Regulating an Innovative Industry*

Steven Callander† Hongyi Li‡

October 11, 2024

Abstract

Innovations bring many benefits to society, but they can also bring harm. We study the problem of a regulator deciding whether to approve an innovation where information about the impact of the innovation is held within the firms that are developing it. We show that competition for the innovation undermines the regulator's ability to extract the information she needs to make good policy. As the number of firms increases and the expected benefit of the innovation grows, the probability that the regulator is persuaded to approve an innovation decreases. This tension between competition and communication reverses Arrow's famous "replacement effect." Thus, in regulated markets, more competition can lead to fewer innovations making it to market. We explore how this tension can be mitigated, but not eliminated, by political and market design.

*For helpful comments, we thank Takuo Sugaya, Dana Foarta, Peter Buisseret, Spencer Pantoja, Phedias Theophanous, Dave Baron, Al Roth, German Gieczewski, Nolan McCarty, Kevin Bryan, Oğuzhan Çelebi, Alessandro Lizzeri, Dan Carpenter, Robert Akerlof, Heski Bar-Isaac, Ali Lazrak, and particularly our discussant, Joshua Gans, and the audience of the NBER Economics of AI conference.
†Stanford University, Stanford Graduate School of Business; sjc@stanford.edu.
‡University of New South Wales, School of Economics; hongyi@hongyi.li.

1 Introduction

New ideas and new technologies are the lifeblood of society, driving social progress and economic development. Yet not all innovations are positive. The changes wrought by an innovation can bring harm to society, whether through the technology itself or by undermining existing social practices. The internet, for instance, has opened up the flow of information across society in many positive ways, but the emergence of social media, particularly its impact on teenage mental health, is less clear. Or consider the development of Artificial Intelligence (AI) and the prospect of machines advancing to the point that humans lose control.[1]

The possibility of harm creates a role for regulation. But regulating an innovative industry is hard. How is a regulator to know which innovations are beneficial to society and which are harmful? Designing and implementing policies in the face of such uncertainty makes it difficult for a regulator to extract the benefits of innovation while avoiding the harmful costs. This challenge is particularly relevant in an age when the technical demands to master new technologies are beyond the skills of almost all policymakers.

Fortunately, regulators need not fly blind. They can consult experts with deep knowledge of the underlying technologies and better information about their likely impact. These experts, however, often work for the firms that are developing and, therefore, hoping to profit from the innovations. The conflict of interest this generates makes extracting the information of experts far from straightforward.

This problem is particularly challenging in innovative markets as they are typically competitive. Market competition creates a wedge between market outcomes and outcomes for individual firms and this changes how firms lobby. OpenAI lobbies not for regulations that favor AI per se, but for regulations that favor OpenAI in market competition with Google, Apple, and the bevy of start-ups trying to capture the market. Competition requires the regulator to disentangle the often conflicting advice she receives from firms to determine what is in the interests of the market and society at large.

In this paper we develop a model to study this regulatory problem and focus on the role of market competition. The regulator must decide whether to approve an innovation that has the potential to benefit society or cause harm. If it is approved, firms compete in the market for the innovation, potentially disrupting the incumbent firm. The regulator can solicit advice from the firms—the firms can lobby the regulator—where the firms possess better knowledge of the impact of the new technology on society but are motivated by profit for themselves. To isolate the effect of competition, we consider a setting in which market and social outcomes

[1] The development of nuclear weapons provides a striking historical example. Edward Teller famously speculated that the chain reaction in an atomic explosion created a small but non-zero probability the atmosphere would ignite and destroy the earth. President Roosevelt chose to approve the use of the atomic bomb despite this (and other) not insignificant risk.

are aligned so that the only misalignment of firms with the regulator is that between market outcomes and the outcomes of individual firms.[2] We show that this misalignment generates a novel failure of regulatory decision making.

In standard analyses of markets, competition improves outcomes. The more firms that enter a market, the better the outcome for consumers. In the regulation of an innovative industry, the opposite can be true. Competition undermines outcomes in innovative industries because it undermines the regulator's ability to extract the information she needs.

In regulated markets, a fundamental tension emerges between competition and communication. Although an increase in the number of innovative firms increases the potential of the market, and increases the sources of expertise that are available to the regulator, the regulator learns less from lobbying and the probability that she is persuaded to approve the innovation decreases the more competitive is the market. We show that this dark side of competition can dominate such that few innovations make it to market despite the benefits they bring. In fact, for a regulator motivated primarily by avoiding harm, the probability of approving an innovation converges to zero as competition increases.

The logic of this result derives from how the regulator's incentives align differently with the different firms. The regulator's incentive aligns more closely with the incumbent firm as both benefit from the status quo. Thus, both want the innovation only if it improves on this outcome. In contrast, entrants to the market possess a very different incentive structure. The entrants gain nothing under the status quo and want the innovation to be approved whenever they benefit from it in the market, regardless of the impact it has on society. This differential alignment of interests means that whereas the incumbent firm can communicate productively with the regulator, communication with the entrants is much more limited, and in most situations it cannot benefit the entrants at all.[3]

The alignment between incumbent and regulator is imperfect, however. The incumbent's interests are not the market's interests, and the incumbent will recommend the innovation be approved only if the innovation is successful *and* the incumbent is the one that captures the benefit of the innovation. Precisely because an innovation can disrupt the market, the incumbent may not inform the regulator of its benefits. The probability that the incumbent is disrupted increases in the number of firms innovating. Thus, the more competition there will be in the market, the higher is the expected benefit from of the innovation, but the less likely it is that the incumbent communicates the benefits to the regulator.

Our result on competition reverses the conclusion of the famous *Arrow replacement effect.* Arrow (1962) posited that competition enhances innovation because an entrant replaces—or,

[2]This distinguishes us from the classic work of Baron and Myerson (1982) that studies monopoly and the conflict between market and social outcomes.

[3]This reflects an equilibrium selection that we discuss momentarily.

in the modern parlance, disrupts—the incumbent. Entrants have zero outside option and so they have more to gain and greater incentive to innovate. In a regulated market, it is precisely because the entrant's outside option is zero that the regulator cannot trust the information that the entrant provides. Entrants may have greater incentive to innovate, but if they cannot persuade the regulator to approve their innovations, they cannot capture the benefit. In a regulated market, competition and innovation are in tension.

This leads to the question of what can be done. One option is to stifle competition for the innovation. Without the threat of disruption, the incumbent could communicate freely with the regulator, but this comes at the cost of perpetual monopoly. Another option is to eliminate the regulator. This solves the communication problem by eliminating it, but leaves unchecked the risk that the innovation causes harm to society. The challenge is to obtain the benefit of market competition and regulation simultaneously.

In the second half of the paper we take up this challenge through market and regulatory design. We offer several adaptations that extract information from the market to ease the tension between competition and regulation and allow the benefits of both, at least in part, to be captured. The adaptations we pursue represent three distinct informational approaches based on the following questions: Can the incumbent be induced to reveal its information more often? Can the incumbent's information be made less important? Can the incumbent's information be changed?

We operationalize these information effects by altering the structure of market competition. The first adaptation reduces competition by delaying it, and shows how this can induce the incumbent to reveal its information more frequently without losing the benefit of competition altogether. The second adaptation increases competition, although in the market without the innovation rather than for the innovation itself. We show that this lessens the importance of the incumbent's information and improves the information flow about the innovation. The third adaptation allows for takeovers, changing the incumbent's information by changing the incumbent itself. The adaptations all improve policymaking, showing how market and regulatory design can alleviate the problems that arise when politics and innovative industries intersect.

Our results rest on a particular approach to lobbying. We suppose that firms lobby in their self-interest, supporting the policy that benefits them the most, and that if there is no policy that benefits the firm, the firm does not lobby at all. In effect, this selects what we refer to as a *regular* equilibrium to the lobbying game. Our setting with many firms corresponds to a game of multisender cheap talk. These games permit many equilibria. This includes babbling and equilibria in which the firms all coordinate on the truth and fully reveal their information to the regulator. These equilibria leave the firms indifferent over lobbying or not. We view it as reasonable in such situations to suppose that firms would instead simply not engage in the lobbying process, and this premise selects the regular equilibrium.

For concreteness, we take the decision of a regulator as our leading example. The ideas underlying our analysis are relevant to policymakers of all forms, whether legislators, members of the executive branch, or independent agencies. Indeed, the mechanism we identify is relevant beyond politics to decision making in firms and any organization in which insiders with superior information can exploit the uncertainty of decision makers to stifle innovation and change.

Innovation and Regulation in Practice

The model we develop describes many innovations observed in practice but not all. It captures innovations whose consequences are difficult to reverse and for which the cost of a mistake can be significant. Examples include nuclear power that risks the possibility of meltdowns and radiation poisoning, or the emergence of AI and the possibility of an artificial general intelligence (AGI) that supersedes humans, and even broad-based social programs like universal basic income that have long-term general equilibrium effects.

The model does not capture innovations of sufficiently small scale and with clearly identified harms such that a liability system restrains firm behavior or for which small-scale experiments are possible, such as with pharmaceutical development. Even with pharmaceutical regulation, however, general equilibrium effects beyond the power of random controlled trials to detect are possible. For example, the society-wide epidemic of Oxycontin occurred even though the drug passed all clinical trials and, to this day, is successful in the narrow medicinal context for which it was designed. (We return to this example in more detail in Section 5.) One may think of the domain of our model as the innovations that fit within Pandora's Box — innovations that once opened cannot be easily contained.[4]

The equilibrium behavior we identify resonates with practice and provides a new logic for the undue success of incumbent firms in lobbying to block the entry of new technologies. One example that is particularly illuminating is the development of FM radio in the 1930's (Wu, 2011). The incumbent firm, RCA, lobbied the FCC, arguing that FM broadcasts would interfere established AM radio networks as well as the more promising but nascent technology of television. The FCC cited these untestable fears in placing onerous restrictions on the use of FM radio that delayed for decades the widespread roll-out of the technology. This handicapped the firms entering the market to the benefit of the incumbent.

FM technology did ultimately make it to market, but the delay meant society missed out on the superior FM technology for well over a decade, and not coincidentally, when it did make it to market, the incumbent firm RCA was able to capture the benefit. Our model is

[4]Our focus on "Pandora's Box" innovations is in line with Gans (2024)'s prescription for regulators to be circumspect about enabling AI adoption if changes are irreversible and contained experiments are not feasible.

able to explain FM radio's regulatory experience as well as that of other technologies with a similar trajectory. This includes the particularly troubling examples of technologies that never made it to market because of an incumbent firm's lobbying. We do this without appealing to corruption or regulatory capture (Stigler, 1971), or to trust or a longstanding and ongoing relationship between the incumbent and the regulator, or to a revolving door. Rather, we show why the regulator will rationally listen to *only* the incumbent firm, why outcomes like this are pervasive, and what can be done so that regulators can make better use of the information held by firms.

Related Literature

In studying the regulation of a *competitive market*, our work is complementary to Baron and Myerson's (1982) seminal model of regulating a *monopolist*. Baron and Myerson (1982, p.912) study situations in which market competition is completely absent because "there are no other producers capable of supplying the product efficiently." We identify a tension created by market competition itself, showing how the presence of other producers can make regulation more, not less, difficult.

Our different emphasis leads us to study different preferences for the regulator. In studying monopoly, the misalignment that arises in Baron and Myerson (1982) is that between market and social outcomes. In competitive markets a second misalignment opens up, that between market outcomes and the outcomes for individual firms. We focus on this misalignment to most clearly illustrate the impact of market competition on regulation.

We also differ from Baron and Myerson (1982) and the literature that followed in how we approach the regulatory problem. Baron and Myerson (1982) takes a mechanism design approach, assuming commitment on behalf of the regulator and allowing transfers. Our approach is with neither commitment nor transfers, and we allow only cheap talk communication between the firms and the regulator.

The relationship between market structure and innovation has long been a focus in economics. It has been described as the second-most studied question in all of industrial organization and led to a growing literature on how government policy can foster innovation (Bryan and Williams, 2021). What has received less attention is the strategic behavior of policymakers. Whereas the literature focuses on the incentive to innovate, we take innovation as given and ask whether it will be approved by a regulator and make it to market. Our conclusion aligns with Schumpeter's (1942) conclusion that monopoly power enhances innovation although our mechanism is completely different to his. In fact, the ingredients for our result are closer to those of Arrow (1962) even though he obtains the opposite conclusion. This shows how the addition of strategic policymakers and asymmetric information can upend the standard relationship between competition and innovation.

Recent work by Acemoglu and Lensman (2024) and Gans (2024) share our interest in how to regulate an innovation with potential for societal harm. They consider settings where a regulator chooses the rate at which an innovation is adopted and directly observes the consequences of the innovation over time. In doing so, they abstract completely from the key elements of our analysis: asymmetric information, lobbying, and how market competition distorts strategic lobbying by firms and regulatory decision making.

Our model fits with a recent series of papers that connect market power and political power (Callander et al., 2022, 2023; Cowgill et al., 2023). These papers are complete information models of quid-pro-quo lobbying (i.e., the buying of favors). We introduce asymmetric information into the relationship between markets and politics.

The large literature on informational lobbying in political science, beginning with Gilligan and Krehbiel (1987), has situated expertise within the government, thereby setting aside the interaction between markets and politics.[5] Two exceptions are McCarty (2017) and Carpenter and Ting (2007).[6] These papers model a single firm lobbying a regulator whereas we focus on many firms and the interaction of market competition and political decision making.[7]

2 The Model

Setting: There are n firms $1, 2, \ldots, n$ (he) and a regulator (she). The regulator has to choose between the status-quo policy, $p = 0$, in which a new technological innovation is prohibited, and a set of k pro-innovation policies, $p \in \{1, 2, 3, ..., k\}$ that approve the innovation in some form.

For $k = 1$ the regulator's decision is simply to approve or not the new innovation. When $k > 1$ the regulator decides not only whether to approve the innovation but how to regulate it. She may, for example, allow unfettered versus restricted use, apply some "guardrails" against misuse, or allow some or all applications of the innovation. We say informally that the innovation is *approved* if the regulator chooses any policy $p \geq 1$.

Technology: The firms develop products to exploit the innovation. The quality of firm j's product under policy p is denoted by $q_j(p)$. For a pro-innovation policy $p \geq 1$, each quality $q_j(p) \geq 0$ is an independent draw from a distribution $F(x)$. We impose minimal restrictions

[5] McCarty (2017, p. 1221) observes, "In this regard, the approach in political science is quite different from that of regulatory economics, where the regulator's extraction of information from the regulated firm is the central problem."

[6] See also McCarty (2013). Many models interpret the lobbyist as a firm or other outside actor, although they model the preferences of the outsider as equivalent to those of another legislator or bureaucrat and do not include any features of a market.

[7] Gilligan and Krehbiel (1989) introduced multiple experts to the lobbying literature in the traditional setting in which expertise is situated within the legislature.

on $F(\cdot)$, assuming only that it is atomless, with strictly positive density $f(x)$ for all $x \in [0, \infty)$, and with finite mean.

The Market: The market is winner-take-all.[8] For each policy p, denote the firm with the highest quality by $w(p)$ and its quality by $q^{\text{win}}(p)$, such that:

$$w(p) = \underset{j}{\text{argmax}}\, q_j(p) \qquad \text{and} \qquad q^{\text{win}}(p) = q_{w(p)}(p).$$

We refer to firm $w(p)$ as the market winner under policy p.

We assume the winner of market competition under the status-quo policy has already been determined as firm 1. Denote its winning status-quo quality by $q_1(0) = q^{\text{SQ}}$. We refer to firm 1 as the *incumbent* and the other firms as *entrants.*

Preferences: The profit of the winning firm under policy p is equal to his quality and all other firms receive zero:

$$u_j(p) = \begin{cases} q_j(p) = q^{\text{win}}(p) & \text{if } j = w(p), \\ 0 & \text{otherwise.} \end{cases}$$

To focus on the effects of competition on communication, we assume that the utility of the regulator is (weakly) increasing in the winning firm's quality, $q^{\text{win}}(p)$. This means the regulator favors innovation but does not care which firm captures the benefit of the innovation. We also assume that the regulator is more averse to harm from an innovation than are the firms, capturing the reality that the harm is most likely an externality borne by the regulator and society rather than the firms themselves. To represent these features simply, we set the regulator's utility to be the social surplus from market competition, minus a penalty $\Delta \geq 0$ which kicks in below some threshold $q^L > 0$. Specifically:

$$U(p) = \begin{cases} q^{\text{win}}(p) & \text{if } q^{\text{win}}(p) \geq q^L, \\ q^{\text{win}}(p) - \Delta & \text{if } q^{\text{win}}(p) < q^L. \end{cases}$$

Formally, an aversion to harm translates into risk aversion on behalf of the regulator, which is a commonly ascribed characteristic of regulators in practice.[9] While a regulator's greater aver-

[8] Winner-take-all is common in markets for new technologies. This assumption simplifies our presentation although it is not essential for our results. See the discussion in Section 5.

[9] Risk aversion in politics can derive from career concerns—the loss of an election or a job following a bad outcome—or be due to the reputational concerns of bureaucracies motivated by funding, recruitment, and credibility, as detailed in Carpenter (2002, 2004). As it is typical in the industrial organization literature to assume firms are risk neutral, any risk aversion on behalf of the regulator ensures she is more risk averse

sion to harm is not essential for our results, the effect we identify becomes more pronounced when $\Delta > 0$, suggesting that the incentives regulators typically encounter in practice make them more vulnerable to the policymaking failure we have identified. We set $q^{\text{SQ}} \geq q^L$ such that the regulator obtains positive utility under the status quo policy.

Given these preferences, the first-best outcome, q^{FB}, is simply the highest quality draw across all firms and all policies, i.e. the best of the market winners:

$$q^{\text{FB}} = \max\{q^{\text{SQ}}, q^{\text{win}}(1), q^{\text{win}}(2), \ldots, q^{\text{win}}(k)\}.$$

Information and Messages: The firms observe the full realization of their own quality and that of all other firms.[10] The regulator knows the distribution of product quality but not the realisations. The firms lobby the regulator by recommending a policy choice or they do not lobby, which we represent by an "empty" message, such that $m \in \{0, 1, 2, 3, ..., k, \varnothing\}$.

Timing:

1. Nature draws the product quality for all n firms for all policies $p \geq 1$, $\boldsymbol{q} = \{q_1(1), \ldots, q_n(1), q_1(2), \ldots, q_n(k)\}$. Firms observe $\boldsymbol{q}$.
2. Firms simultaneously send messages $m_j \in \{0, 1, 2, ..., k, \varnothing\}$.
3. The regulator chooses a policy $p \in \{0, 1, 2, 3, ..., k\}$.

Equilibrium Selection. A firm j's lobbying—or messaging—strategy is *regular* if he sends the empty message $m_j = \varnothing$ when indifferent over all policies, and otherwise recommends the policy that, if implemented, would maximize his profit:

$$m_j^* = \underset{p}{\text{argmax}}\{u_j(p)\}.$$

In all of the equilibria we consider, the entrants use a regular strategy. When the incumbent firm also uses a regular strategy, we refer to this as a *regular equilibrium*.[11] The regulator chooses her optimal policy given the lobbying strategy of firms, denoted by p^*.

than the firms. The risk aversion of the regulator may also be interpreted as a direct representation of the preferences of society.

[10]Alternatively, we could assume firms observe only their own quality without fundamentally changing the logic of our results; see the discussion in Section 5.

[11]It is straightforward to construct off-path beliefs for the regulator that support the regular equilibrium and other equilibria we characterize. We describe these beliefs in full in the appendix.

Discussion of Modeling Assumptions

Lobbying: By focusing on regular strategies, we select an equilibrium where each firm lobbies in a manner consistent with transparent self-interest: (i) he communicates *as if* he expects the regulator to take him at his word in adopting his recommendation, while (ii) he refrains from lobbying whenever he is indifferent between all policies and so cannot possibly benefit from influencing the regulator's decision. In our view, such transparent self-interest seems to be a reasonable description of lobbying behavior.

It is instructive to compare our regular equilibrium to the simple and well-known fully-revealing equilibrium where all firms coordinate on recommending the policy that is the first-best for society. With three or more firms, the regulator can detect (and ignore) any unilaterally deviating firm.[12] In this fully-revealing equilibrium, only those firms that win under some policy can potentially benefit from lobbying. The other firms are indifferent between all policy outcomes; yet they actively lobby, not to their own benefit, but in support of the equilibrium and effectively for the benefit of the winning firm(s). Such equilibrium logic, which holds in some form in most existing models of lobbying with competing experts, runs counter to the spirit of transparent self-interest that motivates our regular equilibrium.[13]

Status quo product quality: One may interpret the dominance of the incumbent firm under the status quo in several ways. It may be that the entrants do not compete under the status quo or that they did compete and were defeated (or, in the dynamic extension of the model to come later, that they haven't yet competed). The interpretation that they did not compete fits novel technologies that are very different from what has come before, such as Amazon competing only online against bricks-and-mortar stores. The interpretation that the firms did compete and lost fits incremental innovations, or innovations that have already come to market in some form but for which new uses emerge. Consider the introduction of generative AI technology to the market for internet search with Google as the incumbent and entrants that include other tech companies that lost the search wars.

Innovation product quality: A feature of our model is that approval of the innovation may lead to a decrease in the incumbent firm's product quality. This can arise for multiple reasons. It may be that the policy change itself disallows the incumbent's status quo product, compelling the incumbent to try something new and possibly developing a less successful product. For example, consider new technology that allows self-driving cars to coordinate their movements almost perfectly. For this coordination to work well, the existing status-quo

[12] The literature has focused almost exclusively on the case of two experts such that the identity of the deviator is not apparent to the receiver (the regulator) whose response must punish both senders.

[13] Our approach is akin to the "coordination-free" approach of Lu (2017, p. 178), who argues against the precise coordination necessary to support full revelation in cheap-talk games. He derives a class of coordination-free equilibria and shows that senders beyond the second add only modest benefit. We show that when the senders are market competitors, more senders can actually lead to worse communication.

product—human-operated cars—cannot share the roads. Thus, policy change that allows the new technology will likely restrict the road access of human-operated cars to some degree.

The new technology may also directly degrade the status quo product even if it is still allowed. This was the argument against FM radio offered by the incumbent firm RCA, that the broadcast of FM signals would interfere with and degrade AM radio signals and the nascent technology of television. Another example is the internet search market following the introduction of generative AI technology, where AI-powered chatbot products such as chatGPT now compete with traditional search engines such as Google (the incumbent) for queries. It is commonly accepted that the increasing use of generative AI has "polluted" the internet with AI-generated web content and, thus, reduced the quality of traditional search-engine results.[14]

3 Results

3.1 Monopoly: Incumbent as Innovator

Consider the situation where the incumbent firm is the only innovator ($n = 1$). As a monopolist, his interest's are the market's interests and perfectly aligned with the regulator. The incumbent recommends the innovation be approved only if his profits increase. As this implies social welfare also increases, the regulator follows the incumbent's recommendation and social welfare is maximized.

Proposition 1 *With only one firm, $n = 1$, a regular equilibrium exists and the regulator follows the incumbent's recommendation, $p^* = m_1^*$. The incumbent recommends the policy for which his product has highest quality, $m_1^* = p^m$, where $p^m = \text{argmax}_p\{q_1(p)\}$.*

The incumbent's strategy for $k = 3$ is depicted in Figure 1. Both policies 2 and 3 increase the incumbent's quality, and he recommends the higher of these. The incumbent would never recommend policy 1 as this leads to lower quality. The regulator follows the incumbent's advice and maximizes her own utility.

Without competition, the regulator is able to learn all that she needs to know to make an optimal decision and the first-best outcome is achieved. Information that is held in the market is transferred to the regulator and used efficiently in policymaking. This establishes a clear benchmark to see the impact of market competition.

[14] Further examples can be found in many fields. One example from drug development is the introduction of methicillin as an antibiotic. The use of methicillin to treat bacterial infections led to the emergence of Methicillin-Resistant Staphylococcus Aureus (MRSA) bacteria that evolved broad-spectrum resistance, not only to methicillin (the entrant) but also to a wider class of existing antibiotics (status quo products).

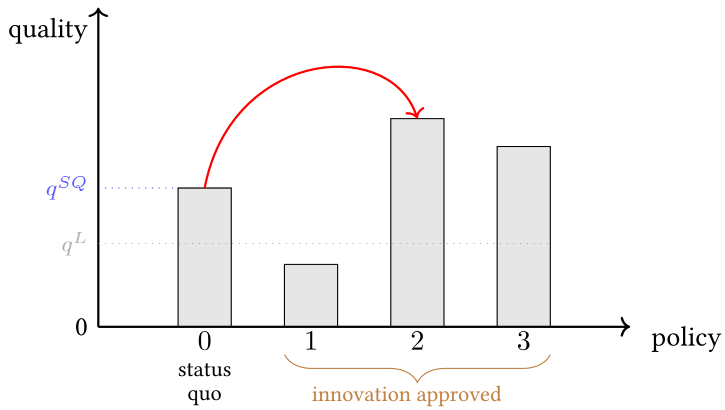

Figure 1: Monopoly: Incumbent as Innovator

3.2 Competition: Many Innovators

Competition brings more innovators to the market. Mechanically, this increases the expected benefit to society. As the number of firms increases, there are more draws of quality and the first-best outcome, q^{FB}, increases in expectation as it is the maximum of those draws. To obtain this outcome, however, the regulator must learn which policy leads to the first-best. As we will see, a tension between competition and communication makes this difficult. We first consider what the regulator would do if there were no lobbying at all.

Competition Without Lobbying.
That the first-best is increasing in n implies that the regulator is more inclined to approve the innovation the more intense is competition. In fact, the regulator's *default* policy—the policy she would implement based on her prior—shifts from the status-quo policy to a pro-innovation policy if competition is intense enough that the expected quality of the winning product overcomes her aversion to harm. The required level of competition is increasing in Δ, the regulator's aversion to harm, and for $\Delta = \infty$ no level of competition will shift the regulator from the status-quo default. For brevity, we first state results in this section for when the regulator either approves or prohibits the innovation ($k = 1$).

Lemma 1 *There exists, for $k = 1$, a threshold level of competition $\hat{n}(\Delta)$ such that the regulator's default is to prohibit innovation if $n \leq \hat{n}(\Delta)$ and to approve innovation if $n \geq \hat{n}(\Delta) + 1$. Moreover, the threshold $\hat{n}(\Delta) \in \{1, 2, \dots\}$ is increasing in Δ and $\hat{n}(\Delta) \to \infty$ as $\Delta \to \infty$.*

The default policy does not affect what the regulator learns from lobbying, but it does affect how the regulator responds to that information.

Competition and Lobbying.
Competition has two significant effects on lobbying. On one hand, it expands the sources of information available to the regulator. On the other hand, it disrupts the alignment between

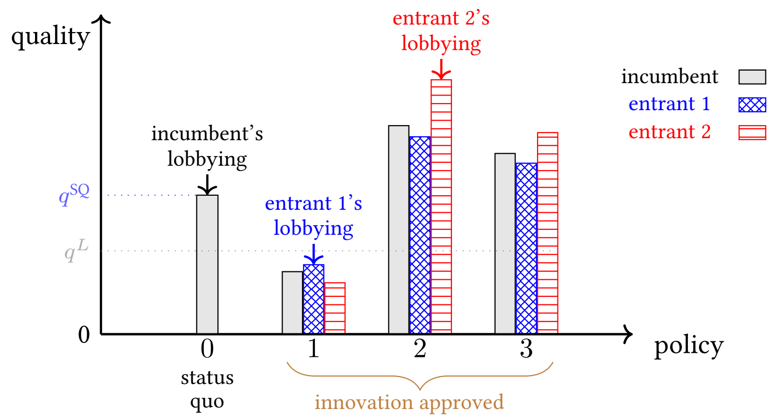

Figure 2: Lobbying with Competition

the interests of individual firms and the overall market. While more firms may engage in lobbying, the interests of any single firm no longer align with what is good for the market as a whole.

To see the misalignment caused by competition, consider the incentives of the incumbent firm. The incumbent lobbies the regulator to approve the innovation if his product quality goes up *and* he wins the market. Otherwise he lobbies for the status quo. Thus, the incumbent continues to lobby for the innovation if it delivers a benefit but now he only does so if he is the one who captures the benefit.

Despite this, the regulator can still rely on the incumbent's information. The incumbent never lobbies for an innovation that causes harm, and every innovation he lobbies for benefits the regulator as well. The cost to the regulator of misalignment with the incumbent is a sin of omission. The incumbent does not tell the regulator about an innovation that would benefit her if it will only lead to the incumbent's disruption, even if the incumbent's product quality—and the regulator's utility—are higher. Figure 2 illustrates with an example: the incumbent achieves his highest product quality under policy 2 but—because he is not the market winner there—recommends the status quo instead.

In contrast to the incumbent, entrants lobby for the innovation too much rather than too little. Entrants sin by commission rather than omission. Because an entrant's outside option is zero, he lobbies for the innovation even when his quality is lower than the status quo and, in particular, lower than the regulator's threshold for harm, q^L. When an entrant lobbies, the regulator learns nothing about the quality of the product, only that the entrant lobbying is the one that will win the market if that policy setting is chosen. As a result, the regulator does not follow the advice of the entrants.

Proposition 2 describes the equilibrium. It establishes that a regular equilibrium exists, and describes how the regulator responds to lobbying given her default policy when each firm's lobbying is conditional on them winning the market.

Proposition 2 *A regular equilibrium exists for $k = 1$. Each firm recommends the policy for which his product has highest quality and he wins the market, $m^*_j = p^c_j$, where $p^c_j = \text{argmax}_p\{q_j(p) \mid w(p) = j\}$.*

- ***(i)*** *For $n \leq \hat{n}(\Delta)$, the regulator follows the incumbent's recommendation.*
- ***(ii)*** *For $n \geq \hat{n}(\Delta)+1$, the regulator approves the innovation, $p^* = 1$, if any firm recommends it ($m^*_j = 1$ for any j), otherwise she prohibits the innovation, $p^* = 0$.*

When the regulator's default is the status-quo (case (i)), she must receive positive information about the innovation to be persuaded to approve it. This can come only from the incumbent and only when he wins the market under the innovation. Thus, the regulator sticks to her status-quo default unless the incumbent lobbies her to approve the innovation. In this case, the regulator's decision always follows the advice of the incumbent.

The situation is reversed when the regulator's default is the pro-innovation policy (case (ii)). In this case, the regulator shifts from her default if she receives negative information about the innovation. Surprisingly, this is possible even by observing the entrants, although only by what they don't say rather than what they do say. If no firm lobbies for the innovation, the regulator learns two things. She learns that it must be the incumbent who wins the market under the innovation—as otherwise an entrant would lobby—and because the incumbent nevertheless lobbies for the status quo policy, the regulator infers that the incumbent's quality and, thus, her own payoff must be higher than under the innovation, and she reverts to the status quo. If an entrant lobbies, the regulator learns nothing about the quality of the innovation, but by not lobbying, the entrants convey valuable negative information to the regulator. This information is credible precisely because the entrants do not benefit from it.

In case (ii), therefore, the regulator follows the incumbent's advice but only conditionally. In particular, the regulator sticks to her default and approves the innovation if an entrant lobbies even though the incumbent lobbies for the status quo. To an outside observer, this gives the appearance of an entrant successfully persuading the regulator to approve an innovation over the objections of the incumbent. However, ascribing such influence to the entrant would be mistaken. The entrant doesn't so much persuade the regulator as he reinforces the regulator's default policy. The entrant's lobbying provides no positive information to the regulator; it is nevertheless effective as it removes the incumbent's capacity to dissuade the regulator from approving the innovation.

The regulator makes different types of policy mistakes depending on which type of firm influences her decision. When the regulator's default is the status-quo, she listens to the incumbent. As the incumbent's lobbying exhibits a sin of omission—lobbying too infrequently for the innovation—the regulator's policy mistake is one of omission too, approving too few innovations. In contrast, when the regulator's default is pro-innovation, she listens to the incumbent only conditionally, sticking to her default whenever an entrant lobbies for it. As the entrant's lobbying exhibits a sin of commission—lobbying too frequently—the regulator's policy mistake is a sin of commission too. In the former case the mistake is unobserved—it resides in the beneficial innovations that never make it to market. In the latter case, the mistake is evident, manifesting in the harm caused by innovations that should not have been approved. In both cases, the regulator's mistakes come from sticking to her default too tightly; she is never persuaded to make a mistake by switching policy.

The relationship between the incumbent and the regulator also changes as the default policy changes. When the default is the status quo, the relationship is Pareto efficient. There is no policy that would make both the regulator and the incumbent better off. In contrast, when the default is pro-innovation, the regulator will approve innovations that cause harm, leaving both her and the incumbent worse off. An observer might be tempted to see this difference as a difference in the degree of regulatory capture across domains. As as our model makes clear, however, the difference is purely a function of the default policy rather than the fundamentals of lobbying and policymaking.

Competition and Innovation in a Regulated Market

Proposition 2 implies that competition is a cursed benefit in a regulated market. This reflects the fundamental trade-off between competition and communication. Competition increases the potential benefit that an innovation can deliver, but because it does so by disrupting the incumbent, the flow of information to the regulator is diminished and the quality of policymaking suffers. The incumbent can only communicate productively with the regulator when he wins the market for the innovation—that is, exactly when he *isn't* disrupted. Consequently, holding the default policy constant, competition degrades the informational content of lobbying and the quality of policymaking precisely because it increases the chance of disruption.

Corollary 1 *For $k = 1$, in the regular equilibrium, the probability that the regulator's policy choice differs from her default because of lobbying strictly decreases on $n \in \{1, \dots, \hat{n}(\Delta)\}$ and on $n \in \{\hat{n}(\Delta) + 1, \dots\}$ as the number of firms n increases.*

The negative effect of competition on policymaking is particularly deleterious when the regulator's default policy is the status quo. In this case, an increase in competition makes society strictly worse off. This is despite the fact that a more competitive market increases the expected quality—and decreases the expected harm—of innovation. Up to the threshold at which the regulator's default policy shifts, more competition not only fails to increase innovation in a regulated market, it actively undermines it.

Corollary 2 *For $k = 1$ and $n \in \{1, \dots, \hat{n}(\Delta)\}$ in the regular equilibrium, as the number of firms n increases,*

(i) *the probability $\Pr[p^* = 1]$ that the innovation is approved strictly decreases,*

(ii) *the regulator's utility $U(p^*)$ strictly decreases.*

In the special case of a regulator motivated primarily by avoiding harm ($\Delta = \infty$): as $n \to \infty$,

(iii) *the probability that the innovation is approved vanishes, $\Pr[p^* = 1] \to 0$, even though the expected quality of the innovation increases without bound, $E[q^{win}(1)] \to \infty$,*

(iv) *the regulator's expected utility converges to the status-quo quality, $E[U(p^*)] \to q^{SQ}$.*

The negative relationship between competition and innovation in Corollary 2 reverses the conclusion of the famous *Arrow replacement effect.* Arrow (1962) showed in a market without regulation that competition induces more innovation because an entrant has greater incentive to innovate. In Arrow's market, the entrant's incentive to innovate is greater than the incumbent's precisely because it starts with nothing, and thus gains the monopolist's profit plus the incremental improvement from the innovation. Our results show that the exact opposite logic holds in a regulated market when the regulator's default is the status quo. Precisely because an entrant starts with nothing, his advice is not credible. Thus, even when innovation is a given, as in our model, the entrant cannot persuade the regulator to approve it and open up the market. The entrant fails to persuade the regulator not despite, but because, of the firm's greater incentive to innovate.

Arrow's effect is nominally restored for competition above the threshold, although this is a mechanical by-product of the assumptions in our model. Indeed, in assuming that the quality of each firm's product is a draw from an unbounded distribution, we assume both that innovation is given and that the winning quality will grow without bound as competition intensifies. In this domain, our result instead shows how competition restrains a regulator's ability to intervene in a market to avoid harmful innovation.[15] Conditional on an innovation being harmful to society, the more competition there is, the less likely the incumbent is the market winner, and therefore, the less likely that the regulator learns about the harm and prohibits the innovation. Thus, while a lot of competition does bring more innovation to market, it bring more harmful innovation as well.

Restricting (and Encouraging) Competition:
A striking implication of Corollary 2 is that if competition were below the threshold $\hat{n}(\Delta)$, the regulator and society would be better off if competition were eliminated altogether. That

[15] Arrow did not allow for this possibility. He assumed implicitly that all innovation is beneficial.

is, if given the choice of any level of competition up to $\hat{n}(\Delta)$, the regulator will choose to restrict competition even to the point of granting monopoly power to the incumbent. We capture this inversion of the standard logic of market competition in the following corollary.

Corollary 3 *Suppose the regulator can restrict market access by choosing any* $1 \leq n \leq \hat{n}(\Delta)$*. Conditional on the regular equilibrium, it is optimal for her to give the incumbent a monopoly position,* $n = 1$*.*

To an outside observer, the regulator's choice to block all competition would give the appearance of regulatory capture. In our model the choice emerges from the regulator's own preference. The regulator chooses to protect the incumbent not *against* but rather *for* the public interest. If the regulator's aversion to harm reflects her own career interests rather than society's interests, then the regulator's choice to protect the incumbent is more problematic. We return to regulatory capture and this interpretational issue in Section 5.1.

Corollary 3 holds only for competition below the threshold, $\hat{n}(\Delta)$. Above the threshold, more competition can be better for the regulator, and for $n \to \infty$ the chance of harm becomes vanishingly small, even as lobbying becomes increasingly uninformative. If intense competition can be induced, it can be the salve to the problem of competition that we identify in our model as it effectively eliminates the risk of harm.[16]

In practice, however, the regulator's ability to influence entry is asymmetric. It is very possible for the regulator to restrict entry, but much harder to induce it. For innovative technologies, there is good reason to believe that entry is neither free nor easy, and that it is due to forces outside the regulator's control.

Two barriers to entry, in particular, come to mind. The first is the classic logic of market entry. Firms must invest in R&D and pay other start-up costs of entry to enter the market. More intense competition lowers the probability that any individual firm wins, and, thus, lowers the incentive to innovate in the first place. This bounds the level of innovative competition that the regulator can achieve.

The second barrier is novel and emerges from our model. Entry into an innovative industry involves a market-wide collective action problem for the entrants. If few or no entrants choose to invest in R&D, then the incumbent wins as the regulator will block the innovation otherwise. Thus, even for small entry costs, it follows that no entry will be equilibrium behavior.[17] To avoid this trap, the entrants must coordinate on innovation. Counter-intuitively, each entrant

[16] However, if there is common uncertainty about product quality under the innovation, then the risk of harm may remain even under intense competition.

[17] To fix ideas, consider an extension where entrants can, before product qualities are realized, choose to pay an entry cost and compete with the incumbent. Then even for small costs, an zero-entry regular equilibrium always exists: firms only enter the market if they know the regulator would listen to them, but the regulator listens to entrants only if sufficiently many firms enter the market. Further, for any positive entry cost, if Δ is sufficiently large, the unique regular equilibrium involves zero entry.

needs others to also innovate to even be able to crack the market open. Their ability to pull this off is beyond the scope of our paper, but would appear to be difficult in practice – particularly given innovation is inherently an uncertain occurrence that can strike at random times. This collective action problem may thus limit the regulator's ability to foster market competition.

Absent barriers to entry, competition can be good for the regulator, and if competition exceeds the threshold at which the regulator's default switches ($n \geq \hat{n}(\Delta) + 1$), more competition can be better. But if barriers to entry have sufficient bite that entry falls below this threshold ($n \leq \hat{n}(\Delta)$), then the regulator is better off *further* limiting competition – even to the extent of granting the incumbent a monopoly.

A Larger Policy Space ($k > 1$)

The equilibrium logic extends to a richer policy space.[18] For a regulator who is motivated primarily by avoiding harm, the extension is immediate.

Proposition 3 *A regular equilibrium exists for $k > 1$ when $\Delta = \infty$. Each firm will recommend, from amongst the policies where his product wins the market, the policy where his product has highest quality: $m_j^* = p_j^c$ where $p_j^c = \text{argmax}_p\{q_j(p) \mid w(p) = j\}$. If a firm does not win the market under any policy, he does not lobby. The regulator follows the incumbent's recommendation, $p^* = m_1^*$.*

For this regulator, the default policy is to prohibit the innovation regardless of the level of competition (as $\hat{n}(\Delta) = \infty$), and the regulator follows the advice of the incumbent firm. An interesting novelty in the larger policy space is that policy can now be inefficient even when the regulator approves an innovation that benefits society. As the incumbent recommends the pro-innovation policy that is most favorable to himself, and the regulator follows his advice, the regulator may miss out on other policy settings that lead to even higher quality products but in which an entrant disrupts the incumbent.

The same logic holds when the regulator's aversion to harm is finite, although the regulator's decision rule becomes more subtle for levels of competition just below the threshold $\hat{n}(\Delta)$. In particular, the regulator may now condition her decision on the number of entrants who lobby for a policy. We omit the details in the interests of space.

4 Regulatory and Market Adaptations

In this section, we consider adaptations of the regulatory and market system to alleviate the policymaking problems identified in the baseline model of the previous section. The three

[18] The threshold $\hat{n}(\Delta)$ continues to apply as the best the regulator can do when she approves the innovation is to select a policy setting at random.

adaptations we develop focus on three distinct approaches to how information flows between the firms and the regulator. First, how can the incumbent be induced to reveal his information more often? Second, how can the incumbent's information be made less important? And, third, how can the incumbent's information be changed to improve outcomes? We address these questions in turn in the subsections that follow.

Throughout this section we allow, unless otherwise specified, for arbitrary $k \geq 1$ and, for simplicity, specialize to a regulator motivated primarily by avoiding harm ($\Delta = \infty$). The first two adaptations are dynamic. To accommodate them, we extend the static model of Section 2 as follows. In each period $t = 1, 2, ...$ the static model recurs. Firms lobby the regulator, the regulator chooses a policy setting, and the firms compete in the market. Product qualities are drawn once at the beginning of play and are persistent. The regulator's stage payoff is $U(p_t)$ and each firm j's stage payoff is $u_j(p_t)$. The firms and regulator have a common discount factor $\delta \in (0, 1)$. We refer to this set of assumptions as the *dynamic model*.

4.1 Adaptation I: Monopoly Protection

The failure of communication hurts the regulator and society, and it also hurts the incumbent firm. The incumbent would like to utilize his information more, but only if he is able to benefit from it. In this section we explore how the incumbent can be induced to reveal his information more often.

Corollary 3 demonstrated one way to do this. By blocking entrants from the market and guaranteeing the incumbent monopoly over the innovation, the regulator is able to open up the channel of communication and gain the information she needs to approve the innovation. The cost is the loss of competition and the higher quality that it would bring.

The adaptation we consider in this section is a middle ground between too much and too little competition. A middle ground not in the number of firms that compete, but in when they compete. Rather than the entrants competing immediately for the innovation when it is improved, we empower the regulator to award the incumbent with a limited period of monopoly protection. Formally, suppose the regulator can award monopoly protection to the incumbent for a fixed number of periods $\tau \in \{0, 1, 2, ...\}$; so only the incumbent is active for periods $t = 1, ..., \tau$, and all n firms are active in subsequent periods $t > \tau$.

A time-limited monopoly presents the incumbent with a potentially delicate choice. He would like to recommend the policy where his quality is highest (p^{m} from Proposition 1). But if he is not the market winner under this policy, he will be disrupted by a higher-quality entrant upon expiry of the monopoly term. The threat of disruption may drive him to instead recommend the best policy in which he wins the market (p_1^{c} from Proposition 2), thereby guaranteeing for himself a perpetual stream of lower profit.

Proposition 4 shows that when the gap between p^{m} and p_1^{c} is high enough, and the length of monopoly protection long enough, the inefficiency of competition can be mitigated. The incumbent recommends p^{m} and the innovation be approved even when he knows it will lead inevitably to his own disruption.

To state this result, we amend the incumbent's lobbying strategy as follows.

Definition 1 *In a τ-regular messaging strategy, the incumbent sends message $m_1^* = p^m$ if*

$$q_1(p^m)(1-\delta^\tau) > q_1(p_1^c),$$

otherwise he recommends p_1^c.

The incumbent recommends p^{m} even if he would ultimately lose the market after the monopoly period ends, as long as the profit $q_1(p^{\mathrm{m}})$ for τ periods exceeds the recurring profit $q_1(p_1^{\mathrm{c}})$ that he would obtain under p_1^{c}. Of course, after the τ periods, the incumbent would like for policy to switch back to the status quo or some other policy in which he wins the market, but by then it is too late and the regulator will no longer follow his advice.

Define a τ-regular equilibrium as one in which the incumbent uses a τ-regular messaging strategy and each entrant uses a regular strategy. We can then state the result.

Proposition 4 *A τ-regular equilibrium exists when the incumbent has τ-periods of monopoly protection. The regulator follows the incumbent's recommendation in the first period and chooses $p_t^* = m_1^*$ for $t = 1, 2, ..., \tau+1$. In period $\tau+2$, the regulator follows the recommendation of the market winner, $w(p_{\tau+1}^*)$, and chooses $p_{\tau+2}^* = m_{w(p_{\tau+1}^*)}^*$ thereafter.*

An interesting feature of equilibrium behavior is that policy may change more than once. The regulator may no longer listen to the incumbent after his monopoly ends, but she will listen to the entrant who enters in period $\tau+1$ and wins the market. The regulator did not listen to this entrant initially as the regulator could not trust his advice. But by entering and establishing a profit level above the regulator's threshold, q^L, the winning entrant is now credible. The winning entrant is, in effect, the new incumbent. Thus, if the victorious entrant prefers a different policy, he lobbies the regulator and she follows his advice, changing policy a second time.[19] An example of the policy dynamic this gives rise to is given in Figure 3, where entrant 2 lobbies credibly for policy 2 after winning under policy 3.

The τ-monopoly mechanism delivers multiple benefits to the regulator. The incumbent may be induced to reveal information he otherwise wouldn't. When he does, society not only benefits from a higher quality innovation for those τ periods, but when the incumbent is disrupted, the quality increases even more with the entrant's better product and, as just

[19] Policy will not change a third time. Additional policy changes are possible if the regulator were to offer a τ-period monopoly to each new incumbent.

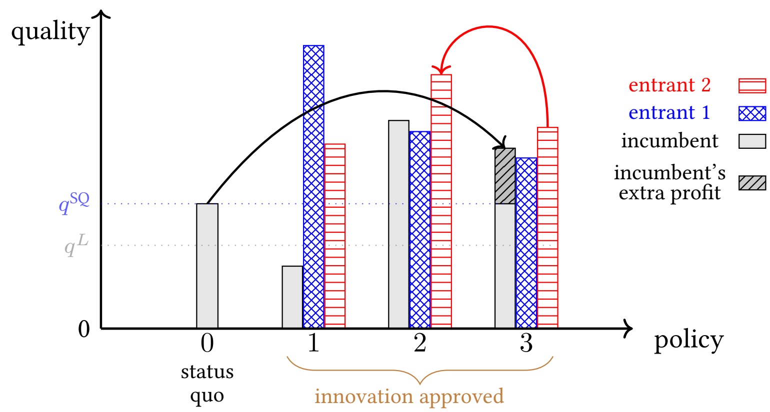

Figure 3: Example of a Policy Dynamic under Temporary Monopoly Protection

discussed, the entrant may use his newly earned credibility to recommend a switch to an even better policy setting.

Despite all of these potential benefits, the τ-monopoly mechanism does not ensure the efficient use of information. The incumbent may recommend a policy that will lead to his ultimate disruption, but he recommends the policy that maximizes his profit for τ periods and not necessarily the policy that maximizes social benefit. The same logic applies after the incumbent is disrupted and the winning entrant uses his information for his own benefit. This is evident in Figure 3 as the regulator would obtain a higher benefit from policy 2 but the credible firm—Entrant 2—recommends a switch from policy 3 to policy 1 and does not reveal the superiority to the regulator of policy 2. That the flow of information to the regulator remains in the control of a single firm, whether the initial incumbent or a disrupting entrant, places an upper bound on the efficiency of policymaking.

The logic of the τ-monopoly mechanism is at once new but familiar. It is, in effect, the logic of patents. It is the deliberate award of monopoly power for a fixed term so as to induce innovation. The novelty here is in how this logic is applied. Rather than rewarding the innovative firm, the award of market power is to a firm that would otherwise block innovation. Counterintuitively, the award of market power is to the loser in the innovation race and not the winner, and the market power removes the informational blockage that the loser creates rather than incentivizing innovation directly.

This discussion leads to the question of the optimal length of protection, τ. The ideal monopoly length provides the incumbent with enough protection to persuade him to reveal some of his information, but no more than that as protection comes at a cost of withholding competition. The perpetual monopoly power of Corollary 3 delivers, in a sense, too much benefit to the incumbent to get him to reveal his information.

The problem for the regulator is that she does not know the firms' qualities and, thus, does not know the length of protection that will induce the incumbent to reveal his information. This presents the regulator with a trade-off: The longer is τ the more likely it is that the incumbent recommends a better policy for society, but the longer it is before the benefits of competition are realized.

To see the elements in the trade-off, we decompose the regulator's expected discounted payoff as a function of τ where, for clarity, we set $k = 1$ such that the regulatory decision is to approve the innovation or not:

$$E\left[\sum_{t=1}^{\infty}\delta^{t-1}U(p_t^*)\right] = \underbrace{\int_0^{\infty}\max\{q, q^{\mathrm{SQ}}\}f(q)dq}_{\text{expected quality under perpetual monopoly}} - \underbrace{\int_{q^{\mathrm{SQ}}}^{\frac{q^{\mathrm{SQ}}}{1-\delta^{\tau}}}(q-q^{\mathrm{SQ}})(1-F^{n-1}(q))f(q)dq}_{\text{loss from incumbent hiding information}}$$
$$+ \underbrace{\int_{\frac{q^{\mathrm{SQ}}}{1-\delta^{\tau}}}^{\infty}\delta^{\tau}(E[Q_{n-1}|q < Q_{n-1}] - q)(1-F^{n-1}(q))f(q)dq}_{\text{gain when incumbent does not win but recommends the innovation anyway}}\ , \qquad (1)$$

where Q_{n-1} is the maximum of $n-1$ independent, F-distributed quality draws.

The first term is the baseline outcome for the regulator if the incumbent is granted monopoly power in perpetuity. The second term is the cost of a finite τ. It consists of the incumbent types that have higher quality under the innovation but do not recommend it be approved as the profit bump is not sufficient to compensate for the disruption that will come after τ periods. The third term is the benefit of a finite τ. This is the extra boost in quality from the entrant who disrupts the incumbent. The gain is the difference between the expected quality of the winning entrant over the quality of the incumbent.

For simplicity, in Proposition 4b we state the properties of optimal τ with a continuous domain. Specifically, let $\tau^* \in \arg\max_{\tau\in[0,\infty)} V(\tau)$ be an optimal monopoly length, where $V(\tau)$ is the regulator's expected discounted payoff (given τ) from Equation (1).

Proposition 4 (b) *For $k = 1$, the regulator's optimal monopoly length τ^* exists, is strictly positive, and is strictly increasing in the discount factor δ.*

Proposition 4b implies that the optimal monopoly length $\tau^* \in T^*$ is positive and finite—in other words, some protection is strictly better than no protection or eternal protection. When τ is required to be finite the optimal monopoly length may not be unique. We show in the proof of Proposition 4b that in this case the result holds for each element in the set of optimal monopoly lengths.

The optimal τ increases in the discount factor δ because of two forces working in concert. First, the incumbent loses more from being disrupted the more patient he is, and thus he requires a longer monopoly period to reveal his information. Second, a patient regulator is willing to tolerate a longer monopoly period in order to induce information revelation from more incumbent types. As a result, it is better for the regulator—and the incumbent—for the regulator to offer a longer period of protection.

The impact of more competition on the optimal monopoly length is less clear. As n increases, the incumbent is more likely to be disrupted by an entrant, and the quality of the winning entrant is expected to be higher. This incentivizes the regulator to reduce τ and bring the benefit of competition forward in time. But doing so will deter the marginal incumbent from revealing his information, and this reticence is now more costly to the regulator. These conflicting forces yield no clear comparative static of the effect of n on the optimal τ. Numerical simulations show that generally the benefit of a shorter τ dominates the cost, and the optimal τ decreases in n, approaching a limiting value that depends on the parameters of the environment. Examples of how τ^* varies with n and δ, given Frèchet-distributed quality shocks, are shown in the panels of Figure 4.

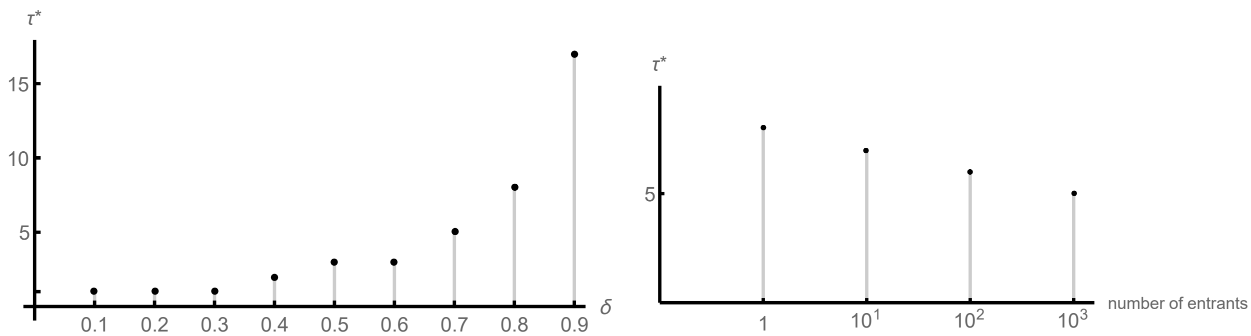

(a) discount factor δ
($n = 5$)

(b) number of entrants
($\delta = 0.8$)

Figure 4: Optimal Protection Length: τ^*
(Frèchet distribution with shape $\alpha = 2$, scale $s = 1$; $q^{\mathrm{SQ}} = 2$)

4.2 Adaptation II: More Competition

Rather than induce the incumbent to speak more frequently, in this section we explore how his information can be made less important. We do this by increasing market competition, specifically by opening up the status quo policy to competition from the entrants. The idea is that the baseline model may suffer not from competition *per se* but from too little competition. Formally, augment the dynamic model as follows. Let each entrant j now receive an independent F-distributed quality draw $q_j(0)$ under the status quo policy $p = 0$, while the incumbent's status-quo quality remains fixed at q^{SQ}, so the market under the status quo policy can potentially be won by an entrant.[20]

[20] Accordingly, amend $\boldsymbol{q}$ in step 1 of the stage game to be $\boldsymbol{q} = \{q_2(0), \ldots, q_n(0), q_1(1), \ldots, q_n(k)\}$.

Competition for the status quo opens up the possibility that firms other than the incumbent will be able to communicate credibly with the regulator. The opening this creates, however, is narrow. Only the entrant who successfully disrupts the incumbent becomes credible. Thus, competition for the status quo cannot increase the number of credible voices, it can only change the identity of which firm is credible.

In fact, competition for the status quo reduces communication between the market and the regulator. As the status quo winner must now overcome $n-1$ competitors, his quality is higher in expectation. Thus, the likelihood that he wins under the innovation at an even higher quality—and, thus, recommends that the innovation be approved—is lower.

This creates a trade-off for the regulator: competition for the status quo increases the winning status quo quality but worsens communication between the regulator and the market. In the baseline model it is straightforward to show that the regulator strictly prefers more competition, effectively giving up on communication and the possibility that the innovation will be approved (see Proposition 5 below).

The contrast between this preference and that in Corollary 1 is illuminating. The difference between the environments is where competition manifests. In Corollary 1 the regulator prefers communication when competition is only for the innovation. Here, competition is for the status quo as well and the regulator prefers competition. In both environments, competition for the innovation harms communication without bringing any compensating benefit (as it is never realized). The regulator prefers competition here as it delivers the best quality draw from n firms, whereas banning competition and communicating with the incumbent delivers only the best of the two quality draws for the incumbent.

In an ideal environment, the regulator could extract the benefit of both competition and communication. Doing so is not possible in the baseline model because the quality draws are assumed to be identical and independent across all firms and policies. Such assumptions are unlikely to hold in practice and are made in the baseline model only for simplicity. To capture richer possibilities, we extend the model further by allowing for the qualities of a firm to be correlated across policy settings.

Correlated product qualities. We formalize these changes in a deliberately simple way and limit attention to a choice to approve the innovation or not ($k = 1$). Given the quality draws described above under the status quo, let the quality of the innovation for any firm $j \geq 1$ be given by:

$$q_j(1) = \begin{cases} s \cdot q_j(0) & \text{with probability } v \\ s \cdot q'_j & \text{with probability } 1-v, \end{cases} \tag{2}$$

where $s \geq 1$, each q'_j is drawn from the distribution $F(\cdot)$, and the randomization between $q_j(0)$ and q'_j is independent across firms. The parameter $v \in [0, 1)$ captures the correlation in firm quality across policies. s is a scale parameter that allows product quality to differ under the innovation and the status quo. $s = 1$ represents the baseline model, and $s > 1$ captures the more common situation in which the innovation improves upon the status quo in expectation.

The regular equilibrium exists in this environment. All firms, including the incumbent, adopt the regular lobbying strategy. However, the regulator no longer listens to the incumbent as she did in the baseline model. Rather, she waits for competition to play out in the first period and then listens thereafter to whichever firm wins the market under the status quo.

Proposition 5 *The regular equilibrium exists when there is competition under the status quo and product qualities are given by (2). For $n > 1$ the regulator ignores all recommendations and implements the status-quo policy in the first period, $p_1^* = 0$. In all subsequent periods $t \geq 2$, the regulator follows the recommendation of the market winner for the status quo policy in the first period, $p_t^* = m_{w(0)}^*$. For $n = 1$ the regulator follows the incumbent's recommendation in the first period, $p_t^* = m_1^*$, and thereafter.*

Policymaking remains inefficient in this equilibrium as the regulator can only listen to the market winner. Although it is likely to not be the incumbent, the new market winner censors his policy recommendation in exactly the same way, recommending the innovation only if quality is higher *and* he wins the market. If another entrant, or even the initial incumbent, were able to win with a higher quality than the new market winner, the regulator misses out on that benefit.

This result reinforces that, even in the baseline model, the incumbent is favored not because he already exists or has a relationship with the regulator, but because he is the current market winner. Proposition 5 shows that once that is lost, the incumbent loses his credibility with the regulator who instead turns to whichever firm disrupts the incumbent and wins the market.

It is important to Proposition 5 that competition under the status quo is not a choice of the regulator. This is why the regulator does not, and in fact cannot, listen to the incumbent's recommendation in the first period. If the incumbent anticipates being disrupted, he will recommend the innovation be approved even if the quality is worse than under the status quo. As the incumbent's information is then no different from that of an entrant, the regulator has no choice but to let competition under the status quo policy play out and exploit the alignment that emerges endogenously between her and whichever firm wins the market.

This raises the question of whether the regulator would block competition for the status quo if she could. As noted above, for the baseline model ($v = 0, s = 1$), the regulator prefers competition to communication. Correlation and a higher expected quality of the innovation ($s > 1$) make this choice less clear.

When the scale parameter s is large, the value of approving the innovation is, in expectation, much higher. The regulator wishes, therefore, to maximize communication, even if it comes at the expense of having to limit competition. If, however, the correlation within firms v is high, competition comes with a smaller cost to communication. With high v, the market winner under the status quo is likely to also be the winner under the innovation and the tension between competition and communication is ameliorated.

Proposition 5b describes the equilibrium trade-off between competition and communication for the regulator. As the exact balance depends on the distributions of quality and other parameters, which we allow to be very general, we state the result only for when the regulator prefers either maximum or minimum competition.

Proposition 5 (b) *Suppose the regulator can choose the level of competition* $n \in \{1, ..., \overline{n}\}$ *when there is competition under the status quo and product qualities are given by (2). The regulator's optimal degree of competition is:*

(i) $n^* = 1$ *if* s *is sufficiently large and* v *is sufficiently small,*

(ii) $n^* = \overline{n}$ *if* v *is sufficiently large or if* s *is sufficiently small.*

In Proposition 5b the regulator can limit the number of firms but cannot control competition within the market itself. An interesting variant on this result is to suppose the regulator's decision to allow competition can be a function of her policy choice. In contrast to what might be the conventional wisdom, the logic of this section suggests that, for the baseline model at least, the regulator's best choice is to encourage competition under the status quo policy but *not* for the innovation. Only in this way can the regulator gain the benefit of competition while keeping open the channel of communication with the market.

This logic also suggests an interesting twist on business strategy. In practice, innovating firms typically enter a market only for the innovation, even if entry under the status quo is possible. The ideas in this section suggest that an innovating firm should enter the market even *without* putting the innovation to work as, by doing so, it may win the market under the status quo and establish credibility with the regulator. This allows him to persuade the regulator to open the market to the innovation when she otherwise may not.

4.3 Adaptation III: Antitrust

In the first two adaptations, a change in the nature of market competition changed the way information is shared with the regulator. In this section we consider changing the competitors themselves. We suppose that the incumbent can make a takeover offer for one or more entrants. When successful, a takeover changes the quality of the incumbent's product and, thus, the information he holds.

We return to the baseline (static) model of Section 2, amending the timing of the game as follows. In between Steps 1 and 2, the incumbent has one opportunity to make a *private* take-over offer for an entrant at price $a_j \geq 0$, and the entrant decides whether to accept. If indifferent, the entrant accepts. The regulator observes a successful takeover, but cannot distinguish between a rejected takeover offer and no offer. After buying firm j, the incumbent's quality under policy p becomes $\hat{q}_1(p) = max\{q_1(p), q_j(p)\}$, the higher of its own and the acquired entrant's quality.

That takeovers are possible does not cause the communication problem of the baseline model to go away, or indeed to even change. However, a causal effect does emerge that runs in the opposite direction. Precisely because there is a communication problem, a takeover is profitable. The incumbent understands that because the regulator listens only to him, he can capture the benefit of a higher product quality, and, at the same time, the entrant knows that he won't be able to.

This is a new logic for takeovers that does not rely on synergies or asymmetric information between the two firms involved. Rather, it is the asymmetry in the ability to communicate with the regulator that makes a takeover profitable. The takeover does not create value, instead it unlocks value that is latent in the market.

Equilibrium behavior is given in Proposition 6. To maximize his profit, the incumbent targets the entrant with the highest quality. Denote the maximum quality by $q^{\text{win}} = \max_p\{q^{\text{win}}(p)\}$, and the global winner w as the firm j that produces this quality for some policy, $q_j(p) = q^{\text{win}}$.

Proposition 6 *There exists a regular equilibrium in which the incumbent acquires the global winner for price zero whenever $w \neq 1$, and the regulator follows the incumbent's recommendation, $p^* = m_1^*$.*

Takeovers benefit the incumbent and they benefit society as well. By aligning product quality with the ability to communicate with the regulator, takeovers bring innovations to market more often. Moreover, it does so efficiently, with the product that comes to market being the highest quality product. Allowing takeovers outperforms the adaptations of the previous sections as it leaves the market to sort out which product comes to market. As the market is where information resides, this is done with maximum efficiency.

A prominent question in the practice of antitrust is whether large incumbents—nowadays predominantly tech companies—should be allowed to acquire start-ups that may eventually compete directly with them. In recent times such acquisitions have gone unchallenged by antitrust authorities, leading to a rigorous debate inside and outside of academia on the costs of such a policy (see, e.g., Bryan and Hovenkamp, 2020). To this debate we add a new dimension. We show that when an industry is regulated in the shadow of antitrust policy, lax antitrust policy can alleviate the informational blockage that emerges between the market and the regulator. Takeovers in our model give life to innovations that otherwise would not make it to the market.

We emphasize that the informational mechanism we identify is only one element in any assessment of the competitive effects of a takeover. In particular, our model sets aside the need for firms to innovate. Were we to add this to the baseline model, the incentive to innovate would be non-existent for the entrants and it would not be helped by the prospect of a takeover by the incumbent. In Proposition 6 the incumbent has *too* much power, driving the acquisition price to zero, and this disincentivizes an entrant from innovating. This suggests that for takeovers to be efficiency enhancing, the bargaining power of the entrants would need to be increased so as to increase the takeover price paid by the incumbent.

It is important to Proposition 6 that the takeover offer is private in the sense that it is unobserved by the regulator. If the regulator were to observe the offer, that information alone may be enough for her to learn that the innovation is safe, and the entrant could reject the offer in anticipation that the policy would change. However, the incumbent could circumvent this inference by making an offer of zero to *all* entrants, regardless of the state, revealing no information and maintaining its leverage over the entrants. The possibilities this opens up suggest interesting strategies in antitrust when a regulator with the power to open a market is watching.

5 Discussion

5.1 Regulatory Capture

In their volume on regulatory capture, Carpenter and Moss (2013) define *strong* and *weak* forms of capture. Weak capture is when the regulator's behavior distorts outcomes from society's interests, and strong capture is when the distortion is so large that society would be better off without the regulator.

Behavior in our model represents neither form of capture, but only to the extent that the regulator's risk aversion reflects society's preferences. If the regulator's risk aversion instead represents a principal-agent problem between her and society, then capture is strong when the number of innovating entrants is large and society prefers the innovation be approved whereas the regulator blocks it. Regardless of the form of capture, or whether there is capture at all, there does exist regulatory *protection*. The incumbent firm benefits from regulation even when the regulator is acting in perfect alignment with society's interests (see Carpenter, 2004).

5.2 Regulating Innovation in Practice

Our model resonates with how innovations are regulated in practice, from society-altering innovations in high tech industries to smaller breakthroughs in more mundane industries. In this section, we discuss a few examples, both historical and current, to illustrate the essential aspects of our model.

Fitting the model to practice. Our model sits on a foundation of several key features: innovation offers the potential of enormous benefit but at the risk of social harm, information about the impact of the innovation is locked up in the firms developing the technologies, the market for the innovation is competitive and, as a result, firms have an incentive to misrepresent their information to regulators. This collection of features—and the challenge they present to regulators—is perhaps nowhere more clearly evident than in the rapid developments of Artificial Intelligence (AI). The benefits of AI are clear, from labor-saving productivity gains, to improvements in the diagnosis and treatment of diseases, to easier global communications, and so on. A much more capable Artificial General Intelligence (AGI), however, has been speculated to portend substantial harm for society and a much darker future for humanity.

The development of AI fits our model also in that it is competitive and that information asymmetry is a key part of the regulatory process. Many firms are developing AI capabilities and products, from the largest tech companies to new start-ups. These firms are lining up to compete in the market and also in the policy domain, with many players eagerly participating in regulatory hearings and policy briefings hoping to shape regulations in their favor. The space for firms to influence regulations is large as most policymakers lack even a basic understanding of the technology and how it will impact society. Whilst it is difficult at this stage to prove that the firms have misrepresented their information to policymakers, the willingness to misrepresent was exposed when the socially-minded board of directors at OpenAI concluded that the CEO was "not consistently candid" in his communications about the power—and potential dangers—of the latest technological advances (OpenAI, 2023).

Regulatory approval and the harm of a failed innovation. A key premise of our model is that innovations can bring harm as well as benefit society. Whilst the potential harm of AI remains on the horizon, examples abound of harmful innovations that *did* receive regulatory approval, highlighting the danger of getting the policy decision on a new technology wrong. A striking recent example of a damaging innovation comes from pharmaceuticals and the experience of Oxycodone. Oxycodone has brought enormous benefit to the world—it appears on the WHO's List of Essential Medicines—but in the form of Oxycontin it has played a central role in the opioid epidemic that has caused enormous harm in the U.S. and elsewhere. As Dr. David Kessler, FDA commissioner at the time Oxycontin was approved, acknowledged, "No doubt it was a mistake. It was certainly one of the worst medical mistakes, a major mistake. ... No doubt [the opioid epidemic] began there." (See Kolodny, 2020; Mitchell, 2018).

The examples of AI and Oxycodone also highlight the "Pandora's box" nature of many harmful innovations. Once Oxycodone was approved and the Opioid epidemic took hold of the US, illicit sources of substitute opioids such as Fentanyl quickly flooded the US market, which rendered irrelevant subsequent regulatory measures to control legal uses of Oxycodone. Similarly, once developed, advances in AI may be co-opted by hostile powers for military purposes or by profit-maximizing firms. More existentially, an Artificial General Intelligence that is able to recursively improve itself may rapidly develop superhuman capabilities and outwit any attempts to regulate or control it.

Regulatory policy that misses the benefit of innovation. Another central feature of our model is that beneficial innovations are not easily identifiable to regulators. The subtlety in matching this feature to practice is that it requires examples of innovations that we don't see, innovations that would have been beneficial had they been permitted. It is easier to identify such examples by looking at innovations that *did* make it to market but in a constrained way such that much of the benefit was missed. One such example is nuclear power in the U.S. and in many other countries. Although the building of nuclear power plants has always been legal, fear of a nuclear accident led to overly-restrictive regulations that made building and operating nuclear plants uneconomical, to the point that the industry stagnated. After an initial burst of construction, no new nuclear plants began construction in the U.S. in the years 1977-2013. This regulatory caution has come at great cost, with an increased reliance on fossil fuels, namely coal, causing enormous harm to health in the short term and to the environment in the long term (Markandya and Wilkinson, 2007).

Firm Strategy Beyond the Model. In characterizing the limits to innovation in a regulated market, the model also provides insight into when and why entrants attempt to circumvent the boundaries of the model. Uber provides a canonical example.

Consistent with our model, Uber, as an entrant was initially unable to persuade regulators to approve its superior ride-sharing innovation because the incumbent taxi industry lobbied against it, emphasizing the possible harm from unlicensed cars and drivers. Within the model, the industry would have then remained at the suboptimal but safe status quo. Famously, however, Uber entered the market in brazen defiance of the regulation. By putting its innovation to work, Uber was able to show regulators that its innovation did, in fact, work. Following the logic of the model, and with the informational blockage we identify resolved, regulators in many markets then approved Uber's innovation.[21]

[21] Our model provides a purely informational explanation for Uber's initial failures and ultimate success in

5.3 Variations and Extensions

We have, so far, extended the baseline model in multiple directions and considered some natural variants. In this section we develop informally several more, focusing on the robustness of the communication problem at the heart of the model and the implications of that problem for broader political and market behavior.

Common Uncertainty. It is realistic that the firms have better information about an innovation than does the regulator, but a stretch to believe, despite their often-extreme confidence, that the firms know outcomes precisely. Adding some common uncertainty about the innovation does not affect the core mechanism driving our results as long as the firms remain better informed than the regulator. Common uncertainty does, however, lead to several additional effects.

When the firms are uncertain about the innovation a wedge forms between the preferences of the incumbent and the regulator due to the regulator's risk aversion. The impact of this wedge depends on the degree of the common uncertainty. The incumbent lobbies for the innovation if he wins the market and the expected outcome is above the status quo outcome, whereas the regulator prefers to reject the innovation if there is significant, or possibly any, chance the outcome falls below her threshold, q^L.

This leads to two interesting effects. First, the regulator will favor an incumbent that has high profit under the status quo policy. This increases the expected quality of the innovation, should the incumbent recommend it be approved, and lowers the probability that the true quality falls below the regulator's threshold q^L.[22] This pattern resonates with the findings on FDA regulation in Carpenter (2004).

Second, there may be a benefit to allowing some competition. Recall that, conditional on recommending the innovation, the incumbent's expected quality is higher when he faces competition. Thus, it may be that a monopolist incumbent is unable to persuade the regulator whereas an incumbent facing competition can. Of course, the more competition there is, the less likely the incumbent wins the market and the lower the probability the innovation makes it to market. In this case, there may be a sweet spot degree of competition for the regulator that allows communication between her and the incumbent whilst not suppressing too much the probability that it happens.

Common uncertainty also leads to more subtle lobbying behavior. If uncertainty is idiosyncratic, a firm may not know whether he will win the market, and this makes him even less likely that an innovation comes to market as the incumbent will be hesitant to recom-

many markets. Other factors were also important, in particular Uber's ability to mobilize its user and driver bases.

[22] Assuming residual uncertainty is independent of realized quality or, at least, not increasing too sharply in it.

mend it be approved. It may also be the case that firms lobby for multiple policies as they trade off the probability of winning the market against the expected profit. This does not necessarily lead to a greater information flow as firms may lobby only for a policy that has lower expected quality if their advantage over the competition is clearer. A similar logic holds if the firms observe their own quality precisely but the quality of the other firms imprecisely. The regulator can then trust the incumbent fully, but the incumbent will be less likely to recommend the innovation and the entrants may lobby strategically for multiple policies.

Non Winner Take All Markets. This leads to contrasting information effects. When more firms profit from the status quo, communication with the regulator increases as more firms are credible. Working against this increase in communication, however, is that the information the regulator receives is fragmented. Each firm makes a recommendation based on their own relative quality and the regulator will be less confident that an innovation in aggregate is beneficial to society. Whether the regulator responds to lobbying to approve the innovation will depend on the number of incumbents recommending the innovation and the nature of competition in the market.

Innovation Over Time. In practice, entrants arrive at a market over time rather than all at once. This has contrasting effects on innovation that can both delay and speed-up innovation when the regulator's risk aversion is less than extreme. As shown in Proposition 2, the regulator will approve an innovation only if there are enough entrants. Thus, if the entrants take time to arrive, her approval will also take time as she waits for their number to accumulate. More surprisingly, the slow arrival of entrants can speed-up innovation by inducing the incumbent to recommend approval even when he knows he will ultimately be disrupted. Following the same logic of temporary monopoly in Proposition 4, the incumbent will trade-off a short-term profit increase versus long-term loss of the market, and will recommend the innovation if he expects disruption to not occur too soon.

Bureaucratic expertise and whistleblower laws. As with Baron and Myerson (1982) and any regulatory model of incomplete information, the regulator benefits if she can develop her own expertise independent of the firms. In most contexts this is difficult for well-understood reasons, and these reasons apply with greater force to new technologies, such as AI, where the market opportunities for expertise are particularly lucrative.

An interesting possibility is whistleblower laws. We have treated each firm as a single entity, whereas in practice expertise is held by a multitude of researchers that make up the firm. Whistleblower laws attempt to drive a wedge between employees and the firm. To create the right incentives, these laws work best in retrospect so that the whistleblower can prove the harm and be compensated in proportion to it. With forward-looking technologies where the harm is only a potential outcome, inducing the right incentives is far from straightforward but worth exploring further.

Informational Collusion. Allowing the incumbent to make a take-it-or-leave-it acquisition offer to the entrants is just one example of what Gans and Stern (2000) refer to as a "collusive agreement" between firms (see also Gans (2017)). Our model identifies a novel benefit of collusive agreements between firms. One firm holds a productive innovation whereas the other holds credibility with the regulator. We showed how a takeover can unlock this value. An interesting set of questions emerge if antitrust is strictly enforced such that takeovers are not allowed. Can the incumbent and a higher-quality entrant find some other way to collude and capture the benefit of the innovation? Can the entrants collude in their lobbying and market strategies so as to create credibility with the regulator? That many innovations are missed because of informational asymmetry opens up new angles on business strategy, both in the market and in the firm's relationship with a regulator.

6 Conclusion

Breakthrough technologies like generative-AI, nuclear power, and gene editing, present regulators with difficult choices. A regulator can approve their use and unleash their capabilities on society, thereby capturing the potential benefits but also opening society up to harm. Or a regulator can hold the new technologies back, keep society safe but deprive us all of the potential benefits. The model we develop in this paper shows the power this dilemma delivers to the incumbent firm. Even when there are many firms competing in the market, all with better information about the innovation than the regulator, the endogenous alignment between the incumbent and the regulator puts the incumbent firm in a position of power.

This conclusion may be a source of optimism if we expect the incumbent firm to prevail under the new technology. However, there is little reason to expect this to be the case. Indeed, for breakthrough technologies, there is good reason to expect that the incumbent is singularly unlikely to be successful with the new technology, particularly when the new technology requires a different way of organizing production, what Gans (2016) refers to as supply-side disruption.

Since at least the time of Schumpeter (1942), the importance of innovation to economic outcomes has occupied a central place in economics. More recently, the role of government policy in fostering innovation has come into focus (Bryan and Williams, 2021). What has received less attention is the strategic behavior of policymakers and how their incentives interact with those of the market. We have shown here how including the policymaker's problem in the analysis can upend standard intuitions about market competition. Exploring this connection more fully is an important pathway to a deeper understanding of innovation in the economy.

References

Acemoglu, D. and T. Lensman (2024). Regulating transformative technologies. *American Economic Review: Insights 6*(3), 359–376.

Arrow, K. J. (1962). Economic welfare and the allocation of resources for invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pp. 609–626. Princeton, NJ: Princeton University Press.

Baron, D. P. and R. B. Myerson (1982). Regulating a monopolist with unknown costs. *Econometrica*, 911–930.

Bryan, K. A. and E. Hovenkamp (2020). Startup acquisitions, error costs, and antitrust policy. *The University of Chicago Law Review 87*(2), 331–356.

Bryan, K. A. and H. L. Williams (2021). Innovation: market failures and public policies. In *Handbook of Industrial Organization*, Volume 5, pp. 281–388. Elsevier.

Callander, S., D. Foarta, and T. Sugaya (2022). Market competition and political influence: An integrated approach. *Econometrica 90*(6), 2723–2753.

Callander, S., D. Foarta, and T. Sugaya (2023). The dynamics of a policy outcome: Market response and bureaucratic enforcement of a policy change. *American Journal of Political Science*, forthcoming.

Carpenter, D. and D. A. Moss (2013). Introduction. In D. Carpenter and D. A. Moss (Eds.), *Preventing regulatory capture: Special interest influence and how to limit it*, pp. 1–22. Cambridge University Press.

Carpenter, D. and M. M. Ting (2007). Regulatory errors with endogenous agendas. *American Journal of Political Science 51*(4), 835–852.

Carpenter, D. P. (2002). *The forging of bureaucratic autonomy: Reputations, networks, and policy innovation in executive agencies, 1862-1928*. Princeton University Press.

Carpenter, D. P. (2004). Protection without capture: Product approval by a politically responsive, learning regulator. *American Political Science Review 98*(4), 613–631.

Cowgill, B., A. Prat, and T. Valletti (2023). Political power and market power.

Gans, J. (2016). *The disruption dilemma*. MIT press.

Gans, J. S. (2017). Negotiating for the market. *Advances in Strategic Management 37*, 3–36.

Gans, J. S. (2024). How learning about harms impacts the optimal rate of artificial intelligence adoption. *NBER Working Paper w32105*.

Gans, J. S. and S. Stern (2000, Winter). Incumbency and R&D incentives: Licensing the gale of creative destruction. *Journal of Economics and Management Strategy 9*(4), 485–511.

Gilligan, T. and K. Krehbiel (1987). Collective decision making and standing committees: An informational rationale for restrictive amendment procedures. *Journal of Law, Economics, and Organization 3*(2), 287–335.

Gilligan, T. W. and K. Krehbiel (1989). Asymmetric information and legislative rules with a heterogeneous committee. *American journal of political science*, 459–490.

Kolodny, A. (2020). How FDA failures contributed to the opioid crisis. *AMA Journal of Ethics 22*(8), E743–750.

Lu, S. E. (2017). Coordination-free equilibria in cheap talk games. *Journal of Economic Theory 168*, 177–208.

Markandya, A. and P. Wilkinson (2007). Electricity generation and health. *The Lancet 370*(9591), P979–990.

McCarty, N. (2013). Complexity, capacity, and capture. In D. Carpenter and D. A. Moss (Eds.), *Preventing regulatory capture: Special interest influence and how to limit it*, pp. 99–123. Cambridge University Press.

McCarty, N. (2017). The regulation and self-regulation of a complex industry. *The Journal of Politics 79*(4), 1220–1236.

Mitchell, J. (2018). How the FDA helped pave the way for an opioid epidemic. *Mississippi Clarion Ledger*. January 26.

OpenAI (2023). OpenAI announces leadership transition. Statement of the board of directors, November 17.

Schumpeter, J. A. (1942). *Capitalism, Socialism, and Democracy.* New York: Harper and Brothers.

Stigler, G. J. (1971, Spring). The theory of economic regulation. *Bell Journal of Economics and Management Science 2*(1), 3–21.

Wu, T. (2011). *The master switch: The rise and fall of information empires.* Vintage.

A Proofs

Specifying Off-Path Beliefs The propositions in the main text specify on-path equilibrium behavior. In each of the proofs, we will also specify off-path beliefs that support these equilibria – specifically, the regulator's beliefs and corresponding decisions following events that are inconsistent with on-path behavior. (We need only consider events that may arise due to unilateral deviations.) Such events occur with zero-probability on-path, and arise either because of deviations from equilibrium play, or because of knife-edge outcomes where two firms have the same quality realizations for two policies, i.e., where $q_j(p) = q_{j'}(p')$ for some $j \neq j'$. Off-path beliefs following such events are constructed as follows. The regulator attributes zero-probability outcomes to "mistakes" by individual players. The regulator believes that (i) mistakes are uninformative, i.e., uncorrelated with the realization of quality shocks; that (ii) the incumbent and entrant are roughly equally likely to make mistakes; that (iii) each firm's mistakes are uniformly distributed over the set of possible recommendations (including the empty recommendation); and that (iv) mistakes are infinitely more likely than knife-edge quality realizations. So, the regulator overwhelmingly favours explanations where mistakes are made over explanations where knife-edge outcomes occur.

Proof of Proposition 1. Given that the incumbent's and regulator's equilibrium strategies, the highest possible quality $\max_p q_1(p)$ is implemented. Both the incumbent's and regulator's utilities are increasing in the quality of the implemented policy, so neither the regulator nor incumbent has any profitable deviations.

Off-Path Beliefs If the incumbent deviates by making an empty recommendation, the regulator infers that the incumbent made a mistake, and responds by implementing the status quo. This is weakly dominated by the regular strategy. ■

Proof of Lemma 1. The regulator's utility from prohibiting the innovation is q^{SQ}. If she approves the innovation, the winning quality is distributed as

$$q^{\max}(1) = \max\{q_1(1), q_2(1), \ldots, q_n(1)\},$$

and the regulator's expected utility is

$$E[q^{\max}(1)] - \Delta \Pr[q^{\max}(1) < q^L].$$

Given that each firm's $p = 1$ quality is i.i.d. with right-unbounded support, (i) $q^{\max}(1)$ is thus increasing (in stochastic dominance) in the number of firms n and $E[q^{\max}(1)]$ increases in n approaches ∞ as $n \to \infty$, and (ii) $\Pr[q^{\max}(1) < q^L]$ decreases in n and vanishes as $n \to \infty$; thus the regulator's expected utility from implementing $p = 1$ is strictly increasing in n and approaches ∞. In comparison, the regulator's utility from implementing the status quo is invariably q^{SQ}. Thus the regulator approves the innovation by default if and only if n exceeds the threshold

$$\hat{n}(\Delta) = \max\left\{n \in \{1, 2, \dots\} : E[q^{\max}(1)] - \Delta \Pr[q^{\max}(1) < q^L] \le q^{\text{SQ}}\right\}$$

Given that $E[q^{\max}(1)] - \Delta \Pr[q^{\max}(1) < q^L]$ is decreasing in Δ and increasing in n, the threshold $\hat{n}(\Delta)$ must be increasing in Δ. Further, for any n, $E[q^{\max}(1)] - \Delta \Pr[q^{\max}(1) < q^L] \le q^{\text{SQ}}$ for sufficiently large Δ; so $\lim_{\Delta\to\infty} \hat{n}(\Delta) = \infty$. ■

Proof of Proposition 2. Suppose that some entrant recommends policy $p = 1$ (in which case, given all firms play regular strategies, the incumbent must have recommended the status quo). Then the regulator's belief about the quality of the winning firm under policy $p = 1$, which is independent of whether an entrant or the incumbent wins under policy $p = 1$, remains unchanged from the default distribution; by Lemma 1, the regulator chooses $p = 0$ if $n \le \hat{n}(\Delta)$ and $p = 1$ otherwise.

It remains to note the two remaining possibilities: either the incumbent recommends the status quo $p = 0$ and nobody recommends $p = 1$, or the incumbent recommends $p = 1$ and nobody recommends the status quo $p = 0$. In both cases, the regulator infers that the incumbent wins under both policies and is recommending the first-best policy; the regulator thus optimally implements the incumbent's recommendation.

Off-Path Beliefs There are three relevant types of off-path outcomes. (i) The incumbent does not make a recommendation, in which case the regulator infers that the incumbent made a mistake. In the case $n \le \hat{n}(\Delta)$, this cannot be a profitable deviation for the incumbent regardless of the regulator's response given that the incumbent's preferred policy would be implemented in equilibrium. If $n \ge \hat{n}(\Delta) + 1$ and an entrant recommends policy $p = 1$, then (whether in equilibrium or if the regulator infers that incumbent made a mistake) the regulator will implement policy $p = 1$; so the incumbent cannot gain from deviating. If $n \ge \hat{n}(\Delta) + 1$ and no entrant recommends policy $p = 1$, then the regulator would follow the incumbent's recommendation in equilibrium; so again the incumbent cannot gain by deviating. (ii) The incumbent and an entrant both recommend policy $p = 1$. The regulator infers that one of the two firms made a mistake. In the case $n \le \hat{n}(\Delta)$, this cannot be a profitable deviation for the incumbent regardless of the regulator's response – given that the incumbent's preferred policy would be implemented in equilibrium anyway. It cannot be a profitable deviation for the entrant because the entrant cannot have won under either policy. If $n \ge \hat{n}(\Delta) + 1$, then the regulator will implement policy $p = 1$ following such an off-path outcome; neither firm would profit from such a deviation. (iii) Two entrants both recommend policy $p = 1$ (and the incumbent recommends the status quo $p = 0$). The regulator infers that one of the two entrants made a mistake. If $n \le \hat{n}(\Delta)$, then the regulator follows the incumbent's status-quo recommendation; neither entrant can profit from such a deviation. If $n \ge \hat{n}(\Delta) + 1$, then the regulator implements policy $p = 1$ anyway, so such a deviation by either entrant has no effect on the outcome. ■

Proof of Corollary 1. Recall that (i) for $n \leq \hat{n}(\Delta)$, the regulator's default decision is to prohibit the innovation; (ii) in equilibrium, the regulator follows the incumbent's recommendation. Thus, lobbying influences the regulator's decision if and only if:

$$q_1(1) = q^{\max}(1) \quad \text{and} \quad q_1(1) > q^{\text{SQ}}.$$

In other words, the probability that lobbying influences the regulator's decision equals

$$\Pr[q_1(1) > q^{\text{SQ}}] \Pr[q_1(1) = q^{\max}(1) | q_1(1) > q^{\text{SQ}}].$$

The first term is constant in n, whereas the second term is decreasing in n; thus their product is strictly decreasing.

For $n \geq \hat{n}(\Delta) + 1$, the regulator's uninformed decision is to choose $p = 1$; while in equilibrium, the regulator chooses $p = 0$ if and only if no firm recommends the innovation. Thus, lobbying influences the regulator's decision if and only if:

$$q_1(1) = q^{\max}(1) \quad \text{and} \quad q_1(1) < q^{\text{SQ}}.$$

Analogously to the case $n \leq \hat{n}(\Delta)$, the probability that lobbying influences the regulator's decision equals

$$\Pr[q_1(1) < q^{\text{SQ}}] \Pr[q_1(1) = q^{\max}(1) | q_1(1) < q^{\text{SQ}}].$$

The first term is constant in n, whereas the second term is decreasing in n; thus their product is strictly decreasing. ■

Proof of Corollary 3. With $n \leq \hat{n}(\Delta)$, the regulator follows the incumbent's recommendation, and the incumbent will recommend policy $p = 1$ if and only if he wins there and is more productive there than under the status quo. So we have

$$U(p^*) = \max\{q^{\text{SQ}}, u_1(1)\}$$

$$\text{where } u_1(1) = \begin{cases} q_1(1) & \text{if } q_1(1) = \max_{j \in \{1,\ldots,n\}} \{q_j(1)\} \\ 0 & \text{otherwise.} \end{cases}$$

Observe that $u_1(1)$ and thus $U(p^*)$ decreases in distribution as n increases.

In the special case $\Delta = \infty$, the regulator follows the incumbent's recommendation for all n. The probability that the innovation is approved is weakly smaller than the probability $1/n$ that the incumbent wins under the pro-innovation policy—thus $\Pr[p^* = 1] \to 0$. Next, notice that the random variable $u_1(1)$ has probability $1/n$ of taking a nonzero value (which occurs if and only if the incumbent wins under the pro-innovative policy), in which case it coincides with $q_1(1)$. Consequently, we can bound $E[u_1(1)]$ as

$$E[u_1(1)] \leq \int_{1-1/n}^{1} G^{-1}(q)dq$$

where $G(x)$ is the cumulative distribution function of $q_1(1)$. The RHS of this expression converges to zero as $n \to \infty$; thus

$$E[U(p^*) \leq q^{\text{SQ}} + E[u_1(1)] \to q^{\text{SQ}} \text{ as } n \to \infty.$$

■

Proof of Proposition 3. *Regulator's Strategy*: Given that the firms play regular strategies, the regulator has to pick from one of four options: (i) follow an entrant's recommendation (if at least one entrant made a recommendation), (ii) adopt a policy $p \geq 1$ that was not recommended by any firm, (iii) follow the incumbent's recommendation, (iv) adopt the status-quo policy (even if it was not recommended by the incumbent).

Suppose the regulator follows entrant j's recommendation. With positive probability, entrant j recommends one of the policies $p \geq 1$ that he wins with quality strictly less than q^L; thus the regulator's expected utility would be $-\infty$. The regulator would be strictly better off always implementing the status-quo and receiving utility q^{SQ}.

Suppose the regulator adopts an unrecommended policy $p \geq 1$. Given that firms play regular strategies, if a policy is unrecommended, then with probability one the incumbent wins under that policy (but prefers the status quo).With positive probability, the incumbent's winning quality under this policy is strictly less than q^L. The regulator would be strictly better off always implementing the status-quo.

Suppose the regulator listens to the incumbent. The incumbent's quality under the recommended policy must be weakly (and sometimes strictly) higher than his status quo quality; thus the regulator is strictly better off listening to the incumbent than adopting the status-quo.

Collecting our observations so far, that the regulator's best response is to listen to the incumbent.

Firm Strategies: By construction, the incumbent's preferred policy is implemented, so the incumbent has no profitable deviations. Given the regulator's and incumbent's strategies, each entrant receives zero regardless of their recommendation; thus they have no profitable deviations either.

Off-Path Beliefs: There are five possible types of off-path outcomes that are consistent with unilateral deviations. (i) The incumbent does not make a recommendation. (ii) The incumbent recommends the same policy as one of the entrants. (iii) Two of the entrants recommend the same policy. (iv) more entrants make distinct recommendations than there are policies. (v) an entrant recommends the status quo.

In the case of (i), the regulator infers that the incumbent made a mistake, and chooses the status-quo decision. In (ii), the regulator infers that either the incumbent or entrant made a mistake, and chooses the status-quo decision. In (iii), (iv), and (v), the regulator infers that an entrant made a mistake, and follows the incumbent's recommendation. ■

Proof of Proposition 4. Fix τ. Let P be the set of all policies, and P_j be the set of policies where firm j is the market winner. Consider the incumbent's first-period recommendation. Given the regulator's decision rule: if the incumbent recommends a policy $p \in P_1$ where he is the market winner, then he will receive his winning payoff $q_1(p)$ in all periods. Otherwise, if the incumbent recommends a policy $p' \in P \setminus P_1$, then in period $\tau + 1$, the winner of that policy will recommend his own preferred policy, which will be implemented in all subsequent periods; thus the incumbent will receive $q_1(p')$ for the first τ periods, and 0 in subsequent periods. It follows that the incumbent's optimal choice in the first period must be one of the following:

1. Recommend p_1^{c} (his highest-quality policy from P_1)
2. Recommend $p_1^{\neg\mathrm{c}}$ (his highest-quality policy from $P \setminus P_1$).

Further, one of these two choices is almost surely strictly optimal. His expected discounted payoff is then

$$V = \begin{cases} \frac{1}{1-\delta} q_1(p_1^{\mathrm{c}}) & \text{if he chooses } p_1^{\mathrm{c}} \\ \frac{1-\delta^\tau}{1-\delta}\, q_1(p_1^{\neg\mathrm{c}}) & \text{if he chooses } p_1^{\neg\mathrm{c}} \end{cases}$$

Notice that $p^{\mathrm{m}} = \arg\max_{p \in \{p_1^{\mathrm{c}}, p_1^{\neg\mathrm{c}}\}} \{q_1(p)\}$. It follows that the τ-regular strategy is optimal for the incumbent given the regulator's strategy. Now consider the regulator's strategy. We say that a policy p is *safe* at a given history if the regulator believes that $q_{w(p)}(p) \geq q^L$ with probability one. The regulator's expected utility from implementing an unsafe policy is $-\infty$, so the regulator only implements safe policies. During the monopoly period, the only safe policies are the status quo and the incumbent's recommendation. The regulator knows that the incumbent's recommendation (if distinct from the status quo) will have higher quality than the status quo, both during and after the monopoly period. It is thus strictly optimal for the regulator to implement the incumbent's recommendation for all rounds $t \leq \tau + 1$. Further, if the winner at period $\tau + 1$ is someone other than the incumbent, then this firm's will recommend his preferred policy in round $\tau+2$; which will (if different from the period-$(\tau+1)$ policy) be weakly (and almost surely strictly) be of higher quality than the period-$(\tau+1)$ policy and thus safe as well. It is thus strictly optimal for the regulator to follow this firm's recommendation for all rounds from $\tau+1$ onwards.

Off-Path Beliefs If the incumbent deviates by making the empty recommendation in the first period, the regulator infers that the incumbent made a mistake and responds by implementing the status-quo policy during the monopoly period. In the continuation equilibrium, the incumbent then recommends his preferred policy p^w in period $\tau+1$, and the regulator implements this policy in all subsequent periods. (This is analogous to the equilibrium outcome from Proposition 2). This outcome is weakly dominated for the incumbent by the on-path outcome where the incumbent recommends p^w in the first period, and thus is not a profitable deviation.

Suppose now that in period $\tau+2$, the period-$\tau+1$ winner deviates by recommending either the status quo or by making an empty recommendation. The regulator infers that the period-$\tau+1$ winner made a mistake, and implements the same policy as in period $\tau+1$. From period $\tau+2$ onwards, the continuation game from period $\tau+1$ is then replayed. This outcome is dominated by the original continuation equilibrium, and thus cannot be a profitable deviation, for the period-$\tau+1$ winner. There are various other potential deviations by the other players; the regulator responds to these by implementing the period-$\tau+1$ winner's recommendation, thus these deviations cannot be profitable. ■

Proof of Proposition 4b. Recall that the optimal monopoly length maximizes the regulator's discounted expected payoff:

$$\tau^* \in \arg\max_{\tau\in(0,\infty)} V(\tau)$$

where

$$V(\tau) = \int_0^\infty \max\{q, q^{\text{SQ}}\} f(q)dq - \int_{q^{\text{SQ}}}^{\frac{q^{\text{SQ}}}{1-\delta^\tau}} (q - q^{\text{SQ}})(1 - F^{n-1}(q))f(q)dq$$
$$+ \int_{\frac{q^{\text{SQ}}}{1-\delta^\tau}}^\infty \delta^\tau (E[Q_{n-1}|q < Q_{n-1}] - q)(1 - F^{n-1}(q))f(q)dq. \quad (3)$$

Let's collect a few observations. First, $V(\tau)$ is continuously differentiable in τ. Second, $V'(\tau) < 0$ for sufficiently large τ. To see why, write

$$V'(\tau) = \frac{d\delta^\tau}{d\tau}(A(\tau) + B(\tau))$$

where

$$A(\tau) = \frac{q^{\text{SQ}}}{(1-\delta^t)^2}\left(q^\star - q^{\text{SQ}} + \delta^t(E[Q_{n-1}|q^\star < Q_{n-1}] - q^\star)\right)(1 - F^{n-1}(q^\star))f\left(q^\star\right) \text{ with } q^\star = \frac{q^{\text{SQ}}}{1-\delta^t},$$

$$B(\tau) = \int_{q^{\text{SQ}}}^{\infty}(E[Q_{n-1}|q < Q_{n-1}] - q)(1 - F^{n-1}(q))f(q)dq.$$

As $t \to \infty$,

$$A(\tau) \to \frac{q^{\text{SQ}}}{(1-\delta^t)^2}\delta^t\left(q^{\text{SQ}} + (E[Q_{n-1}|q^\star < Q_{n-1}] - q^\star)\right)(1 - F^{n-1}(q^\star))f\left(q^\star\right) = \delta^\tau O_{\tau\to\infty}(1),$$

and

$$B(\tau) = \int_{q^{\text{SQ}}}^{\infty}(E[Q_{n-1}|q < Q_{n-1}] - q)(1 - F^{n-1}(q))f(q)dq = O_{\tau\to\infty}(1).$$

Given that $\frac{d\delta^\tau}{d\tau} = \delta^\tau \log\delta < 0$ and that $\lim_{\tau\to\infty} B(\tau) > 0$, it follows that $V'(\tau) < 0$ for sufficiently large τ. Third, $V(\tau)$ is strictly minimized when $\tau = 0$, because

$$V(0) = \int_0^{\infty}\max\{q, q^{\text{SQ}}\}f(q)dq - \int_{q^{\text{SQ}}}^{\infty}(q - q^{\text{SQ}})(1 - F^{n-1}(q))f(q)dq.$$

Collecting these three observations, we infer that the set $T^* = \arg\max_{\tau\in(0,\infty)} V(\tau)$ of optimal monopoly lengths is nonempty and each element of T^* is strictly bounded away from 0 and ∞.

We will now show that T^* is strictly increasing in δ in the sense of weak set order. Notice that, with the substitution $d = \delta^\tau$, we can rewrite (3) as

$$\begin{aligned}\hat{V}(d) = &\int_0^{\infty}\max\{q, q^{\text{SQ}}\}f(q)dq - \int_{q^{\text{SQ}}}^{\frac{q^{\text{SQ}}}{1-d}}(q - q^{\text{SQ}})(1 - F^{n-1}(q))f(q)dq \\ &+ \int_{\frac{q^{\text{SQ}}}{1-d}}^{\infty} d(E[Q_{n-1}|q < Q_{n-1}] - q)(1 - F^{n-1}(q))f(q)dq. \end{aligned} \tag{4}$$

Given $\delta \in (0,1)$, the regulator's optimization problem over τ can be transformed to the equivalent problem

$$T^* = \arg\max_{\tau \in (0,\infty)} V(\tau) \qquad \iff \qquad D^* = \arg\max_{d \in (0,1]} \hat{V}(d)$$
$$\text{where} \quad D^* = \{d^* : d^* = \delta^{\tau^*} \text{ for some } \tau^* \in T^*.\}$$

.

Notice that this transformed problem is invariant to δ. It follows that if δ increases, the set of elements $d^* \in D^*$ remains unchanged while each corresponding element $\tau^* = \frac{\log d^*}{\log \delta}$ increases. That is, T^* is strictly increasing in δ.

So far, we have considered an optimization problem where the regulator chooses any (possibly non-integer) $\tau \in [0, \infty)$. Suppose instead that the regulator chooses from integer-valued monopoly lengths, $\tau \in \{0, 1, 2, ...\}$. Discrete versions of the second and third observations above follow immediately from their continuous counterparts: $V(\tau)$ is strictly decreasing in τ for sufficiently large τ, and $V(\tau)$ is strictly minimized when $\tau = 0$. Thus, an optimal (finite) integer-valued τ^* exists. ■

Proof of Proposition 5. Given the regulator's strategy, the status-quo winner's strategy is optimal: he recommends his preferred policy and has it implemented in all periods $t \geq 2$, while the status-quo policy is implemented regardless of his recommendation in the first period. The other firm's strategies are also best responses: their recommendations have no effect on the chosen policies. Now consider the regulator's strategy. As with the proof of Proposition 4, we say that a policy p is *safe* at a given history if the regulator believes that $q_{w(p)}(p) \geq q^L$ with probability one. The regulator's expected utility from implementing an unsafe policy is $-\infty$, so the regulator only implements safe policies. In the first period, before the regulator learns the identity of the status-quo winner, the only safe policy is the status quo. In subsequent periods, the only safe policies are the status quo and the status-quo winner's recommendation. In particular, the regulator knows that given regular communication, the status-quo winner's recommendation will have (weakly, and sometime strictly) higher quality than the status quo. Thus it is optimal for the regulator to implement the status quo winner's recommendation for all periods $t \geq 2$.

Off-Path Beliefs In the first period, there are three relevant types of off-path outcomes. (i) Two firms make the same recommendation. The regulator infers that one of the two firms made a mistake. (ii) More firms make recommendations than there are policies. The regulator infers that one of the firms made a mistake. (iii) No firms make recommendations. The regulator infers that one of the firms made a mistake. In all three cases, it remains optimal for the regulator to choose the status-quo policy.

For the second period, the regulator's beliefs following off-path outcomes follow those in the baseline model, except with the first-period winner taking the position of the incumbent.

■

Proof of Proposition 5b. Each firm recommends their preferred policy in all periods. In the first period, the regulator chooses the status-quo, which is the only safe policy. In all subsequent periods, the regulator chooses the status-quo winner's recommendation. (On the equilibrium path, this policy and the status quo are the only two safe policies.) The first-period outcome corresponds to the status-quo policy's winning quality:

$$\tilde{q}_n = q_{w(0)}(0) = \max\{q^{SQ}, Q_{n-1}\}.$$

The status-quo winner recommends the alternative, $p = 1$, if he wins there and his $p = 1$ quality is higher than his status-quo quality. The outcome in all periods $t \geq 2$ is thus

$$\tilde{\tilde{q}}_n = \max\{q_{w(0)}(0), s \cdot q_{w(0)}(1) \cdot 1_{q_{w(0)}(1) > \max_{j \neq w(0)}\{q_j(1)\}}\};$$

notice that $E[\tilde{\tilde{q}}_n]$ is continuous in s. Also, the regulator's expected discounted payoff (normalized by factor $1 - \delta$) is

$$E[V_n] = (1 - \delta)E[\tilde{q}_n] + \delta E[\tilde{\tilde{q}}_n].$$

High correlation ($v \to 1$) As $v \to 1$, the probability approaches one that (i) every firm's $p = 1$ quality is exactly factor $s > 1$ larger than his status-quo quality, and thus that (ii) the first-period winner is also the second-period winner, in which case $\tilde{\tilde{q}}_n = s\tilde{q}_n$. Thus as $v \to 1$, we have (pointwise)

$$V_n \to E[(1 - \delta)\tilde{q}_n + \delta s \tilde{q}_n] \to ((1 - \delta) + \delta s)E[\max\{q^{SQ}, Q_{n-1}\}].$$

This expression is strictly increasing in n. It follows that for $n \in \{1, ..., \overline{n}\}$ and for sufficiently large v, V_n must be strictly increasing in n. Thus (for sufficiently large v) the optimal degree of competition is $n = \overline{n}$.

Limited scale ($s \to 1$) Let $\prec$ denote first-order stochastic dominance. If $s = 1$, then

$$\tilde{\tilde{q}}_n \prec \max\{q^{SQ}, Q_{n-1}, q_{w(0)}(1)\} \prec \max\{q^{SQ}, Q_n\};$$

so $E[\tilde{\tilde{q}}_n] < E[\max\{q^{SQ}, Q_n\}]$ for $s = 1$. Given that $E[\tilde{\tilde{q}}_n]$ is continuous in s, it follows that for s sufficiently close to one, for all $n \in \{1, ...\overline{n}\}$, we have $E[\tilde{\tilde{q}}_n] < E[\max\{q^{SQ}, Q_n\}]$ and thus

$$E[\max\{q^{SQ}, Q_{n-1}\}] < E[\underbrace{(1 - \delta)\tilde{q}_n + \delta \tilde{\tilde{q}}_n}_{V_n}] < E[\max\{q^{SQ}, Q_n\}].$$

It follows that $E[V_n]$ is strictly increasing in n within some right-neighbourhood of $s = 1$.

Low correlation, large scale ($v \to 0, s \to \infty$) Denote the probability $\Pr[q'_{w(0)} = \max\{q'_j\}_{j\in\{1,\dots,n\}}]$ that the status-quo-winning firm $w(0)$ also has the highest draw amongst $\{q'_j\}_{j\in\{1,\dots,n\}}$ by $P(n)$. $P(n)$ is independent of s and v, and is strictly less than one and strictly decreasing in n; while $\tilde{q}_n$ is independent of s and v and increasing in n.

As $v \to 0$, $q_{w(0)}(1)$ equals $\hat{q}_1$ with probability approaching one; thus (for given n) $E[q_{w(0)}(1)] = E[\hat{q}_1](1 + o_{v\to 0}(1))$, with the convergence rate being independent of s. Further, as $s \to \infty$ (for given v), the probability that firm $w(0)$ has higher quality under policy $p = 1$ than under the status quo $p = 0$ approaches one; thus the regulator's expected discounted payoff is asymptotically

$$
\begin{aligned}
E[V_n] &= (1-\delta)E[\tilde{q}_n] + \delta E[\tilde{\tilde{q}}_n] \\
&= (1 - \delta + \delta v s + \delta(1 - P(n))(1 - v))E[\tilde{q}_n] \\
&\quad + \delta P(n)(1 - v)E[\max\{q_{w(0)}(0), s q_{w(0)}(1)\}|w(0) = w(1)] \\
&= (1 - \delta + \delta v s + \delta(1 - P(n))(1 - v))E[\tilde{q}_n] + \delta \cdot P(n) \cdot s \cdot E[q_{w(0)}(1)](1 + o_{v\to 0, s\to\infty}(1)) \\
&= \delta \cdot P(n) \cdot s \cdot E[q_{w(0)}(1)](1 + o_{v\to 0, s\to\infty}(1)) \\
&= \delta \cdot P(n) \cdot s \cdot E[\hat{q}_1](1 + o_{v\to 0}(1))(1 + o_{v\to 0, s\to\infty}(1)).
\end{aligned}
$$

Given that $P(n) \cdot s \cdot E[\hat{q}_1]$ is strictly decreasing in n, we can thus select $\overline{s}$ and $\underline{v}$ such that $E[V_n]$ is strictly decreasing in n for $n \in \{1, ..., \overline{n}\}$, $s > \overline{s}$ and $v \in [0, \underline{v})$. That is, for given sufficiently small v and sufficiently large s, the regulator optimally chooses $n = 1$. ■

Proof of Proposition 6. Here, we focus on an equilibrium where the incumbent chooses not to make any offer to an entrant if he is indifferent (as is the case when he is the global winner, $w = 1$).

Suppose the global winner is not the incumbent. It is optimal for the global winner to accept the incumbent's offer: given that the regulator will implement the incumbent's recommendation, the global winner receives a payoff of zero regardless of whether he accepts or rejects the offer. It is also optimal for the incumbent to make an offer of zero to the global winner. Given that the global winner will accept, and the regulator will implement the incumbent's recommendation, the incumbent's payoff equals the global-maximum quality: this is the best possible outcome for the incumbent.

It is optimal for the incumbent to recommend his best policy, given that the regulator will implement his recommendation. It is optimal for the regulator to implement the incumbent's recommendation: the regulator believes that the incumbent recommends the policy that globally maximizes quality (regardless of whether the incumbent makes an acquisition), so there are no profitable deviations for the regulator. ■