$ and $x \mapsto V_x(L, 1)$ is strictly increasing by Lemma 1 (note that we have used the fact that $E_G[T] > 0$ to apply Lemma 1 here). Because $\epsilon \mapsto \bar{V}_{m(L,1)}(L, 1, \epsilon)$ is continuous, it follows that $\bar{V}_{m(L,1)}(L, 1, \epsilon) > 0$ for all $\epsilon > 0$ sufficiently small. But then $\alpha(L, 1) = 1$ by Condition (iii), a contradiction.

Second, if $E_G[T] = 0$, then we can employ a similar argument using the fact that, by part (ii) of lemma 2, $\left. \frac{\partial W_\epsilon(p(L, x))}{\partial \epsilon} \right|_{\epsilon=0}$ is strictly increasing in x and that, by Lemma 2, we have

$$\lim_{\epsilon \rightarrow 0} \frac{\bar{V}_x(L, 1, \epsilon)}{E_{G_\epsilon}[T]} = \lim_{\epsilon \rightarrow 0} \frac{W_\epsilon(p(L, x))}{\epsilon} = \left. \frac{\partial W_\epsilon(p(L, x))}{\partial \epsilon} \right|_{\epsilon=0}$$

■

The remaining results of the paper can then be proved in this model.

⁶¹Suppose for the sake of contradiction that $1 \in \alpha(L, 0)$ and such a sequence does not exist. Then for all $\epsilon > 0$ sufficiently small we have $\bar{V}_{m(L,0)}(L, 1, \epsilon) < 0$ (note that we have used the fact that we subtract $\frac{s}{\gamma}$ from every value function here). Then $\bar{V}_{m(L,0)}(L, 1, \epsilon) < \bar{V}_{m(L,0)}(L, 0, \epsilon) = 0$, which contradicts $1 \in \alpha(L, 0)$ by Condition (iii).