Competition and Product Innovation in Dynamic Oligopoly*

Ronald L. Goettler
University of Rochester

Brett R. Gordon
Columbia University

April 17, 2013

Abstract

We investigate the relationship between competition and innovation using a dynamic oligopoly model that endogenizes both the long-run innovation rate and market structure. We use the model to examine how various determinants of competition, such as product substitutability, entry costs, and innovation spillovers, affect firms' equilibrium strategies for entry, exit, and investment in product quality. We find an inverted-U relationship between product substitutability and innovation: the returns to innovation initially rise for all firms but eventually, as the market approaches a winner-take-all environment, laggards have few profit scraps to fight over and give up pursuit of the leader, knowing he will defend his lead. The increasing portion of the inverted-U reflects changes in firm's investment policy functions, whereas the decreasing portion arises from the industry transiting to states with fewer firms and wider quality gaps. Allowing market structure to be endogenous yields different results compared to extant work that fixes the market structure.

Keywords: competition and innovation, R&D, dynamic oligopoly, inverted-U, spillovers.

JEL Classification: C73, D43, L11, L13, L40

*We appreciate helpful comments from Ron Borkovsky and Wes Hartmann. All remaining errors are our own. Goettler: University of Rochester, Simon School of Business, River Campus Box 270100, Rochester, NY 14627, ronald.goettler@simon.rochester.edu. Gordon: Columbia Business School, 3022 Broadway, New York, NY 10027, brgordon@columbia.edu.

1 Introduction

Innovation drives consumer welfare and firm profitability. Economists and marketers have therefore sought to understand what drives innovation, and in particular, whether competition fosters innovation. Despite this interest, the theoretical literature remains ambiguous, with some advocating a positive relationship (Arrow, 1962; Nickell, 1996), some a negative relationship (Schumpeter, 1942), and others an inverted-U relationship (Scherer, 1967; Aghion et al., 2005). The empirical literature, most recently reviewed by Gilbert (2006), finds varying support for each hypothesis, in part because the relationship depends on industry-specific factors such as product substitutability, entry costs, and the presence of innovation spillovers.

Although innovation incentives increasingly influence antitrust and merger policy (DOJ-FTC, 2010), the academic community has yet to provide policymakers with appropriate tools to evaluate the effect of policy on innovation. Extant research examines innovation incentives holding market structure fixed (e.g., Reinganum, 1983; Doraszelski, 2003; Ofek and Sarvary, 2003; Aghion et al., 2005), exogenously varying market structure (e.g., Spence, 1984; Goettler and Gordon, 2011), or in a static setting (e.g., Vives, 2008). None of these approaches provides what is needed: a dynamic model in which firms with market power endogenously determine market structure and the long-run (or steady-state) innovation rate. We provide such a model and use it to evaluate the effect on innovation of three measures of competition: entry costs, product substitutability, and innovation spillovers.

We choose these measures of competition instead of the number of firms because the endogeneity of market structure renders firm count a poor measure of competition. For example, despite the PC microprocessor industry's high concentration—Intel and AMD control about 95% of the market—competition is fierce as consumers focus on the vertical characteristic of computing speed (Reuters, 2011). Conversely, many brands compete in the apparel industry, but sufficient horizontal differentiation in consumers' tastes for clothing relaxes competitive intensity. We therefore model the endogenous determination of market structure in response to competitive forces, instead of directly changing the number of firms.

In this paper we make two contributions, one methodological and one substantive. First, we develop a dynamic oligopoly model, based on Pakes and McGuire (1994), that yields an endogenous long-run innovation rate that depends on consumer preferences, firms' costs and R&D processes, and the regulatory environment.[1] Our supply side is less restrictive than those used in previous studies, by allowing entry and exit and endogenous quality differences between leaders and laggards. Our demand side uses differentiated-products logit, a standard model used in policy analyses to assess market power and in managerial applications to evaluate firms' strategies. A consequence of such a rich specification is the need to use numerical methods. We therefore inspect policy functions and state transitions to explain our findings.

Second, our primary contribution is to provide comparative statics of the model's equilibrium to examine the effects of product substitutability, entry costs, and innovation spillovers, on market structure and inno-

[1] We define the long-run innovation rate as the average rate at which the industry's frontier quality improves, averaged across the ergodic set of states. In Pakes and McGuire (1994), the long-run innovation rate equals the exogenous rate at which the outside good's quality improves, as discussed in Appendix A.

vation. We show how competitive forces affect industry evolution through two avenues: directly on firms' state-dependent innovation strategies (i.e., investment policy functions) and through the evolution of market structure as firms implement their strategies. We also explore the differential impact of competitive forces on the policy functions of leaders and laggards.[2]

We find a pronounced inverted-U relationship between innovation and product-market competition (PMC), measured as product substitutability. When PMC is low, leaders and laggards have similar qualities because low substitutability implies quality differences have small effects on sales. As PMC increases, business-stealing incentives rise and investment policy functions for the leader and contending laggards shift up, which generates the upward portion of the inverted-U. As PMC continues to increase, the leader further defends its increasingly valuable lead with faster innovation, whereas the laggard's investment policy shifts down, in response to the declining probability of catching the leader. The resulting leader-laggard gap in innovation rates widens the quality gap. The leader therefore evaluates its policy function at states with less competition and accordingly lowers its investments, which generates the downward portion of the inverted-U. In short, the upward portion of the inverted-U results from changes in firms' policy functions whereas the downward portion results from changes in the distribution of realized market structures.

Our model differs from previous models of dynamic oligopoly by allowing for innovation spillovers.[3] The inclusion of spillovers is empirically motivated: non-frontier firms often benefit from innovations at the frontier through reverse engineering or public disclosure (Griliches, 1998; Bloom et al., forthcoming) and laggard firms engage in catch-up behavior resulting in battles for the frontier (Khanna, 1995; Lerner, 1997). We implement spillovers as higher investment efficiencies for laggard firms because advancing the frontier is naturally more difficult than closing the technological gap with the industry leader.[4]

Although significant empirical evidence exists in support of spillovers, their impact on innovation is unclear. Spillovers enable a laggard to innovate faster, but the equilibrium effect depends on the leader's response. The leader might innovate less because the benefits of innovation are shared or the leader might innovate more to escape competition with the laggards (Cohen and Levinthal, 1989). We find the relationship between spillovers and industry innovation depends on the degree of PMC. The industry innovation exhibits a moderate inverted-U relationship with spillovers when product-market competition is strong (i.e., products of similar quality are close substitutes). A higher spillover enables the laggard to close the gap with the leader and the winner-take-all nature of strong product-market competition induces the leader to innovate to defend its lead. However, for a sufficiently large spillover, the laggards innovates more rapidly as the gains

[2]To clarify our terminology, a firm is a leader if its product quality is greater than or equal to all other firms in the industry. Since multiple firms can exist at this industry frontier, there can be multiple lead firms. A laggard is any firm with a quality level strictly below the frontier product quality.

[3]Analyses of dynamic oligopoly with differentiated products include the study of collusive pricing (Fershtman and Pakes, 2000), competition among hospitals (Gowrisankaran and Town, 1997), state-dependent demand (Chan and Seetharaman, 2004), advertising dynamics (Dubé, Hitsch, and Manchanda, 2005), and oblivious equilibrium (Weintraub et al., 2008).

[4]Other work considers different forms of innovation spillovers. Levin and Reiss (1988) incorporate spillovers by allowing a firm's innovation outcome to depend on the investment levels of all firms in the industry. Similarly, Kamien, Muller, and Zang (1992) capture innovation spillovers in the form of research joint ventures. For our purposes, the key notion behind the innovation spillover is that a laggard firm's R&D is more efficient.

to attaining the lead are high. When product-market competition is weak, however, industry innovation eventually declines to zero after the leader gives up. The leader's gains to innovation are lowest in this case and the laggards quickly close the quality gap. Once the laggards are sufficiently close, the leader stops investing. This comparative static therefore demonstrates how the effect of one competitive force can depend on the level of others.

The current paper is distinct from our earlier work in Goettler and Gordon (2011) for three reasons. First, this paper considers endogenous market structure by allowing entry and exit whereas the earlier work only contrasts monopoly outcomes with duopoly outcomes. Second, this paper studies non-durable goods, whereas our earlier work considers durable goods such that a monopolist faces competition with its own sales in past and future periods. Third, we use different assumptions in the two papers to bound the state space. With durable goods, the no-purchase option pertains to using the version of the product bought in a previous period, which evolves endogenously as consumers make upgrade purchases. With non-durable goods, we must address the utility of a truly "outside good" as discussed in section 3.

Our work builds on an extensive literature, which we summarize in the next section. The two closest papers are Ofek and Sarvary (2003) and Aghion et al. (2005). The former uses a dynamic oligopoly model to study how firms' innovation efforts depend on market structure, profit incentives (i.e., PMC), and the leader's innovative advantage. By assuming an exogenous market structure and reduced-form profit functions, Ofek and Sarvary (2003) obtain analytical solutions and find that investment monotonically increases in PMC. We instead allow the number of firms and quality differences to be endogenous and find a pronounced inverted-U relating PMC and investment. Aghion et al. (2005) provides a duopoly model in which innovation and product-market competition have an inverted-U relationship. However, firms in their model invest in process innovations that reduce the marginal costs of production, and a lone lead firm never innovates. In contrast, firms in our model seek to improve quality, and a lead firm invests to fend off competition from encroaching laggards.

Finally, our focus on long-run innovation relates to empirical work measuring the long-term effects of marketing decisions and their impact on firm- and industry-level outcomes (Dutta et al., 1999; Bronnenburg et al., 2008). Firms' capabilities and resources are important sources of innovative and competitive advantage (Wernerfelt, 1984; Day, 1994). In our model, a firm's relative quality determines its current and expected future profit and its innovation capability. Competition for quality leadership generates continuous industry churn where entrants have the ability to supplant leaders and underperforming firms exit. We contribute to these literatures by investigating how changes in fundamental market conditions, such as product substitutability and entry costs, affect equilibrium strategies of leaders and laggards. Managers should be aware of how changes in their industry affect not only their own strategies, but also competitors' decisions and the long-run evolution of the industry.

2 Literature on Competition and Innovation

This section provides a historical perspective of theoretical work on the competition-innovation relationship. As this literature is immense, we direct the reader to surveys by Cohen and Levin (1989) and Gilbert (2006) for further details.

Schumpeter (1942) suggests that, relative to perfectly competitive firms, firms with market power would innovate more due to their ability to extract rents in the post-innovation market. In contrast, Arrow (1962) argues that firms in a competitive market have greater incentives than a monopolist to invest in a cost-reducing patent because the competitive firms, earning zero profits prior to innovating, have more to gain from an innovation.

Rather than comparing the extremes of monopoly and perfect competition, subsequent investigations of Schumpeter's claim consider patent-race games between a fixed number of firms. Loury (1979) and Lee and Wilde (1980) show the incentives to invest in R&D depend on the cost structure of the innovation process. Gilbert and Newbery (1982) demonstrate that an incumbent innovates to deter a lone entrant from innovating to win the race, leading to the increasing dominance of the incumbent. Reinganum (1983) shows that introducing uncertainty into the innovation process can reverse Gilbert and Newbery's (1982) finding, whereas Harris and Vickers (1987) shows that including both uncertainty and strategic interaction as the race unfolds re-establishes the increasing-dominance result.

Later work relaxes four substantive assumptions of the early patent race literature. First, consistent with empirical evidence in Cohen and Levin (1989) and Griliches (1998), researchers relax the assumption that innovation opportunities are exclusive and perfectly appropriable. Vives (2008) relaxes exclusivity by allowing firms to independently achieve cost reductions through R&D and finds the relationship between competition and innovation depends on the measure of competition: innovation increases with product substitutability but decreases with the number of firms. Spence (1984) relaxes perfect appropriability by allowing for innovation spillovers, such that investments by one firm can benefit others. He finds an inverted-U relationship between competition and innovation when spillovers are low and a monotonically decreasing relationship when spillovers are high.

Second, market structure and innovation are both endogenously determined, whereas much work assumes an exogenous market structure. Dasgupta and Stiglitz (1980) argue that both market structure and innovation incentives "depend on more basic ingredients," such as innovation technology, demand conditions, and the degree of appropriability. They endogenize market structure through free entry in an oligopoly model with firms investing in process innovations under Cournot competition, and show higher demand elasticity reduces the number of active firms and increases R&D investment. However, extending their static analysis to a multi-period model requires a winner-take-all assumption and leads to a persistent monopoly in equilibrium. Due to modeling challenges, few recent papers consider the simultaneous determination of market structure and innovation, especially in a dynamic setting.

Third, in models with pre-emption or increasing dominance, the R&D race is effectively decided once

a firm falls behind. A period of intense competition is followed by periods in which the winner innovates unencumbered. Empirical evidence suggests firms often engage in catch-up behavior leading to extended battles for the frontier (Khanna, 1995; Lerner, 1997). Doraszelski (2003) therefore allows a firm's cumulative stock of R&D expenditures to increase the probability of successful innovation, which relaxes the assumption in earlier studies that past R&D investments have no effect on current R&D decisions and outcomes. The more realistic investment process makes the model analytically intractable. Simulations of the equilibrium strategies reveal action-reaction behavior: a laggard with a sufficiently large knowledge stock strives to catch up to the leader.

Fourth, most models in this literature focus on process innovations that reduce firms' marginal costs of production. Examining firms' product innovation strategies is more complicated because differentiated-product markets require more complex models of firms and consumers. Greenstein and Ramey (1998) and Chen and Schwartz (2010) both reassess Arrow's (1962) results under product innovation, and their findings demonstrate that such conclusions depend on the nature of product differentiation (horizontal vs. vertical), the magnitude of innovations ("drastic" vs. "nondrastic"), and whether barriers to entry exist.

We focus on product-quality innovations because they are critical for marketing strategies related to product differentiation and customer segmentation. As in modern empirical analyses, we derive aggregate demand from a consumer utility model. This approach permits diverse market structures with varying numbers of firms and relative qualities and yields different results than models using reduced-form profit functions for leaders and laggards. Firms climb a quality ladder and investment spillovers enable laggards to contend for the lead, generating action-reaction dynamics.[5] Firms' entry, exit, and innovation policies generate an endogenously evolving market structure.

In summary, this review demonstrates the need for a model that reflects important characteristics of real-world R&D competition. Analyzing a model closer to those used in empirical research should help others map our general results into more specific industry settings.

3 Model

We model an infinite-horizon dynamic oligopoly in a differentiated-products market in which each firm sells a non-durable good and invests to improve its future quality. Firms make entry, exit, price, and investment decisions to maximize their expected discounted profits. Investment is a dynamic control because investment sacrifices current profits for potentially higher future profits. Exogenous improvements in the outside good generate industry-wide depreciation shocks.

The sequence of events is as follows:

1. Incumbent firms compete in the product market. Nash equilibrium prices are determined independently

[5] Work in endogenous growth theory also considers quality-ladder models (e.g., Grossman and Helpman, 1991; Aghion and Howitt, 1992; Aghion et al., 2001) to study topics ranging from competition in international trade to optimal policies for national economic growth.

of the dynamic aspects of the industry because prices only affect current profits.

2. Each incumbent observes its scrap value and implements its investment and exit strategies, while the potential entrant simultaneously observes its entry cost and implements its entry decision.

3. Incumbents' innovation outcomes and the industry-wide depreciation shock are realized.

In the subsections that follow, we describe the consumers' decision problem, the firms' decision problem, the equilibrium, and contrast our model with Pakes and McGuire (1994).

3.1 Consumers

Consumers solve a static discrete-choice utility maximization in each period t. Define $\nu_j \in (\ldots, -1, 0, 1, \ldots)$ as the index of the quality of good j in the current period (t subscripts omitted). Consumer i receives utility from good $j = 0, 1, \ldots, J$ according to Cobb-Douglas preferences:

$$\tilde{u}_{ij} = \gamma \nu_j + \log(y - p_j) + \sigma_\varepsilon \varepsilon_{ij} \ , \tag{1}$$

where γ is the marginal utility of quality, y is income, p_j is the price, and ε_{ij} is an idiosyncratic preference shock.

The scale parameter σ_ε, typically normalized to unity, controls the degree of horizontal differentiation in tastes and facilitates the comparative static in product-market competition. The outside good, indexed by $j = 0$, has price $p_0 = 0$.

To express utility in terms of relative quality, we subtract $\gamma\nu_0$ from each option:

$$u_{ij} = \gamma \omega_j + \log(y - p_j) + \sigma_\varepsilon \varepsilon_{ij} \ , \tag{2}$$

where $\omega_j = \nu_j - \nu_0$ is quality relative to the outside good.

The specification above allows for vertical differentiation via ω and horizontal differentiation enters through ε_{ij}. The model can be extended to include heterogeneity in consumer preferences over product quality γ or variation in consumer incomes y. Using a demand model as rich as those in empirical studies such as Berry et al. (1995), however, would be computationally challenging because they require expanding the state space to accommodate multiple product characteristics.

3.2 Firms

In each period, incumbent firms choose whether to exit or to invest in product quality, and a potential entrant chooses whether to enter. We incorporate random scrap values to ensure existence of symmetric equilibria in pure strategies, as suggested in Doraszelski and Satterthwaite (2010).

States and Transitions A firm's state is the pair (ω, s) where $\omega \in \Omega \subseteq \mathbb{Z}$ is the firm's own quality and

$s \in S \subseteq \mathbb{Z}_+^{\mathbb{Z}}$ is a vector with an element s_i that denotes the number of firms at quality level $i \in \mathbb{Z}$. The market contains $J(s) = \sum_i s_i$ active firms. We define $\bar{\omega} = \max(\omega_1, \ldots, \omega_J)$ as the quality of the frontier firm(s) relative to the outside good.

Firms choose a level $x_j \in \mathbb{R}_+$ to invest in R&D, and the industry's collective investment policies determine the laws of motion for (ω, s). The outcome of a particular firm's investment decision $\tau_j \in \{0, 1\}$ is probabilistic and stochastically increasing in x_j. Firms implement the quality improvements that become available with a successful innovation in the next period.[6] We specify the probability of a successful investment as

$$f(\tau_j = 1 | x_j, \omega_j, s) = \frac{a(\omega_j, s) x_j}{1 + a(\omega_j, s) x_j} \ , \tag{3}$$

where, to allow for investment spillovers, the investment efficiency $a(\omega, s)$ is specified as

$$a(\omega_j, s) = a_0 (1 + a_1 \mathcal{I}(\omega_j < \bar{\omega})) \ , \tag{4}$$

where $\mathcal{I}$ is an indicator function. Our choice for $f(\tau | x, \omega, s)$ yields closed-form solutions for optimal investment and satisfies the unique investment choice (UIC) criterion in Doraszelski and Satterthwaite (2010).[7]

The specification of spillovers in $a(\omega_j, s)$ has the benefit of being parsimonious and capturing the notion that spillovers from the leader are the most important. If firms have the ability to reverse engineer another firms' technology, they would presumably choose to reverse engineer the most advanced technology. We refer to $a_1 > 0$ as the level of spillovers.

Let $\tau_0 = 1$ indicate an improvement in the quality of the outside good ν_0 which reduces firms' relative qualities. This "depreciation" shock occurs with exogenous probability $\delta \in [0, 1]$. If firms innovate more rapidly than δ, the space of firms' relative qualities Ω becomes unbounded. We therefore restrict relative qualities: $\omega_j \leq \bar{\omega}_{max}$ for all j. Thus, a firm's relative product quality in the next period is

$$\omega_j' = \omega_j + \tau_j - I_\omega \tag{5}$$

where $I_\omega = \max\left(\tau_0, \{\tau_j \cdot \mathcal{I}(\omega_j = \bar{\omega}_{max})\}_j^J\right)$. The indicator I_ω takes a value of one if either (a) an industry-wide depreciation shock occurs ($\tau_0 = 1$) or (b) if a firm at $\bar{\omega}_{max}$ has a successful innovation. In either event, a firm's next-period relative quality, ω_j', shifts down by one unit. Choosing $\bar{\omega}_{max}$ to be sufficiently large ensures the bound restriction does not affect equilibrium firm behavior.[8]

Static Profits Incumbent firms set prices and compete for demand from consumers of mass M, denoted in millions. Assuming the idiosyncratic shocks are distributed independently and identically type I extreme

[6] Borkovsky (2012) studies a more realistic innovation process in a dynamic quality-ladder model where firms time the release of new innovations and can stockpile successful innovations.

[7] Doraszelski and Satterthwaite (2010) define an investment transition function, such as $f(\tau | x, \omega, s)$, as being unique investment choice (UIC) admissible if the function leads to a unique investment choice for the firm.

[8] We choose the bound such that the outside good's share is less than .001, so that firms' profits would not improve much if the outside good were even further behind.

value, the market share for firm j is

$$\psi_j(\omega, s; p) = \frac{\exp\{(\gamma\omega_j - \log(y - p_j))/\sigma_\varepsilon\}}{1 + \sum_{k\in\{1,\ldots,J(s)\}} \exp\{(\gamma\omega_k - \log(y - p_k))/\sigma_\varepsilon)\}} \ . \tag{6}$$

A unique Nash equilibrium exists in prices where firm j receives static period profits of $\pi_j(\omega, s) = M\psi_j(\omega, s; p)(p_j - mc)$, where p is the vector of $J(s)$ prices and mc is the constant marginal cost of production across firms (Caplin and Nalebuff, 1991). Since prices only affect static profits, we solve for equilibrium static profit, $\pi_j(\omega, s)$, and solve the dynamic game in which firms choose investment, entry, and exit given $\pi_j(\omega, s)$.

Incumbents At the start of each period, each incumbent privately observes its random scrap value ϕ_j drawn from a uniform distribution $F(\cdot)$ with support $\left[\phi^L, \phi^H\right]$.[9] Incumbent firms then decide whether to stay in the industry and invest in product quality or to exit after collecting static profits and the scrap value. An incumbent exits whenever its randomly drawn scrap value exceeds its continuation value of staying in the industry, which enables us to express the exit policy prior to observing the scrap value as a probability.

Let $e[\omega]$ be a vector with one in the ω spot and zero elsewhere, and $\hat{s} = s - e[\omega]$ be the configuration of competitors' qualities. The incumbent makes decisions given beliefs $q^\omega(\hat{s}'|s)$ over the number and qualities of competitors in the next period. Prior to observing ϕ_j, the exit decision rule is the probability $\xi_j \in [0, 1]$ that the firm exits in the industry. An incumbent maximizes its expected discounted profits, yielding the Bellman equation

$$V(\omega, s) = \pi_j(\omega, s) + \max_{\xi_j\in[0,1],\ x_j\geq 0} \left\{\xi_j \mathbb{E}\left[\phi_j | \phi_j \geq F^{-1}(1 - \xi_j)\right] + (1 - \xi_j)\left\{-x_j + \beta W_j(\omega, s, x_j)\right\}\right\}, \tag{7}$$

where x_j is denoted in units of the market size (millions of consumers). The expectation in the second term is the expected scrap value the firm collects conditional on exiting. The continuation value, $W_j(\omega, s, x_j)$, integrates over the firm's own investment outcome τ_j, the outside good's outcome τ_0, and other firms' qualities $\hat{s}'$,

$$W_j(\omega, s, x_j) = \sum_{\tau_j, \tau_0, \hat{s}'} V_j(\omega_j + \tau_j - I_\omega, \hat{s}' + e[\omega_j + \tau_j - I_\omega]) f(\tau_j | x_j, \omega, s) q^\omega(\hat{s}'|s)\, \rho(\tau_0) \tag{8}$$

The first-order condition for the exit policy function $\xi_j^*(\omega, s)$ is:

$$-F^{-1}(\xi_j^*(\omega, s)) + \{-x_j + \beta W_j(\omega, s, x_j)\} = 0 \ . \tag{9}$$

Differentiating equation (7) with respect to investment and rearranging terms yields the first-order condition

[9] The distribution of scrap values could be a function of ω.

for the investment policy function $x_j^*(\omega, s)$:

$$x_j^*(\omega, s) - \left(\frac{a_j(\omega, s)}{1 - (\beta a_j(\omega, s)(EW_j^+(\omega, s) - EW_j^-(\omega, s)))^{-1/2}} - a_j(\omega, s) \right)^{-1} = 0 \ . \tag{10}$$

$EW_j^+(\omega, s)$ and $EW_j^-(\omega, s)$ are the firm's expected continuation values conditional on positive and negative innovation outcomes, respectively. We obtain EW_j^+ and EW_j^- by integrating over competitors' innovation outcomes, the potential entrant's action, and the exogenous depreciation shock due to δ.

Entrants At the start of each period, the potential entrant observes its random entry cost x_e drawn from a uniform distribution $F_e(\cdot)$ with support $x_e \in \left[x_e^L, x_e^H\right]$. If x_e is less than the expected discounted net cash flows $\beta W_e(s)$, the entrant pays x_e and begins operation in the next period at quality level $\omega_e - I_\omega$. Rather than fixing the value of ω_e, as in Pakes and McGuire (1994), we let the relative quality of an entrant be drawn from a discrete distribution $G_{\omega_e}(\cdot)$ with support $\{0, 1, \ldots, \bar{\omega}_{max}\}$. This specification allows entrants to potentially leap to the frontier of the industry.[10]

Prior to observing x_e, the entrant's decision rule can be represented as the probability $\xi_e \in [0, 1]$ that the firm enters the industry currently in state s. The optimal entry policy maximizes the expected discounted continuation payoff net of expected entry costs conditional on entry, yielding the Bellman equation:

$$V_e(s) = \max_{\xi_e \in [0,1]} \xi_e \left\{ -\mathbb{E}\left[x_e | x_e \leq F_e^{-1}(\xi_e)\right] + \beta W_e(s) \right\} \ , \tag{11}$$

where the expectation term is the expected entry cost paid conditional on the entrant choosing to enter. The continuation value for the entrant $W_e(s)$ integrates over other firms' qualities $\hat{s}'$, the outside good's improvement τ_0, and the distribution of entry qualities $\omega_e \sim G_{\omega_e}(\cdot)$,

$$W_e(s) = \sum_{\hat{s}', \tau_0, \omega_e} V\left(\omega_e - I_\omega, \hat{s}' + e\left[\omega_e - I_\omega\right]\right) q^\omega\left(\hat{s}' | s\right) \rho(\tau_0) G_{\omega_e}(\omega_e) \ . \tag{12}$$

Taking the derivative of the entrant's maximization problem yields the first-order condition for the entry policy function $\xi_e^*(s)$ as

$$-F_e^{-1}(\xi_e^*(s)) + \beta W_e(s) = 0 \ . \tag{13}$$

3.3 Equilibrium

We focus on pure-strategy symmetric Markov-Perfect Nash Equilibria (MPNE), which Doraszelski and Satterthwaite (2010) establish are guaranteed to exist. Firms choose their optimal policies based on consistent expectations about the distribution of competitors' actions, leading them to accurately estimate the transition kernel $q^\omega\left(\hat{s}' | s\right)$ for competitors' states. Firms' beliefs are rational in the sense that $q^\omega\left(\hat{s}' | s\right)$ can be

[10] Iskhakov, Rust, and Schjerning (2013) examines leapfrogging behavior in the context of a dynamic duopoly model with cost-reducing investments.

derived from the transition probabilities resulting from the equilibrium policies of other firms.

Besanko et al. (2010) and Borkovsky et al. (2012) document multiple equilibria in dynamic oligopoly models based on Ericson and Pakes (1995). Borkovsky et al. (2012) find that, although multiple equilibria exist, they are similar in their simulated outcomes. The differences mostly relate to policies at particular states that are rarely, if ever, visited.

We refine the set of equilibria to only consider the limit of the finitely repeated game. We use backwards induction to solve for equilibrium in a T-period game in which terminal values are zero and then let $T \to \infty$. As such, we solve for equilibrium strategies within each state at each iteration, as opposed to merely playing best responses from the previous iteration. For each T, we solve the system of $2J(s)+1$ first-order conditions corresponding to the $J(s)$ firms' investment and exit decisions (equations 10 and 9) and the potential entrant's decision (equation 13), which are sufficient for equilibrium at industry state s.[11] This refinement yields a unique equilibrium if the sub-game within each state at each iteration has a unique equilibrium (i.e., if the model exhibits state-wise uniqueness as defined in Besanko et al., 2010). We have numerically verified state-wise uniqueness at various states and iterations of our solution algorithm.

3.4 Comparison to Pakes & McGuire (1994)

Pakes and McGuire (1994, hereafter PM) numerically solve the differentiated-products version of the Ericson and Pakes (1995) framework in which firms invest to move up a quality ladder. As detailed in Appendix A, the long-run industry innovation rate in PM is exogenous because of their approach to bounding the state space. Firms in PM stop innovating when their product quality is sufficiently ahead of the outside good, regardless of competitors' qualities. The long-run innovation rate is therefore the outside good's exogenous rate of improvement. This property limits the use of PM for policymakers and managerial guidance because neck-and-neck firms in the real world frequently strive for industry leadership, even when the outside good is weak competition.

Conceptually, our model restricts the degree to which the outside good can be inferior to frontier products, rather than distorting innovation incentives to firms at the technology frontier, as in PM. Because frontier firms generate most of the industry's sales, profits, and surplus, assumptions regarding the outside good are more innocuous than assumptions limiting firms' benefits from innovation. Moreover, we allow the outside good to be sufficiently inferior that the restriction has no appreciable effect on equilibrium policies.

In Appendix A, we compare industry outcomes in our model with those in PM for various δ. As δ increases, prices and industry profits increase in PM. In our model, the increased competition from a more rapidly improving outside good lowers prices and profits, as one would expect. More importantly, industry innovation in our model depends on consumers preferences, firms technologies, and the regulatory environment and is largely independent of δ. Our model is better suited for evaluating candidate policies than are models with exogenous industry innovation because policy usually aims to enhance welfare and innovation

[11] We include a complementary slackness condition due to the non-negativity constraint on investment.

Table 1: Parameter values used in comparative statics

	x_e^L	x_e^H	ϕ^L	ϕ^H	PMC	a_1	a_0	M	mc	β	γ	y	δ	ω_e
Baseline	22	24	21	22	.78	$\{.2, .6\}$	.5	5	5	.925	.5	15	0	4
PMC					.3–1.9									
Spillover						0–6								
Entry	16–40		15	16										

drives welfare.

4 Comparative Statics for Competition and Innovation

We repeatedly solve and simulate the model to obtain comparative statics in the degree of product-market competition (PMC, as measured by $1/\sigma_\varepsilon$), the degree of innovation spillovers, and in entry costs. To evaluate the comparative statics, we plot industry outcomes averaged over simulations starting with one firm at $\omega_1 = 7$ and a second at $\omega_2 = \omega_e = 4$. We simulate each industry for 100 periods (100 years) and 10,000 repetitions, such that the standard errors on the reported averages are negligible. The primary outcome of interest is the industry's innovation rate, calculated as the share of periods in which the frontier advances.[12]

The first row of Table 1 summarizes the parameters in our baseline specification and subsequent rows indicate deviations from this baseline to generate each comparative static. We choose the range of the scrap value distribution to be just below the lowest entry cost we consider. In the baseline model, the static profits in the initial period are 29.5 and 8.5, respectively, for the leader and laggard. Entry costs between 22 and 24 are therefore about three-fourths of the leader's initial period profit and almost triple the laggard's initial profit.[13] Setting y to 15 and marginal cost to 5 yields a maximum markup of 200%. The discount factor $\beta = .925$ such that we interpret each period as one year.[14] In our baseline we specify $G_{\omega_e}(\cdot)$ as a point mass on $\omega_e = 4$ and for robustness we consider a positive probability of entering at the frontier.

Our choices of parameters in the baseline and which parameters to vary for the comparative statics reflect two considerations. First, we desire interesting equilibria with entry, exit, and innovation.[15] Second, we focus

[12]The averages we report in the simulations are not necessarily from the recurrent class of states. The innovation rate we report corresponds to the steady-state innovation rate if the initial state falls in the recurrent class, otherwise the measure we report is simply the average over the first 100 years.

[13]To relate these costs to a particular industry, consider that Intel's 2012 4th-quarter net income was \$2.5 billion. Since cost estimates of semiconductor fabrication plants are typically around \$3 billion, our range of entry costs seems at least plausible. For estimates of fabrication plants costs, see `http://finance.yahoo.com/news/Construction-Of-Chip-twst-2711924876.html`, `http://topics.nytimes.com/top/news/business/companies/taiwan-semiconductor-manufacturing-company-ltd/index.html`, and `http://www.optessa.com/industries_semi.htm`, all accessed on 2.8.13.

[14]Many of our model's parameters could be reasonably chosen, via calibration or formal estimation, using industry data. The R&D efficiency and spillover parameters, a_0 and a_1, could be estimated using data on R&D expenditures and product innovation outcomes. Quality preferences, γ, could be estimated using standard demand data or calibrated based on data from similar industries. Income can be chosen based on Census information or knowledge of the relevant consumer demographics. Under log utility, the maximum amount any firm can charge is y.

[15]For some parameterizations, we can find equilibria with absorbing states in which no firms invest. These equilibria typically either have a single leader at the frontier with multiple laggards below the entry threshold or all firms at the frontier. The laggards serve to deter potential entrants, and the leader prefers not to invest because it could induce the laggards to exit, prompting entry by new firms at a higher quality level. These parameterizations seem unrealistic and so we avoid them, in part by having scrap values that induce severe laggards to exit.

on parameters that managers or policymakers can influence. For example, policymakers can affect entry costs through increased regulatory requirements or by enacting trade barriers. Given the benefit to incumbents of higher entry costs, firms may petition policymakers to take such actions, such as the steel industry's successful bid to raise import tariffs (The Economist, 2002) or longstanding import tariffs and quotas in the sugar industry (The New Yorker, 2006). Regarding spillovers, firms often form R&D alliances. In the semiconductor industry, SEMATECH is a research joint venture that enables firms to share developments in chip technology and fabrication (Hof, 2011). Similar research consortia exist in the pharmaceutical industry (Fortune, 2007) and in the automobile industry (The New Yorker, 2007). Although alliances are endogenous, these choices occur at an earlier stage than the price and quality competition we consider.[16]

To focus on competition among inside firms, we assume the innovation rate of the outside good is zero ($\delta = 0$). With $\delta = 0$, the share of the outside good quickly goes to zero in all simulations. Since we are focusing on long-run industry outcomes, we ignore the outside good by assuming it does not exist (or equivalently has an extremely low utility). We retain $\nu_0 = 0$ for the purposes of defining relative qualities. With $\delta = 0$, the frontier is always at the highest of the 7 rungs on the quality ladder. In section 4.1 we provide an analysis with $\delta > 0$ to assess the robustness of our results based on $\delta = 0$.

We focus on two levels of spillovers, $a_1 = .2$ and $a_1 = .6$. Allowing for spillovers strengthens the action-reaction behavior noted in Lerner (1997). The spillover comparative static in section 4.2 considers an expanded range of $a_1 \in [0, 6]$.

Before presenting our comparative statics, we explore the equilibrium behavior of the model via the firms' policy functions. Fundamentally, equilibrium investment levels depend on the gains in static profits due to innovation. The static profit gains illustrate the immediate effect of innovation on profits, while the long-run benefits depend on the market evolution. Figure 1 plots the change in static profits at different points in the state space when a leader or laggard successfully innovates. The middle column provides these results in the baseline PMC, and the left and right columns depict the profit gains when PMC is low and high, respectively. The upper panels show that as the leader faces increased competition from the laggard, the static profit gains to innovation follow an inverted-U shape. The level of the leader's gains and their sensitivity to the laggard's quality both increase as PMC increases. The laggard's profit gains monotonically increase as it approaches the leader and are higher for higher PMC only when its quality is near the frontier.

To facilitate plotting the firms' policy functions, we restrict attention to states with two incumbents. Figure 2 plots the policy functions of the leader and laggard under different values for the spillover ($a_1 = \{.2, .6\}$), product-market competition ($1/\sigma_\varepsilon = \{.78, 1.56\}$), and entry costs ($x_e^L = \{18, 22\}$ and $x_e^H = x_e^L + 2$).[17] Entry costs increase from the top row to the middle row, resulting in a downward shift in each firm's innovation rate because the firms are more insulated from potential entrants, with the effect being of similar

[16] See Song (2011) for a dynamic oligopoly model of research joint ventures.

[17] PMC = .78 corresponds to $\frac{1}{\pi/\sqrt{6}}$, the inverse of the standard deviation of a logit error term. A PMC of 1.56 is therefore equivalent to doubling the coefficients on all terms in the utility function while maintaining the standard variance of the logit error.

magnitude for both levels of the spillover parameter. The effect of the spillover parameter is most easily seen by comparing the two plots in the middle row. The higher spillover (of the right column) increases the laggard's innovation when it is not too far behind the leader. The higher spillover, however, slightly decreases the innovation rate of the leader because the laggard now finds it easier to catch up.

Increasing the degree of PMC raises consumers' sensitivity to quality differences. Leaders therefore increase innovation to defend their quality advantage when that advantage is threatened. This behavior is evident by comparing leader innovation across the bottom two rows of Figure 2. The laggard's investment declines to zero when it falls four steps behind the leader, after which it will coast with no investment until it exits. This zero investment is due to both the low gains to static profits and the low probability of catching the leader given the high incentives for the leader to defend its advantage when PMC is high.

Differences in simulated outcomes reflect these changes in policy functions, as well as changes in which states are encountered during simulation of the equilibrium. In discussing the comparative statics, we decompose the results according to these two sources.

4.1 Varying the Degree of Product-Market Competition

In our model, product-market competition (PMC) is determined by the variance of consumers' idiosyncratic utility shocks. This form of horizontal product-differentiation can be viewed as capturing exogenous elements of differentiation such as variation in consumers' beliefs concerning unknown product quality.[18] We leave to future research the case of endogenous horizontal differentiation when consumers have heterogeneous preferences over observable characteristics chosen by firms. Figure 3 plots the industry outcomes as we vary PMC (i.e., $1/\sigma_\epsilon$), and observation 1 summarizes our main findings.

Observation 1 *As PMC increases,*

i. Industry innovation exhibits a pronounced inverted-U.

ii. Higher spillovers extend the positive relationship between PMC and innovation to higher levels of PMC (i.e., shifts the peak of the inverted-U to the right).

The inverted-U shape is pronounced for all spillover parameter values and is robust to other parameter changes and a wider range of values for $1/\sigma_\epsilon$.[19] Much empirical evidence supports an inverted-U relationship.

[18] In some markets, PMC evolves over time. When idiosyncratic shocks are perception errors about unknown product quality, PMC intensifies as consumers learn firms' qualities. Consider the search engine market, born in the 1990's with the sequential entry of Excite, Yahoo!, WebCrawler, Lycos, Infoseek, Altavista, Inktomi, AskJeeves, Google, MSN, Overture, and Alltheweb. Consumer uncertainty was initially high regarding search engine quality because most people lacked experience in the domain and evaluating the quality of a given query was difficult. Over time, the search engines refined their algorithms and consumers gained general experience with web-based search technologies. Four years after entering, Google led the U.S. search query market with a 29.2% market share and now maintains its market dominance with a 65.6% share, according to comScore. Google's profits have soared as its competitors struggle (PCWorld 2010, Forbes 2011).

[19] For values of $1/\sigma_\varepsilon < 0.3$, the industry continues to fill up with an increasing number of firms (as illustrated on the left-hand side of Figure 3.d) and industry innovation continues to drop for reasons explained in this section. For values of $1/\sigma_\varepsilon > 1.9$, entry continues to fall (as illustrated on the right-hand side of Figure 3.d) and the industry moves closer to a persistent monopoly with industry innovation continuing to decline.

Building on Scherer (1967), Levin, Cohen and Mowery (1985) combine FTC business unit data with an extensive cross-industry survey, and find results consistent with an inverted-U even after controlling for technological opportunity and appropriability. Aghion et al. (2005) provide empirical support of an inverted-U between markups (i.e., PMC) and citation-weighted patent counts (i.e., innovation) using a panel of firms in the United Kingdom. Using detailed data following policy reforms in the European Union and UK, Aghion et al. (2009) offer empirical evidence that suggests an inverted-U shaped relationship between PMC and innovation.

In explaining the inverted-U, we focus on the spillover of .2 case. As PMC increases, consumers' choices are driven more by quality differences than by idiosyncratic shocks, which moves the market towards a winner-take-all environment. A direct effect of higher PMC is larger static-profit gains for innovation by leaders or laggards near the leader, as evidenced by the upward shifts as PMC increases across the columns in Figure 1. Because the laggard is indeed near the leader when PMC is low (Figure 3.f), both the leader and laggard increase innovation as PMC initially rises, which generates the left half of the inverted-U.

Both firms' profits increase (Figure 3.i) as the higher PMC enables them to steal share from the many firms with inferior products. The number of firms declines rapidly from an average of more than six to less than three (Figure 3.d), at which point further increases in PMC begin to decrease Firm 2's profits since it loses more share to the leader than it steals from the few remaining inferior firms.

The widening gap between leader and laggard profits induces the leader to vigorously defend his lead, as indicated by his policy function in Figure 4.a. For PMC below 1.0 the leader invests more when tied for the lead at $\omega_1 = 7$ than when the laggard is at ω_2 of 5 or 6, but for PMC above 1.0, the long-run profit advantage is sufficiently high that the leader invests heavily to preserve its marginal lead. The laggard responds to this increased defense by reducing its R&D efforts (Figure 4.b): its investments become flat when at ω_2 of 6 and decline sharply when at lower qualities. As a consequence, the laggard falls further behind (Figure 3.f) which enables the leader to safely decrease innovation without jeopardizing its lucrative position at the frontier. These responses to yet higher PMC generate the right-half of the inverted-U. Increasing the spillover to .6 helps the laggard "stay in the game" longer for higher PMC, but eventually the same story plays out yielding the pronounced inverted-U.

When PMC is low, laggards compete more to steal share from other laggards than with the goal of becoming the leader. That is, competing for profit scraps provides incentives for laggards to innovate. In contrast, with high PMC profit scraps are minimal and innovation by laggards is primarily an effort to achieve industry leadership. With low PMC, Firm 2 is the most successful of the laggards competing for the many profit scraps. But when PMC is high, few scraps are left by the leader, thereby reducing incentives for laggards to innovate.

This comparative static shows that increasing PMC lowers innovation incentives through two avenues: laggards have fewer profit scraps to fight over and they know the leader will respond to any attempt to close the quality gap with increasing vigor. Next we discuss the relationship between our findings and the

literature and evaluate the robustness of the results.

To evaluate the importance of endogenous market structure, we present in Figure 5 the PMC comparative static under different exogenous market structures. The top rows of Figures 3 and 5 are the same except that in the latter each line presents outcomes assuming a fixed market structure with two, three, or four firms. In contrast to the inverted-U when the market structure is endogenous, PMC and innovation exhibit a decreasing relationship when the number of firms is held fixed, regardless of the number. The higher the PMC, the sooner the laggards give up their pursuit. Moreover, without entry, the leader ceases innovating once its lead is sufficiently large.[20] This behavior is evident when comparing Figures 5.d to 5.e, where higher PMC (panel (d)) leads the laggard to decrease investment more rapidly as it falls behind the leader.

An inverted-U with exogenous market structure does arise, however, under the assumption that laggard firms never lag the leader by more than a few quality steps (e.g., 2, 3, or 4 steps). Such an assumption, which can be viewed as a stronger investment spillover, ensures competition from laggard firms except when PMC is high, thereby yielding an inverted-U. In essence, competition from a laggard that is never more than, say, three steps behind the leader replaces competition from the entrant who enters three steps behind the frontier when entry is allowed.

Although Ofek and Sarvary (2003) do not discuss product-market competition per se, increases in PMC would correspond to a widening gap between leader and follower period profits ($\pi_l - \pi_f$ in their notation). They show (in equation A5) that firms' investments monotonically increase in $\pi_l - \pi_f$, whereas we find investment exhibits an inverted-U shape. One reason for the different result is that the market structure in Ofek and Sarvary (2003) is fixed as $\pi_l - \pi_f$ varies. In our model, fewer firms enter as PMC increases, which reduces competition for the scraps, and, consequently, laggards' innovation incentives. Another reason is that in Ofek and Sarvary (2003) all laggards are identical but in our model can differ. With multiple laggard types, the leader endogenously responds to higher PMC by innovating fast enough that laggards fall further behind and reduce the intensity of their pursuit.[21] Modifying our model to match the assumption in Ofek and Sarvary (2003) that firms are either tied or differ by one quality unit, we obtain the positive relationship between innovation and competition. However, once firms can differ by more than one quality unit, the relationship is negative.

As a robustness check, we also consider different assumptions regarding the outside good's innovation rate. Figure 6 plots the industry's innovation rate under the PMC comparative static with $\delta = \{0, 0.4, 0.6, 0.8\}$.

[20] The bottom row of Figure 5 plots the policy functions to show the absorbing states when the laggard is sufficiently far behind the leader (who is always at 30 for these plots). Figure 5.d and 5.e depicts policy functions for duopolies with high and low PMC, respectively. Figure 5.f considers a triopoly to show that the presence of the third firm (at $\omega = 14$) only has a small effect on the other two firms' innovation policies, by comparing panel (e) to (f). Note that the leader's higher innovation for $\omega_2 < 14$ results from the third firm being at $\omega_3 = 14$. When both laggards in Figure 5.f are tied at $\omega_2 = \omega_3 = 14$ or lower, neither invests. Even with low PMC, the profit scraps are insufficient motivation to invest when they are far behind the leader.

[21] To draw a comparison with the results in Aghion et al. (2005), Goettler and Gordon (2011) conduct a comparative static in PMC in a nondurable version of their model in which the laggard is at most one step behind the leader. Figure 12 of Goettler and Gordon (2011) shows that industry innovation *increases* in PMC and they claim the same result would hold even if the maximum quality gap between the firms is widened. That claim is correct, but only for moderate increases in the maximum quality gap. If the maximum gap is sufficiently wide that a laggard eventually gives up (as in the current paper), then innovation is decreasing in PMC because the point of giving up is reached more quickly the higher is PMC.

Introducing the outside good has two effects: it creates more competition in the product space and it increases the quality of entrants' products. We plot outcomes with an outside good that never innovates (i.e., $\delta = 0$) to isolate the former effect. The lines corresponding to no outside good are the same as those in Figure 3.a. The figure demonstrates that the inverted-U persists when $\delta \leq .6$. With $\delta = 0.8$, the outside good innovates so rapidly that firms innovate in response to competition from the outside good, not from each other. Thus, our main findings hold as long as the outside good does not advance too quickly. The comparative static results in the other sections are also unchanged for moderate values of δ.

Appendix B presents additional robustness checks that vary the investment efficiency (a_0), the innovation spillover parameter (a_1), the functional form of innovation spillovers $a(\omega_j, s)$, the density of entrant's quality (G_{ω_e}), and consumers' utility u_{ij}. In general we find the inverted-U relationship is robust to changes in all of these parameters except for the entrant's quality. Specifically, we consider the family of G_{ω_e} with two mass points, $\omega_e = 4$ and $\omega_e = \bar{\omega}$, and let κ_e denote the probability of entering at $\bar{\omega}$. If $\kappa_e = 0.1$, the leader is continually challenged and industry innovation declines little as PMC increases. Raising this probability further has little impact. In Aghion et al. (1995), the downward portion of the inverted-U is due to the increased frequency of encountering states in which the leader does not invest. In our model, the declining leader investment is due to the increased frequency of states with little competition from laggards or entrants. Even a small probability of entrants leaping to the frontier reduces the frequency of these states and raises industry innovation.

4.2 Varying the Degree of Investment Spillovers

R&D spillovers enable firms to innovate more easily when catching up to the leader than when advancing the technology frontier as the leader. Spillovers arise in the real world because laggards often have access to technologies from the leader through reverse engineering, personnel turnover, and public disclosure. Even patented innovations do not fully protect a firm's ability to appropriate its discoveries (Levin et al., 1987). A review by Griliches (1998) concludes that broad empirical support for spillovers exists and finds significant cross-industry variation in the magnitude of spillovers. More recently, using a large firm-level panel data set, Bloom et al. (forthcoming) provide evidence of innovation spillovers in the computer hardware, pharmaceuticals, and telecommunications equipment industries.

However, the theoretical impact of spillovers on the relationship between innovation and competition relationship is ambiguous. The direct effect of higher spillovers is that laggards innovate faster because their investments are more efficient. The equilibrium effect depends on the leader's response to the laggards' faster innovation. Spillovers could reduce the leader's incentive to innovate because firms share the innovation benefits. Conversely, spillovers could increase a leader's innovation in an effort to stay ahead of the laggards (Cohen & Levinthal, 1989).

Whereas the previous comparative static reports results for two spillover levels, we now more fully explore the impact of spillovers on equilibrium outcomes. We vary the spillover parameter a_1 from 0 to 6. Figure 7

plots the comparative static in the spillover parameter for PMC values of .31, .78, and 1.56 to show how the incentives to innovate stemming from price competition interact with those from spillovers. Observation 2 summarizes the main results.

Observation 2 *As the innovation spillover increases,*

i. For low PMC, industry innovation declines and eventually drops close to zero once the leader gives up.

ii. For moderate PMC, innovation by the leader declines and industry innovation has a slight inverted-U.

iii. For high PMC, industry innovation exhibits a moderate inverted-U.

Increasing the spillover lowers laggards' innovation costs which leads them to innovate faster (Figure 7.c). The leader's response to its shrinking advantage depends on PMC. The leader's innovation (Figure 7.b) declines with the spillover level when PMC is .31 or .78 and increases when PMC is 1.56. These responses are consistent with the effect of increasing laggard quality on the leader's static profit gains from innovation: the static gains decline slightly with the low and baseline PMC (Figure 1.a and 1.b, respectively) and increase sharply with the high PMC (Figure 1.c). For any PMC level, the laggard's static profit gains from innovation, in Figure 1, increase as its quality approaches the frontier.

As a_1 approaches infinity, the leader's incentive to innovate converges to a single period of higher profits because, with probability one, laggards eliminate the leader's advantage from that innovation in one period. This single period of higher profits is insufficient when PMC is low.

The low PMC case illustrates a useful boundary condition: for high enough spillovers the leader ceases to innovate and industry innovation drops to zero. This behavior specifically arises in the low PMC case because the leader's static profit gains from innovation are the lowest (Figure 1.a). As the spillover rises from zero, the laggards close their quality gap with the leader (Figure 7.f). When the quality gap becomes sufficiently small, the benefits of remaining in the lead no longer outweigh the costs of fighting off the encroaching laggards, and so the leader stops investing.

As noted in the PMC comparative static, high PMC yields a winner-take-all market, giving the leader such a strong incentive to maintain its lead that the laggards are deterred from mounting much of an attack. When PMC is high and the spillover is low, the leader commands over 85 percent of the market (Figure 7.h) and maintains a large quality lead (Figure 7.f). However, once the spillover exceeds about 1.5, the laggards can inexpensively innovate and close the quality gap with the leader (Figure 7.f).

Regardless of PMC levels, a negative relationship between innovation and spillover levels reflects diminishing incentives for leaders to innovate when laggards can more easily catch up. This force dominates for all spillover levels when PMC is low, yielding the monotonically negative relationship. When PMC is high, this force dominates only when the spillover is high, yielding the inverted-U.

Appendix B reports a series of robustness checks on the relationship between spillovers and industry innovation. Our results are consistent across many alternative parameter values and specifications. For low

PMC, the leader stops innovating for different values of the innovation efficiency (a_0) and the leap probability (κ_e). For moderate and high ranges of PMC, the results are similar except when the entrant can leap to directly compete with the leader. In these cases, setting $\kappa_e > 0$ raises industry innovation at low spillover values.

This relationship between spillovers and innovation contrasts with the pointed and pronounced inverted-U relationship in Figure 10 of Goettler and Gordon (2011). Decreasing the spillover from its estimated value, firms fight to remain competitive knowing that if one firm falls behind, it will cede the industry to the leader and reside forever at the lowest quality level. As the spillover decreases, the likelihood rises that a laggard will cede the industry, spurring the leader to innovate more rapidly to avoid becoming the laggard. At a spillover just below 40 percent of its estimated value, pursuit is too costly for the laggard who quickly cedes the market, allowing the leader to cut innovation as its lead grows. The laggard persists at the lower bound of product quality because Goettler and Gordon (2011) does not allow the firm to exit and no firms can enter to take up the fight with the leader. In this paper we do not find a strong inverted-U because the severe laggard exits and is replaced by an entrant closer to the leader. The leader therefore must continue to innovate to defend its position.

4.3 Varying Entry Costs

The number of firms in an industry depends on firms' strategies, given market conditions. One such condition is the entry costs associated with commencing operation. From a strategic perspective, entry costs serve as barriers that help protect incumbent firms from competition from newcomers (Han et al., 2001). Entry costs can be exogenous, such as building a factory or a distribution channel. Sutton (1991) argues that some entry costs are endogenous, examples of which include overcoming incumbents' brand equity (Aaker, 1991) or consumer switching costs (Karakaya and Stahl, 1989; Klemperer, 1995).

We perform a comparative static that varies the lower bound of the entry cost distribution x_e^L from 40 to 16, while setting $x_e^H = x_e^L + 2$ and fixing the scrap value distribution to be Uniform(15,16). The net entry cost, the difference between the realized entry cost and the expected discounted scrap value, is close to zero when $x_e^L = 16$. Entry costs above 40 often yield a monopoly.

Figure 8 plots several simulated industry outcomes, with the result on industry innovation in panel (a). Observation 3 below summarizes our main findings.

Observation 3 *As entry costs fall,*

i. Industry innovation exhibits a slight U-shape.

ii. The laggard firm increases innovation to deter entry, until entry costs become too low for entry to be deterred.

We find that lowering entry costs, and consequently increasing the number of firms, yields a slight U-shape in industry innovation. The mildly declining leader and industry innovation rates are driven by the

laggard's response to the increased threat of entry. Lower entry costs have no effect until $x_e^L < 34$, after which the laggard's innovation rate rises to deter entry, as shown in Figure 8.c. The laggard's investment policy and the entrant's policy establish that this rise in laggard innovation is indeed a response to the threat of entry: the laggard at $\omega_2 = 4$ (the dashed dotted line in Figure 9.b) increases innovation as x_e^L falls below 34, which is when the probability of entry becomes nonzero for a potential entrant facing a laggard at $\omega_2 = 3$ (upper portion of Figure 9.c). That is, the laggard increases innovation to avoid dropping to $\omega_2 = 3$ and possibly triggering entry. The laggard engages in entry-deterring behavior, as opposed to the leader, because the laggard is competing for the profit scraps against the potential entrant. A consequence of the laggard's entry deterrence is that the leader advantage (Figure 8.f) declines from an average of 2.4 to 1.5 as x_e^L falls from 34 to 20. In Figure 1.b, the static profit gains to innovation decline as the leader's moderate advantage shrinks, thereby reducing the leader's innovation.

However, the innovation rate of the industry and lead firm rises when the entry cost moves from 20 to 16. Since laggards can no longer deter entry, as seen in the upper portion of Figure 9.c for $x_e^L < 20$, their innovation rates drop and they fall further behind the leader. In addition, the increased number of firms causes the laggards to compete for fewer profit scraps, further reducing their innovation incentives. The leader's growing advantage increases its static gains from innovation (Figure 1.e), thereby prompting it to innovate faster.

Note that as entry costs decline, the industry innovation is slightly higher than the leader's innovation because the industry advances when any of the firms at the frontier innovates.[22] The frequency that firms are tied at the frontier increases (Figure 8.e) until $x_e^L = 20$. The industry benefits from having more firms at the frontier to push industry-wide innovation forward despite the frontier firm's declining innovation rate.

This comparative static highlights the complex relationship between competition and innovation, and demonstrates how changes in laggards' investment choices influence the leader's investments. One reason the effect is not pronounced is that lowering entry costs has no direct effect on the static profit gains to innovation or on the innovation process. This relative lack of responsiveness of industry innovation to entry costs has empirical support. For example, using detailed plant-level data from Mexico on the manufacturing industry, Teshima (2010) shows that product R&D is unresponsive to changes in import tariffs.[23]

5 Conclusion

We investigate the relationship between competition and innovation in a dynamic oligopoly with endogenous innovation and R&D spillovers. We use the model to examine how different competitive forces, such as product substitutability, entry costs, and spillovers, shape an industry's evolution. In our comparative

[22] The industry innovation rate at an arbitrary state s equals one minus the probability that all frontier firms fail to innovate: $1 - (1 - f(\tau = 1 | \bar{x}(s)))^{\sum_{j=1}^{J(s)} I(\omega_j = \bar{\omega})}$ where $\bar{x}(s)$ is the investment by each firm at the frontier and the exponent gives the number of firms at the frontier.

[23] Teshima (2010) does, however, find that process R&D (to lower costs) increases as import tariffs fall.

statics, we disentangle the channels through which changes in these competitive forces operate, in terms of the effects on firms' policy functions or the industry states visited during simulation. Decomposing the results in this way is critical to separate firms' equilibrium responses from variation in industry structure. For example, we show that greater product substitutability raises the leader's innovation policy function but also reduces the qualities and number of competing firms.

To conduct our analysis, we provide an alternative approach to bounding the state space in Pakes and McGuire (1994), the quality-ladder model from the framework of Ericson and Pakes (1995). PM's approach to bounding the state space restricts the behavior of frontier firms such that the exogenous innovation rate of the outside good solely determines the long-run industry innovation rate. We relax this restriction to yield a model which endogenizes both market structure and innovation, allowing us to compare the innovation rate across different market conditions.

Our results demonstrate that the relationship between competition and innovation depends on a variety of factors that are often ignored to obtain closed-form solutions. As noted by Soberman and Gatignon (2005), "researchers have often oversimplified the relationships between competitive dynamics and market evolution... This seems to be an important cause of contradictory or insignificant results." Our findings highlight that the relationship between competition and innovation hinges on the nature of competition. We summarize our key results below:

- Varying the degree of PMC results in an inverted-U relationship with industry innovation. As PMC initially increases, firms innovate more to steal share. Eventually, higher PMC moves the industry toward a winner-take-all market in which the leader defends its advantage sufficiently to deter challengers. If entrants can leap to the frontier, however, the right-hand side of the inverted-U is never realized because the leader is continually challenged.

- If PMC is strong, industry innovation exhibits a mild inverted-U in R&D spillovers. Spillovers allow the laggard to compete more easily with the leader, eventually allowing the laggard to close the quality gap and reducing the leader's innovation incentives. With weak PMC, however, industry innovation declines with increasing spillovers and eventually ceases.

- Lowering entry costs results in a slight U-shaped relationship with industry innovation. Entry deterring investment by laggard firms spurs the lead firms to reduce investments until entry costs are low enough that laggard firms can no longer deter entry.

These comparative statics fill an important gap in the literature. The previous theory literature, reviewed in Section 2, offers conflicting results, and the empirical literature finds support for increasing, decreasing, and inverted-U relationships between competition and innovation. One reason for the conflicting empirical results, as Dasgupta and Stiglitz (1980) highlight, is that innovation and market structure are simultaneously determined, leading cross-sectional analyses to incorrectly infer causal relationships. We provide a model that endogenizes both innovation and market structure and use it to illustrate causal relationships between

industry innovation and competitive forces, such as product substitutability, innovation spillovers, and entry costs.

A Comparison to Pakes & McGuire (1994)

The key difference between our model and PM is the way product quality enters a consumer's utility function. PM sets utility for consumer i from good $j \in (1, \ldots, J)$ as

$$u_{ij} = g(\nu_j - \nu_0) - p_j + \varepsilon_{ij} = g(\omega_j) - p_j + \sigma_\varepsilon \varepsilon_{ij} \tag{14}$$

where $g(\cdot)$ is an increasing and concave function of relative quality. We set σ_ε to its standard value of $\sqrt{6}/\pi$. PM specifies that

$$g(\omega) = \begin{cases} \lambda + \gamma\omega, & \text{if } \omega \leq \omega^* \\ \lambda + \gamma\omega^* + \ln(2 - \exp(\gamma(\omega^* - \omega))), & \text{otherwise,} \end{cases} \tag{15}$$

where λ is a shift parameter that determines the competitiveness of the outside good and $\gamma > 0$ rescales the quality ladder. Hence, $g(\cdot)$ is linear for low values of relative quality and concave for high ones, and utility is a weakly concave function of relative quality.

To facilitate direct comparison with PM, we modify our model to have constant marginal utility for money such that income drops out, yielding the utility function:

$$u_{ij} = \lambda + \gamma\omega_j - p_j + \varepsilon_{ij} \ . \tag{16}$$

The outside good's utility in both PM and our model is $u_{i0} = \varepsilon_{i0}$.

The PM discrete-choice model uses a non-standard normalization: rather than subtracting the mean utility of the outside good from the utility of each choice, PM subtracts the quality of the outside good from the quality of each choice and converts relative quality to utils using a concave function. Accordingly, the dynamic game in PM cannot be derived from a model in which consumers have preferences directly over absolute quality. In our model, $g(\cdot)$ is linear (and hence omitted), and we use the standard normalization to convert from absolute qualities to relative qualities (see section 3.1).

The linearity of our quality index in utility does allow for concave preferences over quality itself. For example, the absolute index ν could be in units of log-quality, in which case innovations represent proportional, not additive, increments in quality.

Consequences of concave $g(\cdot)$ The concave $g(\omega)$ in PM has two important implications for the industry's equilibrium behavior: an exogenous long-run innovation rate and distorted utility rankings of inside products. We discuss each of these implications and then compare industry outcomes in the two models using a comparative static in the outside good's rate of innovation.

Any numerical solution to the dynamic quality-ladder game requires relative qualities be bounded. Both PM and our model ensure a lower bound by providing a scrappage value such that firms exit when their relative quality gets sufficiently low. We provide an upper bound by assuming the outside good improves when a firm with relative quality $\bar{\omega}_{max}$ improves its absolute quality. By choosing $\bar{\omega}_{max}$ sufficiently high, this assumption has no effect on firms' equilibrium policies.

PM creates an upper bound by specifying consumer preferences such that the benefit of higher relative quality ω_j quickly goes to zero once ω_j exceeds some threshold ω^*, regardless of competitors' qualities. The most significant consequence of this bounding approach is that δ, the exogenous innovation rate of the outside good, solely determines the industry's long-run innovation rate. The upper panel of Figure 10 plots $g(\omega)$ as specified in PM. Consumer utility is linear for $\omega \leq \omega^*$ and nearly flat for $\omega > \omega^*$. This kink in utility implies that consumers effectively place no value on product quality improvements above ω^*.[24] In principle, one could choose ω^* high enough that firms never reach it over some finite horizon of interest. When interested in innovation over any medium or long run, however, any such ω^* would be too high for computational feasibility.

To assess the implications of a concave $g(\omega)$, first consider the case with an outside good of fixed quality ($\delta = 0$). A monopolist would stop investing shortly after surpassing ω^* as consumers barely value the improvements and investment is costly. Competing firms, even if neck-and-neck, would also stop investing shortly after surpassing ω^* since the concave $g(\omega)$ compresses utility differences between firms above ω^*. Innovation in the long run would be zero since all firms would stop investing shortly after reaching ω^*.

When $\delta > 0$, improvements in the outside good bring firms back below ω^*, thereby restoring the profit incentive to innovate. Long-run innovation in PM is therefore determined solely by the rate of improvement in the outside good.

As δ increases, firms invest at higher values of ω because the concave investment technology function, $f(\tau|x, \omega, s)$, implies a given long-run innovation rate is attained at lower cost when investments are spread out across periods. For $\omega > \omega^*$, the incentive to innovate is driven by the desire to smooth investments across periods, rather than by the desire to reap higher current profits. That is, rather than a sudden increase in investment when ω falls below ω^*, firms will invest at ω slightly higher than ω^* despite the absence of immediate profit gains.

These consequences of concave $g(\omega)$ may have little impact on the equilibrium if firms rarely reach $\omega > \omega^*$. The lower panel of Figure 10, however, shows that firms tend to be above ω^*, unless δ is very high. For δ ranging from 0 to .9, we plot the percentage of simulated firms whose ω exceeds ω^*. This percentage exceeds .9 for $\delta < .45$ and only falls below .5 for $\delta > .85$. Simulated industry outcomes, such as innovation, profits, consumer surplus, and social surplus, are therefore heavily influenced by the concave $g(\omega)$. Conceptually, a model in which the outside good is the primary driver of industry behavior fails to endogenously determine

[24] PM set $\lambda = -4$, $\gamma = 3$, and report $\omega^* = 12$. However, the GAUSS code for their model (`http://www.economics.harvard.edu/faculty/pakes/program`), uses a value of 12 for the point of concavity *after* scaling by 3 and shifting by -4. The corresponding value for ω^* on the ω grid $(0, 1, 2, \ldots)$ is 16/3.

Table 2: Parameter values used in comparison with PM

x_e^L	x_e^H	ϕ^L	ϕ^H	M	mc	β	λ	γ	a_0	a_1	δ	ω_e	$\bar{\omega}_{max}$
.5	1	.1	.5	5	5	.925	-4	3	.5	0	$\{0, .1, \ldots, .9\}$	4	14

the main outcomes of interest.

The second implication of the concave $g(\cdot)$ is consumers' utility rankings of inside goods depend on the outside good's quality. For example, a consumer indifferent between two products that differ in quality will strictly prefer the higher quality product if the outside alternative improves. Consequently, the relative market shares depend on the outside good's absolute quality: as the outside good improves, holding fixed the inside goods' qualities, the relative market share of the higher-quality product increases. Standard discrete choice models do not have this property.

Comparing Outcomes To compare the equilibrium implications of the PM model and the linear model, we perform a comparative static that varies the exogenous rate of innovation in the outside good δ from 0 to .9. For each δ, we simulate the industry for 100 periods and average the results over 10,000 simulations. We parameterize the models to be similar to the baseline specification in PM, as summarized in Table 2. We bound the outside good's quality to be within 14 steps of the frontier. This lower bound on the outside good is large enough to ensure the outside good's market share is tiny when a firm is at $\bar{\omega}_{max}$. We do not allow for innovation spillovers when comparing our model to PM ($a_0 = 0$). In each simulation, the maximum number of firms is seven, and the simulation starts with one firm at the entry point and another firm two steps ahead.

Figure 11 plots the results for both models. In panel (a), industry innovation (i.e., the average rate of improvement in the frontier product) is exactly δ in PM, as expected. The innovation rate in the linear model exceeds δ and is relatively insensitive to its value because equilibrium innovation is primarily determined by competition between firms, not competition with the outside good. The difference between the two models declines as δ rises: when $\delta = 0$, innovation is zero for PM and about .9 in the linear model, and when $\delta = .9$, both models have innovation rates around .9. Similarly, the gaps across models in markups and the leader's share shrink as δ increases, but do not entirely disappear.

Panel (b) shows the average number of firms active in the industry. In PM, the industry reaches the maximum number of allowed firms when $\delta < .4$. Increasing the number of firms allowed indeed increases the number of firms active, but has little effect on other outcomes. The number of firms in PM is sensitive to δ, decreasing to about 2.5 when $\delta = .8$. In contrast, varying δ has essentially no effect on the number of firms in the linear model.

Panels (c) and (d) show that as δ increases, markups and the leader's share increase in PM and decline in the linear model. In both models, firms face greater competition from the outside good as it innovates faster. But in PM, a sharp reduction in the number of firms offsets the increase in competition from the outside

good, and the net change is less competition resulting in higher markups. This reduction in the number of firms is also responsible for the increase in industry profits in PM, reported in panel (e), as δ increases beyond .4.[25] Panel (f) plots consumer surplus, which in our model is consistently more than twice as large as that found in PM because of the higher innovation rates in our model and despite the higher markups.[26]

In summary, by relaxing restrictions on the innovation incentives of lead firms, our model produces strikingly different outcomes from those in PM.

B Robustness Checks

We compute a large number of comparative statics to explore the robustness of our findings. Specifically, we vary the innovation efficiency (a_0), the innovation spillover (a_1), the probability distribution over entrant's quality (G_{ω_e}), and the degree of product-market competition (PMC). For G_{ω_e}, we consider a distribution with two mass points, such that the entrant enters at the frontier $\omega_e = \bar{\omega}$ with probability κ_e and enters at quality $\omega_e < \bar{\omega}$ with probability $(1 - \kappa_e)$. We refer to κ_e as the probability that the entrant "leaps" to the frontier. In unreported results, we found that varying $\omega_e \in \{3, 4, 5\}$ has little effect on all of the outcomes.

We also consider the effect of changing certain parametric assumptions in our model. First, we implement a different form for innovation spillovers. The baseline form, which appears in equation (4), increases laggard's investment efficiency by a_1 independent of the degree to which the laggard is behind the lead firm. We modify this equation such that a laggard's investment efficiency is a linear function of the quality gap: $\tilde{a}(\omega_j, s) = a_0(1 + a_1(\bar{\omega} - \omega_j))$. Second, we alter the form of the consumer's utility function. Instead of the quasi-linear form in equation (2), we use a linear version defined as $\tilde{u}_{ij} = \gamma\omega_j - p_j + \sigma_\epsilon \varepsilon_{ij}$.

In total, we calculated the equilibrium in 32,832 dynamic oligopoly models. We did not investigate the robustness of our results on entry costs because of the relatively little variation we found between entry costs and industry innovation. For the results in section 4, we ensured the simulated industry outcomes never reached the maximum number of allowed firms by increasing the allowed maximum as needed. In this set of extended comparative statics, due to the number of equilibria, we restrict the number of firms to not exceed nine. When the industry reaches this maximum, relaxing this constraint tends to increase industry innovation by 0.01 to 0.03. We indicate such points in these results with dots.

Product-Market Competition Figure 12 presents a collection of comparative statics between PMC and innovation. The solid line plots the industry innovation rate and the dotted line plots the share of periods in which monopoly arises. In the baseline model PMC= 0.78 and the comparative static in section 4.1 varies PMC from 0.3 to 1.9. In all of these robustness checks we vary PMC from 0.3 to 3.1.

[25] The positive relationship between δ and profits in PM would extend to values lower than .4 if we were to increase the maximum number of firms, which becomes binding at $\delta = .4$ in these simulations.

[26] For our model, we can calculate an alternative measure of consumer surplus without imposing $\bar{\omega}_{max}$. To compute this alternative measure, we simulate the model using policy functions obtained from the model that imposes $\bar{\omega}_{max}$ but track and use the absolute quality grid when computing consumer surplus in each period. When equilibrium in our model never yields a monopoly, the equilibrium policies are good approximations to the policies when $\bar{\omega}_{max}$ is relaxed. When the constraint is relaxed, the difference is much greater, particularly when δ is low.

Columns one to three vary the base innovation efficiency (a_0), the innovation spillover (a_1), and the probability an entrant enters at the frontier quality (κ_e). In the baseline specification in Section 4.1, we set $a_0 = .5$, $a_1 = \{.2, .6\}$, and $\kappa_e = 0$. The top three rows set $a_0 = .5$, whereas the bottom three rows set $a_0 = .25$. Moving horizontally from column one to three varies the spillover $a_1 = \{0, 1, 2\}$. Moving vertically from row one to three varies the leap probability from $\kappa_e = \{0, .1, .2\}$, with the same pattern present in rows four to six. The fourth and fifth columns in Figure 12 consider a linear spillover $\tilde{a}(\omega_j, s)$ and linear utility $\tilde{u}_{ij}$, respectively. Both of these columns fix the spillover at $a_1 = 1$ and vary a_0 and κ_e as in the first three columns.

Looking across rows one and four reveals an inverted-U relationship between PMC and the industry's innovation rate, consistent with the earlier results in Section 4.1. Although the relationship with a linear utility function more closely resembles a spike, the intuition underlying this relationship is the same. Moving down from rows one and four increases the leap probability κ_e to 0.1 and 0.2. A positive leap probability enables entrants to immediately challenge the leader, resulting in higher industry innovation as PMC rises. Thus, the probability of leaping moderates the decline in innovation with high PMC because lead firms still innovate due to stronger threat of entrants.

Innovation Spillovers Figure 13 presents extended comparative statics illustrating the relationship between innovation spillovers and innovation. In the baseline model the spillover parameter is set to $a_1 = \{.2, .6\}$ and the comparative static in section 4.2 varies a_1 over $[0, 6]$. In this robustness check we vary a_1 over $[0, 3]$ since the trends in Figure 7 are steady after $a_1 = 3$ and the smaller range reduces the number of equilibria to compute.

Conditional on a particular level of PMC, our results in section 4.2 are broadly consistent with those in this expanded comparative static. The first column of Figure 13 fixes PMC at .47 and displays the same "giving up" behavior we observe for the low PMC value in Figure 7. The second column, which fixes PMC at 1.09, displays a relatively flat industry innovation rate, a result which lies between the low PMC of the moderate PMC of .78 and the high PMC of 1.56 in Figure 7. In the third column, which fixed PMC at 2.03, innovation increases over the full spillover range. Similar to the findings earlier for PMC, a positive leap probability raises industry innovation at low spillover values and induces the leader to give up at lower spillover levels in the alternative specification columns four and five.

References

[1] Aaker, D. (1991), *Managing Brand Equity*, Free Press.

[2] Aghion, P. and P. Howitt (1992), "A Model of Growth Through Creative Destruction," *Econometrica*, 60(2), 323–351.

[3] Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005), "Competition and Innovation: an Inverted-U Relationship," *The Quarterly J. of Economics*, 120(2), 701–728.

[4] Aghion, P., R. Blundell, R. Griffith, P. Howitt, and S. Prantl (2009), "The Effects of Entry on Incumbent Innovation and Productivity," *Review of Economics and Statistics*, 91(1), 20–32.

[5] Arrow, K. (1962), "Economic Welfare and the Allocation of Resources for Invention," in R. Nelson, ed., *The Rate and Direction of Inventive Activity* Princeton, NJ, Princeton University Press.

[6] Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010), "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics," *Econometrica*, 78, 453–508.

[7] Bloom, N., M. Schankerman, and J. Van Reenen (forthcoming), "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*.

[8] Borkovsky, R. (2012), "The Timing of Version Releases: A Dynamic Duopoly Model," working paper, U. of Toronto.

[9] Borkovsky, R., Doraszelski, U., and Y. Kryukov (2012), "A Dynamic Quality Ladder Model with Entry and Exit: Exploring the Equilibrium Correspondence Using the Homotopy Method," *Quantitative Marketing and Economics*, 10(2), 197–229.

[10] Berry, S., J. Levinsohn, and A. Pakes (1995), "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841–890.

[11] Bronnenberg, B., J. P. Dubé, C. Mela, P. Albuquerque, T. Erdem, B. R. Gordon, D. Hanssens, G. Hitsch, H. Hong, and B. Sun, (2008) "Measuring Long-Run Marketing Effects and Their Implications for Long-Run Marketing Decisions," *Marketing Letters*, 2008, 19, 367–382.

[12] Caplin, A. and B. Nalebuff (1991), "Aggregation and Imperfect Competition: On the Existence of Equilibrium," *Econometrica*, 59, 25–59.

[13] Chan, T. Y. and P. B. Seetharaman (2004), " Estimating Dynamic Pricing Decisions in Oligopolistic Markets: An Empirical Approach Using Micro- and Macro-level Data," working paper, Washington University in St. Louis.

[14] Chen, Y. and M. Schwartz (2010), "Product Innovation Incentives: Monopoly vs. Competition," working paper, Georgetown University.

[15] Cohen, W. M. (1995), "Empirical Studies of Innovative Activity," In: Stoneman, P. (ed.), Handbook of the economics of innovation and technological change. Oxford: Blackwell Publishers.

[16] Cohen, W. M. and R. Levin (1989), "Empirical Studies of Innovation and Market Structure," in R. Schmalensee and R. Willig, eds., *Handbook of Industrial Organization, Volume 2*, North-Holland, Amsterdam, 1059–1107.

[17] Cohen, W. M. and D. A. Levinthal (1989), "The Implications of Spillovers for R&D Investment and Welfare: A New Perspective," in: A. Link and K. Smith, eds., Advances in Applied Micro-Economics,. vol. 5: The factors affecting technological change. Greenwich, Conn.: JAI Press.

[18] Dasgupta, P. and J. Stiglitz (1980), "Industrial Structure and the Nature of Innovative Activity," *Economic J.*, 90(358), 266–293.

[19] Day, G. (1994), "The Capabilities of Market-Driven Organizations," *J. of Marketing*, 58, 37–52.

[20] Department of Justice and Federal Trade Commission (2010), "Horizontal Merger Guidelines," accessed on 5/30/2011 from http://ftc.gov/os/2010/08/100819hmg.pdf.

[21] Doraszelski, U. (2003), "An R&D Race with Knowledge Accumulation," *RAND J. Econ.*, 34(1), 19–41.

[22] Doraszelski, U. and M. Satterthwaite (2010), "Computable Markov-Perfect Industry Dynamics" *RAND J. of Economics*, 41(2), 215–243.

[23] Dubé, J.P., G. Hitsch, and P. Manchanda (2005), "An Empirical Model of Advertising Dynamics," *Quantitative Marketing and Economics*, 3, 107–144.

[24] Dutta, S., O. Narasimhan, and S. Rajiv (1999), "Success in High-Technology Markets: Is Marketing Capability Critical?" *Marketing Science*, 18(4), 547–568.

[25] Harris, C. and J. Vickers (1987), "Racing with Uncertainty," *Rev. of Econ. Studies*, 54(1), 1–21.

[26] Ericson, S. and A. Pakes (1995), "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Rev. of Econ. Studies*, 62(1), 53–82.

[27] Fershtman, C. and A. Pakes (2000) "A dynamic oligopoly with collusion and price wars," *RAND J. of Economics*, 31(2), 207–236.

[28] Gilbert, R. (2006), "Looking for Mr. Schumpeter: Where are We in the Competition-Innovation Debate?" in A. Jaffe, J. Lerner, and S. Stern, eds., *Innovation Policy and the Economy, Vol. 6*.

[29] Gilbert, R. and D. M. Newbery (1982), "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review*, 72, 514–526.

[30] Goettler, R. L. and B. R. Gordon (2011), "Does AMD Spur Intel to Innovate More?," *Journal of Political Economy*, 119(6), 1141–1200.

[31] Gowrisankaran, G. and R.J. Town (1997), "Dynamic Equilibrium in the Hospital Industry," *J. of Economics and Management Strategy*, 6(1), 45–74.

[32] Greenstein, S. and G. Ramey (1998), "Market Structure, Innovation, and Vertical Product Differentiation," *Intl. J. of Industrial Organization*, 16, 285–311.

[33] Griliches, Z. (1998), "The Search for R&D Spillovers," in R&D and Productivity: The Econometric Evidence), University of Chicago Press.

[34] Grossman, G. M. and E. Helpman (1991), "Quality Ladders in the Theory of Growth," *Rev. of Econ. Studies*, 58, 43–61.

[35] Han, J., N. Kim, & H. Kim (2001), "Entry Barriers: A Dull-, One-, or Two-Edged Sword for Incumbents? Unraveling the Paradox from a Contingency Perspective," *J. of Mkting*, 65(1), 1–14.

[36] Hof, R. D. (2011), "Lessons from Sematech," *Technology Review*, July 25, Published by MIT.

[37] Iskhakov, F., J. Rust, B. Schjerning (2013), "The Dynamics of Bertrand Price Competition with Cost-Reducing Investments," working paper, Georgetown University.

[38] Kamien, M. I., E. Muller, and I. Zang (1992), "Research Joint Ventures and R&D Cartels," *American Economic Review*, 82(5), 1293-1306.

[39] Karakaya, F. and M. J. Stahl (1989), "Barriers to Entry and Market Entry Decisions in Consumer and Industrial Goods Markets," *J. of Marketing*, 53(2), 80–91.

[40] Khanna, T. and M. Iansiti (1997), "Firm Asymmetries and Sequential R&D: Theory and Evidence from the Mainframe Computer Industry," *Management Science*, 43, 405-421.

[41] Klemperer, P. (1995), "Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade," *Rev. of Econ. Studies*, 62(4), 515–539.

[42] Lee, T. and L. Wilde (1980), "Market Structure and Innovation: A Reformulation," *Quarterly J. of Economics*, 94, 429–436.

[43] Lerner, J. (1997), "An Empirical Exploration of a Technology Race," *RAND J. Econ.*, 28, 228-247.

[44] Levin, R. C., W. M. Cohen, and D. C. Mowery (1985), "R&D Appropriability, Opportunity, and Market Structure: New Evidence on Some Schumpeterian Hypotheses," *American Economic Rev.*, 75(2), 20–24.

[45] Levin, R. C., A. K. Klevorick, R. R. Nelson, and S. G. Winter (1987), "Appropriating the Returns from Industrial Research and Development," *Brookings Papers on Economic Activity*, 3, 783–831.

[46] Levin, R. C. and P. .C. Reiss (1998), "Cost-Reducing and Demand-Creating R&D with Spillovers," *RAND J. Econ.*, 19(4), 538-556.

[47] Loury, G. (1979), "Market Structure and Innovation," *Quarterly J. of Economics*, 93, 395–410.

[48] New Yorker, The (2006), "Deal Sweeteners" by James Surowiecki, November 27.

[49] New Yorker, The (2007), "Running on Fumes" by Elizabeth Kolbert, November 5.

[50] Nickell, S. (1996), "Competition and Corporate Performance," *Journal of Political Economy*, 104, 724-746.

[51] Ofek, E. and M. Sarvary (2003), "R&D, Marketing, and the Success of Next-Generation Products," *Marketing Science*, 22(3), 355–370.

[52] Pakes, A. and P. McGuire (1994), "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *RAND J. of Economics*, 25(4), 555–589.

[53] Reinganum, J. (1983), "Uncertain Innovation and the Persistence of Monopoly," *American Economic Review*, 73(4), 741–748.

[54] Reuters (2011), "Goldman cuts Intel to sell as supply glut looms, shares fall," May 19.

[55] Scherer, F. (1967), "Market Structure and the Employment of Scientists and Engineers," *American Economic Review*, 57, 524-531.

[56] Schumpeter, J (1942), *Capitalism, Socialism, and Democracy*, New York, Harper.

[57] Spence, M. (1984), "Cost Reduction, Competition, and Industry Performance," *Econometrica*, 52(1), 101-122.

[58] Soberman, D. and H. Gatignon (2005), "Research Issues at the Boundary of Competitive Dynamics and Marketing Evolution," *Marketing Science*, 24(1), 165–174.

[59] Song, M. (2011), "A Dynamic Analysis of Cooperative Research in the Semiconductor Industry," *International Economic Review*, 52(4), 1157-1177.

[60] Sutton, J. (1991), *Sunk Cost and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. Cambridge: MIT Press.

[61] Teshima, K. (2010), "Import Competition and Innovation at the Plant Level: Evidence from Mexico," Working paper, Instituto Tecnologico Autonomo de Mexico.

[62] Vives, X. (2008), "Innovation and Competitive Pressure," *J. of Indust. Economics*, 56(3), 419–469.

[63] Weintraub, G. Y. , C. L. Benkard, B. Van Roy (2008), "Markov Perfect Industry Dynamics with Many Firms," *Econometrica*, 76(6), 1375-1411.

[64] Wernerfelt, B. (1984), "A Resource Based View of the Firm," *Strat. Mngmnt. J.*, 5(2), 171–180.

Figure 1: Gains to Static Profits when the Leader or Firm 2 Innovates, Holding other Firms Fixed: The top row reports the gain in static profits for the leader (at $\omega_1 = 7$) and firm 2, as a function of ω_2 (x-axis) and the number of other firms at $\omega = 2$ (curve labels). In the bottom row, the x-axis and labels are reversed to provide an alternative view. Product-market competition varies across the columns as labeled from .47 to .78 to 1.87. To relate these gains to R&D costs, a leader investing $x = 2$ innovates with probability 1/2.

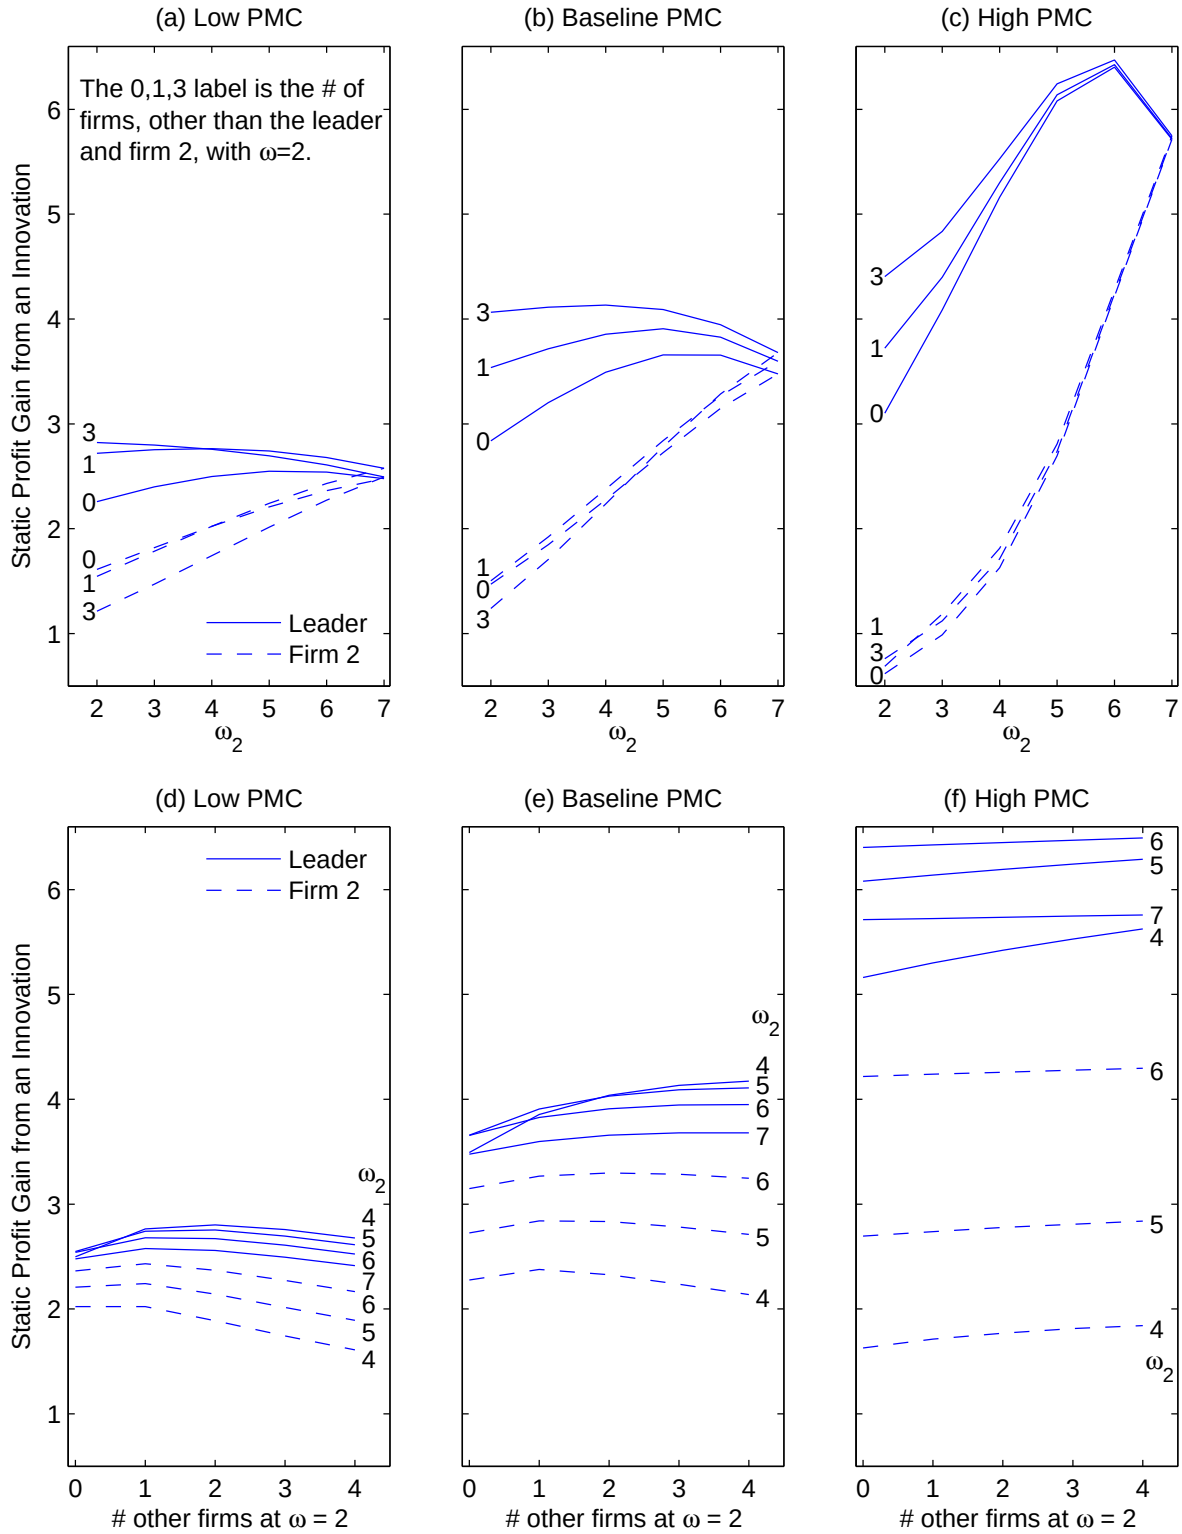

Figure 2: Policy Functions: Each panel presents the policy functions in a doupoly, with the degree of spillovers, PMC, and entry costs varying over the panels. The vertical axes are probabilities of entry or innovation. The horizontal axis is the number of quality steps the laggard is behind the leader.

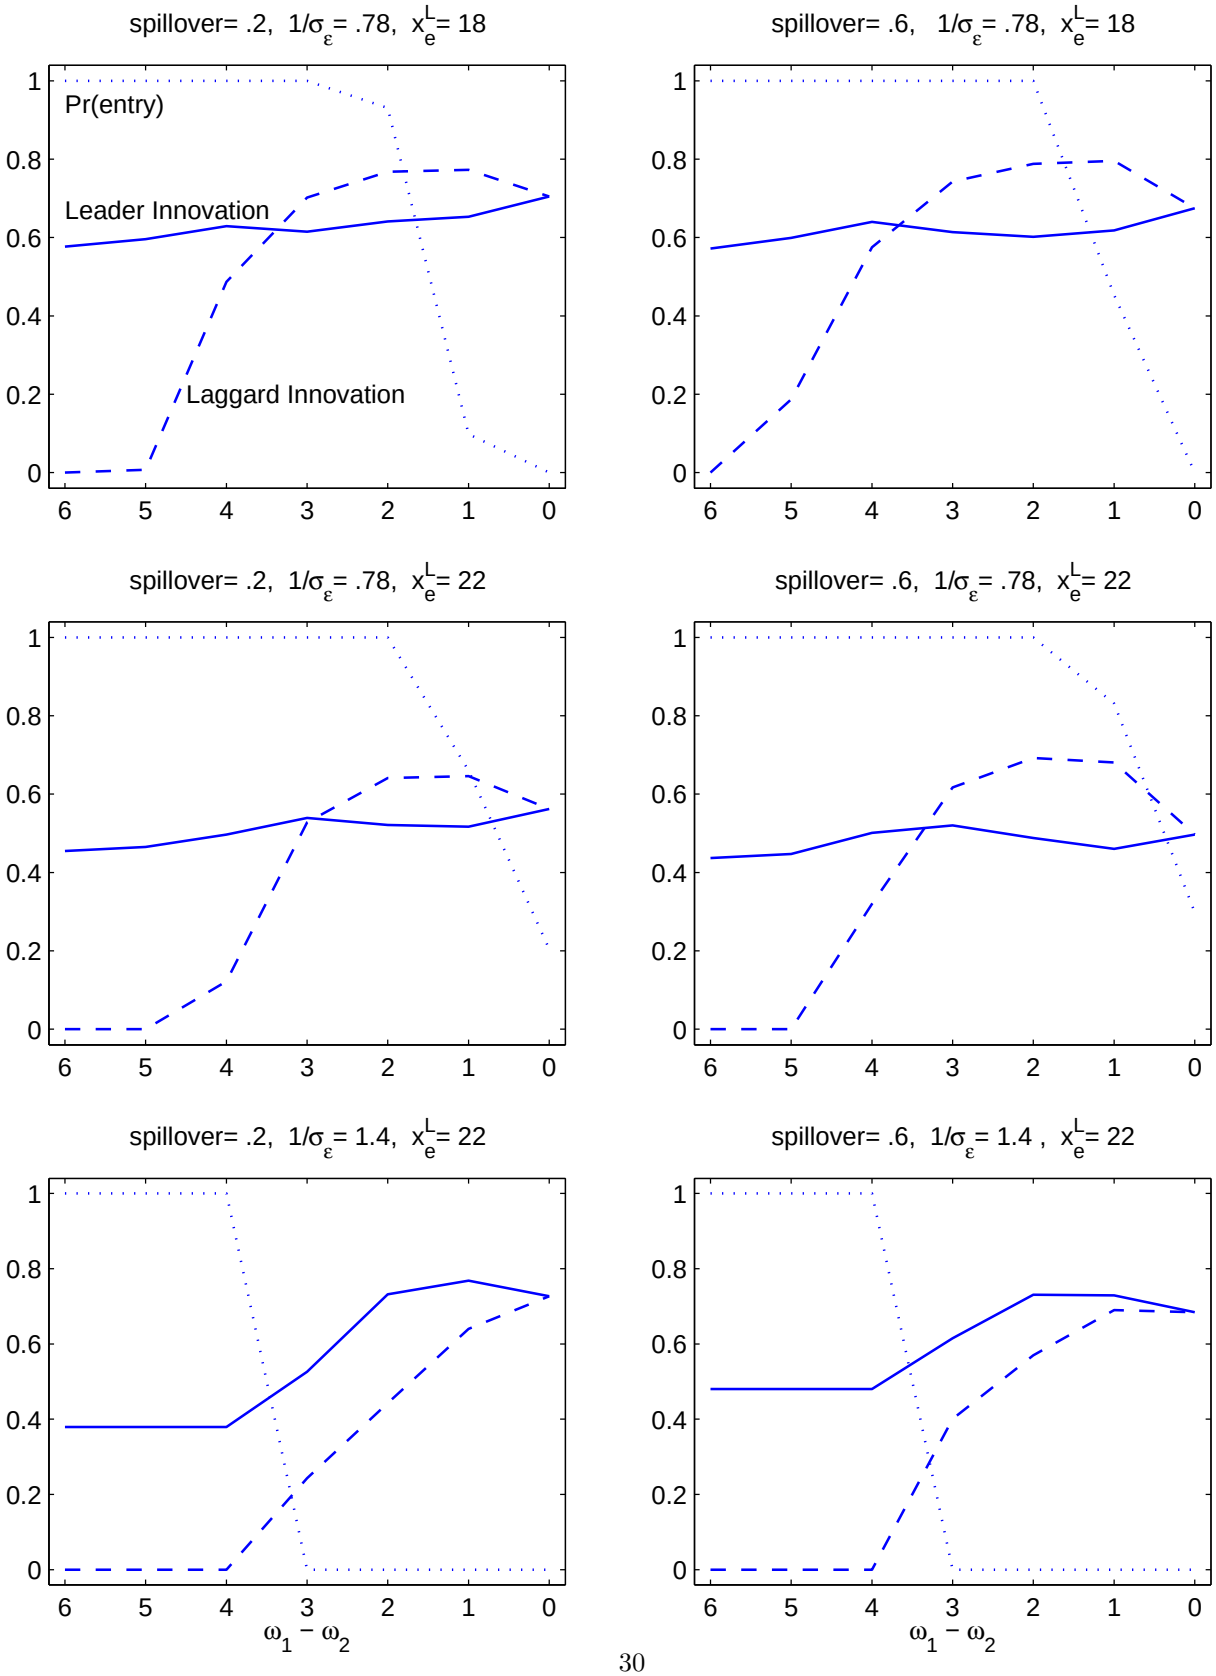

Figure 3: Comparative Static: Varying Product-Market Competition: Vertical axes below correspond to each panel's title and all horizontal axes vary the degree of product-market competition, $1/\sigma_\epsilon$.

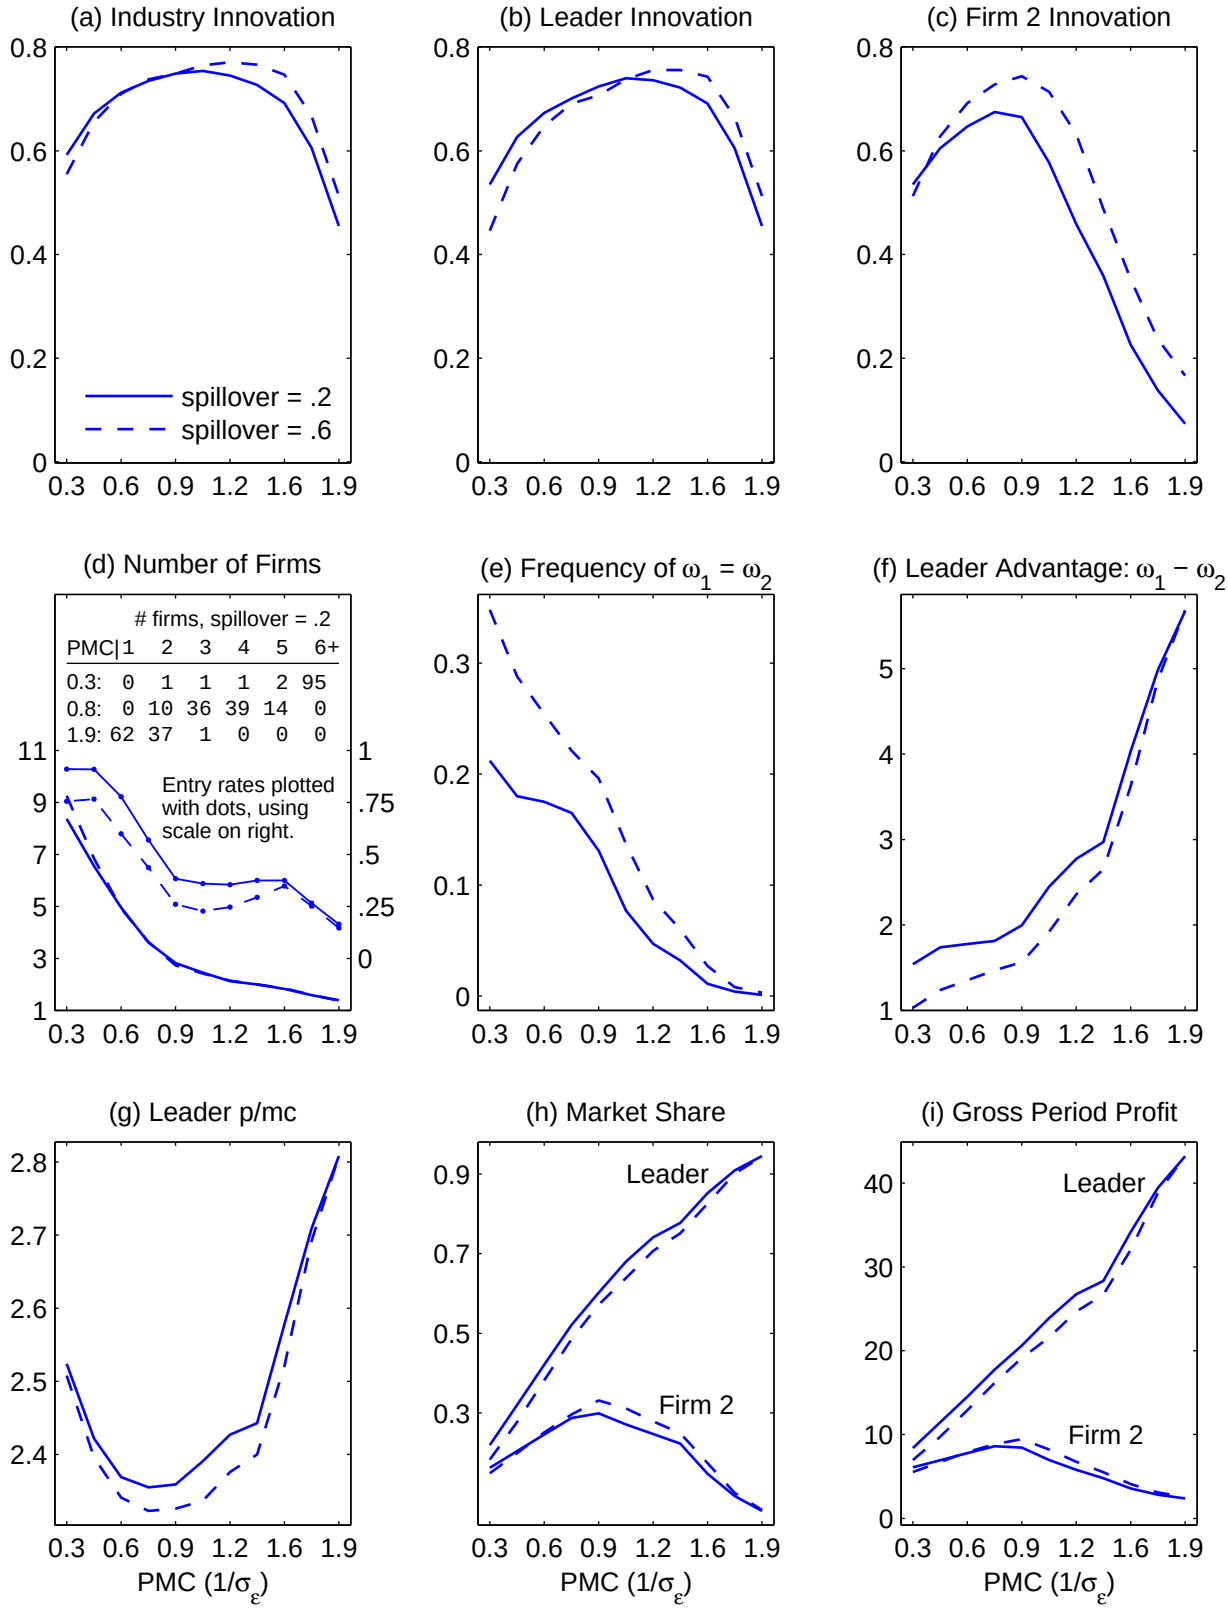

Figure 4: Policy Functions for the Comparative Static in PMC: Panels (a), (b), and (c) provide policy functions at seven states as product-market competition varies along the x-axis. The states have two incumbents with ω_1 always at 7 and ω_2 varying from 1 to 7, as labeled. Panels (d), (e), and (f) plot policies at three states with $\omega_1 = 7$ and either 1, 2, or 3 firms at $\omega_e = 4$, as labeled. The spillover is .2.

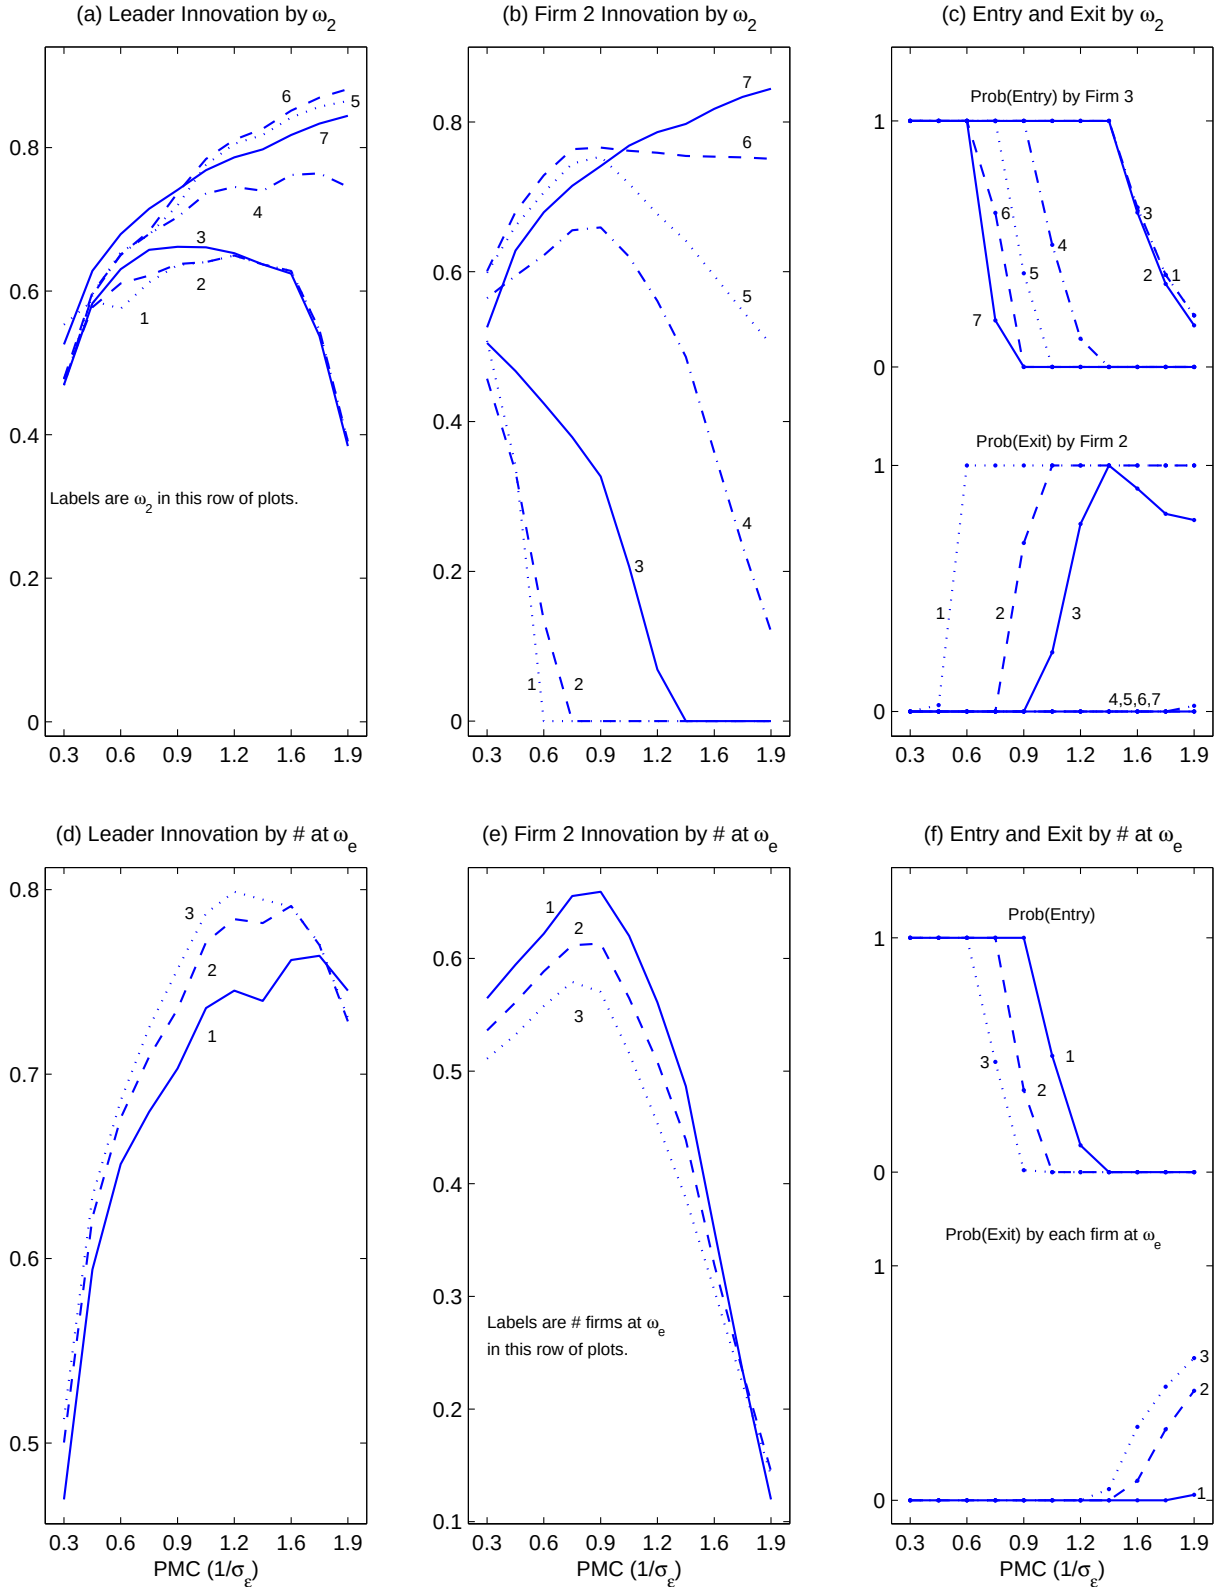

Figure 5: Comparative Static: Varying Product-Market Competition with Exogenous Market Structure: Vertical axes below correspond to each panel's title. The horizontal axes in the top row vary the degree of product-market competition and in the bottom row vary the quality of the laggard firm (where the lead firm is $\omega_1 = 30$). In the top row, each curve corresponds to an exogenous market structure with either two, three, or four firms. In the bottom row, panels (d) and (e) display the innovation rate policies assuming a fixed duopoly with high and low PMC, respectively. Panel (f) presents the innovation policies for the leader and a laggard in a triopoly when the third firm's quality is $\omega_3 = 14$.

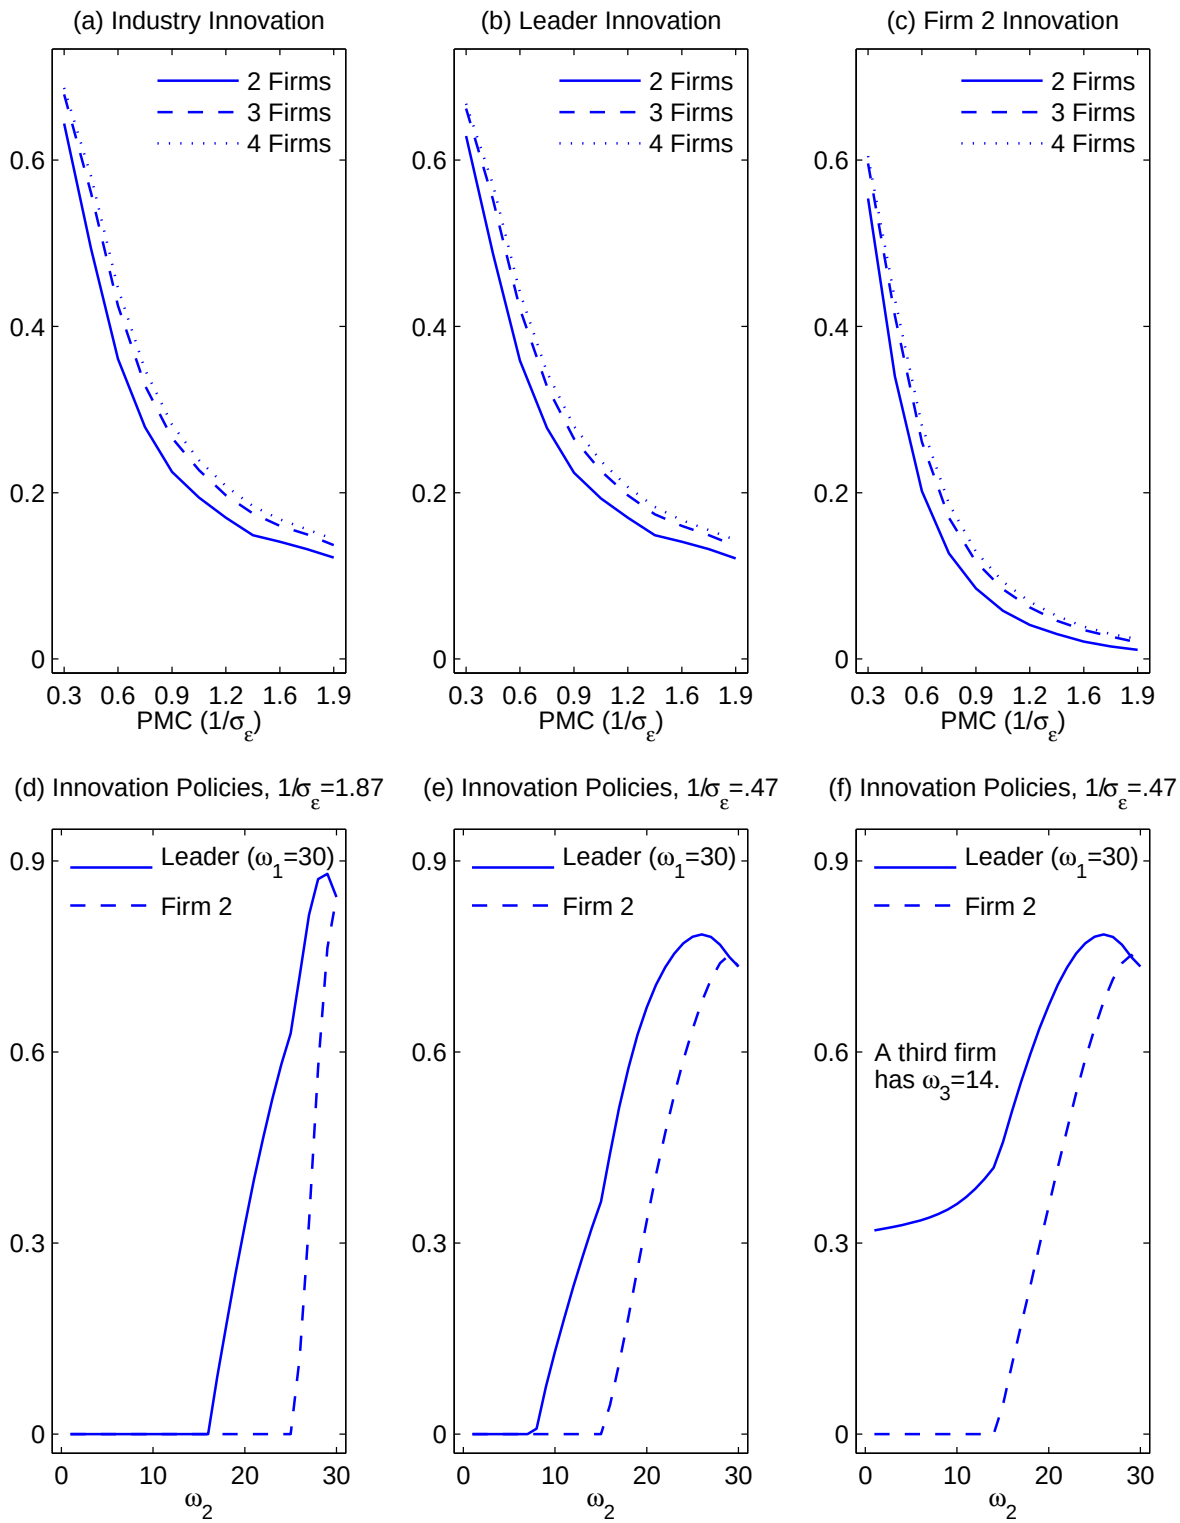

Figure 6: Comparative Static: Varying Product-Market Competition and Outside Good's Innovation Rate: The upper plot reports industry innovation rates under different assumptions concerning δ, the outside good's innovation rate, while varying the degree of product-market competition on the x-axis. The baseline with "No Outside Good" is identical to the outcome in Figure 3.a with spillover = .2. The next curve presents innovation when an outside good exists but never advanced ($\delta = 0$). The remaining curves depict innovation for positive values of δ. The lower plot reports the same outcomes when spillover = .6.

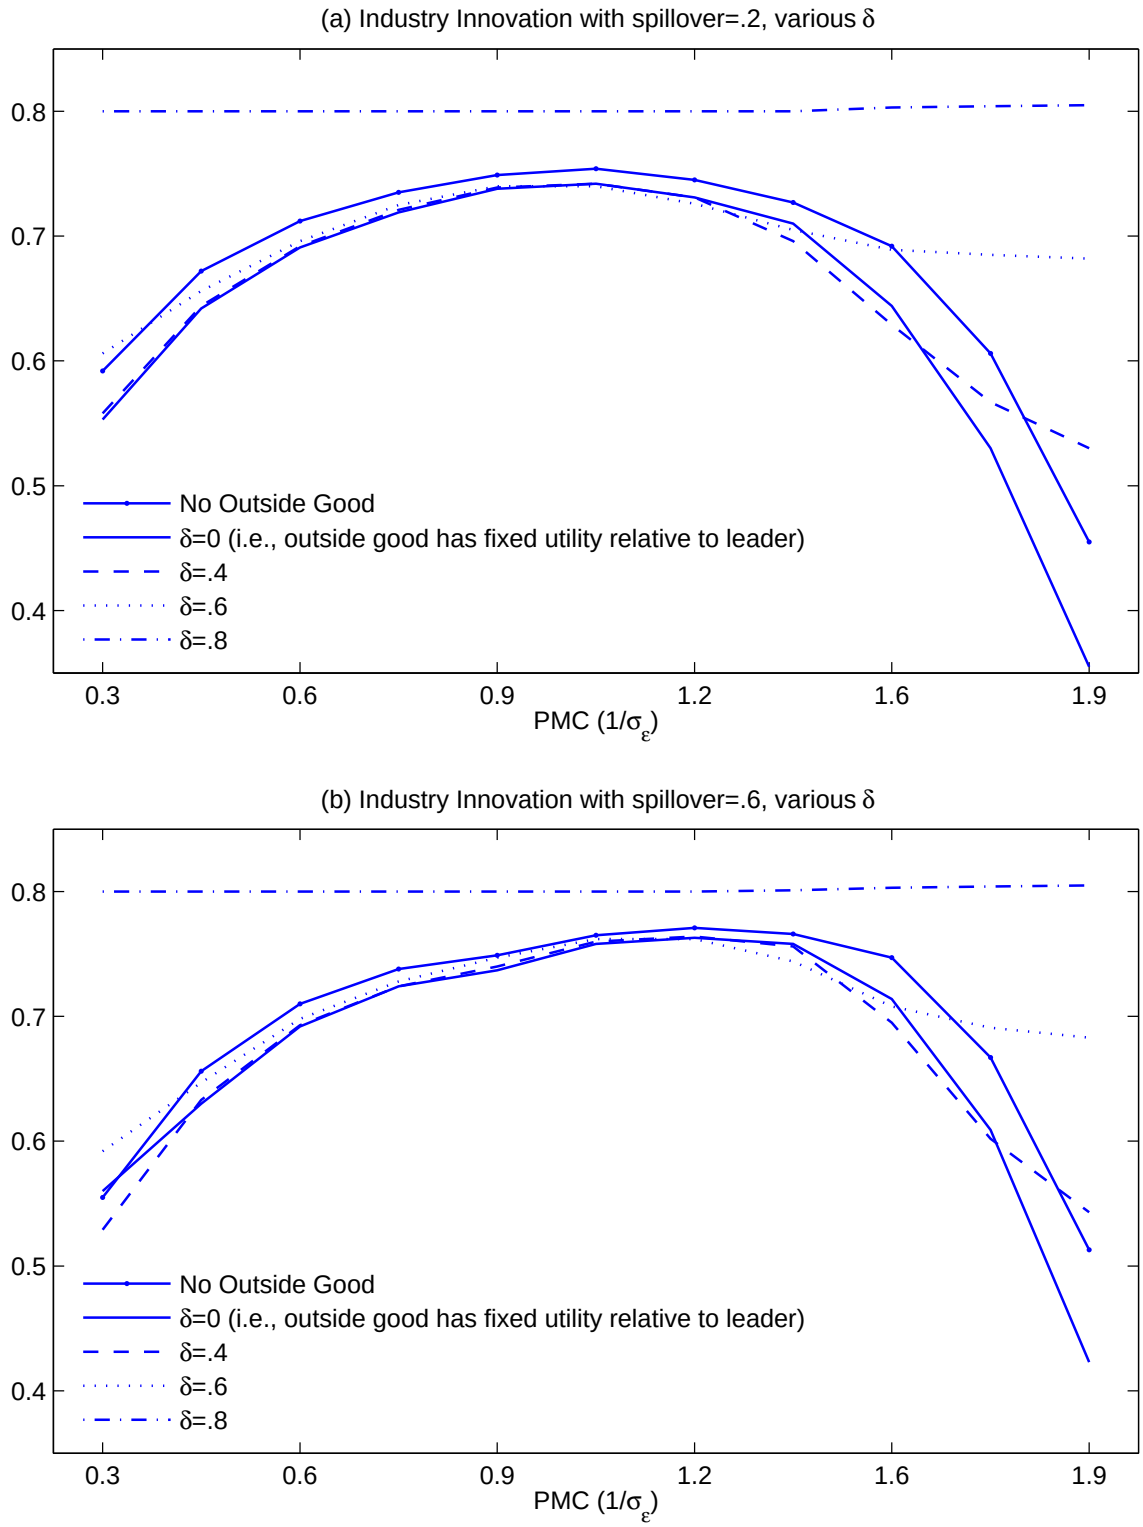

Figure 7: Comparative Static: Varying the Investment Spillover: Vertical axes below correspond to each panel's title and all horizontal axes vary the innovation spillover parameter a_1.

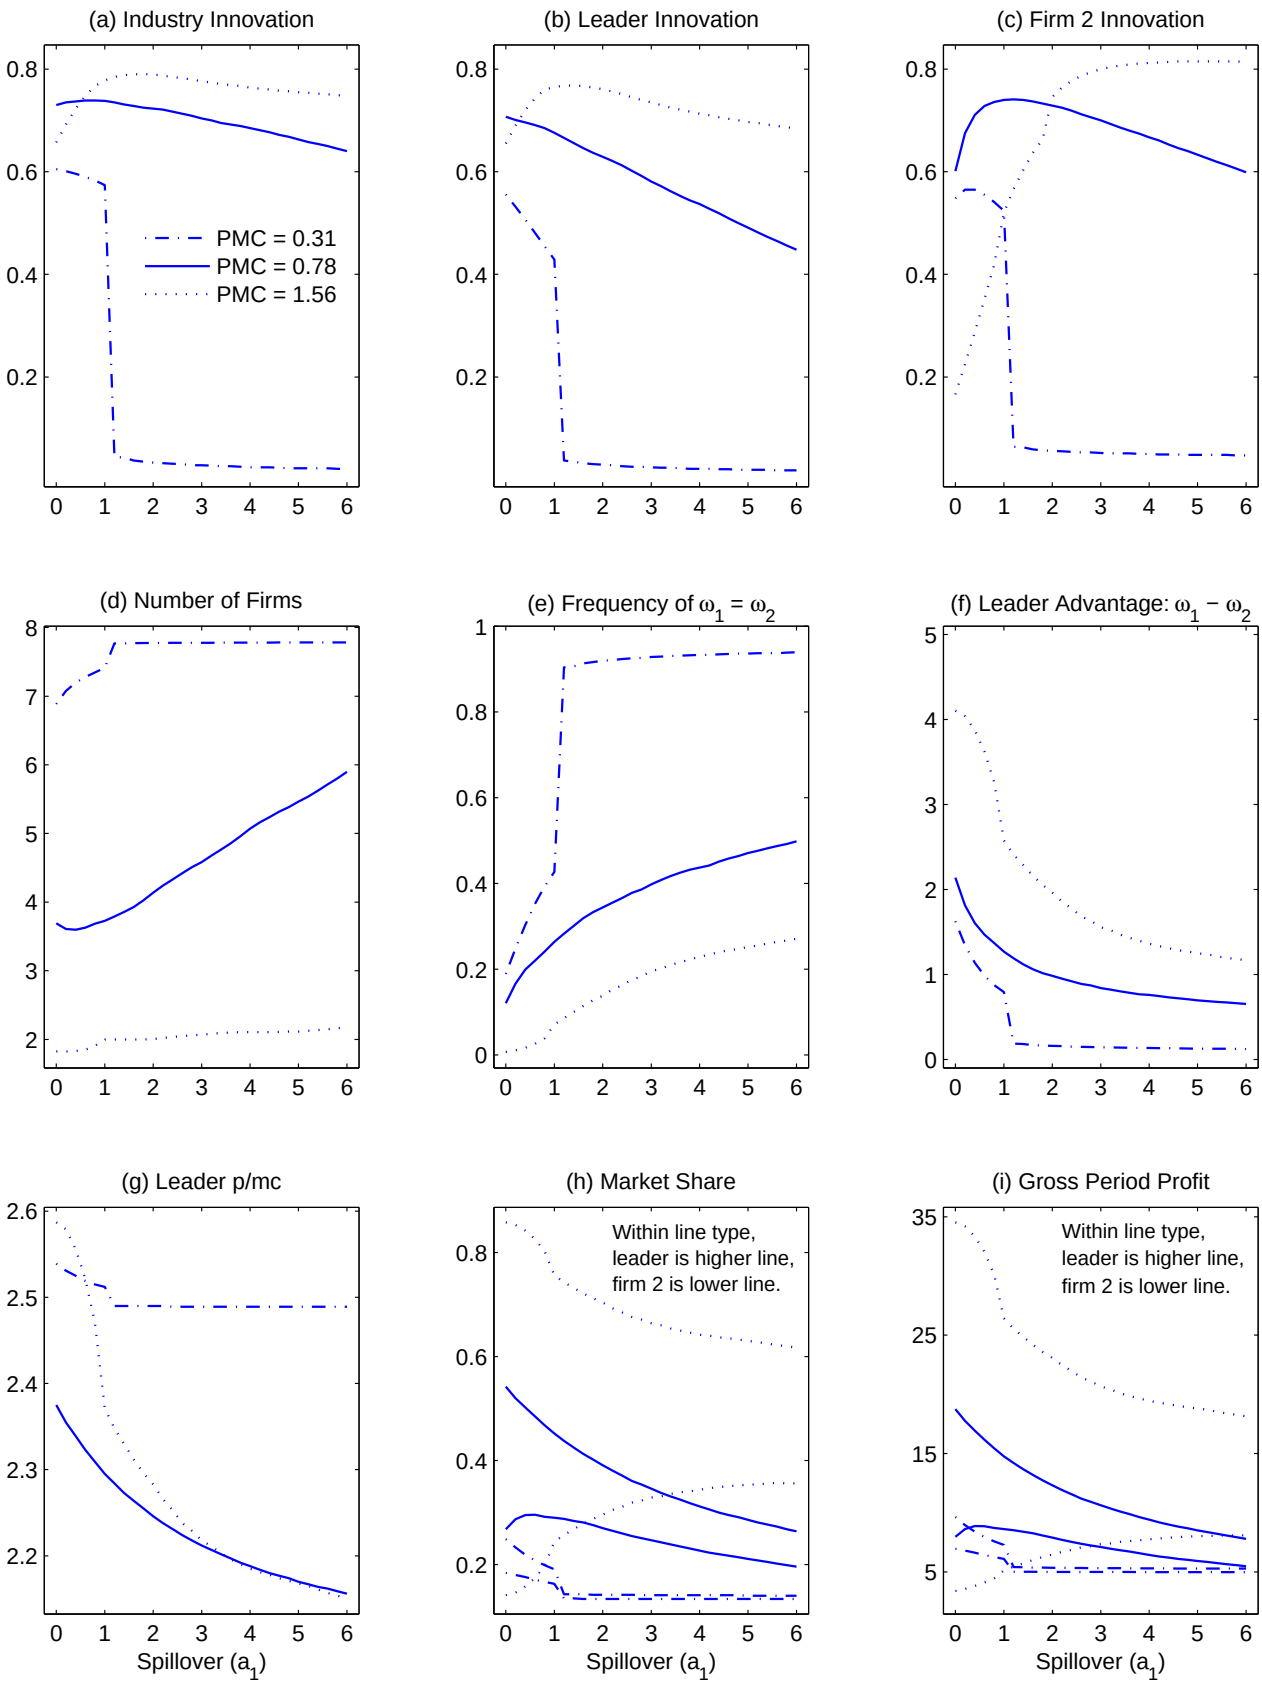

Figure 8: Comparative Static: Varying Entry Costs: Vertical axes below correspond to each panel's title and all horizontal axes vary x_e^L, the lower bound of the entry cost distribution.

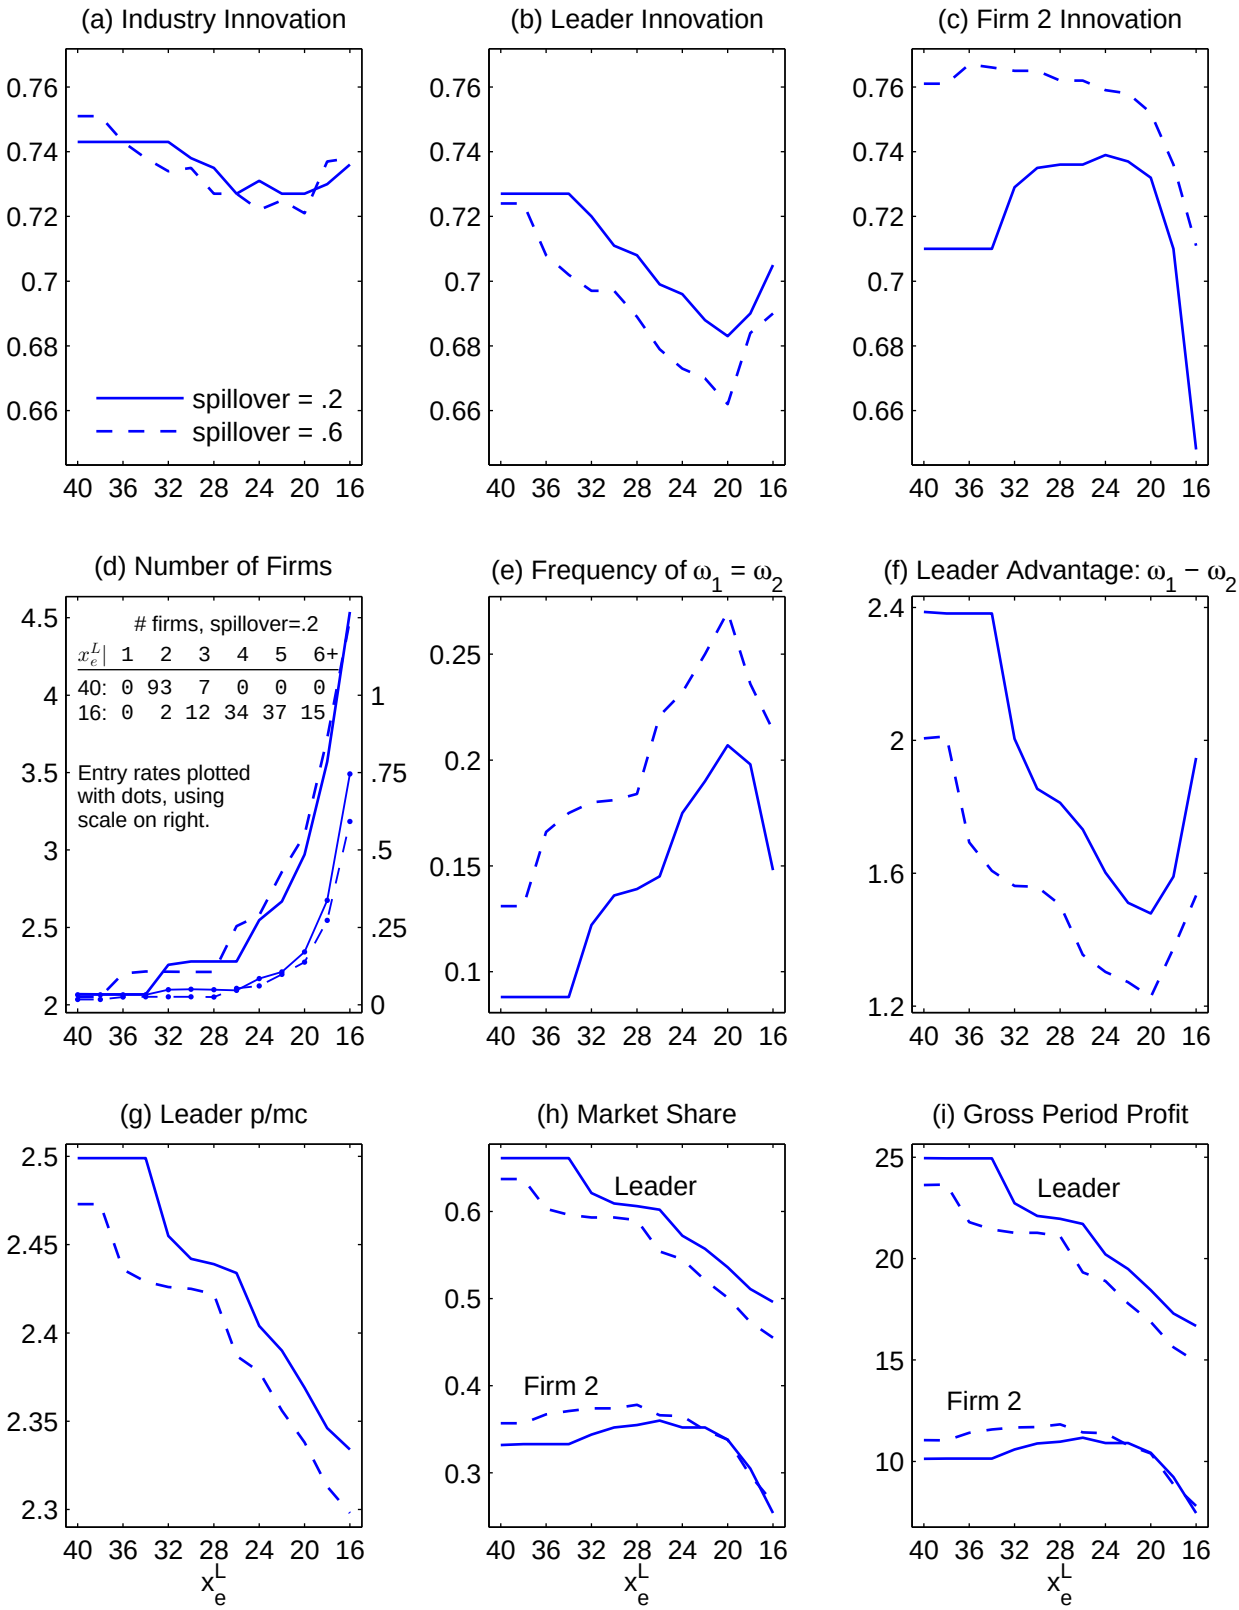

Figure 9: Policy Functions for the Comparative Static in the Entry Cost: Panels (a), (b), and (c) plot policies at seven states as the entry cost varies along the x-axis. The states have two incumbents with ω_1 always at 7 and ω_2 varying from 1 to 7, as labeled. Panels (d), (e), and (f) plot policies at three states with $\omega_1 = 7$ and either 1, 2, or 3 firms at $\omega_e = 4$, as labeled. The spillover is .2.

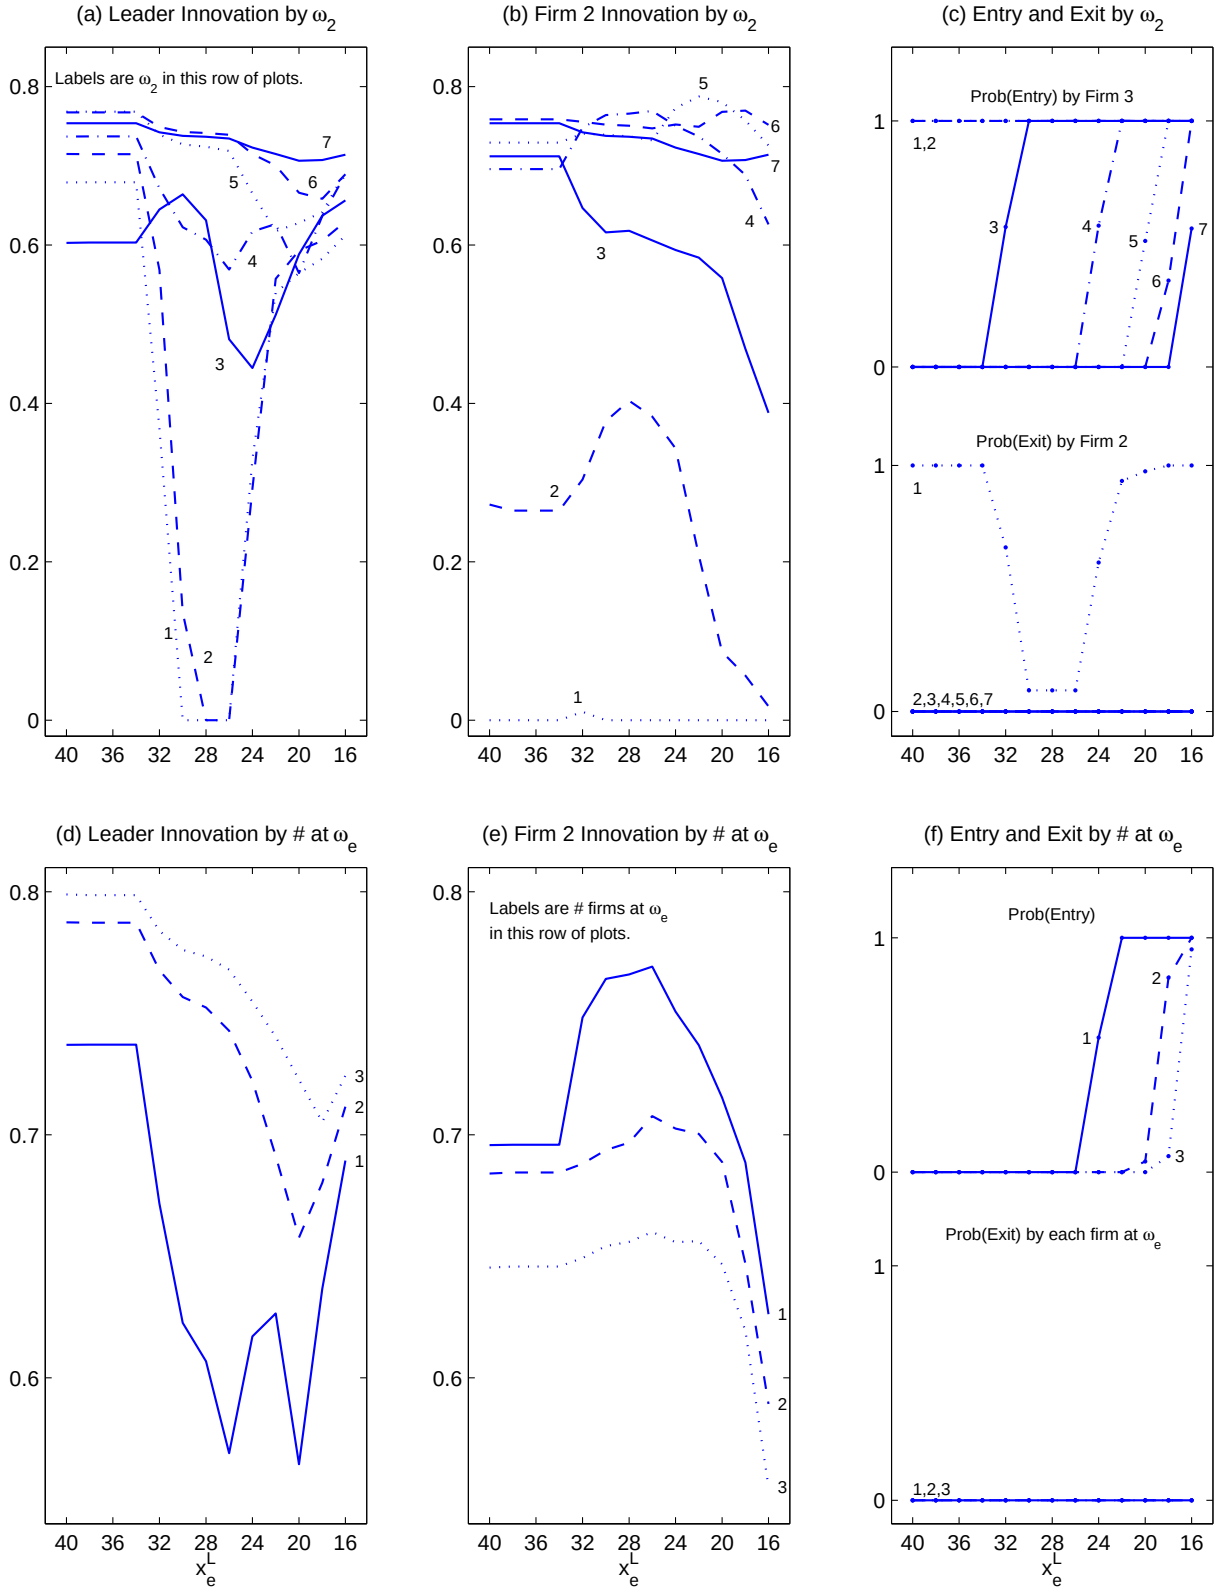

Figure 10: Plot of $g(\omega)$ and share of firms above ω^* in Pakes & McGuire (1994)

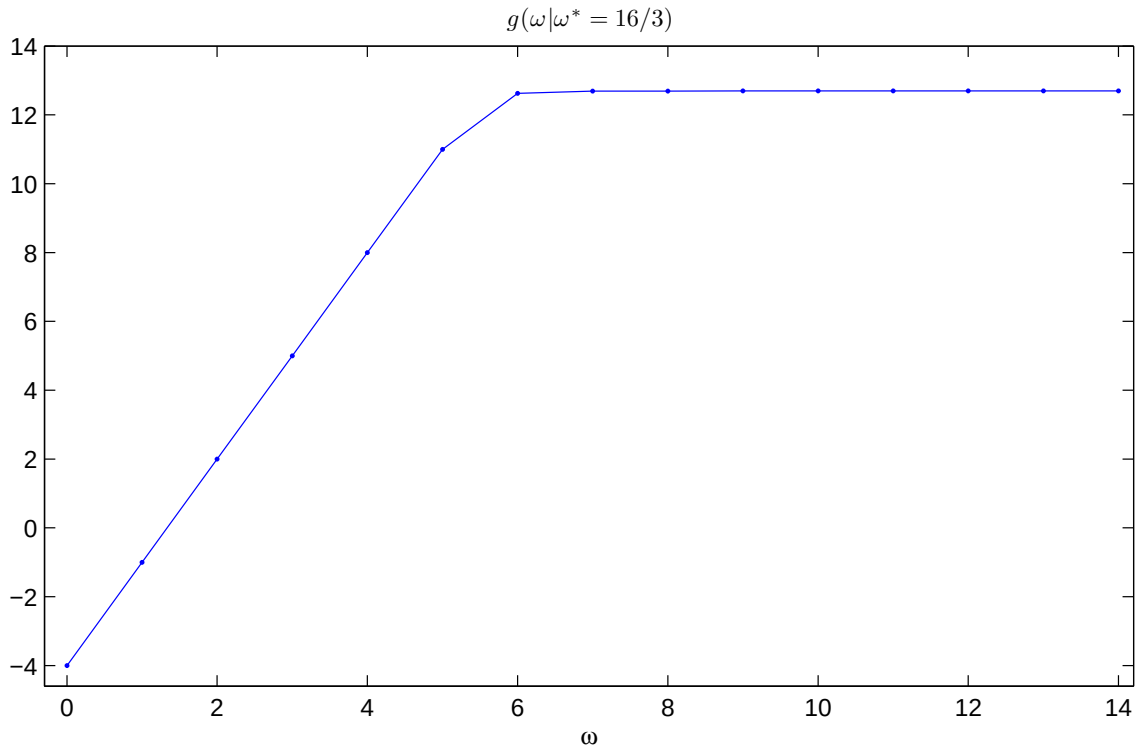

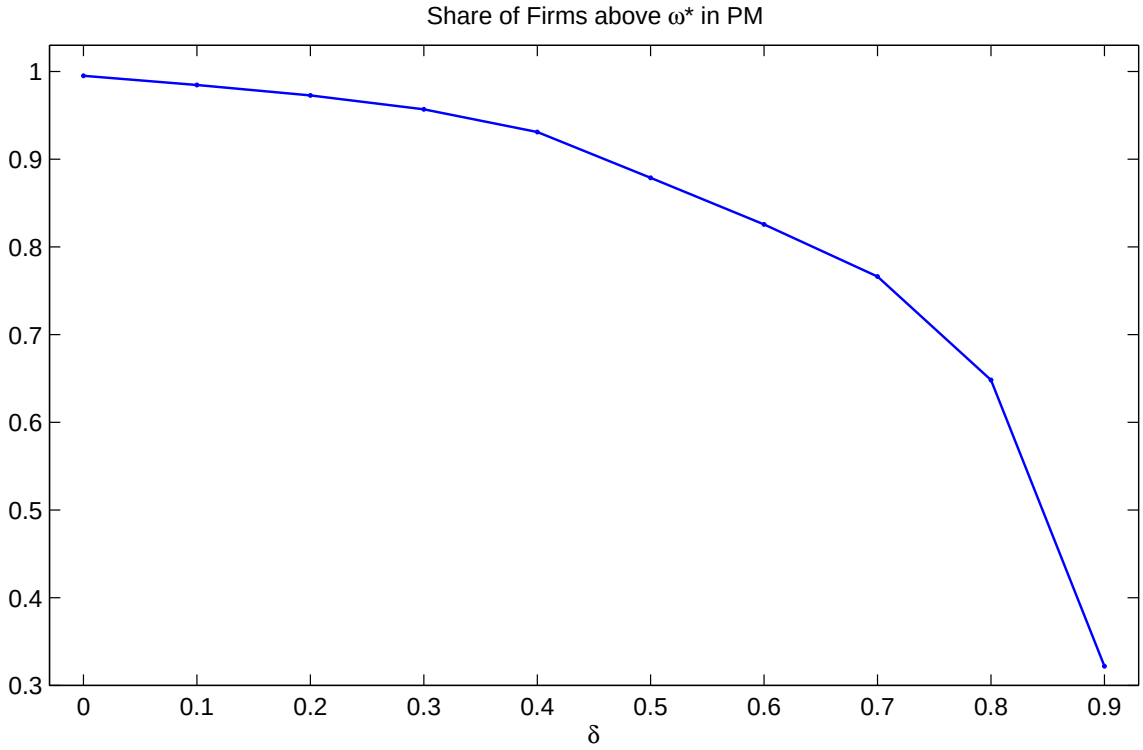

Figure 11: Comparison to Pakes & McGuire (1994) While Varying δ: Our demand model is labeled "linear" because its utility function is linear in some function of absolute product quality (e.g., quality itself or log quality). Pakes-McGuire uses a utility function that is concave in product quality measured relative to the outside good. The y-axis of each panel corresponds to its title. Industry innovation is measured as the share of periods in which the industry's frontier quality improves. The x-axis in each panel is δ, the rate the outside good improves, as labeled on the bottom row.

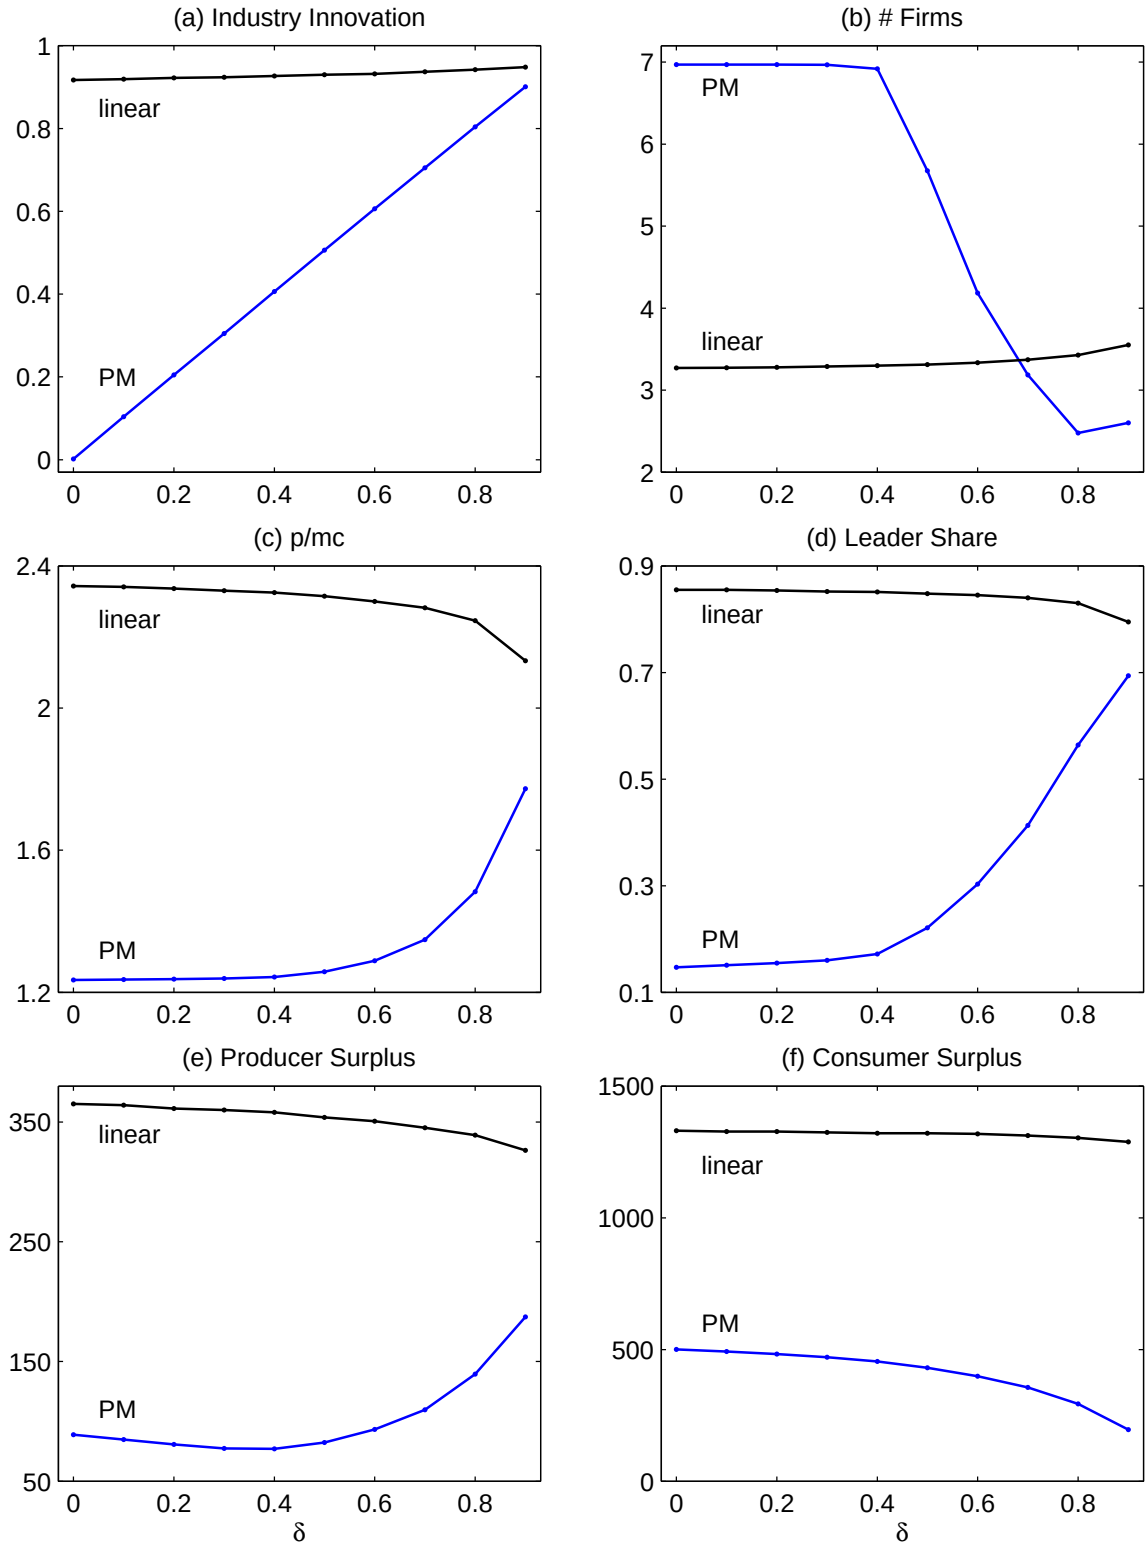

Figure 12: Extended Comparative Statics in PMC: See Appendix B for details.

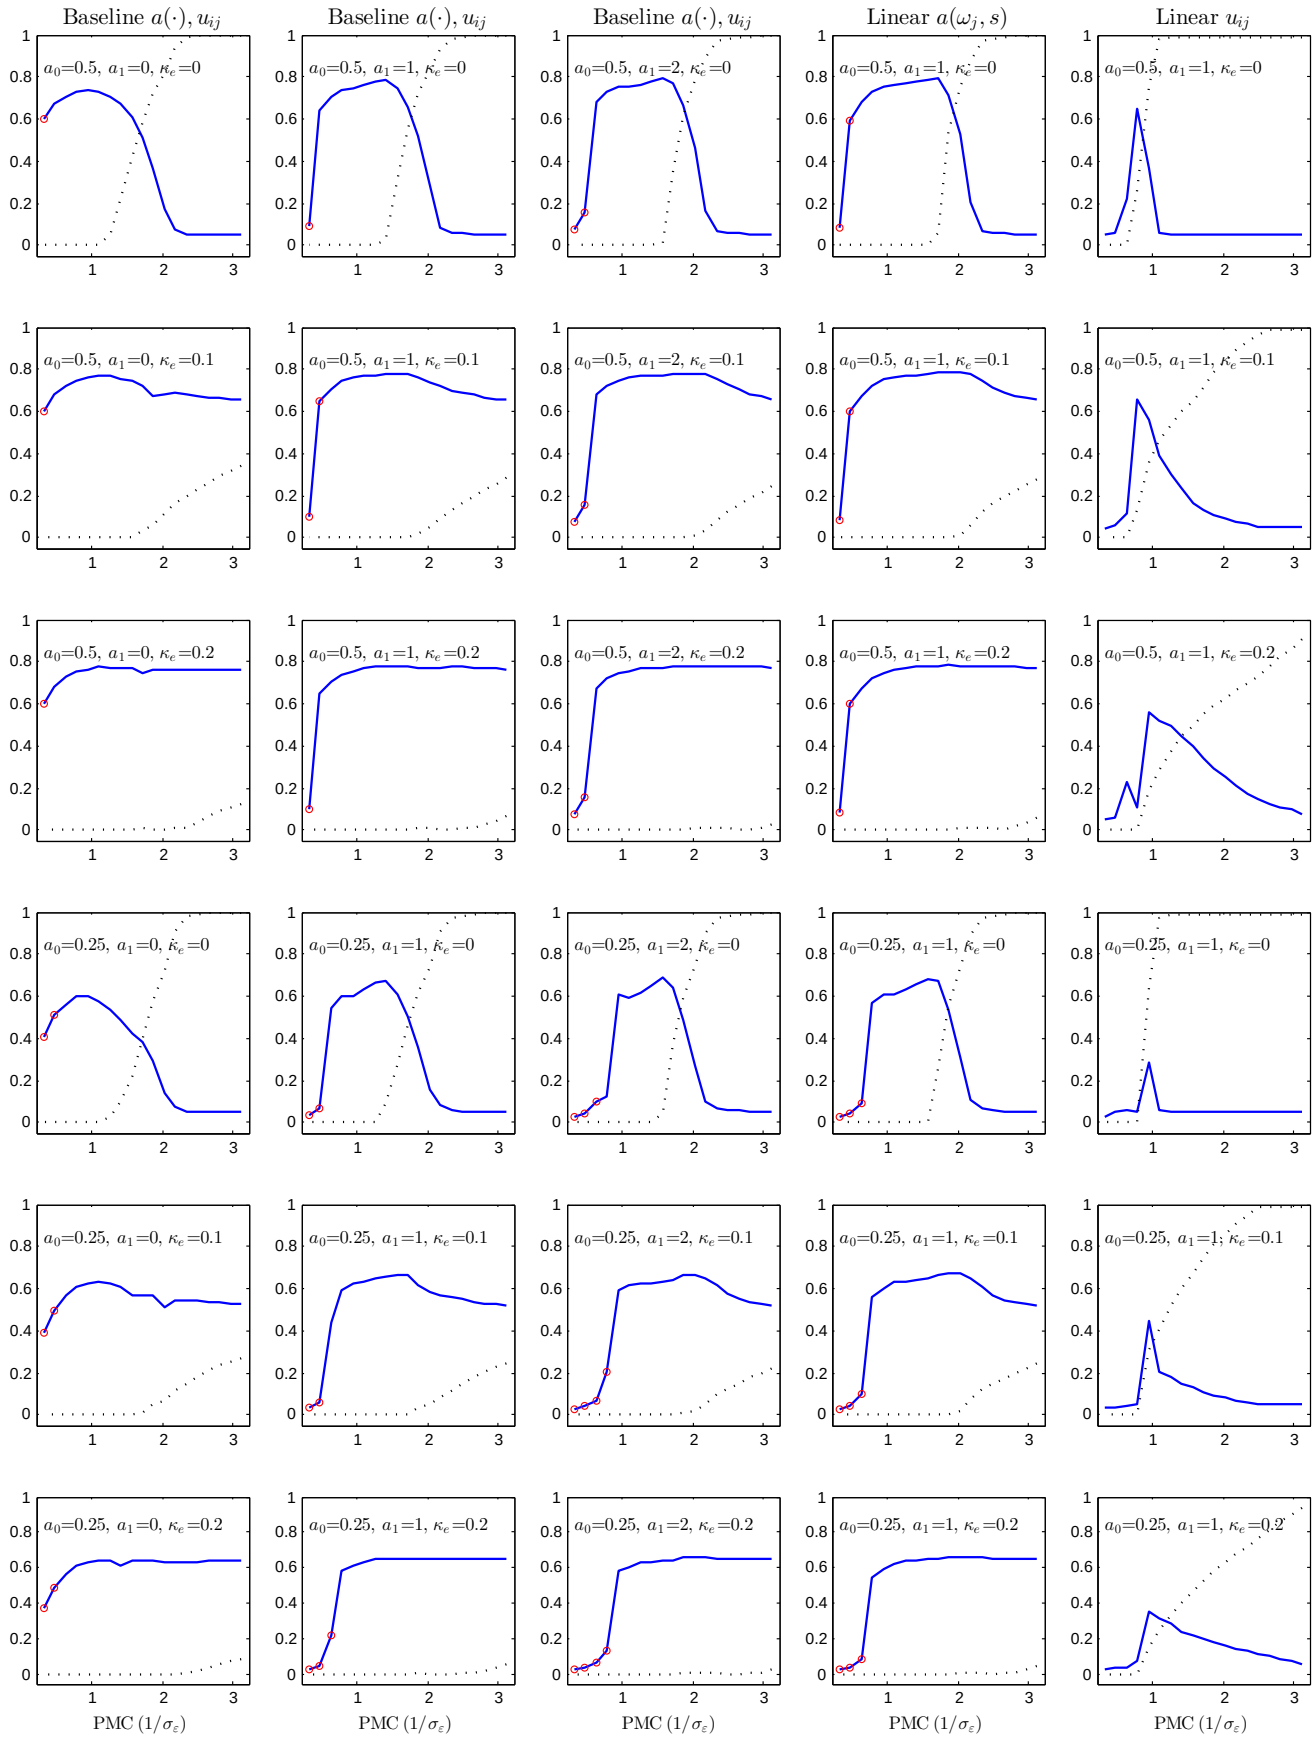
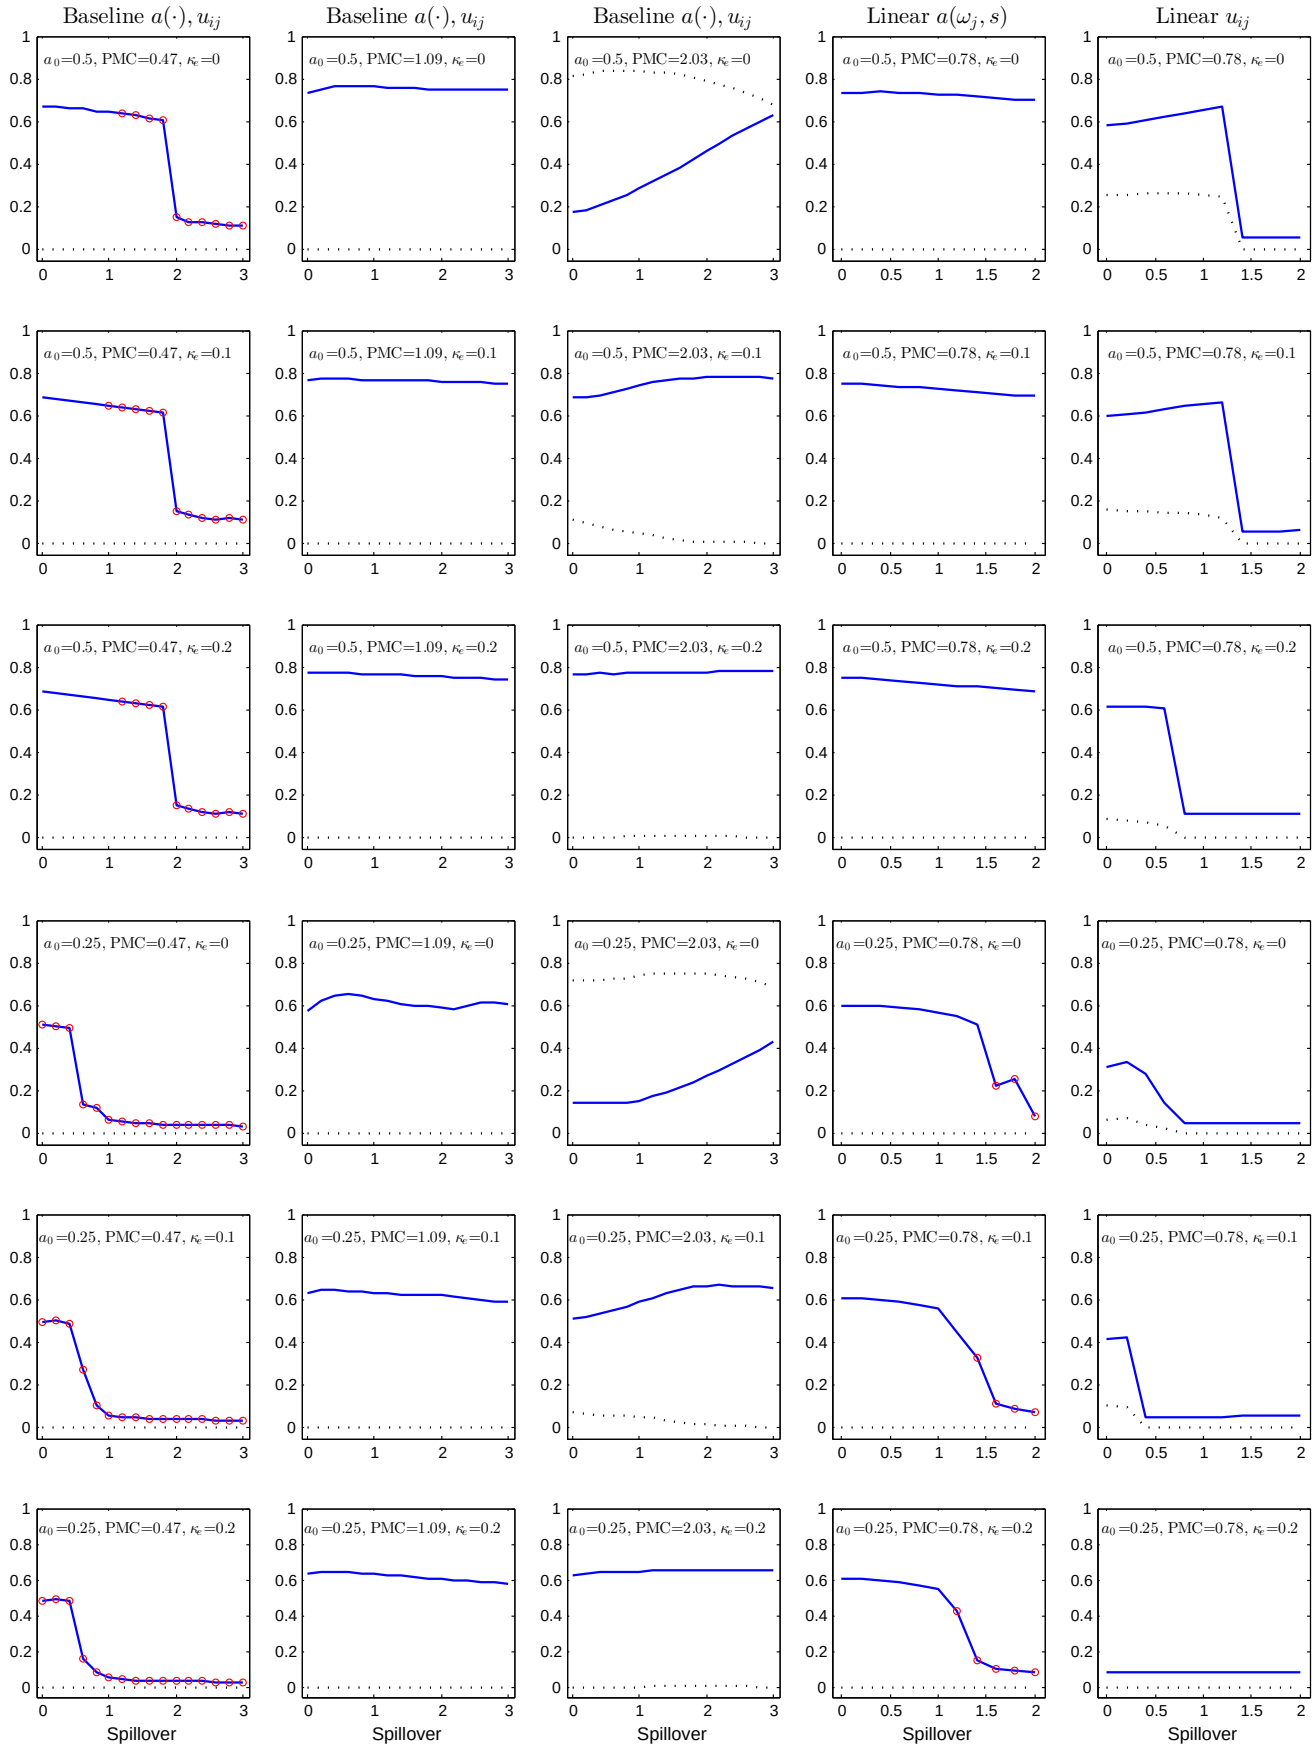
Figure 13: Extended Comparative Statics in Spillover: See Appendix B for details.