Acquisitions, innovation and the entrenchment of monopoly*

Vincenzo Denicolò[§] and Michele Polo[§§]
[§] University of Bologna and CEPR
[§§] Bocconi University, GREEN and IGIER

February 16, 2023

Abstract

We analyze a dynamic model of repeated innovation where inventors may either be acquired by an incumbent or else resist takeover and challenge for leadership. In the short run, acquisitions spur innovation because of the invention-for-buyout effect. In the longer run, however, they may stifle it because of a countervailing effect, the entrenchment of monopoly. The latter occurs when the incumbent's dominance depends on past levels of activity and is therefore reinforced by recurrent acquisitions. We show that if the entrenchment effect is sufficiently strong, forward-looking policymakers should prohibit acquisitions in anticipation of the long-run negative impact on innovation.

Keywords: Acquisitions; Innovation; Market power; Invention-for buyout; Entrenchment-of-monopoly

*E-mail addresses: vincenzo.denicolo@unibo.it, michele.polo@unibocconi.it. We thank, without implicating, Giacomo Calzolari, Igor Letina, Matthew Mitchell, Marc Möller, Massimo Motta, Fausto Panunzi, Patrick Rey, Thom Schmitz, Piercarlo Zanchettin and seminar participants at Amsterdam, Oxford, Paris, the 2021 CRESSE Conference, the 2022 MACCI Conference and the 5th UNIBG Industrial Organization Winter Symposium for useful comments and discussions.

1 Introduction

Technology giants often acquire innovative start-ups with high growth potential.[1] Today, many such acquisitions escape antitrust scrutiny, or are cleared, because merger control focuses on the size of the firms at the time of the takeover, and these target firms are often still small.[2] In the last decade, however, this established approach has increasingly been called into question. Critics argue that it is ill suited to innovative industries, where the acquisition of small entrants may impede Schumpeterian competition (i.e., the replacement of market leaders by new firms) and stifle innovation.[3] Advocates of the permissive policy counter that the prospect of such buyouts heightens the incentive to innovate for small enterprises that lack the assets required to effectively bring their innovations to the market – the so-called *invention-for-buyout* effect.[4]

The entrenchment of monopoly. To contribute to this debate, we propose a Schumpeterian model of repeated innovation and acquisition in which acquisitions have both pro- and anti-competitive consequences. The former stem from the invention-for-buyout mechanism, while the latter derive from a mechanism that we define as *entrenchment of monopoly.*

The entrenchment-of-monopoly effect occurs when an acquisition increases the incumbent's market dominance, its competitive advantage over potential challengers. This effect may follow from different specific mechanisms, but in general tends to arise when the incumbent's strength depends on its past levels of activity. This may be due to such factors as consumer inertia, dynamic economies of scale, exclusive access to more and better data, and the like, all of which imply that acquisitions may strengthen the incumbent's market dominance in the future by increasing its

[1] Recent examples include Facebook's acquisition of WhatsApp and Istagram, Google's acquisition of Youtube and Waze, and Microsoft's acquisition of Linkedin. These prominent cases are just the tip of the iceberg. Focusing only on the "big five", Motta and Peitz (2021) report 42 acquisitions by Amazon, 33 by Apple, 21 by Facebook, 48 by Google, and 53 by Microsoft in the period 2015-2020.

[2] Wallmann (2019) refers to acquisitions that slip under the radar of antitrust agencies as "stealth consolidation."

[3] See e.g. Cremer et al. (2019), Furman et al. (2019), Scott Morton et al. (2019).

[4] After Rasmusen's (1988) "entry-for-buyout."

size today. This obstructs the entry of future inventors, reducing their incentives to innovate. And this is the case even if future inventors are acquired in their turn, because the entrenchment of monopoly worsens their outside options and therefore reduces the share of innovative rents they can negotiate in the bargaining with the incumbent over the price of acquisition.

The results. Our main result is that the effects of acquisitions depend on the time horizon. In the short run, they heighten the incentive to innovate because of the invention-for-buyout effect. In the longer run, however, if the entrenchment effect is strong enough they reduce both the rate of innovation and consumers' surplus. In other words, the buyout effect dominates in the short run, but the entrenchment effect may prevail in the long run.

We also show that the optimal policy on acquisitions may be state dependent: permissive as long as market dominance is weak and restrictive once repeated acquisitions have made it too strong.

Policy implications. Our results carry significant implications for policy. Methodologically, they imply that acquisitions should not be assessed singly, in isolation. This myopic approach, which in our model would produce a lenient policy, is generally sub-optimal. Instead, forward-looking policymakers should consider the cumulative dynamic effects of alternative policy rules.[5]

On substantive grounds, our analysis offers a theory of harm that can be used to block acquisitions that might otherwise go unchallenged. In particular, prohibiting acquisitions benefits consumers if policymakers are sufficiently patient and the entrenchment effect sufficiently strong. The analysis also clarifies the role of other factors, such as the inventors' bargaining power *vis-à-vis* the incumbent, or the speed with which innovations are imitated.

The literature. Although the risk of entrenchment of monopoly is often cited in

[5] This marks an important difference with respect to Nocke and Whinston (2010), where a myopic merger policy based on consumer surplus as a welfare criterion is optimal also in a dynamic setting where a series of mergers may be proposed over time. The reason for this difference is that the entrenchment effect makes our model intrinsically more dynamic: acquisitions affect not only the set of active firms but also future demand, degree of market dominance, and the incentives to innovate.

the acquisition policy debate,[6] to the best of our knowledge this paper is the first formal analysis of this possibility.[7] Previous research on the impact of acquisitions on innovation[8] has either focused on static models of isolated innovations or else posited that the degree of market dominance is time-invariant. In these settings, the entrenchment effect cannot arise.

Models of isolated innovations are the simplest analytical setting in which the invention-for-buyout effect can be demonstrated: see, for instance, Mason and Weeds (2013), Phillips and Zhdanov (2013), and Letina et al. (2020).[9] Static models have also been used to uncover various adverse effects of acquisitions. In an important contribution, Cunningham et al. (2021) have shown, both theoretically and empirically, the profitability of "killer acquisitions," in which the new owner suppresses one or more research projects initiated by the takeover target in order to prevent the cannibalization of its own market. In a similar vein, Kamepalli et al. (2020) suggest the possibility of a "kill zone," where entrants, whose innovations would challenge the incumbent's dominance, are discouraged by the threat of an aggressive reaction. Our analysis abstracts from these effects.

Static models have also shown that acquisitions can have an impact not only on the rate but also on the direction of technological progress. In particular, acquisitions can affect the diversity of research projects (Letina et al., 2020), whether innovators target substitutes or complements of the incumbent's product (Shelegia and Motta, 2021; Dijk, Moraga-Gonzàlez and Motchenkova, 2021), and whether they target the product of the market leader or of the follower (Bryan and Hovenkamp, 2020).

The present paper, instead, forms part of the strand of the literature analyzing antitrust policy in dynamic models of repeated innovation. The pioneering contribution here is Segal and Whinston (2007). Below, we discuss the differences with

[6] See, for instance, Scott Morton et al. (2019) and Bryan and Hovenkamp (2020).

[7] In recent independent work, Fons-Rosen et al. (2022) have developed an endogenous growth model that bears some resemblance to ours. However, their model cannot be solved analytically and thus they resort to numerical analysis.

[8] There is also an extensive literature on the impact of mergers on innovation: see Bourreau et al. (2021) for an excellent synthesis. Unlike that on the acquisition of start-ups, this literature studies mergers that take place before investment in R&D is chosen.

[9] There can also be further positive effects: for example, acquisitions may relax the inventors' financial constraints, as in Fumagalli et al. (2021).

their model at some length; for now, suffices it to say that they do not consider acquisitions and assume that the degree of market dominance is constant over time.

This latter assumption is also made by Cabral (2018, 2021). He distinguishes between incremental and radical innovations. For incremental innovations, the invention-for-buyout effect implies that acquisitions spur innovation. Radical innovations, however, are different: the buyout effect is nil, insofar as these innovations would not be transferred to the incumbent anyway. Still, acquisitions are not neutral because innovators may choose which type of innovation to target. When acquisitions are permitted, incremental innovations may therefore crowd out radical ones. Clearly, this crowding-out mechanism, which may reduce the overall rate of innovation, is different from the entrenchment of monopoly.

Another examination of the invention-for-buyout mechanism is Katz (2021). He notes that acquisitions increase the entrant's payoff; however, the incentive to innovate is determined by the rate with which the payoff increases with the size of the innovation, which in principle may either increase or decrease with acquisitions. This observation also applies to our model, where, however, the buyout effect in itself implies that acquisitions would increase both the level and the slope of the inventor's profit.

While we can solve our model in closed form and derive our results analytically, other papers have used numerical analysis to study a richer industry dynamics. Most existing computational dynamic models however do not feature the entrenchment effect. The results they produce are driven, essentially, by the buyout effect.[10]

Structure of the paper. In the next section, we outline a tractable model of repeated innovation and acquisitions where the entrenchment of monopoly is due to consumer inertia. Section 3 derives the equilibrium. Section 4 examines the effects of acquisitions on the rate of innovation. Sections 5 and 6 analyze the optimal acqui-

[10] In particular, Hollenbeck (2019) examines the trade-off between the static allocative effects of acquisitions, which are always negative, and the dynamic effects *via* level of innovation, which in his model are always positive due to the buyout effect. As a result, he finds that acquisitions are welfare-reducing in the short run but can be positive in the long run. Mermelstein et al. (2020) consider a model where entry may be inefficient because of economies of scale in production and investment. The invention-for-buyout effect implies that acquisitions facilitate entry, but the inefficiency of entry implies that a restrictive policy may be optimal.

sition policy when antitrust authorities adopt non-contingent and state-contingent policy rules, respectively. Section 7 discusses a different mechanism that may underpin the entrenchment effect. Section 8 summarizes and concludes. The proofs are set out in online Appendix A.

2 The model

We propose a tractable model of repeated innovation, where in the absence of acquisitions, incumbents would be systematically replaced by new innovators. The possibility of acquiring these challengers may however lead to the persistence of monopoly.

The model is tailored to industries where the ability to innovate is diffused, so that it is unlikely that the same firm will innovate repeatedly, or that the successful innovator can be identified before the innovation is developed. Once the innovation has been developed, on the other hand, the incumbent can identify the potential challenger precisely.[11]

These features are embedded in a fully specified model of the industry, whose other components are kept as simple as possible in order to obtain closed-form solutions.

2.1 Demand

In each period $t = 1, 2, ...$, a vertically differentiated product is purchased by a mass of infinitely-lived homogeneous consumers (normalized to 1), who may demand either 0 or 1 unit. The net utility from one unit of a product of quality q^i purchased at price p^i is

$$U_t^i = q_t^i - p_t^i, \tag{1}$$

[11] In reality, incumbents often acquire inventors when they are still so small that it is difficult to gauge whether they actually pose a threat for leadership. They may therefore acquire a large number of start ups, in the hope of catching the few ones that have the potential to become the new market leaders before they get too strong.

where the willingness to pay for quality is also normalized to 1. The utility of not purchasing is normalized to 0 (these normalizations entail no loss of generality).

2.2 Innovation, entry and market structure

In each period $t = 1, 2, ..,$ one outsider, drawn at random from a number of potential innovators, gets an idea for improving the existing technology q_{t-1}. This outsider then becomes the period-t inventor by developing the idea into an innovation, i.e., a product of quality $q_t > q_{t-1}$.[12] The inventor chooses the magnitude of innovation $\Delta_t = q_t - q_{t-1}$ so as to maximize its profits. The cost of raising quality by Δ_t, $C(\Delta_t)$, is independent of the current level q_{t-1} and is quadratic in Δ_t. With another innocuous normalization, we can write:

$$C(\Delta_t) = \frac{1}{2}\Delta_t^2. \tag{2}$$

After developing its invention, the inventor firm q_t enters the market. In the absence of acquisitions, it is the technological leader but faces competition from the incumbent (i.e., inventor q_{t-1}). In period $t+1$, inventor q_t becomes the new incumbent and competes with inventor q_{t+1}.

As time passes, inventions can be imitated by a competitive fringe. We assume that the innovation is used exclusively by the inventor for two periods and afterwards can be imitated freely. (For example, the invention might be protected by a patent that lasts for two periods.) As a consequence, in the absence of acquisitions inventor q_{t-2} is absorbed by the competitive fringe in period t. Thus, in each period t there are three types of firm: an entrant (E), which supplies a product of quality $q_t^E = q_t$, an incumbent (I) with quality $q_t^I = q_{t-1}$, and a competitive fringe (F) with $q_t^F = q_{t-2}$.[13]

[12] With a modest abuse of notation, we denote by q_t both the quality and the identity of inventor.

[13] One can allow imitation to be faster, say because intellectual property protection is imperfect. For example, continuing to assume that invention q_t is fully protected in period t and can be imitated freely in period $t+2$, the innovation could be imitated partially in period $t+1$. In this case, the competitive fringe's quality would be $q_t^F = q_{t-2} + \vartheta(q_{t-1} - q_{t-2})$, where the parameter ϑ is an index of the speed of imitation, or an inverse index of the strength of intellectual property protection. When $\vartheta = 1$, the competitive fringe imitates the innovation in just one period, whereas the baseline case in which it takes two periods is re-obtained for $\vartheta = 0$. With this more general formulation, the only change in our formulas is that the discount factor δ is replaced by $\delta(1-\vartheta)$.

The unit production cost c is independent of quality and is normalized to 0.

2.3 Market dominance

Following Segal and Whinston (2007), we assume that although it is technological laggard, the incumbent may have acquired some other competitive advantage, owing, for instance, to such factors as consumer inertia, intertemporal network externalities, dynamic economies of scale, or exclusive access to more and better data. Conversely, entrants may face entry hurdles; for example, some consumers may be unwilling to try new products, or may not even be aware of their existence.

We differ from Segal and Whinston (2007), who assume that the incumbent's dominance is time-invariant, by allowing it to change over time as a function of the industry's past history. In particular, in our model dominance is strengthened by acquisitions, which can therefore lead to the entrenchment of market power.

To be specific, we assume that demand is not entirely contestable:[14] in each period, a fraction μ_t of consumers are "captive" and cannot purchase from the new entrant; the remaining $1 - \mu_t$ consumers are "free" and can buy the new product.[15] The size of the incumbent's captive consumer base, μ_t, is therefore an index of the *degree of market dominance.*

The innovator forms its captive consumer base in the first period of its life cycle, by captivating a fixed fraction κ of its customers. It can then exploit this consumer base in the next period, when it becomes the new incumbent. Thus, the number of captive consumers evolves over time according to the following equation (the superscript NA stands for "no acquisition"):

$$\mu_t^{NA} = \kappa(1 - \mu_{t-1})x_{t-1}^E, \tag{3}$$

Thus, the parameter δ that we use throughout the model can be thought of as capturing both the private rate of time preference and the strength of intellectual property protection.

[14] This is a fairly common assumption in the analysis of exclusionary conduct: see, for instance, Ide and Montero (2020), Oertel and Schmutzler (2021), and the literature cited therein.

[15] Captive consumers can always purchase from the fringe. This limits their exploitation by the incumbent and ensures the stationarity of the model.

where $1-\mu_{t-1}$ is the number of free consumers in period $t-1$, who represent the potential buyers from the q_{t-1} innovator, x_{t-1}^E is the fraction of such free consumers who actually buy from the innovator, and κ is the fraction of the innovator's customers who become loyal (or captive).

The key property of this formulation is that the incumbent's degree of market dominance in period t, μ_t, depends on its past sales, $(1-\mu_{t-1})x_{t-1}^E$. Equation (3), which captures this broad idea in a specific way, may be interpreted literally, or taken as a metaphor for various possible reasons why the dependence may arise.

The literal interpretation may be justified as follows. Suppose that consumers face a cost of switching to the latest product (they have to learn how to use it, say, or have to conduct a search to learn of its existence). These learning costs are heterogeneous. For a fraction κ of consumers, they are sufficiently high that it is worth paying them only to move ahead by two quality steps.[16] As a consequence, a high-cost consumer who purchased the state-of the art product in period $t-1$, q_{t-1}, would not be willing to switch to product q_t in period t. In other words, such a consumer would be captive in period t. (In the absence of acquisitions, he will return to be a free consumer in period $t+1$). For the remaining fraction $1-\kappa$ of consumers, learning costs are negligible, so they can always purchase all the products on offer, including the latest. In this interpretation, it is learning costs that create inertia in market shares.

2.4 Acquisitions

If acquisitions are allowed, the incumbent may take over the inventor after it has fully developed its new product. The merged entity resulting from the acquisition is denoted by M. The entrant furnishes M with its new technology q_t, which is ready for use without incurring in any further development cost,[17] while the incumbent brings its exclusive control over the old technology q_{t-1} and its captive consumer

[16] Consumers do not need to be posited permanently high- or low-cost: the parameter κ may be interpreted as the probability of a consumer's being high-cost in a given period. Any level of correlation across periods is consistent with our formulation.

[17] This distinguishes ours from models of killer acquisitions.

base.

We assume that consumers who are captive to firm M can purchase any product from it, including the newest one.[18] If they do, they may remain captive for several periods in a row. Specifically, we assume that firm M's captive consumer base in period $t+1$ comprises a fraction κ of the free consumers that it served in period t and a fraction ξ, possibly greater than κ, of the already captive ones. As a result, in period $t+1$ the merged firm will have (the superscript A stands for "acquisition"):

$$\mu_{t+1}^{A} = \kappa(1-\mu_t)x_t^{F,M} + \xi\mu_t x_t^{C,M} \tag{4}$$

captive consumers, where $x_t^{F,M}$ denotes the fraction of free consumers and $x_t^{C,M}$ the fraction of captive consumers served by firm M in period t.[19] Since the merged entity can serve more consumers than the entrant alone, it can build a larger captive consumer base for the next period. This is how acquisitions increase market dominance in our model, creating what we call the entrenchment-of-monopoly effect.

We assume that acquisitions will take place whenever they are jointly profitable. The acquisition price paid by the incumbent, P_t,[20] determines the division of the surplus among the two parties. The price is determined by a simple bargaining process, where one of the two firms is picked up randomly to make a take-it-or-leave-it offer to the other. We denote by α the probability that the entrant makes the offer and the incumbent receives it; with probability $1-\alpha$, these roles are reversed. Thus, α is the share of the bargaining surplus obtained on average by the entrant – a measure of its bargaining power.[21]

[18] In the learning cost interpretation of the model, one may imagine that as long as consumers purchase from the same firm, they bear no more learning cost. In other words, the merged entity can remove the factors that would otherwise prevent captive consumers from purchasing the new product q_t. For example, the merged entity may ensure backward compatibility with product q_{t-1}, it may facilitate the transition by providing the same usage modes, or it may guarantee a seamless transfer of data to the new service. See Kamepalli et al. (2020) for a similar assumption.

[19] In the learning-cost interpretation of the model, this assumption can be justified on the ground that consumers who purchased from the same firm repeatedly in the past may have undermined their ability to learn or search.

[20] The assumption that the incumbent acquires the entrant, and not the other way around, is just an accounting convention. Nothing would change if the roles were reversed.

[21] We adopt a strategic approach to the bargaining process in order to avoid mixing notions from cooperative and non-cooperative game theory. In any case, many different bargaining solutions

2.5 Timing and Payoffs

We consider an infinite horizon in discrete time. Each period t is divided into three stages. In the first stage (*ex ante*), the inventor is randomly selected and chooses the innovation size, Δ_t. In the second (*interim*), the incumbent and the entrant bargain over the acquisition price and, if they reach an agreement, the acquisition occurs. In the third stage (*ex post*), firms compete in prices. This sequence is repeated in every period $t = 1, 2, ...$ If acquisitions are prohibited, the second stage is absent.

Firms are risk neutral and maximize intertemporal profits, future values being discounted by the common discount factor $\delta < 1$. Total discounted profits as of time t are denoted by Π_t^i, current profits by π_t^i, with $i \in \{E, I, M\}$.

2.6 Equilibrium

We analyze the Markov perfect equilibria of this game. Under our assumptions, at the beginning of each period t, the payoff-relevant variables are μ_t, q_{t-1} and q_{t-2}. At the *interim* stage, i.e. once the entrant has chosen the size of the innovation Δ_t, they also include q_t.

3 The acquisition game

We find the model's equilibrium under the assumption that acquisitions are always permitted. Since the merged entity can replicate any behavior of both entrant and incumbent, acquisitions are weakly profitable. In fact, we shall presently show that they are always strictly profitable. This implies that acquisitions will always take place in equilibrium.

To ensure a perfect equilibrium, we start from the pricing subgames and proceed backwardly to the bargaining process and choice of innovation size.

---

would lead to the same expected outcome as our non-cooperative assumptions.

3.1 Pricing subgames

We assume that firms cannot distinguish between captive and non captive consumers, so price discrimination is not practicable.

On path. We begin from the pricing subgame that is actually played on the equilibrium path, i.e., the one starting after the acquisition.

The merged entity's only competitor is the fringe, which supplies the best freely available quality, $q_t^F = q_{t-2}$, and prices it at cost, $p_t^F = 0$. The equilibrium strategy of the merged firm is given by the following lemma.[22]

Lemma 1 *The merged entity supplies only one product of quality $q_t^M = q_t$. It serves all consumers ($x_t^{FM} = x_t^{CM} = 1$) at price $p_t^M = q_t - q_{t-2} = \Delta_t + \Delta_{t-1}$, reaping a profit of*

$$\pi_t^M = \Delta_t + \Delta_{t-1}. \tag{5}$$

The intuition is simple. The competitive fringe does not sell any output in equilibrium but exerts competitive pressure by providing an outside option to consumers. The merged entity undercuts the fringe in utility space, charging a price equal to the value of the quality differential.

Note that the presence of the competitive fringe prevents prices and profits from increasing without limit even though the quality level continues to rise over time. From an economic point of view, this guarantees that all benefits from technological progress eventually accrue to consumers; from an analytical point of view, it guarantees the stationarity of the equilibrium.

Note also that the merged entity always uses the state-of-the-art technology q_t; the model excludes "killer acquisitions."

Off path. Next, we characterize the price equilibrium that arises, out of the equilibrium path, if the incumbent does not acquire the entrant.[23] In this case, there

[22]To simplify the presentation, we adopt the following tie-breaking rule: when a consumer or a firm is indifferent among different actions, it chooses the one that maximizes aggregate profits. This assumption captures the idea that the stronger firm could shave the price marginally in order to overcome the indifference.

[23]The same equilibrium arises also on the equilibrium path if acquisitions are prohibited (see footnote 27 below).

are two active firms besides the fringe. In the baseline specification, we assume that these two firms price sequentially, with the incumbent acting as the price leader. As we shall see, this implies that acquisitions do not affect consumer surplus for a given state of the technology. For our purposes, this is a conservative property that biases the analysis against prohibiting acquisitions.[24] Alternative timings are considered later.

Lemma 2 *If the incumbent acts as price leader, it serves all captive consumers and the entrant serves all free consumers ($x_t^E = 1$). The incumbent prices at $p_t^I = q_{t-1} - q_{t-2} = \Delta_{t-1}$ and obtains a profit of*

$$\pi_t^I(\mu_t) = \mu_t \Delta_{t-1}. \tag{6}$$

The entrant's equilibrium price is $p_t^E = q_t - q_{t-2} = \Delta_t + \Delta_{t-1}$, so the profit it earns in the first period of its life cycle is

$$\pi_t^E(\mu_t) = (1 - \mu_t)(\Delta_t + \Delta_{t-1}). \tag{7}$$

When the incumbent acts as price leader, both the incumbent and the entrant slightly undercut the competitive fringe in utility space, and the entrant also slightly undercuts the incumbent. Consequently, consumers get the same net utility from any firm they may buy from.[25] The incumbent does not compete more aggressively for the free consumers because it anticipates that it would be underpriced by the entrant.

Implications. Lemmas 1 and 2 carry several significant implications. First, the per-period profit functions π_t^i are additively separable in the quality steps Δ_t and Δ_{t-1}. This property is central to the following analysis, as it allows for a closed-form

[24] Sequential pricing also facilitates the analysis by guaranteeing the existence of a pure-strategy equilibrium. With simultaneous moves, there is generally no pure-strategy pricing equilibrium. Intuitively, the existence of captive consumers is analogous to a capacity constraint, as the entrant cannot supply more than $(1 - \mu_t)$ units.

[25] Note, incidentally, that this implies that free and captive consumers obtain the same net utility in equilibrium. As a consequence, consumers have no incentive not to buy the state-of-the-art version of the product in order to avoid becoming captive.

solution.

Second, the lemmas imply the following:

Corollary 1 *Acquisitions are always strictly profitable.*

Acquisitions are profitable for two reasons. From the static viewpoint, they facilitate the diffusion of the innovation: the state-of-the-art product is sold not only to the free but also to the captive consumers. From the dynamic viewpoint, they increase the fraction of captive consumers that the merged entity can exploit in the next period.

Third, the lemmas imply:

Corollary 2 *Either with or without acquisitions, consumers obtain exactly the surplus guaranteed to them by the fringe:*

$$CS_t = q_{t-2}. \tag{8}$$

From the consumer viewpoint, therefore, acquisitions matter only to the extent that they affect innovation, which determines their future surplus. As noted, this is a conservative property for our purposes.

Finally, Lemmas 1 and 2 prove that in equilibrium the entrant serves all free consumers ($x_t^E = 1$), and the merged entity all consumers ($x_t^{F,M} = x_t^{C,M} = 1$). Therefore, with no acquisitions the share of captive consumers evolves over time according to:

$$\mu_t^{NA} = \kappa(1 - \mu_{t-1}). \tag{9}$$

With acquisitions, the dynamics of μ_t is:

$$\mu_t^A = \kappa(1 - \mu_{t-1}) + \xi\mu_{t-1}. \tag{10}$$

3.2 The acquisition price

Proceeding with our backward induction, consider next the bargaining over the acquisition price. Firms are forward looking and correctly anticipate all the future consequences of their choices. Since entrants are systematically acquired, the acquisition price must coincide with the entrant's value function (gross of the innovation cost). This is determined simultaneously with the value functions for the incumbent and the merged entity, as we shall see presently.

To proceed, it is important to keep in mind that in a Markov perfect equilibrium, the value functions depend only on the payoff-relevant variables. From the foregoing, it appears that profits depend on μ_t and the quality differentials Δ's. Thus, the period-t payoff-relevant variables are $\{\mu_t, \Delta_{t-1}\}$ at the *ex ante* stage and $\{\mu_t, \Delta_{t-1}, \Delta_t\}$ at the *interim* stage. Accordingly, we denote by $V_t^i(\mu_t, \Delta_{t-1})$ the firms' *ex ante* value functions, and by $v_t^i(\mu_t, \Delta_{t-1}, \Delta_t)$ the *interim* functions, for $i \in \{E, I, M\}$. These value functions must satisfy the following conditions (to simplify the notation, we suppress the dependence of the *interim* value functions on the relevant variables when this does not create confusion):

$$v_t^M = \pi_t^M + \delta V_{t+1}^I(\mu_{t+1}^A, \Delta_t) \tag{11}$$

$$v_t^E = (1-\alpha)\left[\pi_t^E + \delta V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)\right] + \alpha\left(v_t^M - \pi_t^I\right) \tag{12}$$

$$v_t^I = v_t^M - v_t^E. \tag{13}$$

Equation (11) says that the merged entity obtains profits π_t^M in period t and then becomes the new incumbent with μ_{t+1}^A captive consumers, yielding a continuation value of $\delta V_{t+1}^I(\mu_{t+1}^A, \Delta_t)$. According to (12), the acquisition price (which, as noted, coincides with the entrant's value function) equals the entrant's disagreement payoff plus a fraction α of the bargaining surplus. The entrant's disagreement payoff is equal to the current profit if it resists the takeover, π_t^E, plus the continuation value, $\delta V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)$. The "one-shot deviation principle" implies that the continuation value must be calculated on the expectation that even if there was no acquisition in period t, entrant q_t, once it has become the new incumbent in period $t+1$, will acquire entrant q_{t+1}. At that point, however, it will have only μ_{t+1}^{NA} captive consumers. As for the period-t incumbent, its disagreement payoff is simply π_t^I, given that with no agreement it would exit the market in the next period. The bargaining surplus is therefore $v_t^M - \left[\pi_t^E + \delta V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t) + \pi_t^I\right]$, from which there follows condition (12). The final condition says that the value of being the incumbent must be equal to the value of the merged entity minus the acquisition price. In other words, the acquisition does not change the sum of the firms' values because the extra-profits created by the merger are already included in the forward-looking valuation of the firms.

The system of equilibrium conditions (11)-(13) cannot be solved for the *interim* value functions yet, because it also involves the *ex ante* value functions $V_{t+1}^I(\mu_{t+1}^A, \Delta_t)$ and $V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)$, which depend on Δ_t. In turn, Δ_t might potentially depend on the future values Δ_{t+1}, Δ_{t+2} etc. To proceed, we must therefore consider the optimal choice, *ex ante*, of the size of the innovation.

3.3 The innovation size

Plainly, the equilibrium innovation size must satisfy the following condition:

$$\Delta_t^A(\mu_t, \Delta_{t-1}) = \arg\max_{\Delta_t} \left[v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) - \frac{1}{2}\Delta_t^2 \right]. \tag{14}$$

In a Markov perfect equilibrium, this optimal choice of Δ_t is anticipated by all players, furnishing a link between the *ex ante* and *interim* value functions. That is, the *ex ante* value must be equal to the *interim* value calculated at the optimal innovation size:

$$V_t^i(\mu_t, \Delta_{t-1}) = v_t^i\left[\mu_t, \Delta_{t-1}, \Delta_t^A(\mu_t, \Delta_{t-1})\right] \qquad \text{for } i \in \{E, I, M\}. \tag{15}$$

This completes the set of conditions that must all hold simultaneously in equilibrium.

3.4 Equilibrium

It is easy to see that the set of Markov perfect equilibria coincides with the set of solutions to the system of equilibrium conditions (11)-(15), given the profit functions (5), (6) and (7).

The solution can be calculated explicitly thanks to a key simplifying property of the model, which we noted above: the profit functions π_t^i are additively separable in Δ_t and Δ_{t-1}. This separability implies that while the value function $v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)$ depends on Δ_{t-1}, the marginal value of increasing the innovation size, $\frac{\partial v_t^E(\mu_t,\Delta_{t-1},\Delta_t)}{\partial \Delta_t}$, does not. Therefore, the optimal innovation size in period t, $\Delta_t^A(\mu_t, \Delta_{t-1})$, is independent of Δ_{t-1}, and in turn Δ_t does not affect the future values Δ_{t+1}, $\Delta_{t+2}, ...$, in spite of the forward-looking nature of system (11)-(15).

These properties of the model imply that the derivative $\frac{\partial v_t^E(\mu_t,\Delta_{t-1},\Delta_t)}{\partial \Delta_t}$ can be calculated even without full knowledge of the value function $v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)$, allowing a two-stage solution. In the first stage, we calculate the derivative and find the equilibrium innovation size for any value of μ_t, $\Delta_t^A(\mu_t)$. With this function in hand, in the second stage we determine the value function $v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)$ by a guess-and-verify method. Further details are provided in the proof of Lemma 3 in the Appendix. Applying this procedure, we get:

Lemma 3 *In the baseline model, the equilibrium innovation size depends on the share of captive consumers* μ_t *and is*

$$\Delta_t^A(\mu_t) = (1+\delta\kappa)(1-\mu_t) + \alpha(1+\delta\xi)\mu_t. \tag{16}$$

The ex ante value functions are

$$V_t^E(\mu_t, \Delta_{t-1}) = \phi_0 + \phi_1\mu_t + \phi_2\mu_t^2 + (1-\mu_t)\Delta_{t-1} \tag{17}$$

$$V_t^I(\mu_t, \Delta_{t-1}) = \varphi_0 + \varphi_1\mu_t + \varphi_2\mu_t^2 + \Delta_{t-1} \tag{18}$$

The coefficients ϕ_n and φ_n, for $n = 0, 1, 2$, depend on the exogenous parameters α,

δ, κ and ξ and are reported in the Appendix.[26] Given $V_t^E(\mu_t, \Delta_{t-1})$ and $V_t^I(\mu_t, \Delta_{t-1})$, one can easily recover $V_t^M(\mu_t, \Delta_{t-1})$ and the *interim* value functions $v_t^i(\mu_t, \Delta_t, \Delta_{t-1})$, $i = M, I, E$, from conditions (11)-(15).

4 Acquisitions and innovation

In this section, we analyze the impact of acquisitions on innovation. We show that prohibiting acquisitions always reduces the equilibrium size of innovation in the short run but may increase it in the long run if the entrenchment effect is large enough.

4.1 Benchmark: no acquisitions

To proceed, we determine the innovation size when acquisitions are prohibited. In this case, the entrant's payoff is

$$\Pi_t^{E,NA} = \pi_t^E(\mu_t) + \delta\pi_{t+1}^I(\mu_{t+1}^{NA}), \tag{19}$$

where the profit functions are the same as in Lemma 2.[27] The equilibrium innovation size with no acquisitions then is $\Delta_t^{NA}(\mu_t) = \arg\max_{\Delta_t}\left[\Pi_t^{E,NA} - \frac{1}{2}\Delta_t^2\right]$. Simple calculations lead to the following:

Lemma 4 *If acquisitions are always prohibited, the equilibrium size of innovation is*

$$\Delta_t^{NA}(\mu_t) = (1 + \delta\kappa)(1 - \mu_t). \tag{20}$$

4.2 Innovation and market dominance

We start by showing that market dominance always has an adverse effect on innovation.

[26] The Appendix also verifies that $V_t^I(\mu_t, \Delta_{t-1})$ increases with μ_t, a property that we used in the derivation of the pricing equilibria.

[27] This is not self-evident, because firms are forward looking, and the entrant's continuation value is different with and without acquisitions. With them, the continuation value is $\delta V^I(\mu_{t+1}^A, \Delta_t)$; without, it is $\delta\pi_{t+1}^I(\mu_{t+1}^{NA})$. However, the proof of Lemma 2 shows that all that matters is that the continuation value is non-decreasing in μ_{t+1}, which is true in both cases.

Proposition 1 *The equilibrium size of innovation Δ_t is a decreasing function of the degree of market dominance μ_t both when acquisitions are permitted (eq. (16)) and when they are prohibited (eq. (20)).*

To understand why Proposition 1 holds, consider first the case of no acquisitions. Innovator q_t's marginal benefit from increasing its innovation size is the increase in the discounted sum of its profits in the two stages of its life cycle. Inspection of the profit functions (7) and (6) shows that the marginal profit is equal to the number of free consumers $1 - \mu_t$ in the first period, and to the number of captive consumers $\mu_{t+1}^{NA} = k(1 - \mu_t)$ in the second. Both decrease with μ_t.

With acquisitions, the mechanism is analogous. The q_t innovator's outside option when bargaining on the acquisition price is the profit $\pi_t^E(\mu_t)$ that it would obtain if it resisted the takeover plus the continuation value $\delta V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)$. An increase in μ_t decreases both $\pi_t^E(\mu_t)$ and μ_{t+1}^{NA}, and hence $V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)$. This implies that a higher μ_t worsens the innovator's outside option and thus reduces its incentives to invest even if it expects to be bought out.

4.3 The short run

Comparing (16) and (20) one immediately obtains:

Proposition 2 *Provided that $\frac{1+\delta\kappa}{1+\delta\xi} > \alpha > 0$ and $\mu_t > 0$, prohibiting acquisitions reduces the equilibrium size of innovation in period t:*

$$\Delta_t^A(\mu_t) > \Delta_t^{NA}(\mu_t). \tag{21}$$

This result reflects the invention-for-buyout effect. Intuitively, the innovation is more valuable in the hands of the incumbent, which can supply the state-of-the-art product not only to the free but also to the captive consumers. By transferring the new technology to the incumbent, acquisitions create a surplus, a share of which, when $\alpha > 0$, goes to the inventors. The prospect of being bought out thus increases the value of the innovation to forward-looking inventors, hence their incentives to innovate.[28] The greater the entrant's bargaining power α, the stronger this invention-for-buyout effect.

4.4 The long run

However, acquisitions also affect the dynamics of μ_t. Starting from an arbitrary μ_t, if acquisitions are permitted μ will converge to its steady state level

$$\bar{\mu}^A = \frac{\kappa}{1+\kappa-\xi}, \tag{22}$$

whereas if acquisitions are prohibited the steady state is:[29]

$$\bar{\mu}^{NA} = \frac{\kappa}{1+\kappa}. \tag{23}$$

Clearly, $\bar{\mu}^A > \bar{\mu}^{NA}$. This inequality reflects the entrenchment of monopoly due to acquisitions. The strength of the entrenchment effect can be measured by the percentage increase in the long-run degree of market dominance:

$$\frac{\bar{\mu}^A - \bar{\mu}^{NA}}{\bar{\mu}^{NA}} = \frac{\xi}{1+\kappa-\xi}. \tag{24}$$

When $\xi = \kappa$, the entrenchment effect is simply ξ. In general, the effect increases with ξ, which can therefore be regarded as the entrenchment parameter.

In the steady state, if acquisitions are always prohibited the level of innovation is

$$\Delta^{NA}(\bar{\mu}^{NA}) = \frac{1+\delta\kappa}{1+\kappa}. \tag{25}$$

If acquisitions are always permitted, on the other hand, it is

$$\Delta^{A}(\bar{\mu}^{A}) = \frac{(1+\delta\kappa)(1-\xi)+\alpha\kappa(1+\delta\xi)}{1+\kappa-\xi}. \tag{26}$$

[28] This is not a foregone conclusion, however. The incentive to innovate is not determined by the impact of acquisitions on the inventor's profit, but by the marginal profitability of the innovation size. Proposition 2 guarantees that in our model the marginal and total effects go hand in hand. See Katz (2021) for a model where this property does not necessarily hold.

[29] While both μ_t^A and μ_t^{NA} converge to their respective steady state levels, the dynamics of the former is monotonic, that of the latter oscillatory.

Comparing (25) and (26), it appears that if the entrenchment effect is sufficiently strong, the positive short-run effect of acquisitions on innovation may be reversed in the long run.

Proposition 3 *In the steady state, prohibiting acquisitions increases the equilibrium size of innovation if*

$$\xi > \frac{\alpha\,(1+\kappa)}{1+\delta\kappa-\alpha\delta(1+\kappa)}. \tag{27}$$

Intuitively, the long-run effect of acquisitions is the sum of two components: the difference between Δ_t^A and Δ_t^{NA} for any given μ_t, and the difference between $\bar{\mu}^A$ and $\bar{\mu}^{NA}$. The first component reflects the buyout effect and is positive. The second component reflects the entrenchment effect and is negative. Condition (27) determines when the second component prevails over the first one.

The condition simplifies considerably in the special case $\kappa = \xi$, when it reduces to:

$$\xi > \frac{\alpha}{1-\alpha}. \tag{28}$$

Intuitively, the entrenchment parameter ξ must be large and the entrant's bargaining power α, which determines the magnitude of the invention-for-buyout effect, must be small.

When $\xi > \kappa$, other factors come into play. Prohibiting acquisitions is more likely to raise the long run level of innovation the lower the private discount factor (and hence the lower the speed of imitation, or the stronger the protection of intellectual property), and the lower the fraction of free consumers that are turned into captive.

The effects that we have identified in this section are depicted in Figure 1, with the case of acquisitions being permitted in red, prohibited in blue. First, both with and without acquisitions, the rate of innovation decreases as market dominance increases, reducing innovators' ability to appropriate the returns from their innovations (Proposition 1). Second, for any given level of dominance, the rate of innovation is higher when acquisitions are permitted (Proposition 2), reflecting the invention-for-buyout effect. Third, the long-run degree of dominance is higher if acquisitions are

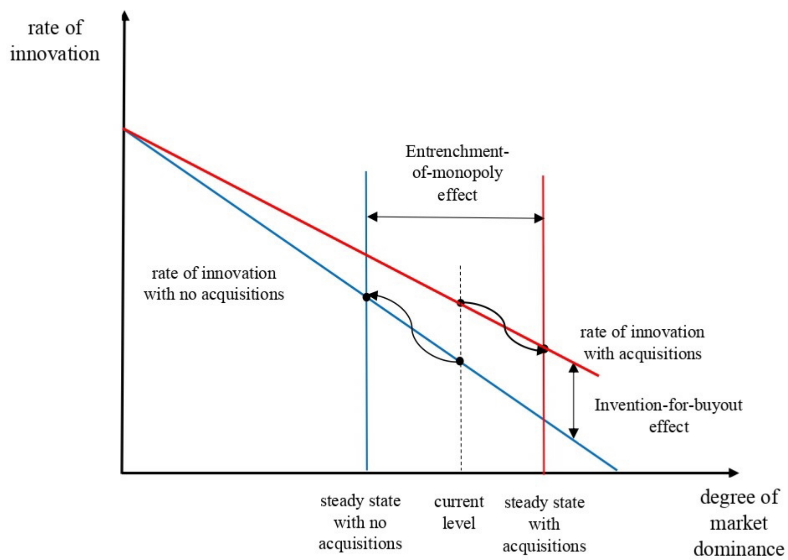

Figure 1: The decreasing lines represent the equilibrium rate of innovation when acquisitions are permitted (red) or prohibited (blue). The distance between the two lines measures the invention-for-buyout effect. The vertical lines represent the long-run degree of market dominance, which is higher when acquisitions are pemitted because of the entrenchment effect. The curley arrows represent the process of convergence to the steady state, starting from the current level of market dominance.

permitted than if they are prohibited, reflecting the entrenchment-of-monopoly effect.

4.5 Transitory dynamics

Our model is tractable enough to allow explicit calculation of the equilibrium dynamics of the innovation size Δ_{t+n} for $n = 1, 2, ...$, starting from an arbitrary μ_t. When acquisitions are permitted, the degree of market power evolves over time as follows:

$$\mu_{t+n}^{A} = \frac{\kappa}{1+\kappa-\xi} + \left(\mu_t - \frac{\kappa}{1+\kappa-\xi}\right)(\xi-\kappa)^n, \tag{29}$$

and thus the level of innovation is:

$$\begin{aligned}\Delta_{t+n}^{A} \quad = \quad & \frac{1+\kappa\left[\alpha(1+\delta\xi)+\delta(1-\xi)\right]-\xi}{1+\kappa-\xi} + \\ & -\left[1+\delta\kappa-\alpha(1+\delta\xi)\right]\left(\mu_t - \frac{\kappa}{1+\kappa-\xi}\right)(\xi-\kappa)^n.\end{aligned} \tag{30}$$

When, on the contrary, acquisitions are prohibited, we have:

$$\mu_{t+n}^{NA} = \frac{\kappa}{1+\kappa} + \left(\mu_t - \frac{\kappa}{1+\kappa}\right)(-\kappa)^n \tag{31}$$

and

$$\Delta_{t+n}^{NA} = \frac{1+\delta\kappa}{1+\kappa} - (1+\delta\kappa)\left(\mu_t - \frac{\kappa}{1+\kappa}\right)(-\kappa)^n. \tag{32}$$

Figure 2 illustrates the dynamics of Δ_{t+n}, starting from an arbitrary $\mu_t \in \left(\bar{\mu}^{NA}, \bar{\mu}^{A}\right)$. Consider a shift from a lenient policy (the red curve) to a restrictive one (the blue curve). Immediately with the policy change, the level of innovation drops, as the invention-for-buyout effect vanishes. In subsequent periods, however, the share of captive consumers μ_{t+n} shrinks, reducing the degree of market dominance, with a positive effect on the entrant's innovative effort, which increases over time. In the counterfactual where acquisitions are permitted, on the other hand, μ_{t+n} increases towards its steady state level $\bar{\mu}^{A}$. The figure illustrates the case where condition (27) holds. In this case, at some point in time the innovation size with

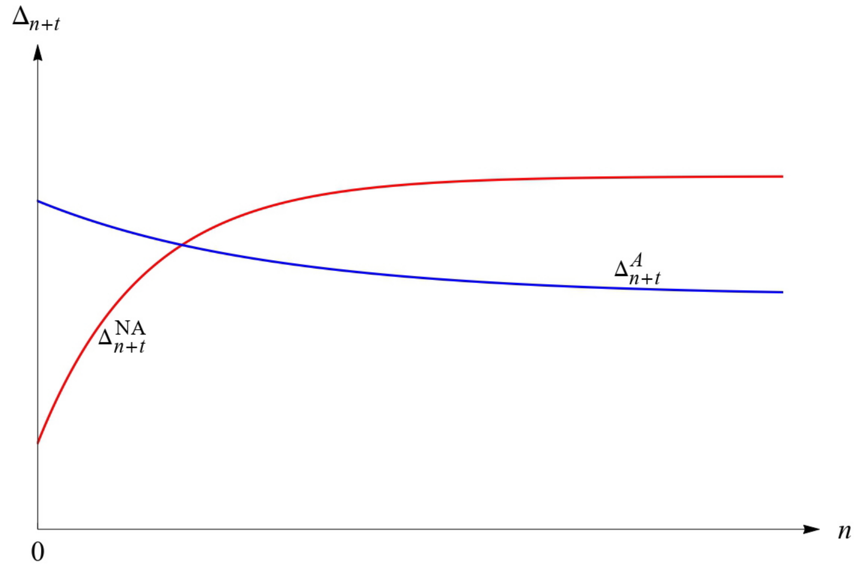

Figure 2: The dynamics of equilibrium innovation size when acquisitions are permitted and prohibited. The figure depicts a continuous time approximation of the discrete dynamics, which eliminates the oscillations that may be exhibited by the discrete dynamics. The picture has been drawn for $\kappa = \frac{3}{4}$, $\xi = \frac{19}{20}$, $\delta = \frac{19}{20}$, $\alpha = \frac{9}{20}$ and $\mu_t = \frac{1}{2}$.

acquisitions banned becomes larger than if acquisitions continued to be permitted.

5 Policy rules

We now analyze the optimal antitrust policy in our model, assuming that the agencies take consumer surplus as their objective and discount future values by the social discount factor δ_S.[30] In view of (8), the policymaker's objective function as of period t is

$$\begin{aligned}\sum_{n=0}^{\infty} CS_{t+n}\delta_S^n &= \sum_{n=0}^{\infty} \delta_S^n q_{t+n-2} \\ &= \frac{1}{1-\delta_S}\left(q_{t-2} + \delta_S q_{t-1} + \delta_S^2 \sum_{s=0}^{\infty} \delta_S^n \Delta_{t+n}\right). \qquad (33)\end{aligned}$$

[30] The social discount factor δ_S is generally greater than the private discount factor δ, because benevolent policymakers ought to be more farsighted than private firms, and because δ may reflect not only the private rate of time preference but also the speed of imitation (see footnote 13). However, our formulas would continue to hold even if $\delta_S < \delta$.

The first two terms inside brackets are pre-determined, so the objective function effectively reduces to

$$W_t = \sum_{n=0}^{\infty} \delta_S^n \Delta_{t+n}. \tag{34}$$

Thus, social welfare comparisons boil down to the comparison of the discounted sum of current and future sizes of innovation.

5.1 Fixed v. contingent policies

In principle, the optimal acquisition policy may vary over time. In particular, if the antitrust authorities can observe the state of the industry μ_t, they may condition acquisition policy on it.[31] In this subsection, however, we show that in our baseline model it is never optimal to take advantage of this possibility.

We start this analysis with a preliminary result. So far, we have considered only the case where acquisitions are always permitted or always prohibited. Under a state-dependent policy, however, acquisitions may be permitted in certain periods and prohibited in others, as the degree of market dominance varies. A convenient simplifying property of the model is that the size of innovation in period t in fact depends only on whether acquisitions are permitted or prohibited in that period, not in subsequent periods.

Lemma 5 *If acquisitions are permitted in period t, then $\Delta_t(\mu_t) = \Delta_t^A(\mu_t)$ irrespective of acquisition policy in all subsequent periods. Likewise, if acquisitions are prohibited in period t, then $\Delta_t(\mu_t) = \Delta_t^{NA}(\mu_t)$ irrespective of acquisition policy in all subsequent periods.*

This is not a foregone conclusion, as firms are forward looking, and future acquisition policy affects the continuation values in the dynamic game. The conclusion nonetheless holds, because the profit functions π_t^i are additively separable in Δ_{t-1} and Δ_t. This property is inherited by the value functions, which implies that changes

[31] Strictly speaking, the state also includes Δ_{t-1}. However, this variable does not impact on future innovation and affects consumer surplus in an additive, separable way; hence, it cannot affect the optimal policy.

in future acquisition policy cause parallel shifts in $v_t^E(\Delta_t)$ that do not affect the incentive to innovate.[32]

Armed with Lemma 5, we can now show that in the baseline model, the possibility of conditioning acquisition policy on the degree of market dominance μ_t is in fact valueless.

Proposition 4 *In the baseline model, the optimal acquisition policy does not depend on* μ_t.

This result relies on the fact that $\Delta_t^A(\mu_t)$ and $\Delta_t^{NA}(\mu_t)$ are linear functions of μ_t, and that μ_{t+1} is a linear function of μ_t. This linearity entails a special property: while the short-run gains and the long-run losses from permitting acquisitions both depend on μ_t,[33] the sign of the sum of the two does not. This property is quite special, and later we shall analyze an extension of the baseline model where it no longer holds.

5.2 Optimal fixed policy rules

In view of Proposition 4, in the rest of this section we focus on non-contingent policy rules, assuming that acquisitions are either always approved or always prohibited. The choice is made once and for all in a generic period t.

Let us compare the two policy regimes for an arbitrary initial μ_t. If acquisitions are always permitted, using (30)and (29) social welfare becomes:

$$
\begin{aligned}
W_t^A(\mu_t) &= \Delta_t^A(\mu_t) + \sum_{n=1}^{\infty} \delta_S^n \Delta_{t+n}^A(\kappa) \\
&= \frac{(1+\delta\kappa)(1-\xi\delta_S) - \alpha\kappa\delta_S(1+\delta\xi)}{(1-\delta_S)\left[1-\delta_S(\xi-\kappa)\right]} - \frac{1-\alpha+\delta\kappa-\alpha\delta\xi}{1-\delta_S(\xi-\kappa)}\mu_t. \qquad (35)
\end{aligned}
$$

If instead acquisitions are always prohibited, using (32) and (31) social welfare be-

[32] Lemma 5 therefore continues to hold in the variant of the model considered in the next section, where the profit functions are still additively separable in Δ_t and Δ_{t-1}.

[33] The short-run gain is the difference $\Delta_t^A(\mu_t) - \Delta_t^{NA}(\mu_t)$, whereas the long-run loss is the difference between $\sum_{n=1}^{\infty} \delta_S^n \Delta_{t+n}$ starting from $\mu_{t+1}^A(\mu_t)$ and $\mu_{t+1}^{NA}(\mu_t)$, respectively.

comes:

$$
\begin{aligned}
W_t^{NA}(\mu_t) &= \Delta_t^{NA}(\mu_t) + \sum_{n=1}^{\infty} \delta_S^n \Delta_{t+n}^{NA}(\mu_{t+n}^{NA}) \\
&= \frac{1+\delta\kappa}{(1-\delta_S)(1+\delta_S\kappa)} - \frac{1+\delta\kappa}{1+\delta_S\kappa}\mu_t.
\end{aligned} \tag{36}
$$

Comparing $W_t^{NA}(\mu_t)$ and $W_t^{A}(\mu_t)$, we get:

Proposition 5 *Prohibiting acquisitions increases social welfare if and only if*

$$
\xi > \frac{\alpha(1+\delta_S\kappa)}{(1-\alpha)\delta_S\delta\kappa + (\delta_S - \alpha\delta)}. \tag{37}
$$

The condition simplifies considerably in the case $\xi = \kappa$, when it reduces to:

$$
\xi > \frac{\alpha}{(1-\alpha)\,\delta_S}. \tag{38}
$$

The effects of ξ and α are the same as in Proposition 3, and for the same reasons. That is, prohibiting acquisitions is the more likely to be optimal, the higher the entrenchment-of-monopoly parameter ξ and the lower the invention-for-buyout parameter α. Furthermore, prohibiting acquisitions is the more likely to be optimal, the higher the social discount factor δ_S. This makes intuitive sense: in our model, prohibiting acquisitions is socially costly in the short run but may bring about long-run benefits. It is therefore logical that a restrictive policy may be optimal only if the policymaker is sufficiently farsighted. When $\delta_S \to 1$ condition (37) collapses to (27): the weight of the transitory dynamics in the social welfare calculation becomes negligible, so the welfare comparison depends only on the steady-state levels of innovation.

When $\xi > \kappa$, two more parameters come into play. Prohibiting acquisition is the less likely to be optimal, the lower the private discount factor δ, and the lower the fraction of free consumers that are turned into captive κ. Since a higher discount factor δ also captures the possibility of slower imitation, as discussed in footnote 13, Proposition 5 suggests that in our model acquisition policy and patent policy may be interconnected: when entrants are better protected against imitation, acquisition policy should be more lenient, while weaker patent protection calls for stricter antitrust rules.

6 State-dependent policy

In this section, we consider a variant of the model where the optimal acquisition policy may vary over time.[34]

As noted, the fact that a non-contingent policy rule is optimal in the baseline model relies on the linearity of the profit functions in the degree of market power μ_t. To allow for the possibility that the optimal policy may be state-dependent, we therefore consider a variant in which the entrant's profit π_t^E is a non-linear function of μ_t. Specifically, we assume that if the acquisition does not take place, the entrant, rather than the incumbent, acts as price leader. This is enough to invalidate the conclusion of Proposition 4.

Reversing the order of moves in the pricing game changes the pricing equilibrium when there is no acquisition. (In the case of acquisitions, nothing changes.) The incumbent's price and profit are the same as in the baseline model, but the entrant's profit in the first period of its life cycle becomes :[35]

$$\pi_t^E(\mu_t) = (1 - \mu_t)\left(\Delta_t + \mu_t \Delta_{t-1}\right). \tag{39}$$

The profit is lower than in the baseline, because the entrant has to cut its price down to the point where the incumbent's incentive to compete for the free consumers is

[34] Omitted details and proofs for this section may be found in online Appendix B.

[35] Under the assumption that the dynamics of market dominance is given directly by (9) and (10), the case of simultaneous moves would produce the same results as when the entrant is price leader. The reason for this is that firms would price myopically, as the impact of acquisitions on future market dominance would not depend on current output levels. Now, in a static pricing game of simultaneous moves, the equilibrium generally involves mixed strategies, but in any case each firm obtains the same payoff as it would if it acted as price leader (as in Kreps and Scheinkman, 1983). Therefore, the entrant's profit would be (39), and the incumbent's profit would be (6). Retaining the original assumptions (3) and (4), however, the simultaneous-move pricing game becomes untractable.

eliminated – a form of limit pricing.[36] However, the marginal impact of Δ_t on π_t^E is the same as in the baseline model, so if acquisitions are prohibited the equilibrium level of innovation does not change. If they are permitted, on the other hand, the new level of innovation is:

$$\Delta_t^A = 1 - (1-\alpha)\mu_t + \delta \left[1 - (1-\xi) \frac{1 + (1-\alpha)(1-\kappa)\xi}{1+\kappa-\xi} \mu_t \right]. \tag{40}$$

Comparing (40) and (20), one sees that acquisitions always spur innovation in the short run. However, as in the baseline model, acquisitions may stifle innovation in the long run. The condition is again that the entrenchment-of-monopoly effect be sufficiently strong; formally, $\xi > \tilde{\xi}(\alpha, \delta, \kappa)$, where the new critical threshold $\tilde{\xi}$ is an increasing function of α, κ and δ.

6.1 Cut-off policy rules

We continue to assume that the authorities maximize consumer welfare. If the acquisition takes place in period t, the consumer surplus is $CS_t^A = q_{t-2}$, as in the baseline model. However, if there is no acquisition the surplus now becomes:

$$CS_t^{NA} = q_{t-2} + (1-\mu_t)^2 \Delta_{t-1}. \tag{41}$$

The additional term, which now arises because firms compete more aggressively, represents the static allocative gains from prohibition.

The discounted consumer surplus is now:

$$\sum_{n=0}^{\infty} \delta_S^n CS_{t+n} = \frac{q_{t-2} + \delta_S q_{t-1} + \delta_S^2 \sum\limits_{s=0}^{\infty} \delta_S^n \Delta_{t+n}}{1-\delta_S} + \sum_{n=0}^{\infty} \mathbf{1}_{t,NA} \left(1 - \mu_{t+n}\right)^2 \delta_S^n \Delta_{t+n}, \tag{42}$$

where $\mathbf{1}_{t,NA}$ is an indicator function equal to 1 if acquisitions are prohibited in period t and to 0 if they are permitted.

For simplicity, we focus on the limiting case $\delta_S \to 1$, allowing us to abstract

[36] If the incumbent acts as price leader, it will not even try to compete for the free consumer in the anticipation that it would be outpriced by the entrant anyway.

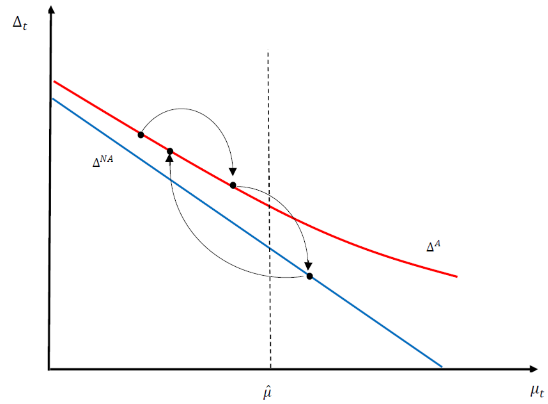

Figure 3: A non-stationary cycle of length $\ell = 2$.

from the second term in the above expression. Indeed, the relative weight in the social welfare calculation of this term, which captures the static effect of prohibition, becomes negligible as δ_S approaches 1. Intuitively, the static allocative costs of acquisitions are transitory, whereas the effects on innovation size are permanent. As $\delta_S \to 1$, welfare comparisons rest uniquely on how acquisition policy affects the long-run level of innovation.

6.1.1 Acquisition cycles

For tractability, in the rest of this section we focus on a class of simple policy rules, where the policy-maker permits acquisitions as long as $\mu_t < \hat{\mu}$ and prohibits them when $\mu_t \geq \hat{\mu}$ for some cut-off value $\hat{\mu}$.[37] The variable $\hat{\mu}$ may be interpreted as the degree of lenience of acquisition policy.

Such cut-off policies produce cycles in which the industry oscillates between periods where market dominance is low, and acquisitions are permitted, and periods where it is high, and they are prohibited. As long as acquisitions are permitted the degree of market dominance μ_t increases, until it crosses the threshold $\hat{\mu}$. When it does, the acquisition is prohibited, μ_t jumps down, and a new cycle starts. Figure 3 illustrates.

[37] We believe that the optimal policy belongs to this class, even though we have not been able to prove this conjecture.

It is convenient to define the length of a cycle, ℓ, as the number of consecutive periods in which acquisitions are permitted between two successive prohibitions. Cycles are degenerate, and policy is effectively state independent, when $\ell = 0$ (acquisitions are always prohibited) or $\ell = \infty$ (acquisitions are always permitted).

Starting from an arbitrary μ_t, the cycles generated by a given cut-off $\hat{\mu}$ are, generically, non stationary, as the one shown in Figure 3. However, there exist limit cycles where the system follows exactly the same steps over and over again. At the end of a limit cycle of length ℓ (and period $\ell + 1$), the value of μ_t is

$$\mu^H(\ell) = \frac{\kappa \sum_{n=0}^{\ell} (\xi - \kappa)^n}{1 + \kappa (\xi - \kappa)^\ell}, \tag{43}$$

and at the beginning of the cycle it is $\mu^L(\ell) = \kappa \left[1 - \mu^H(\ell)\right]$.

These limit cycles are stable attractors.

Lemma 6 *Under a cut-off policy rule, the degree of market dominance μ_t converges to a limit cycle. The length of the limit cycle is* 0 *if*

$$\hat{\mu} \leq \mu^H(0) = \bar{\mu}^{NA},$$

it is $\ell + 1$ if

$$\mu^H(\ell) < \hat{\mu} \leq \mu^H(\ell + 1), \tag{44}$$

and it is infinite if

$$\hat{\mu} > \lim_{\ell \to \infty} \mu^H(\ell) = \bar{\mu}^A.$$

Lemma 6 says that the length of the limit cycles is a stepwise increasing function of the cut-off $\hat{\mu}$ and thus it depends on policy lenience. Intuitively, the greater $\hat{\mu}$, the longer it takes, starting from any given $\mu_t < \hat{\mu}$, to pass the threshold. Furthermore, the greater $\hat{\mu}$, the lower μ_t at the beginning of the next cycle.

Asymptotically, a cut-off policy is fully characterized by the length ℓ of the limit cycle which the industry converges to. If $\hat{\mu}$ changes but remains within one and the same interval of Lemma 6, the change in $\hat{\mu}$ may affect the industry's transitory dynamics but not the long-run level of innovation.

6.1.2 Optimal cycles

We are interested in determining whether acquisition cycles may be optimal, and if so, when. Let us start from the case $\kappa = \xi$. In this case, there is only one non-degenerate limit cycle, which has length $\ell = 1$ and period 2; its initial and final point are, respectively, $\mu^L = \kappa(1-\kappa)$ and $\mu^H = \kappa$.

The following proposition says that inducing such a cycle may indeed be the optimal policy when the policy-maker is nearly indifferent between always permitting and always prohibiting acquisitions. This is true, in particular, when α and δ are not too large.

Proposition 6 *If $\kappa = \xi$, provided that α and δ are not too large, there exists a neighborhood of $\tilde{\xi}(\alpha, \delta)$ such that when ξ lies in that set, the optimal acquisition policy entails a period-2 cycle.*

The analytical characterization of the region of parameter values where such a neighborhood exists is unmanageable, but the region can be identified by numerical methods.

Note that any value of $\hat{\mu}$ in the interval $[\kappa(1-\kappa), \kappa]$ leads to the same limit cycle, so the asymptotic dynamics of the industry does not depend on the exact value of $\hat{\mu}$ within that interval. Nevertheless, the choice of $\hat{\mu}$ requires care, in that when $\hat{\mu}$ is close to κ, the merged firm might opt for self-restraint in order to retain the possibility of additional acquisitions in the future. From this standpoint, the merged firm should voluntarily limit its sales so as to prevent μ_t from crossing the threshold $\hat{\mu}$.[38] In particular, when $\mu_t = \hat{\mu} < \kappa$ the firm should serve only a fraction $x^M = \frac{\hat{\mu}}{\kappa}$ of its potential demand, so that $\mu_{t+1} = \kappa \frac{\hat{\mu}}{\kappa} = \hat{\mu}$. When $\hat{\mu}$ is close to κ, the amount of profit foregone by restricting output in this way would be small, and the strategy would therefore be profitable.

[38] The merged firm cannot raise its price because of the competitive pressure from the fringe. An alternative way not to cross the threshold is to resist takeover, but it is easy to see that this strategy is always dominated by merging and then restraining output.

One may wonder whether it might be optimal to induce such self-restraint by the merged firm. In fact, it is not.[39] Therefore, the policy-maker must be careful not to set $\hat{\mu}$ too high. But if $\hat{\mu}$ is set just barely above $\kappa(1-\kappa)$, it is easy to see that the merged entity would definitely serve its entire demand.[40] Therefore, when $\hat{\mu}$ is marginally higher that $\kappa(1-\kappa)$ there will be no self-restraint, and the industry will shuttle between the high- and low-dominance states.

Turning to the general case $\xi > \kappa$, we have analyzed the optimal length of the limit cycles by means of numerical calculations. Figure 4 illustrates an example where, depending on the value of ξ, cycles of different lengths may be optimal. Numerical calculations show that the optimal length of the limit cycle – which as noted is an index of the lenience of acquisition policy – is (i) a stepwise decreasing function of the entrenchment-of-monopoly parameter ξ, (ii) a stepwise increasing function of the entrant's bargaining power α, (iii) a stepwise decreasing function of the private discount factor δ, and (iv) a stepwise decreasing function of the capture rate κ.

7 Absorptive capacity

So far, we have focused on a specific source of market dominance, namely, consumer inertia. However, there might be other reasons why acquisitions may create, or enhance, a competitive advantage for incumbents, in spite of the lower quality of their products. To demonstrate the robustness of our mechanism, in this section we analyze a toy model of absorptive capacity, in which acquisitions affects incumbents' incentives to imitate the entrants.

[39] This is proved in Appendix B.

[40] The reason is that to avoid crossing the threshold, the merged entity would have to serve only $1-\kappa$ consumers. In this case, however, its profit would be lower than the aggregate profit of the incumbent and the entrant in the absence of acquisition. In other words, acquisitions would no longer be profitable, so retaining the option to future acquisitions would be valueless.

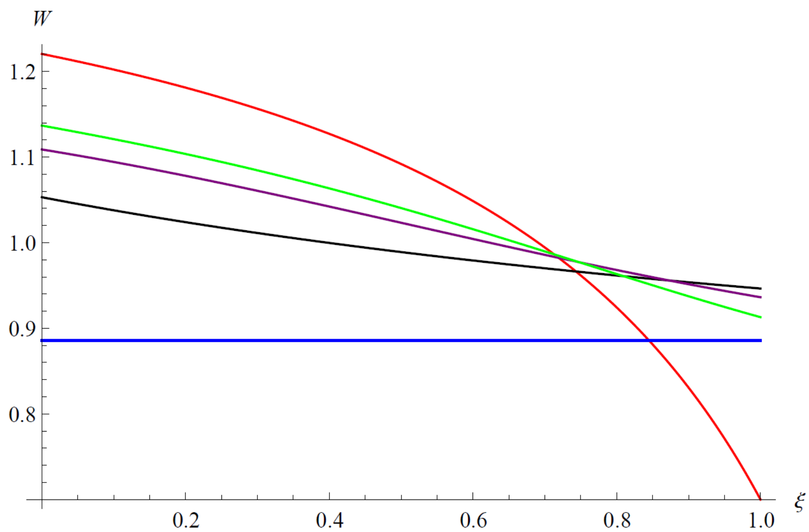

Figure 4: Social welfare under different lengths of the limit cycles. For low values of ξ, the optimal policy is to always permit acquisitions ($\ell = \infty$, the red curve). As ξ increases, limit cycles of finite length become optimal. The figure represents the cases of limit cycles of length $\ell = 3$ (the green curve), $\ell = 2$ (the purple curve), and $\ell = 1$ (the black curve). The blue curve is the case where acquisitions are always prohibited. The figure has been drawn for $\alpha = 0.1$, $\delta = 0.6$ and $\kappa = 0.4$. For these values, permanent prohibition is never optimal.

7.1 Assumptions

Consumers' utility function is still given by (1), and firms' types and dynamics are as in the baseline model. However, all consumers now are free, eliminating the consumer inertia mechanism that we have focused on so far. The new mechanism relies instead on the incumbent's ability to imitate the entrant's product.

Specifically, we now assume that if the incumbent makes an investment in "absorptive capacity," it can supply a product of quality

$$q_t^I = q_{t-1} + \phi \Delta_t \tag{45}$$

rather than q_{t-1}, where $\phi \in [0, 1]$. The cost of the investment is positive but arbitrarily small. The new mechanism of entrenchment relies on the fact that the investment is profitable only if the incumbent acquires the entrants.

We start with a stationary version of the model. After explaining the new mechanism, we shall briefly discuss a more dynamic setting.

7.2 Equilibrium

Since all consumers are free, the pricing equilibrium no longer depends on the order of moves. The q_t-entrant always serves all consumers, charging a price cost margin of $q_t^E - q_t^I$ and obtaining a profit of $\pi_t^E = q_t^E - q_t^I$. Therefore, the entrant's profit π_t^E is Δ_t if the incumbent does not invest in absorptive capacity but is $(1-\phi)\,\Delta_t$ if it does. As for the incumbent, it always obtains $\pi_t^I = 0$. With acquisitions, on the other hand, the merged entity's profit is $\pi_t^M = \Delta_t + \Delta_{t-1}$, as in the baseline model.

It appears that when acquisitions are prohibited, incumbents have no incentive to invest in absorptive capacity, as they would make zero profits anyway. Therefore, the entrant's discounted profit over its life cycle is $\Pi_t^E = \Delta_t$. In view of (2), this immediately implies that:

$$\Delta_t^{NA} = 1. \tag{46}$$

When acquisitions are permitted, on the other hand, incumbents gain from the investment in absorptive capacity, as it worsens the entrants' outside option and thus allows incumbents to extract a larger share of the rents from the innovation. Since the investment cost is arbitrarily small, the merged firm will therefore always invest in absorptive capacity.

As in the baseline model, the incentive to innovate with acquisitions is determined by the value functions v_t^i.[41] These are now determined by the following conditions:

$$v_t^M = \Delta_t + \Delta_{t-1} + \delta v_{t+1}^I \tag{47}$$

$$v_t^E = (1-\alpha)\left[(1-\phi)\,\Delta_t + \delta v_{t+1}^I\right] + \alpha v_t^M \tag{48}$$

$$v_t^I = v_t^M - v_t^E. \tag{49}$$

Condition (47) says that the value of the merged firm is its current profit $\pi_t^M = \Delta_t + \Delta_{t-1}$ plus the continuation value, which is the discounted value of being the

[41] In this static version of the model, there are no state variables so the value functions reduce to constants, and it is no longer necessary to distinguish between *ex ante* and *interim* values.

incumbent in period $t+1$. Condition (48) says that the entrant's payoff is a weighted combination of its disagreement payoff and the bargaining surplus, with weights given by its bargaining power α. The disagreement payoff is the sum of the current profit, which is $\pi_t^E = (1-\phi)\,\Delta_t$, plus the continuation value. (The disagreement payoff of the current incumbent is nil, as with no acquisition this incumbent makes zero profits and is then absorbed by the competitive fringe.) The joint payoff in case of agreement, on the other hand, is v_t^M, so the bargaining surplus is $v_t^M - \left[(1-\phi)\,\Delta_t + \delta v_{t+1}^I\right]$. Condition (49) says that acquisitions do not change the firms' aggregate value, as in the baseline.

Solving the system (47)-(49) yields:

$$v_t^E = \left[1-(1-\alpha)(\phi-\delta)\right]\Delta_t + \alpha\Delta_{t-1} + (1-\alpha)\,\delta\phi\Delta_{t+1}. \tag{50}$$

The corresponding equilibrium level of innovation is

$$\Delta_t^A = 1-(1-\alpha)(\phi-\delta)\,. \tag{51}$$

In this stationary version of the model, Δ_t^A is constant over time. The level of innovation is a decreasing function of ϕ, which represents the new entrenchment parameter.

Comparing (46) and (51), one sees immediately that acquisitions stifle innovation if

$$\phi > \delta. \tag{52}$$

Once again, the intuition is that the entrenchment effect must be strong enough to outweigh the invention-for-buyout effect. Like in the baseline model, prohibiting acquisitions is more likely to raise the level of innovation the lower the private discount factor δ. This simple model therefore yields the same basic insight as the baseline model.

7.3 Dynamics

One may obtain a more fully dynamic model by assuming that the merged firm's absorptive capacity varies over time. For example, suppose that if the firm enters period t with an absorptive capacity of ϕ_{t-1} and invests in that period, its ability to imitate the entrant becomes:

$$\phi_t = \vartheta \left(\bar{\phi} - \phi_{t-1} \right), \tag{53}$$

where $\bar{\phi}$ is the maximum feasible level of absorptive capacity and $\vartheta \in [0, 1]$ is the rate at which ϕ_t is accumulated. This captures the idea that the longer a firm has been the industry leader, the better it can appropriate the results of outsiders' innovative activity.

In this case, the merged entity will invest in all periods if acquisitions are permitted. As a result, its absorptive capacity will gradually increase, converging towards its steady state level $\bar{\phi}$. Assuming that $\phi_0 < \delta < \bar{\phi}$, it then follows that acquisitions increase innovation in the short run but decrease it in the longer run, as in the baseline model.

8 Conclusion

We have analyzed a tractable model of repeated innovation, where incumbents may either compete with innovative entrants or else acquire them. Acquisitions have both positive and negative effects on innovation. The former stems from the invention-for-buyout mechanism: inventors earn more by transferring their innovations to the incumbent than by exploiting them themselves, so their incentive to innovate is greater when such technology transfers are permitted. The negative effect, on the other hand, derives from the entrenchment of monopoly due to acquisitions. When these are permitted, that is to say, incumbents come to enjoy a higher degree of market dominance, which in turn reduces the entrants' incentive to innovate.

We have shown that the invention-for-buyout effect always prevails in the short run but can be outweighed in the long run by the entrenchment effect. As a result, if policymakers are sufficiently farsighted and the entrenchment effect is sufficiently strong, prohibiting acquisitions may be the optimal policy. In some cases, the optimal policy may be state-dependent. In other words, it may be best to permit acquisitions as long as market dominance is weak and prohibit them once repeated acquisitions have made it too strong.

Throughout our analysis, we have assumed that the sole objective of the antitrust authorities is consumer welfare. It is sometimes contended that this narrow focus is responsible for the leniency of antitrust policy. However, we have shown that if the authorities are forward looking and consider the cumulative dynamic effects of different policy rules, taking consumer surplus as the welfare criterion may justify a restrictive policy on acquisitions.

Our results imply that the small size of the target firm should not provide a shield against antitrust scrutiny. The critical policy variable should not be the size of the takeover target but the incumbent's degree of market dominance. To the extent that in innovative industries this correlates with the size of the incumbent, then it is the latter's size that should matter in the antitrust assessment.

We believe that the mechanism highlighted in this paper is quite general, but we have carefully chosen specific assumptions that allow for a closed-form solution. Extending the model therefore requires care. In particular, the special property that allows us to calculate the value functions explicitly is that the profit functions are linear in the current and past innovation sizes, Δ_t and Δ_{t-1}. If this property fails, one must either resort to numerical solutions, or else restrict attention to a two-period model where the first period is a stylized representation of the short run and the second period, the long run. These alternative approaches, while being perhaps less elegant, may allow one to analyze richer variants of the model.

References

[1] Bryan, K. A. and Hovenkamp, E. (2020a). Antitrust limits on startup acquisitions. Review of Industrial Organization.

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

[3] Bourreau, M., Jullien, B., and Lefouili, Y. (2021). Mergers and demand-enhancing innovation.

[4] Cabral, L. (2018). Standing on the shoulders of dwarfs: Dominant firms and innovation incentives, CEPR DP 13115.

[5] Cabral, L. (2021). Merger policy in digital industries. Information Economics and Policy, 54, 100866.

[6] Cunningham, C., Ederer, F., and Ma, S. (2021). Killer acquisitions. Journal of Political Economy, 129(3), 649-702.

[7] Fons-Rosen, C., Roldan-Blanco, P., & Schmitz, T. (2022). The Aggregate Effects of Acquisitions on Innovation and Economic Growth. Available at SSRN 3785485.

[8] Fumagalli, C., Motta, M., and Tarantino, E. (2021). Shelving or developing? The acquisition of potential competitors under financial constraints. Mimeo.

[9] Hollenbeck, B. (2020). "Horizontal mergers and innovation in concentrated industries," Quantitative Marketing and Economics, 18(1): 1-37.

[10] Ide, E., and Montero, J. P. (2020). A Theory of Must-Haves, mimeo.

[11] Kamepalli, S. K., Rajan, R., and Zingales, L. (2020). Kill zone. Stigler Center for the Study of the Economy and the State.

[12] Katz, M. L. (2021). Big Tech mergers: Innovation, competition for the market, and the acquisition of emerging competitors. Information Economics and Policy, 54, 100883.

[13] Katz, M. L. (2020). Big-Tech Mergers: Schumpeterian Competition, Innovation, and the Treatment of Emerging Competitors.

[14] Kreps, D. M., and Scheinkman, J. A. (1983). Quantity precommitment and Bertrand competition yield Cournot outcomes. The Bell Journal of Economics, 326-337.

[15] Letina, I., Schmutzler, A., and Seibel, R. (2020). Killer acquisitions and beyond: policy effects on innovation strategies. University of Zurich, Department of Economics, Working Paper, (358).

[16] Mason, R., and Weeds, H. (2013). Merger policy, entry, and entrepreneurship. European Economic Review, 57, 23-38.

[17] Mermelstein, B., Nocke, V., Satterthwaite, M. A., & Whinston, M. D. (2020). Internal versus external growth in industries with scale economies: A computational model of optimal merger policy. Journal of Political Economy, 128(1), 301-341.

[18] Motta, M., and Peitz, M. (2021). Big tech mergers. Information Economics and Policy, 54, 100868.

[19] Nocke, V., and Whinston, M. D. (2010). Dynamic merger review. Journal of Political Economy, 118(6), 1200-1251.

[20] Oertel, C., and Schmutzler, A. (2021). Challenging the incumbent: Entry in markets with captive consumers and taste heterogeneity. University of Zurich, Department of Economics, Working Paper, (386).

[21] Phillips, G. M., and Zhdanov, A. (2013). R&D and the incentives from merger and acquisition activity. The Review of Financial Studies, 26(1), 34-78.

[22] Rasmusen, E. (1988). Entry for buyout. The Journal of Industrial Economics, 281-299.

[23] Scott-Morton, F., Bouvier, P., Ezrachi, A., Jullien, B., Katz, R., Kimmelman, G., Melamed, D., and Morgenstern, J. (2019). Committee for the Study of Digital Platforms. Report. Stigler Center for the Study of the Economy and the State.

[24] Segal, I., & Whinston, M. D. (2007). Antitrust in innovative industries. American Economic Review, 97(5), 1703-1730.

[25] Shelegia, S., and Motta, M. (2021). The "kill zone": Copying, acquisition and start-ups' direction of innovation.

[26] Spulber, D. F. (2013). How do competitive pressures affect incentives to innovate when there is a market for inventions?. Journal of Political Economy, 121(6), 1007-1054.

[27] Stein, J. C. (1997). Waves of creative destruction: Firm-specific learning-by-doing and the dynamics of innovation. The Review of Economic Studies, 64(2), 265-288.

[28] Wollmann, T. G. (2019). Stealth consolidation: Evidence from an amendment to the Hart-Scott-Rodino Act. American Economic Review: Insights, 1(1), 77-94.

For online publication
Appendix A
Proofs

Proof of Lemma 1. The merged entity's objective function is

$$\Pi_t^M = \pi_t^M + \delta V_{t+1}^I(\mu_{t+1}^A, \Delta_t),$$

where $\pi_t^M = x^M p_t^M$ is the current period profit, δ is the discount factor, and $\delta V_{t+1}^I(\mu_{t+1}^A, \Delta_t)$ is the continuation value, i.e., the discounted value of being the incumbent in the next period with $\mu_{t+1}^A = \kappa x_t^M$ captive consumers. Note that $x^{F,M} = x^{C,M} = x^M$ as the firm cannot price discriminate.

To begin with, assume that the merged entity prices myopically, i.e., $\delta = 0$. Since all consumers are effectively identical in this case, there is no incentive to price discriminate by supplying different quality levels. Thus, the dominant firm will supply only the highest quality, q_t. Since the competitive fringe guarantees to all consumers an outside option of $U_t^F = q_{t-2}$, the merged entity must match this utility level:

$$U_t^M = q_t^M - p_t^M = U_t^F,$$

with a tiny price discount to break the indifference, if necessary. Therefore, $p_t^M = \Delta_t + \Delta_{t-1}$. In this myopic equilibrium, $x_t^M = 1$.

Next suppose that $\delta > 0$. It is intuitive (and we shall confirm below) that the continuation value $V_{t+1}^I(\mu_{t+1}^A, \Delta_t)$ is a non-decreasing function of μ_{t+1}^A, which is in turn a non-decreasing function of x_t^M. Therefore, a forward-looking firm would have an incentive to further reduce the price so as to increase x_t^M if possible. But since x_t^M is already equal to 1, the myopic price remains optimal also for a forward-looking firm. ■

Proof of Lemma 2. Plainly, all firms supply the highest quality level that they control: $q_t^E = q_t$, $q_t^I = q_{t-1}$, and $q_t^F = q_{t-2}$, and the fringe prices at marginal cost (i.e., 0). The incumbent and the entrant, on the other hand, price so as to maximize their respective profits:

$$\pi_t^I = \mu_t x_t^I p_t^I$$

and

$$\Pi_t^E = \pi_t^E(\mu_t) + \delta V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)$$

where $\pi_t^E(\mu_t) = (1-\mu_t)x_t^E p_t^E$ denotes the entrant's profit in period t, and $V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)$ is the value of being the incumbent in the next period with $\mu_{t+1}^{NA} = \kappa(1-\mu_t)x_t^E$ captive consumers. The incumbent, which is due to exit in the next period, prices myopically. A forward-looking entrant, in contrast, must keep into account the impact of its current price on the number of captive consumers that it will inherit in the second period of its life-cycle, as this affects the profits that it will earn in its capacity as the new incumbent.

To begin with, however, suppose that the entrant prices myopically ($\delta = 0$). Given the behavior of the fringe, consider the entrant's best response to p_t^I. Free consumers choose to purchase from the entrant if $U_t^E > \max\left\{U_t^I, U_t^F\right\}$, that is, if $p_t^E < \min\left\{p_t^I + \Delta_t, \Delta_t + \Delta_{t-1}\right\}$. Therefore, the entrant's best response is

$$p_t^E(p_t^I) = \begin{cases} p_t^I + \Delta_t & \text{if } p_t^I \leq \Delta_{t-1}, \\ \\ \Delta_t + \Delta_{t-1} & \text{if } p_t^I > \Delta_{t-1}. \end{cases}$$

Next, consider the incumbent's strategy as a price leader. The incumbent makes no sales if $p_t^I > \Delta_{t-1}$. On the other hand, it anticipates that if it reduces the price below Δ_{t-1}, it would always be undercut by the entrant and would therefore serve only the captive consumers anyway. Therefore, the incumbent must price exactly at Δ_{t-1} (with a tiny discount to break the captive consumers' indifference, if necessary). By doing so, it gets a profit of $\pi_t^I = \mu_t \Delta_{t-1}$. In response, the entrant prices at $p_t^E = \Delta_t + \Delta_{t-1}$ (again with with a tiny discount if necessary) and will serve all free consumers.

If $\delta > 0$, so that the entrant is forward looking, it would have a further incentive to reduce the price to increase x_t^E if that were possible, as the continuation value V_{t+1}^I is increasing in μ_{t+1}. However, x_t^E is already equal to 1, so the myopic price remains optimal also for a forward-looking firm. ■

Proof of Corollary 1. With no acquisition, the firms' aggregate payoff is

$$(1-\mu_t)(\Delta_t + \Delta_{t-1}) + \mu_t \Delta_{t-1} + \delta V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t).$$

If the incumbent acquires the entrant, in contrast, the aggregate payoff becomes

$$(\Delta_t + \Delta_{t-1}) + \delta V_{t+1}^I(\mu_{t+1}^A, \Delta_t).$$

The lemma then immediately follows by comparing the above expressions, keeping in mind that V_{t+1}^I increases with the fraction of captive consumers μ_{t+1}, and that $\mu_{t+1}^A = \kappa \geq \mu_{t+1}^{NA} = \kappa(1-\mu_t)$. ■

Proof of Lemma 3. From the optimization problem (14) it appears that the equilibrium innovation size depends only on the derivative of $v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)$ with respect to Δ_t, which is the marginal profitability of increasing the size of the innovation. To calculate the derivative, let us substitute (11) into (12), obtaining

$$\begin{aligned} v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) &= (1-\alpha)\left[\pi_t^E + \delta V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t)\right] + \alpha[\pi_t^M + \delta V_t^I(\mu_{t+1}^A, \Delta_t) - \pi_t^I] \\ &= (1-\alpha)\pi_t^E + \alpha(\pi_t^M - \pi_t^I) + \delta\left[(1-\alpha)V_{t+1}^I(\mu_{t+1}^{NA}, \Delta_t) + \alpha V_t^I(\mu_{t+1}^A, \Delta_t)\right]. \end{aligned}$$

The on-path continuation value is

$$\begin{aligned} V_{t+1}^I(\mu_{t+1}^A, \Delta_t) &= v_{t+1}^I(\mu_{t+1}^A, \Delta_t, \Delta_{t+1}) \\ &= v_{t+1}^M(\mu_{t+1}^A, \Delta_t, \Delta_{t+1}) - v_{t+1}^E(\mu_{t+1}^A, \Delta_t, \Delta_{t+1}) \\ &= \pi_{t+1}^M + \delta V_{t+2}^I(\mu_{t+2}^{A,A}, \Delta_{t+1}) + \\ &\quad -(1-\alpha)\left[\pi_{t+1}^E + \delta V_{t+2}^I(\mu_{t+2}^{A,A}, \Delta_{t+1})] - \alpha[v_{t+1}^M(\mu_{t+1}^A, \Delta_t, \Delta_{t+1}) - \pi_{t+1}^I\right], \end{aligned}$$

where Δ_{t+1} is the innovation size in period $t+1$, which is correctly anticipated in period t, and $\mu^{A,A}_{t+2}$ is the fraction of captive consumers in period $t+2$ if the entrant is acquired both in period t and in period $t+1$. Likewise, the off-path continuation value is

$$\begin{aligned} V^I_{t+1}(\mu^{NA}_{t+1}, \Delta_t) &= v^I_{t+1}(\mu^{NA}_{t+1}, \Delta_t, \Delta_{t+1}) \\ &= v^M_{t+1}(\mu^{NA}_{t+1}, \Delta_t, \Delta_{t+1}) - v^E_{t+1}(\mu^{NA}_{t+1}, \Delta_t, \Delta_{t+1}) \\ &= \pi^M_{t+1} + \delta V^I_{t+2}(\mu^{NA,A}_{t+2}, \Delta_{t+1}) + \\ &\quad -(1-\alpha)\left[\pi^E_{t+1} + \delta V^I_{t+2}(\mu^{NA,A}_{t+2}, \Delta_{t+1})\right] + \alpha[v^M_{t+1}(\mu^A_{t+1}, \Delta_t, \Delta_{t+1}) - \pi^I_{t+1}\Big] . \end{aligned}$$

where $\mu^{NA,A}_{t+2}$ is the fraction of captive consumers in period $t+2$ if the entrant is not acquired in period t but is acquired in period $t+1$. (It follows from the one-shot deviation principle that this is indeed the relevant value of μ.)

Next, note that all current-period profit functions

$$\begin{aligned} \pi^M_t &= \Delta_t + \Delta_{t-1} \\ \pi^I_t(\mu_t) &= \mu_t \Delta_{t-1} \\ \pi^E_t(\mu_t) &= (1-\mu_t)(\Delta_t + \Delta_{t-1}) \end{aligned}$$

are additively separable in Δ_{t-1} and Δ_t, and that all other terms in the expression for $v^E_t(\mu_t, \Delta_{t-1}, \Delta_t)$ do not depend on Δ_{t-1}. This implies that $v^E_t(\mu_t, \Delta_{t-1}, \Delta_t)$ depends on Δ_{t-1} in an additively separable way and that, as a result, the optimal choice of Δ_t does not depend on Δ_{t-1}.

Since a similar argument applies to all subsequent periods, it follows that Δ_{t+1} does not depend on Δ_t, and the same is true of Δ_{t+2}, Δ_{t+3} etc. These future values depend only on μ_t. In particular

$$\begin{aligned} V^I_{t+1}(\mu^A_{t+1}, \Delta_t) &= (1-\alpha)\left[\pi^M_{t+1}(\mu^A_{t+1}) - \pi^E_{t+1}(\mu^A_{t+1})\right] + \alpha \pi^I_{t+1}(\mu^A_{t+1}) + \\ &\quad + \text{terms that depend only on } \mu_t \end{aligned}$$

and

$$V^I_{t+1}(\mu^{NA}_{t+1}, \Delta_t) = (1-\alpha)\left[\pi^M_{t+1}(\mu^{NA}_{t+1}) - \pi^E_{t+1}(\mu^{NA}_{t+1})\right] + \alpha\pi^I_{t+1}(\mu^{NA}_{t+1}) + $$
$$+ \text{ terms that depend only on } \mu_t$$

Thus, we have

$$\begin{aligned} v^E_t(\mu_t, \Delta_{t-1}, \Delta_t) &= (1-\alpha)\pi^E_t(\mu_t) + \alpha\left[\pi^M_t - \pi^I_t(\mu_t)\right]) + \\ &+\delta\left\{ (1-\alpha)\left[\pi^M_{t+1} - \pi^E_{t+1}(\mu^{NA}_{t+1})\right] + \alpha\pi^I_{t+1}(\mu^{NA}_{t+1}) + \right. \\ &\left. +\alpha\left[\pi^M_{t+1} - \pi^E_{t+1}(\mu^A_{t+1})\right] + \alpha\pi^I_{t+1}(\mu^A_{t+1}) \right\} \\ &+ \text{ terms that depend only on } \mu_t \end{aligned}$$

Collecting all terms that depend on Δ_{t-1} and Δ_t, we finally have

$$\begin{aligned} v^E_t(\mu_t, \Delta_{t-1}, \Delta_t) &= (1-\mu_t)\left[\alpha + (1-\alpha)\mu_t\right]\Delta_{t-1} + \\ &+ \left[(1+\delta\kappa)(1-\mu_t) + \alpha\mu_t(1+\delta\xi)\right]\Delta_t + \\ &+ \text{ terms that depend only on } \mu_t \end{aligned} \tag{54}$$

From this expression, (16) follows immediately, proving the first part of the lemma.

With the equilibrium innovation size at hand, we can now determine the equilibrium value function, and hence the acquisition price. To this end, we make a guess on the functional form of the value functions and find them by the method of undetermined coefficients. Since Δ^A_t is linear in μ_t and (54) shows that the expression for the value function $v^E_t(\mu_t, \Delta_{t-1}, \Delta_t)$ involves the product $\mu_t \times \Delta_t$, we conjecture that the *ex ante* value functions are polynomials of degree 2:

$$V^E_t(\mu_t, \Delta_{t-1}) = (1-\mu_t)\Delta_{t-1} + \phi_0 + \phi_1\mu_t + \phi_2\mu_t^2 \tag{55}$$
$$V^M_t(\mu_t, \Delta_{t-1}) = \Delta_{t-1} + \varphi_0 + \varphi_1\mu_t + \varphi_2\mu_t^2. \tag{56}$$

Given $V^E_t(\mu_t, \Delta_{t-1})$ and $V^M_t(\mu_t, \Delta_{t-1})$, we have $V^I_t(\mu_t, \Delta_{t-1}) = V^M_t(\mu_t, \Delta_{t-1}) -$

$V_t^E(\mu_t, \Delta_{t-1})$. We then identify the coefficients $\phi_0, \phi_1, \phi_2, \varphi_0, \varphi_1$ and φ_2 by imposing the condition that (55)-(56) must be identically satisfied. Here we report the solution for the special case $\xi = \kappa$ (the general solution can be found in a mathematical appendix available from the authors upon request):

$$\begin{aligned}
\phi_0 &= \frac{(2-\alpha)(1+\delta\kappa)^2 \{1+(1-\alpha)\delta[1-\alpha(1-\kappa)]\kappa\}}{[1+(1-\alpha)\delta][1-(1-\alpha)\delta\kappa][1+(1-\alpha)\delta\kappa^2]} \\
\phi_1 &= -\frac{(1-\alpha)(1+\delta\kappa)^2 \left[3-\alpha-(1-\alpha)\delta\kappa+(1-\alpha^2)\delta\kappa^2+(1-\alpha)^2\delta^2\kappa^3\right]}{[1-(1-\alpha)\delta\kappa][1+(1-\alpha)\delta\kappa^2]} \\
\phi_2 &= \frac{(1-\alpha)^2(1+\delta\kappa)^2}{[1+(1-\alpha)\delta\kappa^2]}
\end{aligned}$$

and

$$\begin{aligned}
\varphi_0 &= \frac{(1+\delta\kappa)^2 \{1+(1-\alpha)\delta[1-\alpha(1-\kappa)]\kappa\}}{[1+(1-\alpha)\delta][1-(1-\alpha)\delta\kappa][1+(1-\alpha)\delta\kappa^2]} \\
\varphi_1 &= -(1-\alpha)(1+\delta\kappa)^2 \\
\varphi_2 &= 0.
\end{aligned}$$

This proves the second part of the lemma.

It is simple to verify that $V_t^I(\mu_t, \Delta_{t-1})$ is increasing in μ_t – a property the was used repeatedly in the proof of Lemma 1 and 2. ■

Proof of Proposition 1. From (16) and (20) we have

$$\frac{d\Delta_t^A}{d\mu_t} = -(1+\delta\kappa) + \alpha(1+\delta\xi) \lesseqgtr 0 \Longleftrightarrow \alpha \lesseqgtr \frac{1+\delta\kappa}{1+\delta\xi}$$

and

$$\frac{d\Delta_t^{NA}}{d\mu_t} = -(1+\delta\kappa) < 0. \quad ■$$

Proof of Proposition 2. From (16) and (20) we have:

$$\Delta_t^{NA}(\bar{\mu}^{NA}) - \Delta_t^A(\bar{\mu}^A) = \kappa\frac{[1+\delta\kappa-\alpha\delta(1+\kappa)]\xi-\alpha(1+\kappa)}{(1+\kappa)(1+\kappa-\xi)},$$

whence the result follows immediately.■

Proof of Proposition 3. Using the steady state values (22) and (23) we get

$$\Delta_t^{NA}(\bar{\mu}^{NA}) - \Delta_t^{A}(\bar{\mu}^{A}) = \kappa \frac{\left[1 + \delta\kappa - \alpha\delta\left(1+\kappa\right)\right]\xi - \alpha\left(1+\kappa\right)}{\left(1+\kappa\right)\left(1+\kappa-\xi\right)},$$

whence the result follows immediately. ■

Proof of Lemma 5. Proceeding as in the proof of Lemma 3, it is easy to verify that if acquisitions are permitted in period t, the derivative of the acquisition price $P_t = v_t^E(\mu_t, \Delta_{t-1}, \Delta_t)$ with respect to Δ_t may depend only on whether acquisitions are prohibited or permitted in period $t+1$.

We already know that if acquisitions are permitted in period $t+1$, we have

$$\begin{aligned} v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) &= (1-\mu_t)\Delta_{t-1} + \\ &+ \left(1+\delta\kappa\right)\left[(1-\mu_t) + \alpha\mu_t\right]\Delta_t + \text{ other terms} \end{aligned}$$

where the other terms depend on expected values $\Delta_{t+1}, \Delta_{t+2}, \Delta_{t+3}$ etc., and thus depend only on μ_t.

If, on the other hand, acquisitions are prohibited in period $t+1$, the period-t acquisition price is

$$v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) = (1-\alpha)\left[\pi_t^E + \delta\pi_{t+1}^I(\mu_{t+1}^{NA})\right] + \alpha[\pi_t^M + \delta\pi_t^I(\mu_{t+1}^A) - \pi_t^I].$$

That is, the continuation value reduces to the profits that the merged firm will reap in its capacity as the next-period incumbent. Simple calculations yield

$$v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) = (1-\mu_t)\Delta_{t-1} + \left(1+\delta\kappa\right)\left[(1-\mu_t) + \alpha\mu_t\right]\Delta_t.$$

This immediately implies that if acquisitions are permitted in period t, $\Delta_t^A(\mu_t)$ does not depend on whether acquisitions are permitted or prohibited in the subsequent periods.

Next suppose that acquisitions are prohibited in period t. If acquisitions are prohibited also in period $t+1$, the entrant's profit is

$$\begin{aligned}
\Pi_t^{E,NA} &= \pi_t^E(\mu_t) + \delta \pi_{t+1}^I(\mu_{t+1}^{NA}) \\
&= (1-\mu_t)(\Delta_{t-1} + \Delta_t) + \delta\kappa(1-\mu_t)\Delta_t,
\end{aligned}$$

whence (20) follows. If, on the other hand, acquisitions are permitted in period $t+1$, the entrant's profit is

$$\begin{aligned}
\Pi_t^{E,NA} &= \pi_t^E(\mu_t) + \delta \left[v_{t+1}^M(\mu_{t+1}^{NA}, \Delta_t, \Delta_{t+1}) - v_{t+1}^E(\mu_{t+1}^A, \Delta_t, \Delta_{t+1}) \right] \\
&= \pi_t^E(\mu_t) + \delta \left\{ \pi_{t+1}^I(\mu_{t+1}^{NA}) + (1-\alpha) \left[\pi_{t+1}^M - \pi_{t+1}^I(\mu_{t+1}^{NA}) - \pi_{t+1}^E(\mu_{t+1}^{NA}) \right] \right\} + \\
&\quad + \text{other terms that depend only on } \mu_t.
\end{aligned}$$

Plugging the equilibrium profits into this formula, one obtains

$$\begin{aligned}
\Pi_t^{E,NA} &= (1-\mu_t)(\Delta_{t-1} + \Delta_t) + \delta \left[\alpha \mu_{t+1}^{NA} \Delta_t + (1-\alpha)\mu_{t+1}^{NA} (\Delta_t + \Delta_{t+1}) \right] \\
&= (1-\mu_t)\Delta_{t-1} + \left[(1-\mu_t) + \delta \mu_{t+1}^{NA} \right] \Delta_t + \text{other terms} \\
&= (1-\mu_t)\Delta_{t-1} + \\
&\quad +(1+\delta\kappa)(1-\mu_t)\Delta_t + \\
&\quad +\text{other terms that depend only on } \mu_t
\end{aligned}$$

so the optimal level of innovation is still given by (20). ■

Proof of Proposition 4. We verify whether the non state-contingent policy is optimal by applying the single-deviation principle. To begin with, suppose that condition (37) holds, so that always prohibiting acquisitions is the optimal non-contingent policy. Consider an alternative policy where acquisitions are permitted in period t and then are always prohibited from period $t+1$ onwards. With this policy, social welfare is:

$$W_t^{A,NA}(\mu_t) = \Delta_t^A(\mu_t) + \delta_S W_t^{NA}(\mu_{t+1}^A).$$

Using (16) and (36), this rewrites as:

$$\begin{aligned} W_t^{A,NA}(\mu_t) &= (1+\delta\kappa)(1-\mu_t) + \alpha(1+\xi\delta_s)\mu_t + \\ &\quad + \frac{\delta_S(1+\delta\kappa)\left[1-(1-\delta_S)\kappa(1-\mu_t)+\xi(1-\delta_S)\mu_t\right]}{(1-\delta_S)(1+\delta_S\kappa)}. \end{aligned}$$

Simple algebra shows that $W_t^{A,NA}(\mu_t) \leq W_t^{NA}(\mu_t)$ when (37) holds.

Next suppose that condition (37) fails, so that always permitting acquisitions is the optimal non-contingent policy. Consider an alternative policy where acquisitions are prohibited in period t and then are always permitted from period $t+1$ onwards. With this policy, social welfare is:

$$W_t^{NA,A}(\mu_t) = \Delta_t^{NA}(\mu_t) + \delta_S W_t^A(\mu_{t+1}^{NA}).$$

Using (20) and (35), this rewrites as:

$$W_t^{NA,A}(\mu_t) = \frac{1-\delta_S(\xi-\alpha\kappa)+\delta\kappa\left[1-(1-\alpha)\xi\delta_S\right]}{(1-\delta_S)(1+\delta_S\kappa)}\left[1-(1-\delta_S)\mu_t\right].$$

Simple algebra shows that $W_t^{NA,A}(\mu_t) \leq W_t^A(\mu_t)$ when (37) fails.

Thus, the optimal non-contingent policy always survives one-shot deviations. This confirms that such policy remains optimal even when state-contingent policies are feasible. ■

Proof of Proposition 5. Simple algebra shows that

$$W_t^{NA} - W_t^A = \frac{(1+\delta\kappa)\xi\delta_S - \alpha(1+\kappa\delta_S)(1+\delta\xi)}{(1-\delta_S)(1+\kappa\delta_S)\left[1-\delta_S(\xi-\kappa)\right]}\left[\delta_S(\kappa-\mu_t)+\mu_t\right],$$

whence the result follows immediately. ■

For online publication
Appendix B
Additional material for Section 6

Pricing equilibrium. Reversing the order of moves in the pricing game changes the pricing equilibrium as follows:

Lemma 2′: *If the entrant acts as price leader, the incumbent serves all captive consumers and the entrant all free consumers ($x_t^E = 1$). The incumbent prices at $p_t^I = \Delta_{t-1}$ and obtains a profit of*

$$\pi_t^I(\mu_t) = \mu_t \Delta_{t-1}. \tag{57}$$

The entrant prices at $p_t^E = \Delta_t + \mu_t \Delta_{t-1}$, so the profit it earns in the first period of its life cycle is

$$\pi_t^E(\mu_t) = (1 - \mu_t)(\Delta_t + \mu_t \Delta_{t-1}). \tag{58}$$

Proof. As in the proof of Lemma 2, let us begin by assuming that the entrant prices myopically ($\delta = 0$). Given the behavior of the fringe, consider the incumbent's best response to p_t^E. Captive consumers choose to purchase from the incumbent if $U_t^I > U_t^F$, that is if $p_t^I < \Delta_{t-1}$. Free consumers choose to purchase from the incumbent if $U_t^I > \max\{U_t^E, U_t^F\}$, that is, if $p_t^I < \min\{p_t^E - \Delta_t, \Delta_{t-1}\}$. Therefore, the incumbent must always undercut the fringe by setting a price no higher than $p_t^I = \Delta_{t-1}$, for otherwise it would make no sales. (Likewise, the entrant must always undercut the fringe to make positive sales, so we can restrict attention to the case $p_t^E \le \Delta_t + \Delta_{t-1}$ with no loss of generality.) If the incumbent prices exactly at Δ_{t-1} (with a tiny discount to break the captive consumers' indifference, if necessary), it gets a profit of $\pi_t^I = \mu_t \Delta_{t-1}$. If it further reduces the price to undercut not only the fringe but also the entrant, pricing at $p_t^I = p_t^E - \Delta_t$, it earns $\pi_t^I = p_t^E - \Delta_t$. Therefore, the incumbent's best response is

$$p_t^I(p_t^E) = \begin{cases} \Delta_{t-1} & \text{if } p_t^E \leq \Delta_t + \mu_t \Delta_{t-1}, \\ \\ p_t^E - \Delta_t & \text{if } p_t^E > \Delta_t + \mu_t \Delta_{t-1}. \end{cases}$$

Next, consider the entrant's strategy as a price leader. As noted, the entrant makes no sales if $p_t^E > \Delta_t + \Delta_{t-1}$. It also makes no sales if $p_t^E > \Delta_t + \mu_t \Delta_{t-1}$, as in this case the incumbent would undercut it. Therefore, the entrant's equilibrium price is $p_t^E(\mu_t) = \Delta_t + \mu_t \Delta_{t-1}$. At this price, the entrant serves all free consumers: $x_t^E = 1$.

Finally, suppose that $\delta > 0$ so that the entrant is forward looking. Since the continuation value V_{t+1}^I is increasing in μ_{t+1}, the entrant would have a further incentive to reduce the price to increase x_t^E if that were possible. But since x_t^E is already equal to 1, the myopic price remains optimal also for a forward-looking firm. ■

The incumbent's price and profit are the same as in the baseline model. Here, however, when the entrant acts as price leader it cannot just barely undercut the rivals in utility space by setting $p_t^E = \Delta_t + \Delta_{t-1}$. If the entrant priced this way, the incumbent would have the incentive to cut the price below Δ_{t-1} in order to capture the $(1 - \mu_t)$ free consumers. The entrant therefore has to cut its own price further, down to the point where the incumbent's incentive to compete for the free consumers is eliminated – a form of limit pricing.

Derivation of equation (41). Proceeding as in the proof of Lemma 3, but using the current-period profit functions

$$\begin{aligned} \pi_t^M &= \Delta_t + \Delta_{t-1} \\ \pi_t^I &= \mu_t \Delta_{t-1} \\ \pi_t^E &= (1 - \mu_t)(\Delta_t + \mu_t \Delta_{t-1}) \end{aligned}$$

one obtains

$$
\begin{aligned}
v_t^E(\mu_t, \Delta_{t-1}, \Delta_t) &= (1-\mu_t)\left[\alpha + (1-\alpha)\mu_t\right]\Delta_{t-1} + \\
&+ \{\ \alpha + \delta + (1-\alpha)(1-\mu_t) - (1-\alpha)\delta\left[1-\kappa(1-\mu_t)\right]\left[\alpha + \kappa(1-\alpha)(1-\mu_t)\right] + \\
&-\alpha\delta\left[1-\kappa(1-\mu_t) - \mu_t\xi\right]\left[\alpha + \kappa(1-\alpha)(1-\mu_t) - (1-\alpha)\mu_t\xi\right]\ \}\ \Delta_t \\
&+ \text{terms that depend only on } \mu_t
\end{aligned}
$$

From this expression, (40) follows immediately. The equilibrium value functions v_t^i and V_t^i can be calculated by the method of undetermined coefficients, as in the proof of Lemma 3. They are now polynomials of degree 4 in μ_t; the solution can be found in a mathematical appendix available from the authors upon request.

Proof of Lemma 6. First of all, consider a generic cycle of length ℓ; that is, suppose that there are $\ell - 1$ consecutive periods in which acquisitions are permitted, followed by one period in which they are prohibited. Focusing only on the beginning and the end of such a cycle, the dynamics of μ_t can be described by the following difference equation

$$
\mu_{t+\ell} = \kappa - \kappa^2 \sum_{n=0}^{\ell-2} (\xi - \kappa)^n - \kappa(\xi-\kappa)^{\ell-1}\mu_t.
$$

The characteristic root of this difference equation is, in absolute value,

$$
\left| -\kappa^{\frac{1}{\ell}}(\xi - \kappa) \right| < 1.
$$

This shows that all cycles of length ℓ converge to the limit cycle of length ℓ, irrespective of the initial condition.

Next consider a generic cut-off $\hat{\mu} \in \left[\bar{\mu}^{NA}, \bar{\mu}^{A}\right]$. We show that starting from an arbitrary μ_0, after at most one cycle, all subsequent cycles have length $\ell(\hat{\mu})$. With no loss of generality, assume that $\mu_0 < \hat{\mu}$. (If $\mu_0 \geq \hat{\mu}$, it suffices to apply the argument below starting from μ_1, which will then be lower than $\hat{\mu}$.) The value of μ_t for which acquisitions are first prohibited may then range from $\hat{\mu}$ to $\kappa(1-\hat{\mu}) + \xi\hat{\mu}$. Therefore, the value of μ_t at the beginning of the next acquisition cycle may range from $\kappa(1-\kappa) - \kappa(\xi-\kappa)\hat{\mu}$ to $\kappa(1-\hat{\mu})$. Tedious but simple algebra shows that for all starting points in this interval, μ_t will cross the threshold $\hat{\mu}$ in exactly $\ell(\hat{\mu})$ periods.

Combining the fact that under a cut-off policy rule with cut-off $\hat{\mu}$ all cycles will eventually have length $\ell(\hat{\mu})$, and that all cycles of that length converge to the limit cycle of length $\ell(\hat{\mu})$, it follows that the system converges to the limit cycle of length $\ell(\hat{\mu})$.■

Proof of Proposition 6. First of all, note that when $\xi = \kappa$ the critical value $\tilde{\xi}$ is implicitly defined by the condition

$$\mathcal{H}(\alpha, \delta, \xi) \equiv 1 - (1-\alpha)\xi + \delta\left[1 - \xi(1-\xi) - (1-\alpha)\xi^2(1-\xi)^2\right] - \frac{1+\delta\xi}{1+\xi} = 0$$

Simple algebra shows that

$$\begin{aligned} \left.\frac{\partial \mathcal{H}}{\partial \delta}\right|_{\xi=\tilde{\xi}} &< 0 \\ \left.\frac{\partial \mathcal{H}}{\partial \alpha}\right|_{\xi=\tilde{\xi}} &> 0, \end{aligned}$$

implying that $\tilde{\xi}$ is an increasing function of α and a decreasing function of δ.

Next, note that along a period-2 cycle, μ_t oscillates between $\mu_t = \kappa$, where acquisitions are prohibited, and $\mu_t = \kappa(1-\kappa)$, where acquisitions are permitted. By Lemma 5, the level of innovation then oscillates between $\Delta_t^{NA}(\kappa)$ and $\Delta_t^{A}(\kappa(1-\kappa))$. On average, the level of innovation is

$$\begin{aligned} \Delta^{SC} = {} & \frac{2+\delta(1-\alpha)}{2} - \frac{2-\alpha(1+2\delta)}{2}\kappa + \frac{1-\alpha+\delta\left[1-2\alpha(2-\alpha)\right]}{2}\kappa^2 \\ & - \frac{3-(7-4\alpha)\alpha\delta}{2}\kappa^3 + \delta(1-\alpha)^2\frac{3-\kappa+\kappa^2}{2}\kappa^4 \end{aligned}$$

Next, note that for $\delta = 0$ and $\kappa = \tilde{\kappa}(\alpha, 0) = \frac{\alpha}{1-\alpha}$, we have $\Delta^{NA}\left(\bar{\mu}^{NA}\right) = \Delta^{A}\left(\bar{\mu}^{A}\right) = \Delta^{SC}$. It may then be verified that if one increases δ and at the same time increases κ so that $\tilde{\kappa}(\alpha,\delta)$ remains constant at $\frac{\alpha}{1-\alpha}$, Δ^{SC} increases more rapidly than either $\Delta^{NA}\left(\bar{\mu}^{NA}\right)$ or $\Delta^{A}\left(\bar{\mu}^{A}\right)$ (which in fact, by construction, increase at the same rate). This implies that when $\delta = 0$, there exists a neighborhood of $\tilde{\kappa}(\alpha, 0) = \frac{\alpha}{1-\alpha}$ where the state-contingent policy is optimal. This however requires that $\frac{\alpha}{1-\alpha} < 1$, and hence that $\alpha < 1/2$. By continuity, state-contingent policy is optimal for some values of κ as long as δ and α are not too large. ■

Discussion of self-restraints. Suppose now that the merged firm rations demand so as to keep its degree of market dominance below the threshold $\hat{\mu}$. To this end, it must set an output level of $x^M = \frac{\hat{\mu}}{\kappa}$, earning a profit of $\pi^M = \frac{\hat{\mu}}{\kappa}(\Delta_t + \Delta_{t-1})$. Proceeding as in the proof of Lemma 3, one finds that the corresponding level of innovation is

$$\begin{aligned}\Delta^R &= (1-\alpha)\{1-\delta\kappa[1-\alpha(2-\kappa)]\} + \\ &\quad - \left[1-\alpha+\alpha\delta(1-2\alpha) - \frac{\alpha+\delta(1-\alpha)}{\kappa} - (1-\alpha)(1-2\alpha)\delta\kappa + 2(1-\alpha)^2\delta\kappa^2\right]\hat{\mu} \\ &\quad + (1-\alpha)\left[\alpha+(1-\alpha)\kappa^2\right]\delta\hat{\mu}^2.\end{aligned}$$

Tedious algebra shows that Δ^R is always lower than $\max\left[\Delta^{NA}\left(\bar{\mu}^{NA}\right), \Delta^{A}\left(\bar{\mu}^{A}\right)\right]$. This implies that a state contingent policy that induces the merged firm to ration its demand is never optimal.