The Dynamics of Innovation and Horizontal Differentiation*

Borghan Narajabad
Department of Economics
Rice University

Randal Watson
Department of Economics
University of Texas at Austin

Abstract

We study innovation in a dynamic stochastic discrete-time duopoly with endogenous horizontal differentiation. Innovation takes the form of a quality ladder; horizontal differentiation is Hotelling competition.

---

*A previous version of this paper appeared as Chapter 3 in Narajabad (2007), and was circulated under the title 'A dynamic duopoly with endogenous horizontal and vertical differentiation'. We are grateful for helpful comments from anonymous referees and an Associate Editor, and from Satyajit Chatterjee, Russell Cooper, Dean Corbae, Ken Hendricks and seminar participants at Rice, Texas, Texas A&M, IIOC 2007, ESAM 2007, Monash, Latrobe, Texas Tech and Midwest Macro 2008. *JEL classification:* C61, O31. *Keywords:* dynamic duopoly, Markov-perfect equilibrium, innovation, horizontal differentiation.

We compute Markov-perfect equilibria and study the effects on long-run innovation of changes in taste heterogeneity (transport costs) and firms' costs of relocating products. Innovation rises as the industry's long-run position moves toward products that are permanently co-located in the space of horizontal tastes. A large enough fall in taste heterogeneity will raise long-run innovation, while more costly product relocation lowers innovation if taste heterogeneity is high, and raises it otherwise.

1 Introduction

This paper aims at incorporating the dynamics of product location into the study of competition and innovation. A number of recent theoretical approaches have used infinite-horizon models of dynamic stochastic oligopoly to examine the competition-innovation relationship – see for example Budd, Harris and Vickers (1993), and Aghion et al. (2001). These contributions look at how oligopolies compete on investment in a single dimension of differentiation, i.e., cost reduction or quality improvement. Here we allow firms to invest in multiple dimensions. We study innovation in a discrete-time game of repeated price competition between two firms who may invest in both vertical and horizontal differentiation. Using numerical techniques we elicit the relationship between equilibrium long-run innovation and the primitives of the model, namely the heterogeneity in consumer tastes and firms' costs of product relocations. To add to the 'vertical composition' effects uncovered by previous authors, we find that the effects of the primitives on the industry's horizontal composition – the proportion of time spent in horizontally differentiated states – can also be an important influence on innovation.[1]

Examples of firms' strategic use of horizontal differentiation may be found in many industries. Computer manufacturers change platforms (micropro-

---

[1] Interest in the competition-innovation relationship has a long heritage in economics, stretching back at least to the seminal work of Schumpeter (1942) and Arrow (1962). See Gilbert (2006) and Van Cayseele (1998) for surveys of the theory, and Cohen and Levin (1989) and Symeonidis (1996) for surveys of the empirics. See also Sutton (1998).

cessor architecture), radio stations change formats, retail clothing chains restock with different fashions, and so on. At the same time firms in these industries may invest in lower costs, or improved quality of service. Since these investments are jointly determined with their product locations, one might expect that, for example, changes in the costs of horizontal differentiation would have indirect effects on firms' quality investments. To capture both strategic dimensions we model a repeated Hotelling game between two firms each selling a single perishable good. Stage-game profits depend on firms' relative qualities and product locations, which in turn depend on firms' per-period investment choices. These investments have stochastic outcomes. We restrict firms to locating at either end of the Hotelling line, but investment in location shifts may see firms in any given period operating at the same location (i.e., producing homogeneous products), or at different locations (i.e., producing differentiated products). Rivalry in quality upgrades takes the form of a quality ladder in which, for simplicity, a firm can be in one of three states – one step ahead, neck-and-neck, or one step behind. The size of the step is an exogenous parameter representing the returns to successfully drawing ahead in the quality race.

We seek symmetric pure-strategy Markov-perfect equilibria (MPE's) of the infinite-horizon dynamic game of investment in location choices and quality improvement, in which the industry's current state is the configuration of firms' qualities and locations. Solving the game with the algorithm of Pakes and McGuire (1994, henceforth 'PM'), we study the effect on long-run innovation of changes in the model parameters. Since it is fairly intuitive that innovation increases when we raise the size of the 'quality step', we generally fix this parameter and focus on changes in two other quantities: consumer taste heterogeneity and the costs of location switching.[2]

Taste heterogeneity, or 'product substitutability', is represented by a Hotelling transport-cost parameter. The effects of changes in product substitutability on a duopoly's long-run innovation rate were previously studied in Aghion et al. (2001). Working in continuous time with a CES representative-consumer model of preferences, they identified two countervailing effects of greater substitutability on long-run innovation. The first is the *escape-competition* effect, whereby the increased incremental returns to gaining an advantage over one's rival lead firms to invest more when they are neck-and-neck in quality (or cost) competition. The second is the *vertical composition* effect, whereby, as a result of intense investment competition between neck-and-neck firms, the industry spends more time overall in states of asymmetric quality, where innovation is relatively lower. As products become more substitutable these two effects exert opposing influences on long-run innovation, but Aghion et al. argue that for 'most parameter values' the former effect will dominate, leading to a rise in innovation overall (p. 470).[3]

---

[2] Since innovation rates dwindle to nothing as the quality step goes to zero, we consider quality steps that are not too small.

[3] In Aghion et al. (2001) firms invest in cost reduction rather than quality improvement. However there is a direct correspondence between their results and those of a quality-ladder model, since in either case profits depend solely on the difference between firms' current

One disadvantage of the CES framework is that changes in the elasticity of substitution compound two distinct types of changes that might be observed in real-world industries. Products might become more substitutable because consumer tastes are more homogeneous (lower transport costs), or because firms change positions in the product space. The framework presented below allows us to separate these distinct effects, and to study their interactions.

As a first step in this direction, in a companion paper (Narajabad and Watson 2008) we studied innovation in a continuous-time dynamic duopoly in which a heterogeneous-agent Hotelling model of consumer preferences replaces the CES framework of Aghion et al. (2001). Firms' locations are fixed at opposite ends of the Hotelling line. We showed that the two effects identified by Aghion et al. are still present, and that long-run innovation is always maximized when consumers have homogeneous tastes (i.e., zero transport costs) for the two goods. This latter result arises because homogeneous tastes imply that there is no competition between firms for the marginal consumer. Inframarginal consumers then get relatively less surplus, and more of the overall surplus in the market accrues to the quality leader, causing firms to innovate more because they are fighting over a bigger pie.[4]

In the present work we move to discrete time, and show that when product locations are endogenous a third effect arises, a *horizontal composition*

---

technological states – whether qualities or costs.

[4] A difference that arises in moving to Hotelling preferences is that the innovation rate will exhibit a U shape as tastes become more homogeneous, rather than the inverse U emphasized in Aghion et al. (2001).

effect. Take an industry where consumers have heterogeneous tastes (high transport costs), and where firms initially operate at separate locations in the product space. Reducing transport costs will eventually give the current quality leader an incentive to join the rival's location, in order to nullify the rival's local market power. The rival in turn will have an incentive to shift its product away, but as tastes become more homogeneous the quality leader's incentive will dominate, and over the long run the industry will spend more time in states where products are co-located. We find that in discrete time, as in the continuous-time model, these states are where the innovation rate is highest. Hence with endogenous locations the effect of a large enough fall in taste heterogeneity will be to raise innovation overall.

The effect on innovation of a rise in the switching cost depends on the degree of heterogeneity in consumer tastes.[5] If tastes are rather homogeneous then such a rise will tend to raise the long-run innovation rate, because it eventually pushes the industry into a state of permanent co-location where innovations are frequent. On the other hand if tastes are rather heterogeneous then a rise in the switching cost will tend to reduce the long-run innovation rate, because it eventually pushes the industry into a state of permanent

---

[5] The switching cost could represent the resources required to adjust product attributes or learn about new markets. While we do not explicitly model intellectual property, this parameter might also reflect the cost of gaining and defending a new patent (when a firm differentiates away from its rival), or of overcoming an existing patent (when a firm moves closer to its rival). For simplicity we make the switching costs non-directional – they are held to be the same whether a firm is moving away from, or closer to, its rival.

separation, where the incentives for quality investment are lower. Thus the effects on innovation of policies designed to make it easier (or harder) for firms to enter new product markets may depend crucially on market-specific factors such as the heterogeneity in consumer tastes.

As a qualification to the above findings we note that when product locations are endogenous there is a potential selection effect in the observed relationship between innovation and horizontal differentiation. The urge to avoid intense competition in investment and prices may lead firms to bypass the state where equal-quality firms sell undifferentiated goods. This state is where the industry innovation rate would be highest. Conditioning long-run innovation rates on the location state, we may then find that horizontally differentiated firms innovate more, not less – the point is that such an empirical finding could arise from the endogeneity of product locations, and need not represent a causal relationship.[6]

Analysis of our model is complicated by the existence of multiple equilibria. Numerical searches throw up many such cases – fortunately we were able to reduce almost all such cases to a single prediction by applying a few simple criteria.[7] First, multiple equilibria create no problems for our

---

[6] The pattern of equilibrium state transitions resembles the industry dynamics in Bresnahan and Greenstein (1999), who describe the phenomenon of 'indirect entry' in the computer industry, whereby new platforms first avoid direct competition with existing platforms by operating in uncontested market segments.

[7] We have no analytic results that conclusively demonstrate uniqueness at any point in the parameter space. In fact Besanko, Doraszelski, Kryukov and Satterthwaite (2010)

analysis if they predict the same ergodic distribution. Second, we assume that Pareto-dominated equilibria are not played. Third, we assume that the quality gap between a leader and a laggard is not too small. This has the effect of ruling out 'coordination failures' in investment in the neck-and-neck quality states. After applying these criteria we are left with one small area of multiplicity. This area has a clear intuition which we explain in the appendix by providing a complete solution to a *deterministic* dynamic infinite-horizon game of location choices. As far as we are aware this solution appears for the first time here, and thus may be of independent interest.

Our model is related to the literature on endogenous product locations in spatial price equilibrium. Most of the main contributions in this area (e.g., D'Aspremont, Gabszewicz, and Thisse 1979, De Palma et al. 1985) deal with symmetric (equal-quality) firms. A notable recent exception is Vogel (2008). We are not aware of any previous treatments in this literature of the infinite-horizon dynamic location problem. Caulkins et al. (2007) study dynamic product relocations by a firm seeking to maintain horizontal separation from its rivals, but in that model there is no quality (or cost) investment, and the firm's rivals do not behave strategically. Also related to our approach is

---

have shown that, because of instability, the PM algorithm may fail to find a substantial portion of the equilibria that exist in cases of multiplicity. Thus, where we claim a single, or 'unique' prediction at any given parameter values, this is in the weak sense of a sole equilibrium remaining after a numerical search with various starting values, including random starting values and equilibria observed at nearby points in the parameter space, and after further applying the selection criteria referred to in the text.

the growing literature which uses the PM algorithm to compute equilibria in various forms of dynamic oligopoly. See Doraszelski and Pakes (2007) for a survey and Doraszelski and Satterthwaite (2010) for a discussion of related existence issues.[8] As far as we are aware none of these computational papers deal with endogenous location choices.[9]

Sections 2 through 4 respectively cover the stage game, the dynamic model, and the computation of equilibria. Equilibrium state-conditional investments and their implications for the innovation rate are explained in sections 5 and 6. In section 7 we discuss the case of fixed firm locations in more detail, drawing comparisons with the equivalent continuous-time model. Section 8 discusses the empirical possibility of observing a higher location-conditional innovation rate among separated firms than among co-located firms. Section 9 concludes. An online appendix posted on the journal's website draws on results from a deterministic dynamic model to discuss the

---

[8] Recent examples from this literature include Markovich (2008) and Goettler and Gordon (2009) – as in our work the latter paper allows each firm two control variables.

[9] For recent empirical work on the dynamics of product relocations, see Sweeting (2007) and Bernard, Redding and Schott (2010). There is an extensive empirical literature on the role of location-dependent spillovers in fostering R&D, along with some related theoretical contributions, such as Long and Soubeyran (1998) and Piga and Poyago-Theotoky (2005). These latter papers study the tradeoff between spillovers and competition when firms choose whether or not to locate near rivals. Since we assume location-independent spillovers, our analysis of endogenous horizontal differentiation addresses a different tradeoff – that between the development of local market power and the full exploitation of quality advantages over rivals.

multiplicity problem in more detail.

2 The stage game

The stage game is a standard Hotelling duopoly where two firms $i = 1$ and 2 may locate at either end (possibly the same end) of the unit interval. Let $a_i \in \{0, 1\}$, $i = 1, 2$, denote a firm's location. Consumers are uniformly distributed between 0 and 1 and have unit demands and quadratic transport costs with parameter α.[10] Each firm i produces a single non-durable good at constant marginal cost of c_i. For simplicity let $c_i = 0$.

The firms' goods each have an idiosyncratic quality q_i. A consumer located at $\tau \in [0, 1]$ who purchases a good of quality q at price p from a firm located at $a \in \{0, 1\}$ realizes net utility of $z = -\alpha(a - \tau)^2 + q - p$. Consumers are fully informed about prices and qualities in the current period and purchase from the firm offering the highest net utility. We will assume throughout that $q_i > 3\alpha$ for both firms $i = 1, 2$. This ensures that the market is always covered in equilibrium.

Denote the difference between the qualities of the firms as $\delta \equiv q_2 - q_1$, an element of a finite set Δ of possible such differences. Define $s \in S \equiv \{0, 1\}$ as the industry's horizontal differentiation state (or 'location state'), with $s = 0$ and $s = 1$ respectively denoting 'firms are together' (i.e., $(a_1, a_2) = (0, 0)$ or $(1, 1)$) and 'firms are apart' (i.e., $(a_1, a_2) = (0, 1)$ or $(1, 0)$). Consider the

---

[10] All results to follow continue to hold if transport costs are linear.

Nash equilibrium of the one-shot simultaneous price-setting game for given locations and product qualities. Let π_i denote the equilibrium profits of firm $i = 1, 2$. If the firms are apart ($s = 1$), we have:

a. if $\delta > 3\alpha$, $\pi_1 = 0$, $\pi_2 = \delta - \alpha$.

b. if $-3\alpha \leq \delta \leq 3\alpha$,

$$\pi_1 = \frac{\alpha}{2}[1 - (\frac{\delta}{3\alpha})]^2 \ , \ \pi_2 = \frac{\alpha}{2}[1 + (\frac{\delta}{3\alpha})]^2 \ .$$

c. if $\delta <$ -3α, $\pi_1 = -\delta - \alpha$, $\pi_2 = 0$.

Cases (a) and (c) here represent situations where one firm's quality advantage is sufficient to allow it to sell to the whole market in equilibrium. When firms are located together ($s = 0$) we have the Bertrand outcome for homogeneous goods:

a. if $\delta > 0$, $\pi_1 = 0$, $\pi_2 = \delta$,

b. if $\delta = 0$, $\pi_1 = \pi_2 = 0$,

c. if $\delta < 0$, $\pi_1 = -\delta$, $\pi_2 = 0$.

It can be seen that firm 2 (resp. firm 1) gets a higher payoff in the 'co-located' equilibrium than in the 'separate' equilibrium if and only if $\delta > 3\alpha(2 - \sqrt{3}) \approx 0.804\alpha$ (resp. $\delta < -3\alpha(2 - \sqrt{3})$). Note also that, whether firms' locations are separate or coincident, equilibrium profits depend only on the difference in product qualities δ, not directly on the quality levels themselves. This observation holds because the market is by assumption covered in equilibrium, and it allows us to adopt a 'quality ladder' approach in what follows.[11] Furthermore the symmetric distribution of consumers in the market implies that the horizontal differentiation between firms affects stage-game payoffs only through the location state s. In keeping with the 'payoff-relevant' restriction on state variables in the MPE concept (Maskin and Tirole 1988a,b), we therefore define the industry's state at the start of the stage game to be a vector $(s, \delta) \in S \times \Delta$.

Note that if $s = 1$ and firm 2 advances its quality differential from $\delta = -k < 0$ to $\delta = 0$, it receives incremental profits of:

$$K_{L1} = \begin{cases} \frac{\alpha}{2} & \text{if } k > 3\alpha \text{, or} \\ \frac{k}{6}\left[2 - \frac{k}{3\alpha}\right] & \text{if } k \leq 3\alpha. \end{cases}$$

If $s = 1$ and it advances its quality differential from $\delta = 0$ to $k > 0$ it receives incremental profits of:

$$K_{H1} = \begin{cases} k - \frac{3\alpha}{2} & \text{if } k > 3\alpha \text{, or} \\ \frac{k}{6}\left[2 + \frac{k}{3\alpha}\right] & \text{if } k \leq 3\alpha. \end{cases}$$

For all α and $k > 0$ the incremental profits of the quality leader are decreasing in α, while those of the quality laggard are increasing in α, i.e., $dK_{H1}/d\alpha < 0$

---

[11] The absence of an outside good is a key simplification here. Many applications of the PM algorithm allow for an outside good with stochastically evolving quality, and (in contrast to the present model) assume a decreasing marginal utility of quality in order to keep the quality differential between the inside and outside goods within manageable bounds. Note also that utility is linear in quality, which is necessary for the result that profit depends only on the difference in product qualities, not on their levels.

and $dK_{L1}/d\alpha > 0$. If $s = 0$ then stage-game profits are independent of α and we have $K_{L0} = 0$ and $K_{H0} = k$.

It is useful to define the 'leader's surplus' (LS) as the difference between the maximum and minimum stage-game profits, conditional on s. Thus, where $k > 0$ is the maximum possible value of δ, define $LS_1 \equiv \pi_2(s = 1, \delta = k) - \pi_1(s = 1, \delta = k) = K_{L1} + K_{H1}$, so that:

$$LS_1 = \begin{cases} k - \alpha & \text{if } k > 3\alpha, \text{ or} \\ \frac{2}{3}k & \text{if } k \leq 3\alpha. \end{cases}$$

Also $LS_0 \equiv \pi_2(0, k) - \pi_1(0, k) = K_{L0} + K_{H0}$, which implies $LS_0 = k$ for all α. The leader's surplus is effectively the size of the pie over which firms are fighting in the quality race, conditional on there being no change in s.

3 The dynamic game

Turn then to dynamic extensions of this Hotelling game in which firms may invest to change product locations and improve their future qualities – such investments have stochastic outcomes. Time is discrete and the horizon is infinite. Future profits are discounted at the rate β. We split each discrete time period into two stages: one in which firms simultaneously choose investments in horizontal differentiation, followed by a stage in which they simultaneously invest in quality improvement. Outcomes of investment choices in a given stage are realized and observed at the end of that stage. Let $s_h(t) \in S$ denote the industry's differentiation state at the beginning of the first (horizontal-differentiation) stage in period t, and let $s_v(t)$ denote this state at the beginning of the second (quality-investment) stage. Let $\delta(t) \in \Delta$ denote the industry's quality differential at the beginning of period t – since this will still be the quality differential at the beginning of the second stage in that period there is no need to put a stage subscript on $\delta(t)$. For simplicity we restrict Δ to just three elements: $\Delta \equiv \{-k, 0, +k\}$, where $k > 0$ represents the size of a single step in the transitions of the quality differential δ. At the beginning of any stage in period t the industry's quality differential can thus be in one of three states: firm 1 is one step ahead ($\delta(t) = -k$), the firms are of equal quality ($\delta(t) = 0$), or firm 2 is one step ahead ($\delta(t) = k$). By varying k we vary the value of leadership in the firms' quality competition.[12]

Let $\omega_l(t) \equiv (s_l(t), \delta(t)) \in S \times \Delta$ denote the overall state of the industry at the beginning of stage l in period t, for $l = h$ or v, with initial values $\omega_h(0) \equiv (s_h(0), \delta(0))$. We assume that each firm's quality starts at $q_i(0) > 3\alpha$: since there is no depreciation of quality the market will then be covered in any equilibrium of the Hotelling stage game. Within the first (location-investment) stage in any period t the sequence of moves is as follows:

a. firms observe the current state $\omega_h(t)$ and simultaneously choose investment levels.

b. stage-game payoffs are realized given the current state.

---

[12] Note that the quality step k can be interpreted as a marginal utility of quality in the consumer's utility function z.

c. the investment plans chosen in (a) are implemented, their outcomes (location switches) are realized, and the costs of the plans are incurred.

d. play moves on to the second stage, with the new state $\omega_v(t)$ determined by the realizations of the investment plans implemented in (c).

In the second (quality-investment) stage the same sequence applies, with the exception that there is no step (b), because stage-game payoffs are only received once per period. The current state in step (a) would be $\omega_v(t)$, and after the final step the new state would be $\omega_h(t+1)$. We assume that firms only discount future payoffs across periods, not across stages within a period.

The policy function for each firm is in two parts, corresponding to the two stages in each period. Let $h_i : \Delta \times S \rightarrow R_+$ denote the first part, the dollar value of investment in location changes (or 'switching') by firm i. If i invests h in a given period, the investment is successful with probability

$$A(h) \equiv \frac{\exp(-\gamma)h}{1+\exp(-\gamma)h} \ , \tag{1}$$

where γ is a cost parameter. Success in switching investment means that the firm changes its location from $a_i = 0$ to $a_i = 1$, or vice versa. As γ increases the dollar cost of implementing any given probability of switching success also increases.

Whether switching success translates into a change in the industry's differentiation state s depends on the outcome of the other firm's switching investment. If both succeed in switching then there is no change in s; if one switches and the other does not then s changes from 0 to 1, or vice versa.

We assume that the stochastic outcomes of firms' simultaneous switching investments are independent. Where the differentiation state is represented s for this period and s' for next period, we thus have:

- $s' = s$ with probability $(1 - A(h_1))(1 - A(h_2)) + A(h_1)A(h_2)$.
- $s' = 1 - s$ with probability $A(h_1)(1 - A(h_2)) + (1 - A(h_1))A(h_2)$.

A firm's own quality q_i advances from period to period in discrete increments of k units per step. These increments may be generated firstly by the firm's own investment in quality improvement, which is the second part of its policy function, a function $v_i : \Delta \times S \rightarrow R_+$ indicating a dollar value invested in better quality. Second, if the firm is lagging in the quality race it may also benefit from spillovers from the leader's technology. In particular, we assume that a firm which has a one-step quality advantage cannot realize any further improvements without automatically conferring the same increment on the quality of its rival. Our motivation for this restriction is that the low-quality firm can reverse engineer the products of the high-quality firm, and thus avoid falling too far behind. However there are limits to the efficacy of this reverse engineering, and therefore quality differences do not evaporate instantaneously and may persist over time. When a firm has achieved a quality advantage of $\delta = k$ it may nevertheless continue to invest in quality improvements, not to extend its advantage, but to prevent its lower-quality rival from catching up.

Formally, let $\tilde{q}_i(t)$ be a random variable representing increments in firm

i's quality at time t arising *from its own investment efforts.* For a firm that has invested an amount $v \geq 0$ in quality improvement at t this increment is equal to $+k$ with probability

$$Q(v) \equiv \frac{\exp(-\phi)v}{1+\exp(-\phi)v} \ , \tag{2}$$

and is equal to zero otherwise. Here ϕ parameterizes the cost of implementing any given probability of success in quality improvement. We assume that, conditional on firms' actual quality investments in a given period, the random variables $\tilde{q}_i(t)$ and $\tilde{q}_j(t)$ are independent (and also i.i.d. over time).

If a firm i is not a quality laggard, i.e., if $q_i(t) \geq q_j(t)$, then its realized quality increment in period t conditional on $v_i(t)$ will be $q_i(t+1) - q_i(t) = \tilde{q}_i(t)$. However if the firm is a low-quality producer, $q_i(t) < q_j(t)$, then its quality can be improved not just by its own efforts but, alternatively, by spillovers from the innovations of its rival. That is, $q_i(t+1) - q_i(t) = \tilde{q}_i(t) + \mathbb{1}(\tilde{q}_i(t) = 0)[\tilde{q}_j(t)]$, $i \neq j$, where $\mathbb{1}$ represents the indicator function. Note that this specification allows the laggard i to benefit from positive realizations of either $\tilde{q}_i$ or $\tilde{q}_j$, but not both.

Since we are solving for a stationary equilibrium, we henceforth omit time indices and instead represent next-period values of state values as, e.g., δ'. Let $\theta_{(v_1,v_2)}(\delta, \delta')$ denote the probability of transiting from quality differential δ to δ', given quality investments v_1 and v_2. We then have

$$\begin{aligned}
\theta_{(v_1,v_2)}(k,k) &= 1 - Q(v_1)(1-Q(v_2)) \\
\theta_{(v_1,v_2)}(k,0) &= Q(v_1)(1-Q(v_2)) \ .
\end{aligned}$$

Probabilities for transitions $(-k,-k)$ and $(-k,0)$ are formed symmetrically. We also have

$$\begin{aligned}
\theta_{(v_1,v_2)}(0,k) &= Q(v_2)(1-Q(v_1)) \\
\theta_{(v_1,v_2)}(0,-k) &= Q(v_1)(1-Q(v_2)) \\
\theta_{(v_1,v_2)}(0,0) &= 1-\theta_{(v_1,v_2)}(0,k)-\theta_{(v_1,v_2)}(0,-k)\ ,
\end{aligned}$$

and $\theta_{(v_1,v_2)}(k,-k)=\theta_{(v_1,v_2)}(-k,k)=0$, because firms cannot improve their quality by more than one step per period.

A firm's strategy is a vector $u_i(\omega_h,\omega_v)\equiv(h_i(\omega_h),v_i(\omega_v))$ of (non-negative) real-valued functions of ω_h and ω_v. Firm i's best response to a rival's strategy $u_j(\omega_h,\omega_v)$ will be the policy function generated by the two-part Bellman equation:

$$\begin{aligned}
V_i(\omega_h) &= \max_{h_i\in R_+}\{\pi_i(\omega_h)-h_i+\mathrm{E}_{\omega_v}[W_i(\omega_v)\mid\omega_h,h_i,h_j]\}\ , &(3)\\
W_i(\omega_v) &= \max_{v_i\in R_+}\big\{-v_i+\beta\mathrm{E}_{\omega_h'}[V_i(\omega_h')\mid\omega_v,v_i,v_j]\big\}\ . &(4)
\end{aligned}$$

Here $V_i(\omega_h)$ is firm i's value function in the first 'switching' stage and $W_i(\omega_v)$ is the value function for the second 'quality' stage. As noted above, we assume that there is no discounting between stages, so that the discount factor β only appears in (4). Stage-game payoffs only accrue in the first stage. The dependence of $V_i(\omega_h)$ and $W_i(\omega_v)$ on u_j is suppressed for brevity.

We seek an MPE in symmetric pure strategies. That is, we are looking for a pair of policy functions $u_1(\omega_h,\omega_v)$, $u_2(\omega_h,\omega_v)$ and a pair of value functions $(V_1(\omega_h),V_2(\omega_h))$ such that, for $i,j=1,2$, $i\neq j$:

a. $V_i(\omega_h)$ solves (3), given $u_j(\omega_h, \omega_v)$, where $W_i(\omega_v)$ is defined as in (4);

b. given $u_j(\omega_h, \omega_v)$, the components $h_i(\omega_h)$ and $v_i(\omega_v)$ of $u_i(\omega_h, \omega_v)$ are policy functions generated by $V_i(\omega_h)$ and $W_i(\omega_v)$ respectively; and

c. for every $(\omega_h, \omega_v) = ((s_h, \delta), (s_v, \delta))$, we have $u_2(\omega_h, \omega_v) = u_1(\tilde{\omega_h}, \tilde{\omega_v})$, where $\tilde{\omega_h} = (s_h, -\delta)$, $\tilde{\omega_v} = (s_v, -\delta)$.

Condition (c) is the symmetry requirement: firm 2's chosen action when the differentiation state is s and his quality advantage is $q_2 - q_1$, should be the same as firm 1's chosen action given s and the same quality advantage.

Existence of MPE's in dynamic oligopoly models of the present type (specifically, in the framework introduced by Ericson and Pakes 1995) is extensively discussed in Doraszelski and Satterthwaite (2010), building on Whitt (1980). Their argument, which can be suitably adapted to fit our two-stage framework, has a particular requirement of uniqueness in firms' optimal investment choices. This uniqueness follows here from the concavity in h and v of the success probabilities in (1) and (2). To see this, take for example investment by firm 2 in the second stage in state $\omega_v = (s_v, \delta)$. Let

$$\begin{aligned} x_2(\omega_v; V_2, u_1) \quad \equiv \quad & \max\{0\ ,\ Q(v_1(\omega_v))[V_2(s_v, \delta) - V_2(s_v, \max(\delta - k, -k))] \\ & +(1 - Q(v_1(\omega_v))[V_2(s_v, \min(\delta + k, +k)) - V_2(s_v, \delta)]\}\,. \quad (5) \end{aligned}$$

This quantity may be thought of as firm 2's expected incremental benefit from successful quality investment – the min and max operators inside the braces take account of the upper and lower bounds on the quality differential

δ. Given V_2 and u_1, firm 2's unique optimal quality investment in state ω_v is then:

$$v_2^*(\omega_v; V_2, u_1) = \max\left\{0, \exp(\phi)[\sqrt{\beta x_2(\omega_v; V_2, u_1)\exp(-\phi)} - 1]\right\} . \qquad (6)$$

A firm's optimal switching investment may be characterized similarly.[13]

As outlined above the model has four parameters of interest: k, α, ϕ and γ. (We hold the discount factor β to be fixed at 0.95 throughout.) Alternatively, express the parameters as the vector $\xi \equiv (k, \alpha, \exp(\phi), \exp(\gamma))$. When represented in this form it may be seen that one of these parameters can be normalized, because scaling all four up or down by the same factor results in scaling of the value functions and investments and no changes to equilibrium transition probabilities and ergodic distributions. Hence we normalize ϕ to zero w.l.o.g.

4 Analysis of equilibria

The model outlined above is too complicated to solve analytically. This is obviously not a virtue of the model but it is unavoidable given the different elements that we want to incorporate. Instead we rely on the PM algorithm, adapted to our two-stage framework, to solve for the symmetric pure-strategy

---

[13] Let firm 2's expected incremental benefit from switching be denoted $\overline{x}_2 \equiv \max\{0\ ,\ A(h_1(\omega_h))[W_2(s_h,\delta) - W_2(1-s_h,\delta)] + (1-A(h_1(\omega_h)))[W_2(1-s_h,\delta) - W_2(s_h,\delta)]\}$. Then the expression for the optimal h_2^* is analogous to that in (6), with β deleted, γ replacing ϕ, and $\overline{x}_2$ replacing x_2.

MPE numerically.[14] Given the small numbers of firms and states in our model the computational burden of this algorithm is usually not excessive. When coded in Matlab 7 the algorithm typically takes a few seconds to calculate an equilibrium to a tolerance of 10^{-10} on a PC with a 2.7GHz processor. If an MA(4) dampening procedure (or variation thereof) is used (see Judd 1998) the algorithm is fairly stable and does not often fail to converge. An exception is when switching costs are extremely low and taste heterogeneity is high. In such cases the algorithm often has difficulty converging and we were not always able to find an equilibrium. Our numerical analysis therefore avoids this region of the parameter space, effectively assuming that firms' switching costs exceed some (very low) minimal level.

Figure 1 shows the various types of ergodic set that arise in equilibrium at different parameter values. In this figure the taste heterogeneity parameter α is set at 1. The parameter γ determining the difficulty of switching investment is on the vertical axis. On the horizontal axis is k, the size of the leader's quality advantage.

The equilibria are classified into four types according to the nature of the absorbing differentiation state (or states). At the top of figure 1 is a region of no movement (area I – indicated by bold dots). In this region γ is high –

---

[14] Briefly, we start with initial guesses for the equilibrium value and policy functions, plug these into the RHS of (4), solve, plug into the RHS of (3), solve, and repeat until convergence. The algorithm is a Gauss-Jacobi method: at each iteration we only update the value functions after computing optimal investments at all states.

switching locations is expensive, so no-one ever moves, i.e., $h_i(\cdot) = 0$ for all ω_h, $i = 1, 2$. On the right-hand side of the figure, crosses (area III) represent equilibria where co-location is the only absorbing differentiation state. Here switching costs are moderate and k, the leader's advantage, is relatively large, which encourages a quality leader to move close to the laggard when $s = 1$. (Recall from the discussion in section 2 that a quality leader gets higher stage-game profits when co-located than when separated if and only if $k > 0.804\alpha$.) Since k is high the laggard has little to gain from moving away – indeed if $k \geq 3\alpha$ then the laggard gets zero stage-game profits regardless of where he is located. Then he does not respond to the leader's moves by switching, and as a result the firms will always end up together in the long run.

If switching costs fall to a low enough level then the industry moves from the co-location outcome into area IV, represented by the squares, where there is no absorbing differentiation state. That is, conditional on $s = 0$ or 1, along any equilibrium path there is always a strictly positive probability of reaching the other differentiation state $1 - s$. For example when $s = 1$ the quality leader wants to switch so as to join the laggard and increase its stage-game payoffs. But if he succeeds in doing this, thereby changing the differentiation state to $s = 0$, the laggard may then want to move away. The result (once shifts in the quality differential are also factored in) is endless cycling through all possible states. Finally when the quality differential k is low enough (and the costs of switching are not too high) a quality leader will lose his incentive to co-locate with his rival, since his stage-game payoffs are higher when they are separate. Then both players want to differentiate, regardless of their qualities, and $s = 1$ is the only absorbing differentiation state. In the figure such outcomes are represented by x symbols – area II.

Figure 2 represents the same categories of equilibrium outcome, but with γ fixed and $-\log(\alpha)$ on the vertical axis. We include this figure to show the path taken by the industry through the different equilibrium outcomes as α changes. Increases in $-\log(\alpha)$ represent lower transport costs, i.e., less taste heterogeneity, but note that the figure plots $-\log(\alpha)$ against k/α, not against k.[15] Therefore as α increases from a low level (with k fixed) the industry follows a path along a curve such as that shown, moving to the southwest as α increases. A *ceteris paribus* increase in k shifts the industry downwards, to a lower such curve.

Several problems of multiplicity in this model must be addressed before we can proceed to our analysis of the innovation rate. First, if there are zeroes in investment then a given MPE may have multiple ergodic sets. The case arising from zeroes in switching investment has already been explained as area I in figure 1. Not represented in figure 1 are cases of zeroes in quality investment. For example if $v_i(.) = 0$ in all states, $i = 1, 2$, then each value of δ in Δ is itself an absorbing quality state. Such cases may arise when k, the

---

[15] We plot the figure this way because stage-game profits in any state can be expressed as a function just of α and k/α (see section 2), and for consistency with figure **??** in the appendix, which has δ/α on the horizontal axis.

leader's quality advantage, is small. From the present point of view they are uninteresting, since they trivially imply zero innovation. Therefore we rule them out of the subsequent analysis by restricting attention to values of k that are not too small.

Somewhat thornier questions are raised by multiple equilibria, in particular when these imply multiple ergodic sets. To look for cases of multiplicity we applied the PM algorithm over a grid of 4,225 points in (k, γ, α) space. Initially feeding in eleven different starting values per point (six pre-set, five random), we then took all equilibria found and used them as starting values for further searches at neighboring points, iterating through the parameter space until no new equilibria could be found. This method of search throws up cases of multiplicity falling into four different categories, which we deal with in turn below.

We should note that, while our method takes some care to produce a list of possible equilibria that may be 'reasonable' candidates at any given set of parameter values, we have no analytical results to guarantee that it does in fact completely enumerate the set of all Markov-perfect equilibria. In what follows, claims of 'uniqueness' or 'elimination of multiplicity' should therefore be interpreted in the rather loose sense of 'apparent uniqueness within a set of intuitively-reasonable equilibrium classes that we have been able to find with the PM algorithm'. We particularly have in mind here the limitations of the PM algorithm identified by Besanko et al. (2010). Those authors used homotopy techniques to show cases of multiplicity in dynamic oligopoly where the PM algorithm is unable to compute a substantial portion of the equilibria. (Specifically they find equilibria to which the PM algorithm will never converge because it is unstable at those equilibria.) Such cases may arise in the present application, but we leave a full homotopy-based investigation of these as an interesting topic for future work. Instead we simply identify a few classes of equilibria for which there is clear intuition, and at each point in the parameter grid use continuity arguments to identify which of these classes might be reasonable candidates for equilibria at neighbouring points. Furthermore the random-starting-values part of our search method gives us some probability of also turning up completely unexpected equilibria, not otherwise indicated by observations at nearby parameter values.[16]

Our four different categories of multiplicity are discussed in detail in the appendix, along with a simpler deterministic dynamic model which gives some of the intuition for the multiplicity. That discussion shows that one of the cases is of no consequence here because both equilibria imply the same ergodic distribution over quality states, and a second one is a pathological case that can be eliminated by imposing a Pareto-superior criterion. Of the two remaining cases the first occurs as a result of coordination failures in

---

[16] Of course the results of Besanko et al. (2010) would indicate that even with many random starting values there is effectively zero probability of converging to any of the PM-unstable equilibria (if such exist). Nevertheless it would seem that under mild regularity restrictions one could construct an argument whereby an increased density of random starting values at least raises the probability of catching all the PM-stable equilibria (leaving aside the question of whether PM-stability is a useful economic restriction).

quality investment. Consider for example the case of equal-quality firms located apart ($\omega_v = (1,0)$). Suppose for simplicity that they do not invest in switching. One equilibrium may be for neither firm to invest in quality improvement, in which case $\omega_v = (1,0)$ is an absorbing state – both firms 'cruise', incurring no investment costs and earning stage-game payoffs of $\alpha/2$ per period. There may also be an equilibrium where both firms invest in quality improvement in this state – given that one's rival invests something it will be optimal for the own firm to do likewise. Such cases of multiple equilibria in quality investment are frequently observed in the present model when k is low. As k rises such cases will disappear because firms start to 'defect' from the zero-investment strategy profile. Therefore we rule them out by assuming again that k is not too low – in particular we assume $k \geq 1$ (i.e., $k/\exp(\phi) \geq 1 \Leftrightarrow \log(k) \geq \phi$).

The final case of multiple equilibria involves overlaps between the edges of the four areas in figures 1 and 2. Those figures indicate such cases by superposing the symbols representing the different types of ergodic set. For example in figure 1 at $\gamma = 2$ and $k = 0.8$ there is an equilibrium with no switching at all and an equilibrium where firms end up separated. We found that in many such cases one of the equilibria Pareto-dominated the other. In such cases we restrict our attention to the single Pareto-dominant profile. Some overlaps survived this process of deletion – of these survivors (which are the only ones shown in figures 1 and 2) some occur at low values of k, and will be ruled out by the assumption $k \geq 1$ mentioned previously.

Finally we are left with just one ‘troublespot’ – the overlap between areas II and III, represented by, e.g., the point $\gamma = 0.5$, $k = 2$ in figure 1, where firms end up separated in one equilibrium and co-located in another (and neither outcome is Pareto-dominant). The multiplicity in this region of the parameter space has a clear intuition which is explained in the appendix. In presenting our numerical results below we generally gloss over this region of multiplicity by choosing parameter values that avoid it. Our discussion centers on how the industry transitions from the interior of one area with a ‘unique’ ergodic set to the interior of another region of ‘uniqueness’, without special regard to the details of what happens at the boundary between these regions. Any multiplicity along these boundaries will not affect our main conclusions, as long as it is restricted to a relatively small area of the parameter space.

In summary, subject to the above caveats concerning our search procedure and the stability of the PM algorithm, we argue that we can ensure ‘uniqueness’ of the equilibrium ergodic set over a relevant region of the parameter space by assuming:

a. $k \geq 1$,

b. Pareto-dominated equilibria are not played.

We also assume that $\log(\alpha) < 4 + \gamma$, i.e., that the cost of switching is not too low relative to the heterogeneity in consumer tastes. As noted previously, this keeps us out of a region where the algorithm has trouble converging.

Some of the computations represented in the figures to follow violate this restriction, but they simply represent particular cases where we did succeed in computing equilibria in this region of the parameter space.

5 Equilibrium investments

5.1 Quality investments

Figures 3 and 4 show equilibrium quality investments, represented as state-conditional probabilities of success in quality improvement, conditional on $s = 1$ and $s = 0$. The bottom, middle, and top panels in each figure respectively show success probabilities for the three states $\delta = -k, 0, +k$, as functions of α, for various values of γ, with k fixed at 4.45. These are the probabilities $Q(v_i^*)$ of improvement resulting from the firm's own efforts, defined in (2). For the laggard spillover probabilities are not included.

If the cost of switching γ is high then there is no movement between product locations – firms start out either separated or co-located, and stay that way forever. Quality investment in this case is represented by the solid line in each figure, for $\gamma = 5$. Under permanent homogeneity of products (figure 4) this solid line is constant in α, because stage-game profits are then unaffected by changes in taste heterogeneity. On the other hand if products are permanently differentiated then α does affect the equilibrium investments. Figure 3 shows investment in that case to be monotonic in all three quality states – increasing for $\delta = -k$, and decreasing for the other two states. This result is not general – at least for the laggard and for neck-and-neck qualities we have found examples where their equilibrium success probabilities are U-shaped in α rather than monotonic. However in all the cases that we have examined (with fixed locations) the slope in α of quality investment by the laggard is certainly more positive than for the leader (unless quality investment by the latter is zero everywhere).

This latter observation is a natural consequence of the convexity of stage-game profits in the quality differential δ. When α is low, stage-game profits are quite convex in δ, which tends to foster investment by a leader and suppress investment by a laggard. As α increases profits become less convex in δ, which raises the incentives for investment by the laggard and reduces those for the leader.[17] For a neck-and-neck firm the investment incentives are a weighted average of those faced by firms at the extremes. If positive amounts are invested at all values of δ, then equilibrium investment at $\delta = 0$ has a slope in α somewhere between the slopes for firms at $\delta = -k$ and $\delta = +k$.

Of particular note in figure 3 is the change in firms' investment incentives around α such that $k = 3\alpha$, i.e., around $\log(\alpha) = 0.39$. As noted in section 2, when taste heterogeneity falls below that level the quality leader gets the whole market, and as α falls the leader's surplus LS then *increases*, because the leader is able to extract more surplus from consumers. The figure shows

---

[17] The convexity of π_i can be expressed as $K_{H1} - K_{L1}$. Recall that K_{H1} (resp. K_{L1}) is everywhere decreasing (resp. increasing) in α.

that as $\log(\alpha)$ falls below 0.39 (with $\gamma = 5$), investments by the leader and by neck-and-neck firms both start to increase more steeply. Simultaneously the decline in investment by the laggard is attenuated (and at some parameter values may even begin to increase). Since firms are now fighting for a 'larger pie', the incentives to invest are enhanced.

At lower levels of the switching cost γ there may be some movement between product locations in equilibrium. To understand these cases recall from figure 2 the path traced through the different outcomes as α increases. At low α there is little differentiation between locations – hence quality laggards have little incentive to run away from the leader, and the absorbing state sees both firms co-located ($s = 0$). That is, we are in region III of the figure. On the other hand for high α firms have substantial potential market power if they locate separately. Then the absorbing state sees firms permanently separated. In figure 2 we would be in region II. As soon as firms reach this state $s = 1$ they cease all further location switching – hence quality investments conditional on $s = 1$ do not depend on γ. For any γ the value of α at which this situation of permanent separation is reached can be read from figure 3 – it is that α at which the probability-of-success curve merges with the solid curve.

For intermediate levels of α there may be no absorbing state in equilibrium – the industry may pass through region IV in figure 2. In this region there is enough heterogeneity in tastes to make it worthwhile for a laggard to run away from a co-located leader. At the same time there is not so much heterogeneity that the leader does not want to give chase. Although the relationship between quality investment and γ is not monotone, it appears that separated firms invest more at low γ – where product locations are in continual flux – than if their locations are fixed (as at $\gamma = 5$). Such behavior is consistent with the idea that firms innovate more when they are fighting over a bigger leader's surplus. We have $LS = k$ if firms are co-located, and $LS = 2k/3$ if they are separated. The more time the industry spends in co-location, the larger is the leader's surplus in the quality race, which feeds into higher investments by firms in the separated states. Conversely figure 4 shows that, at the levels of α and γ that induce switching, investment by co-located firms is in all three quality states lower than if locations are fixed.

5.2 Switching investment

Figures 5 and 6 show equilibrium state-conditional probabilities of success in switching investment, for the same parameter values as in figures 3 and 4, conditional on $s = 1$ and $s = 0$ respectively.[18] The middle panels show that firms of equal quality only invest in location switching when co-located. That is, we have $h_i(s_h = 1, \delta = 0) = 0$. Since our numerical search does not

---

[18] If α is high and γ is not too high (say $\log(\alpha) = 2$ and $\gamma = -0.5$) then the state $s = 0$ is off the path of long-run equilibrium, under which firms are permanently separated. As discussed in the appendix, there may be multiple equilibria in this case, distinguished by 'who moves away' in state $s = 0$. The figures plot equilibria where the laggard is the one to move in state $s = 0$.

completely exhaust the parameter space, the generality of this result is not certain. Nevertheless the underlying intuition is clear, since equal-quality firms earn $\pi_i = 0$ when co-located.

Explanations for the top panel in figure 5 and the bottom panel in 6 are also straightforward: leaders switch in order to co-locate with rivals, while laggards switch in order to develop their own product niche. In both cases the investment is naturally decreasing in γ. Investment by the laggard is non-decreasing in α, reflecting his increased returns to local market power when tastes are more heterogeneous. Investment by the leader is highest at intermediate α, reflecting the shape of the profit differential $\pi_2(s = 0, \delta = +k) - \pi_2(s = 1, \delta = +k)$, which is maximized at $\alpha = k/3$, i.e., at $\log(\alpha) = 0.39$ (see Narajabad and Watson 2008). For higher α the leader prefers the current location state ($s = 1$), while at lower α a costly relocation is relatively less attractive, since the products are in any case more homogeneous.[19]

These two figures also show that if γ is sufficiently low then there may be investment in switching by laggards in state $s = 1$, and by leaders in state $s = 0$. Such investments may seem counterintuitive, in light of the

---

[19] It will be noted that when α is very low the leader may cease switching altogether in state $s = 1$ (e.g., if $\gamma = 1$ and $\log(\alpha) < -0.5$), because products are already very homogeneous. That is, travelling far enough to the northeast along the curves in figure 2 eventually takes the industry into a no-movement equilibrium of area I. Since α is very low in this area, there will be little difference between location states in the rates of quality improvement, as may be seen from figures 3 and 4. Hence the innovation rates conditional on $s = 0$ and $s = 1$ will be approximately equal.

incentives for horizontal differentiation discussed above. Closer examination reveals that the explanation lies in firms' incentives to pre-empt rivals' product relocations. For example in figure 5 at $\gamma = -2$ and $\log(\alpha) = 1$ the leader switches locations with around 50% probability. Given that switching is not too expensive, the laggard finds it optimal to forestall closer competition by attempting a product relocation of his own. A similar observation applies for $\gamma = -2$, $\log(\alpha) = 1.5$ in figure 6, where the laggard is moving away with probability in excess of 50%, and the leader attempts to follow him.

6 Innovation with endogenous product locations

Following previous authors, long-run innovation is defined as the average probability of an advance in the industry's frontier technology, under the equilibrium invariant distribution over the three quality differentials $\delta = -k, 0, +k$. An advance in the frontier technology means improvement by the leader, in states where $\delta = -k$ or $+k$, or by either firm in states where $\delta = 0$. The invariant distribution weights the probabilities of such advances by the amount of time the industry spends in each state over the long run.

Figure 7 shows innovation over a range of (α, γ) which covers all four regions in figures 1 and 2. The leader's quality advantage is fixed at $k = 4.45$. Note that if the switching cost is high enough (e.g., γ near 2), then there is no movement in equilibrium, in which case there are two invariant distributions, one for the absorbing state $s = 0$ and one for the absorbing state $s = 1$. In such cases (which correspond to area I in figures 1 and 2), figure 7 for simplicity only shows innovation for the absorbing state $s = 1$.

Two flattish areas in the figure constitute lower and upper bounds for the innovation rate, and respectively represent the cases of permanently separated and permanently co-located firms. Thus the upper area, where α is low and γ is not too high, corresponds to area III in figure 1. The lower area corresponds to areas I and II in that figure – it combines both areas because of the convention noted above of plotting innovation for $s = 1$ when there are two ergodic sets.

Consider then the trajectory of innovation as we vary α with the switching cost fixed at, say, $\gamma = -1$. Start from high taste heterogeneity, which puts us in area II in figure 2, where firms are permanently separated. As α is reduced, long-run innovation is initially fairly flat.[20] Moving to the northeast along the curve in figure 2, the industry eventually enters area IV, where location switching persists in the long run. Since this results in firms spending more time in co-located states, innovation leaves its lower bound and starts to increase. Eventually α is so low that there is little incentive for a laggard to move away from the leader. Then the industry reaches area III, where firms are permanently co-located in the long run and innovation is maximized.

---

[20] Innovation is not in general constant in α in this region. Although it appears flat in the figure, depending on the parameter values it may be monotone increasing, or U-shaped, as α falls.

Although innovation is not necessarily monotonic over its whole trajectory, as long as k is not too small a large enough fall in α raises innovation overall at all the parameter values that we have examined.

Now consider varying the switching cost at fixed levels of taste heterogeneity. From figure 7 it is clear that the effect on innovation will depend on α. Suppose that the switching cost and taste heterogeneity are both fairly low, e.g., $\log(\alpha) = 0.5$, $\gamma = -2$. Then the industry will initially be in area IV in figures 1 and 2, with continual changes in product locations in the long run. As γ rises with α fixed the industry heads for $s = 0$ as an absorbing state, spending longer and longer in co-location. Since the leader's surplus is higher in such states this leads to an overall rise in the innovation rate. On the other hand if taste heterogeneity is high ($\log(\alpha) = 1.25$, say), then an opposite effect is seen. Increases in the switching cost push the industry toward $s = 1$ as an absorbing location state, where long-run innovation is lower than under permanent co-location.

The distinction between these two cases is striking. It suggests that the long-run effects of policies designed to enhance firm mobility depend on, among other factors, the heterogeneity in consumer tastes. Such policies might, for example, include measures designed to lower the costs of entry or adjustment, to facilitate the planning of new stores, to weaken patent protection, and so on. Impacts of such 'γ reducing' measures depend on the degree of heterogeneity in consumer tastes. If heterogeneity is not too high and firms start out producing homogeneous products, then reductions in γ

encourage quality laggards to invest more in horizontal differentiation, leading eventually (if k is high enough) to falls in innovation. If heterogeneity is high and firms start out producing differentiated products, then the opposite is true, as lower switching costs encourage quality leaders to minimize differentiation. On this basis it might be argued that models which, for example, examine the link between patent protection and innovation without regard to horizontal differentiation run the risk of missing important elements of the true relationship between these variables.

Aghion et al. (2001) emphasized a tradeoff between escape-competition and composition effects in explaining the innovation-competition relationship. Composition in that model might be termed 'vertical composition', i.e., the proportion of time that an industry spends in states of asymmetric and equal quality. The preceding discussion suggests that, when locations are endogenous, 'horizontal composition' is also important, i.e., the proportion of time that firms spend separated and co-located. This is borne out in figure 8, which shows industry probabilities of co-location under the equilibrium invariant distribution, for the same parameter values as in figure 7. Clearly there is a close correspondence in the shapes of the curves between the two figures.[21] Through its effect on the returns to quality improvement, the propensity of firms to seek either maximal or minimal differentiation from rivals can have substantial effects on industry innovation rates.

---

[21] Figure 8 maintains the convention used in figure 7, to plot the outcome conditional on $s = 1$ in cases where γ is high and there are two ergodic sets.

7 Innovation with fixed locations

In this section we shut down switching investments and compare innovation under fixed locations with the results from our continuous-time model in Narajabad and Watson (2008). That paper analyzes a three-state quality-ladder framework in which firms compete Hotelling-style from opposite ends of the unit interval.[22] Apart from the timeframe, the framework there corresponds to the present model with the location state fixed at $s = 1$.

A common property of continuous-time models of investment competition is that a firm with the maximum quality advantage ($\delta = +k$) always invests nothing – it waits for a laggard to start catching up before responding with its own innovation activity. In our model we further show that equilibrium investment by neck-and-neck firms is monotonically decreasing in α, and investment by laggards is U-shaped in α. Using these results we prove that in the continuous-time model:

i. long-run innovation is always strictly quasiconvex in α, with interior minimum at $\alpha = k/3$;

ii. long-run innovation is maximized at $\alpha = 0$.

In the discrete-time model, with firms at separate locations, the leader's

---

[22] The paper also analyzes an alternate model in which qualities are fixed but locations are endogenous. The endogenous-quality version of the model is similar to that in Aghion and Howitt (1997), with the exception that it studies Hotelling competition rather than Cournot or Bertrand.

equilibrium quality investment may be non-zero: see figure 3. It is more likely to be non-zero when α is low and k is high, which are the situations in which the leader's profit increment K_{H1} is high. This distinction with the continuous-time model feeds into a different pattern of quality investment by laggards. Such investment may now be monotonic increasing in α, rather than U-shaped. This is because the leader has greater incentives to invest at low α, which in turn dampens the laggard's incentives at those values, by reducing the probability that it will be able to successfully catch up. Hence as α declines quality investment by a laggard may fall monotonically.

Quality investment by neck-and-neck firms in discrete time may be U-shaped in α, rather than monotone decreasing. In contrast with the continuous-time situation, in planning their quality investments such firms must now factor in the probability of falling behind in the next period. The loss in π_i from this eventuality would be K_{L1}. Since this quantity is everywhere increasing in α, quality investment may be increasing in α for some parameters.

These differences mean that property (i) of the continuous-time model may not hold in discrete time. In particular our discrete-time computations suggest that when k is large the long-run innovation rate may be monotonically declining in α. This possibility is driven by the leader's quality investment, which in discrete time makes a non-zero contribution to innovation when k is high, and which is monotonically declining in α. On the other hand all the numerical examples that we have computed suggest that result (ii) still holds in discrete time, if k is large enough. That is, innovation is maximized among undifferentiated firms. When k becomes very small the innovation rate goes to zero at all α, and hence in those cases the innovation-maximizing value of α becomes difficult to calculate with certainty.

8 Empirical implications: innovation conditional on locations

The preceding discussion indicates that if locations are fixed then we should expect more innovation when firms are co-located than when they are separated. (The innovation rate among co-located firms is the same as when they are separated and $\alpha = 0$.) It may be possible to test this prediction empirically, by, for example, combining data on product characteristics with firms' patent activity or R&D expenditures. In any such analysis it would be crucial to first determine whether or not product locations may truly be regarded as exogenous. If locations are endogenous there are reasons why we might expect the opposite: more innovation by *differentiated* firms.

To see the reason for this consider first the probability that firms are of equal quality, under the equilibrium invariant distribution, as a function of α. Set γ equal to -1, say, so that the proportion of time spent at separate locations increases with α (see figure 8). For brevity we omit the relevant figure, but it may be seen that the long-run probability of firms having equal quality is also increasing in α. The reason is that at low α stage-game profits are quite convex in δ, which gives a quality leader a strong incentive to maintain its position, and simultaneously diminishes the laggard's incentives to catch up. Then in the long run the industry spends relatively little time in the state $\delta = 0$. When α is high the opposite is true, because the leader's incentives to defend its quality ranking are somewhat reduced, while the laggard's incentives to catch up are increased.

Regardless of the location state, firms innovate most intensely when $\delta = 0$.[23] Changes in quality-state frequencies as α increases may then produce the location-conditional innovation rates shown in figure 9. The figure shows long-run innovation as a function of $\log(\alpha)$ for $\gamma = -1$ and $k = 4.45$. Along with the unconditional innovation rate (the dashed line), also shown are the innovation rates conditional on separation (the solid line), and co-location (the dotted line). Note that the unconditional innovation rate is a weighted average of these two conditional rates – figure 8 shows the change in the weight as α falls with $\gamma = -1$.

The figure reveals the paradoxical result that, at the intermediate values of α, innovation among co-located firms is *lower* than among separated firms. Upon reflection this result is entirely consistent with the preceding discussion. Equal-quality firms earn zero stage-game profits when co-located, but incur high investment costs in the quality race. Firms in this state therefore invest in horizontal differentiation (figure 6), in an attempt to earn the

---

[23] This can be seen by comparing the top and middle panels of figures 3 and 4. First, at the parameters shown each firm invests more if $\delta = 0$ than if they have the lead. Second, in state $\delta = 0$ a quality improvement by *either* firm counts as an innovation.

higher stage-game profits available in the state ($s = 1,\delta = 0$). It follows that, when firms are actively changing product locations, the industry tends not to spend very long in the state ($s = 0,\delta = 0$), relative to the time it spends in ($s = 1,\delta = 0$). As a consequence the equal-quality high-innovation states are relatively underweighted in the conditional invariant distribution over the states ($s = 0$, $\delta \in \{-k, 0, +k\}$), and relatively overweighted in the conditional invariant distribution over the states ($s = 1$, $\delta \in \{-k, 0, +k\}$). Hence the conditional innovation rate can be higher in the latter states than in the former. This observation highlights the danger in empirical work of conditioning on product locations without taking the industry dynamics into account. A finding of greater innovation by horizontally differentiated firms is not inconsistent with a model in which product locations are endogenous.

9 Conclusion

The preceding analysis has for simplicity set aside a number of interesting questions which are left for future research. These include considerations of endogenous entry, borrowing constraints on firms, and location-dependent spillovers. For purposes of comparison with the existing literature we have focused on the long-run determinants of advances in product quality, which are regarded as beneficial by all consumers. A full welfare analysis would also need to take into account the impact on consumers of changes in firms' horizontal differentiation.

In ongoing work we are expanding further on the empirical implications of our model. Using patent data and measures of industry concentration from the UK, Aghion et al. (2005) report an inverse-U relationship between profits and innovation. It turns out that at certain parameter values our model can reproduce this relationship. We are investigating whether the model can also explain other patterns that have been reported in this literature.

References

Aghion, P., Howitt, P., 1997. Competition and growth with step-by-step innovation: An example. European Economic Review, 41, 771-82.

Aghion, P., Harris, C., Howitt, P., Vickers, J., 2001. Competition, imitation and growth with step-by-step innovation. Review of Economic Studies, 68, 467-92.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., Howitt, P., 2005. Competition and innovation: An inverted-U relationship. Quarterly Journal of Economics, 120, 701-28.

Arrow, K.J., 1962. Economic welfare and the allocation of resources for inventions, in: Nelson, R.R. (Ed.), The rate and direction of inventive activity: Economic and social factors. Princeton University Press, Princeton NJ, pp. 609-25.

Bernard, A.B., Redding, S.J., Schott, P.K., 2010. Multiple-product firms and product switching. American Economic Review, 100, 70-97.

Besanko, D., Doraszelski, U., Kryukov, Y., Satterthwaite, M., 2010. Learning-by-doing, organizational forgetting, and industry dynamics. Econometrica, 78, 453-508.

Bresnahan, T.F., Greenstein, S., 1999. Technological competition and the structure of the computer industry. Journal of Industrial Economics, 48, 1-40.

Budd, C., Harris, C., Vickers, J., 1993. A model of the evolution of duopoly: Does the asymmetry between firms tend to increase or decrease? Review of Economic Studies, 60, 543-73.

Caulkins, J.P., Hartl, R.F., Kort, P.M., Feichtinger, G., 2007. Explaining fashion cycles: Imitators chasing innovators in product space. Journal of Economic Dynamics and Control, 31, 1535-56.

Cohen, W.M., Levin, R.C., 1989. Empirical studies of innovation and market structure, in: Schmalensee, R., Willig, R. (Eds.), Handbook of Industrial Organization, Vol. 2. Elsevier, Amsterdam, pp. 1059-1107.

D'Aspremont, C., Gabszewicz, J.J., Thisse, J.-F., 1979. On Hotelling's "Stability in Competition". Econometrica, 47, 1145-50.

De Palma, A., Ginsburgh, V., Papageorgiou, Y.Y., Thisse, J.-F., 1985. The principle of minimum differentiation holds under sufficient heterogeneity. Econometrica, 53, 767-81.

Doraszelski, U., Pakes, A., 2007. A framework for applied dynamic analysis in IO, in: Armstrong, M., Porter, R. (Eds.), Handbook of Industrial Organization, Vol. 3. Elsevier, Amsterdam, pp. 1887-1966.

Doraszelski, U., Satterthwaite, M., 2010. Computable Markov-perfect industry dynamics. RAND Journal of Economics, 41, 215-43.

Ericson, R., Pakes, A., 1995. Markov-perfect industry dynamics: A framework for empirical work. Review of Economic Studies, 62, 53-82.

Gilbert, R., 2006. Looking for Mr. Schumpeter: Where are we in the competition-innovation debate? NBER Innovation Policy and the Economy, 6, 159-215.

Goettler, R., Gordon, B., 2009. Does AMD spur Intel to innovate more? Mimeo., University of Chicago Booth School of Business.

Judd, K.L., 1998. Numerical methods in economics. MIT Press, Cambridge MA.

Long, N.V., Soubeyran, A., 1998. R&D spillovers and location choice under Cournot rivalry. Pacific Economic Review, 3, 105-19.

Markovich, S., 2008. Snowball: A dynamic oligopoly model with indirect network effects. Journal of Economic Dynamics and Control, 32, 909-38.

Maskin, E., Tirole, J., 1988a. A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. Econometrica, 56, 549-69.

Maskin, E., Tirole, J., 1988b. A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. Econometrica, 56, 571-99.

Narajabad, B., 2007. Essays on dynamic markets with heterogeneous agents. Unpublished PhD dissertation, University of Texas-Austin.

Narajabad, B., Watson, R., 2008. Innovation and horizontal differentiation in a continuous-time Hotelling model. Mimeo., Rice University Dept. of Economics.

Pakes, A., McGuire, P., 1994. Computing Markov-perfect Nash equilibria: numerical implications of a dynamic differentiated product model. RAND Journal of Economics, 25, 555-89.

Piga, C., Poyago-Theotoky, J., 2005. Endogenous R&D spillovers and locational choice. Regional Science and Urban Economics, 35, 127-39.

Schumpeter, J.A., 1942. Capitalism, socialism, and democracy. Harper & Brothers, New York.

Sutton, J., 1998. Technology and market structure. MIT Press, Cambridge MA.

Sweeting, A., 2007. Dynamic product repositioning in differentiated product markets: The case of format switching in the commercial radio industry. NBER working paper no. 13522.

Symeonidis, G., 1996. Innovation, firm size and market structure: Schumpeterian hypotheses and some new themes. OECD Economics working paper 161.

Van Cayseele, P.J.G., 1998. Market structure and innovation: A survey of the last twenty years. De Economist, 146, 391-417.

Vogel, J., 2008. Spatial competition with heterogeneous firms. Journal of Political Economy, 116, 423-66.

Whitt, W., 1980. Representation and approximation of noncooperative sequential games. SIAM Journal of Control and Optimization, 18, 33-48.

Figure 1: Equilibrium outcomes in the model with endogenous qualities and product locations

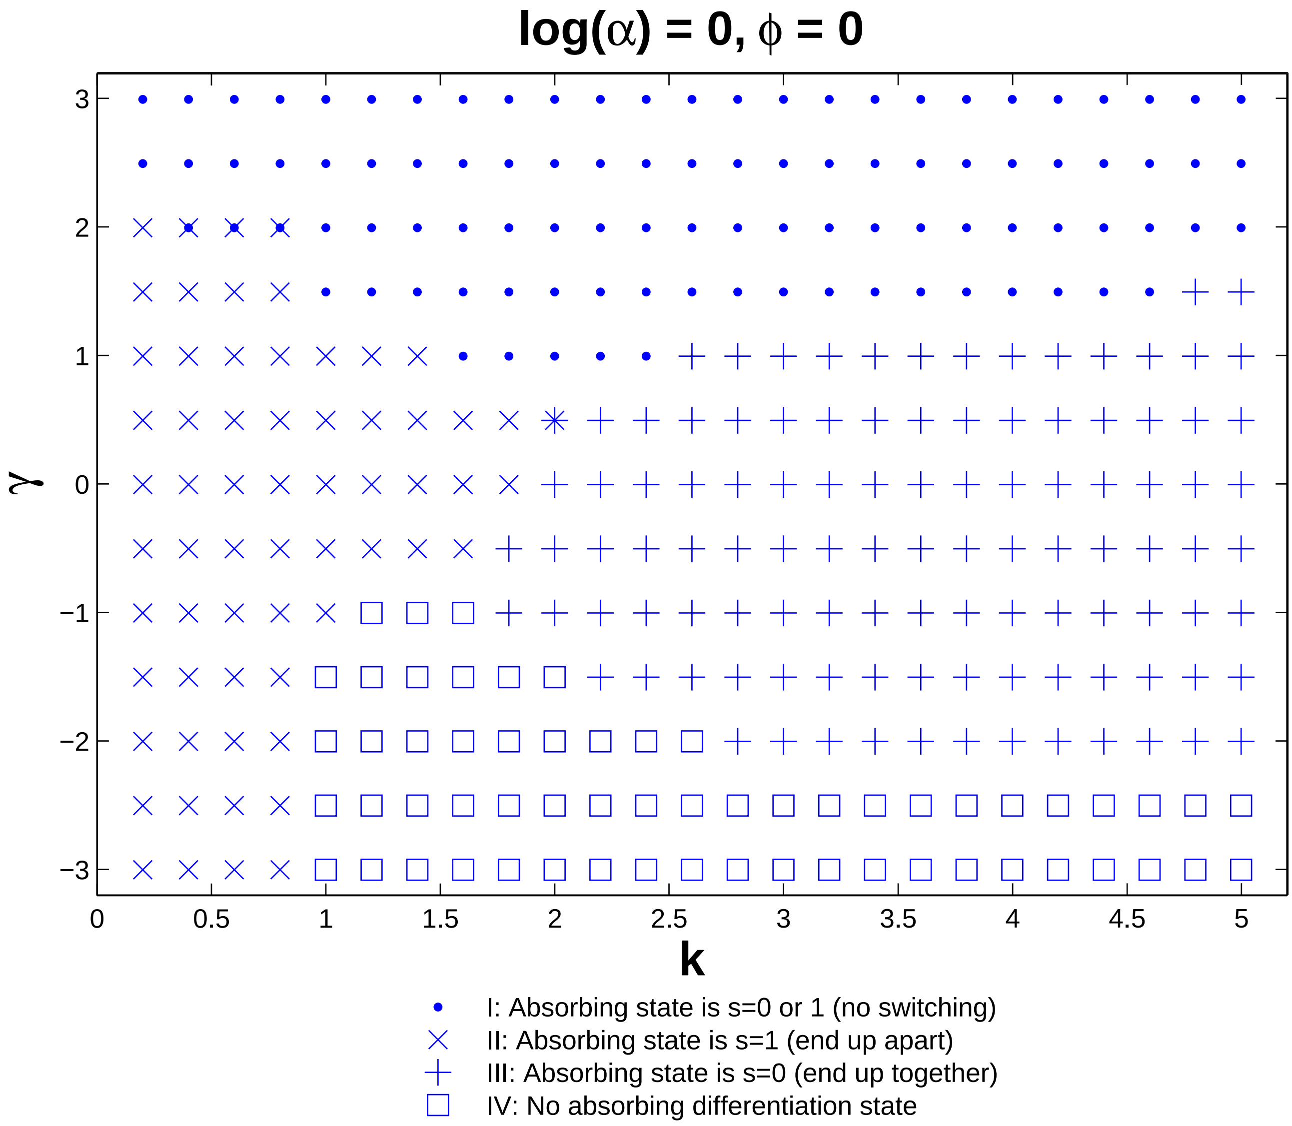

Note: Outcomes for Pareto-dominated equilibria are not shown.

Figure 2: Equilibrium outcomes in the model with endogenous qualities and product locations – effects of changes in α

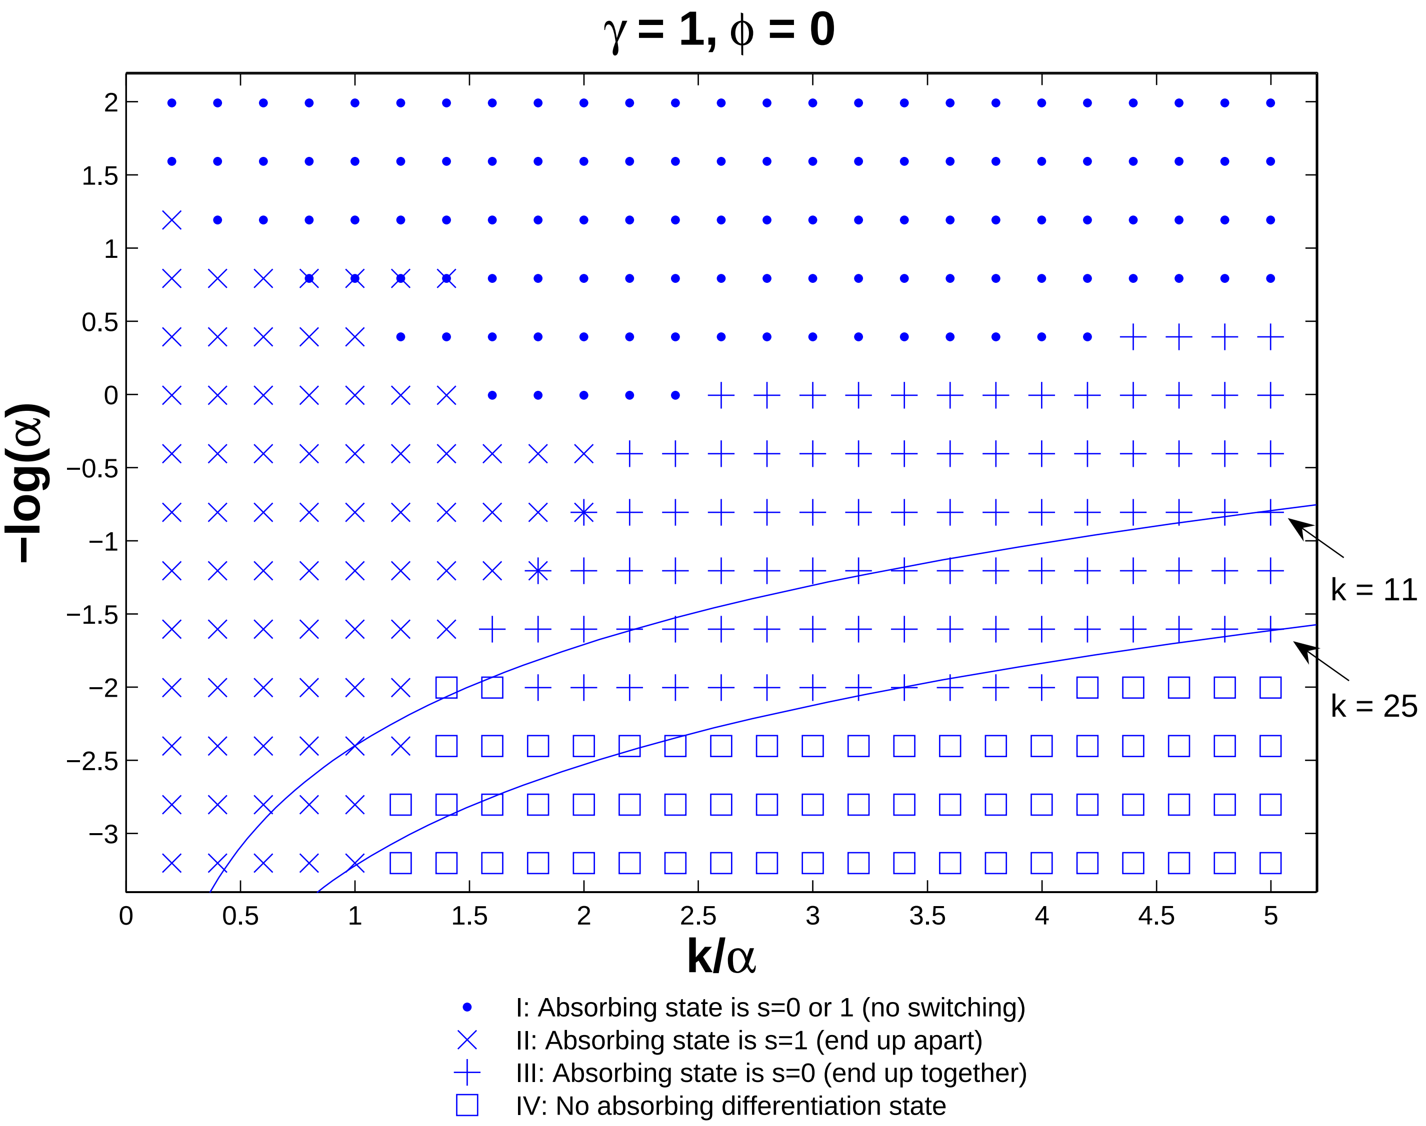

Note: Outcomes for Pareto-dominated equilibria are not shown. Curves show paths as alpha changes with k fixed.

Figure 3: Equilibrium state-conditional probabilities of success in quality investment with endogenous locations, as functions of α, given $s = 1$

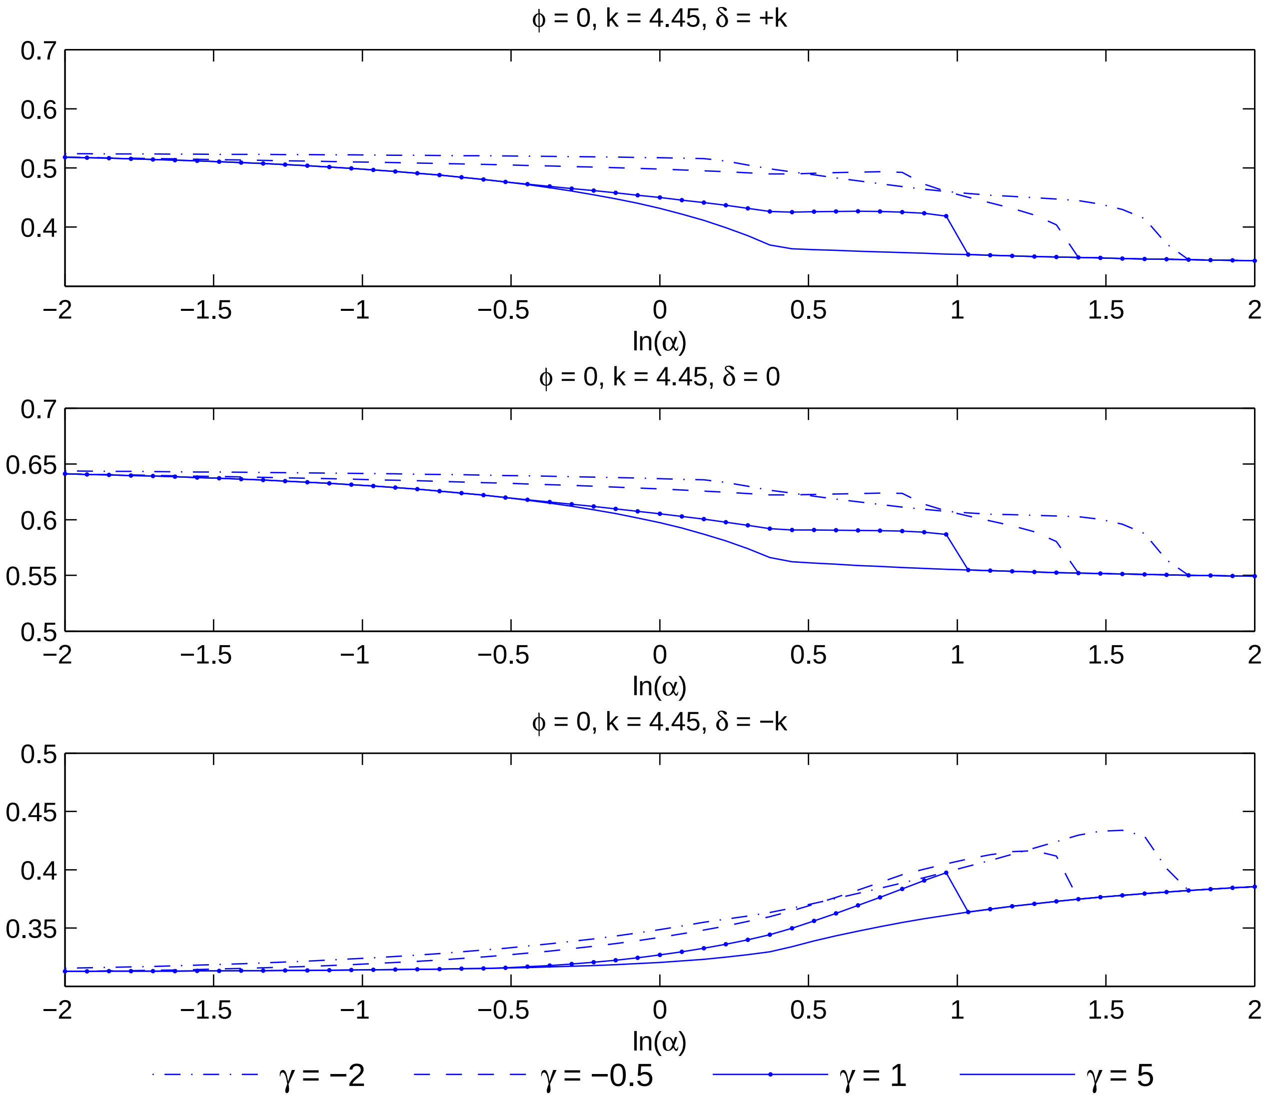

Figure 4: Equilibrium state-conditional probabilities of success in quality investment with endogenous locations, as functions of α, given $s = 0$

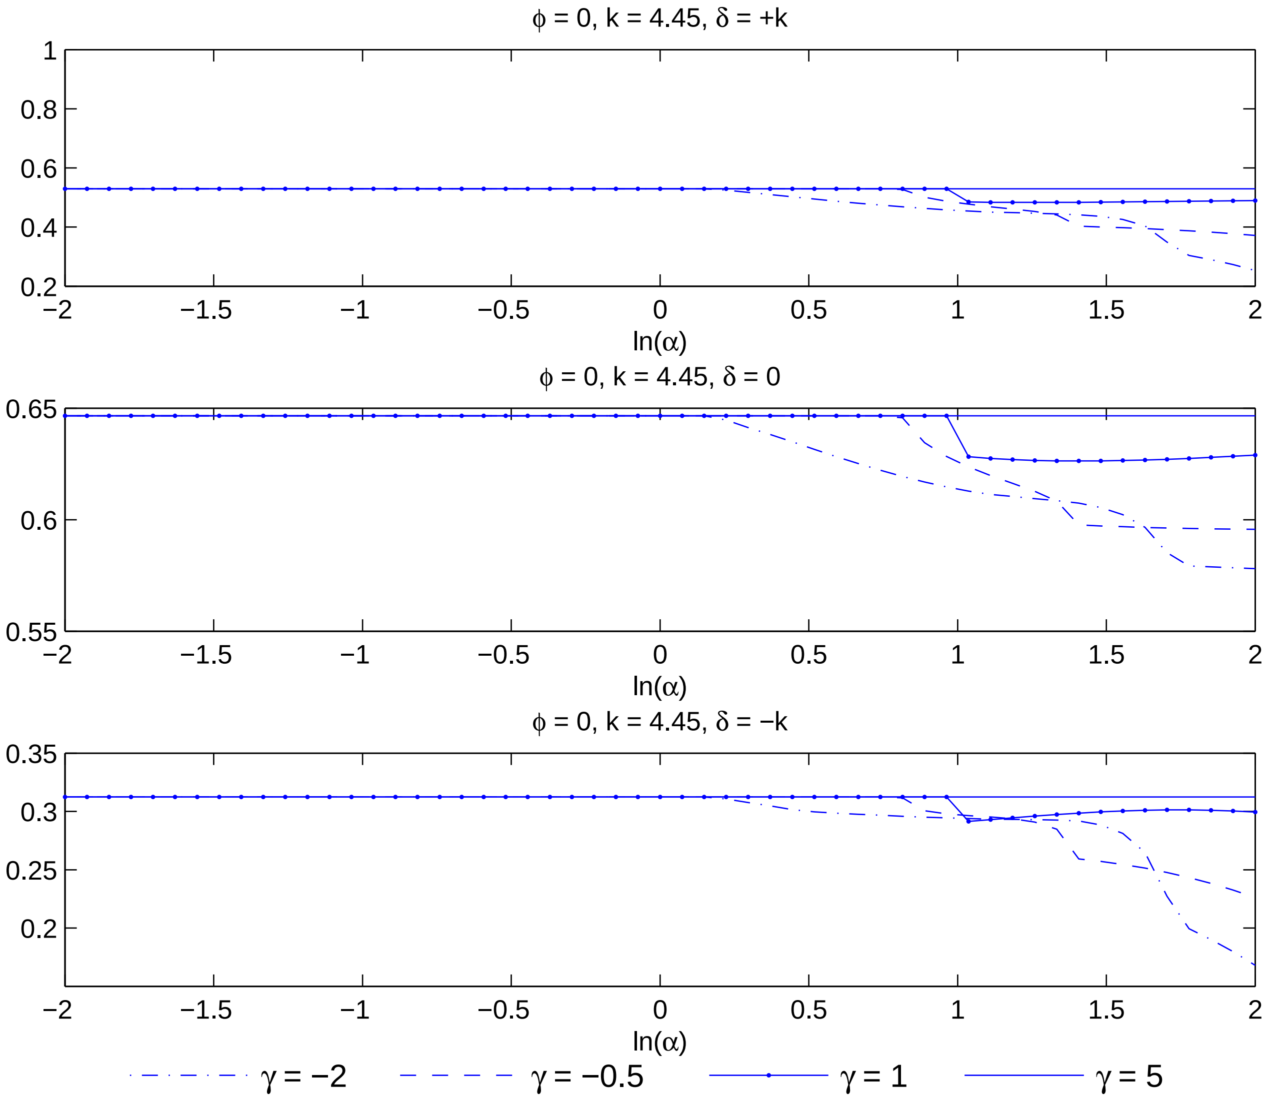

Figure 5: Equilibrium state-conditional probabilities of success in location switching, as functions of α, given $s = 1$

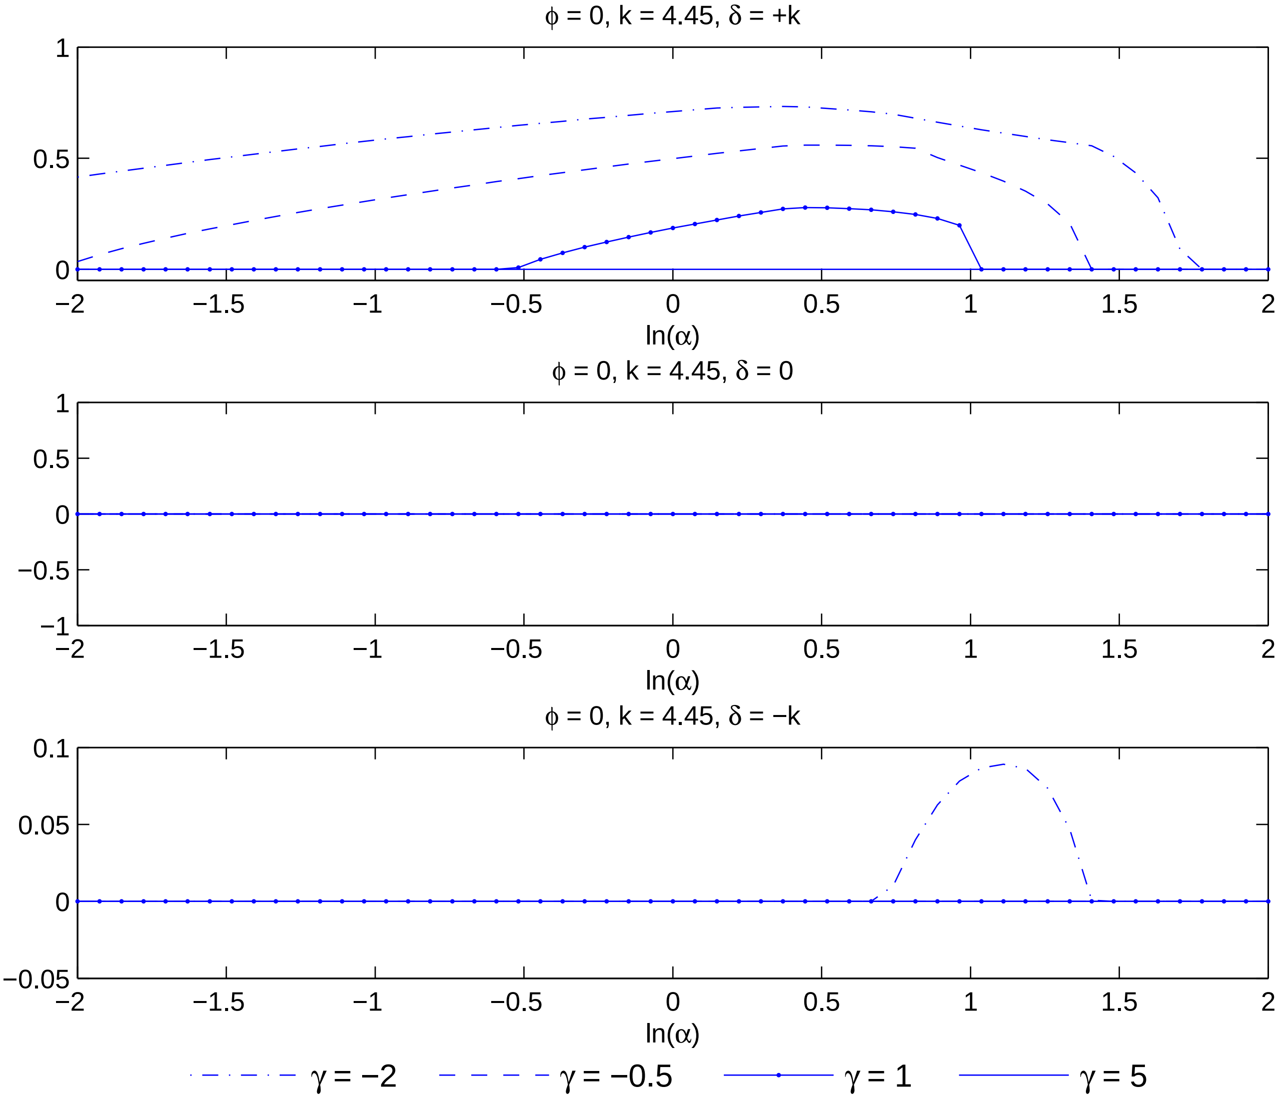

Figure 6: Equilibrium state-conditional probabilities of success in location switching, as functions of α, given $s = 0$

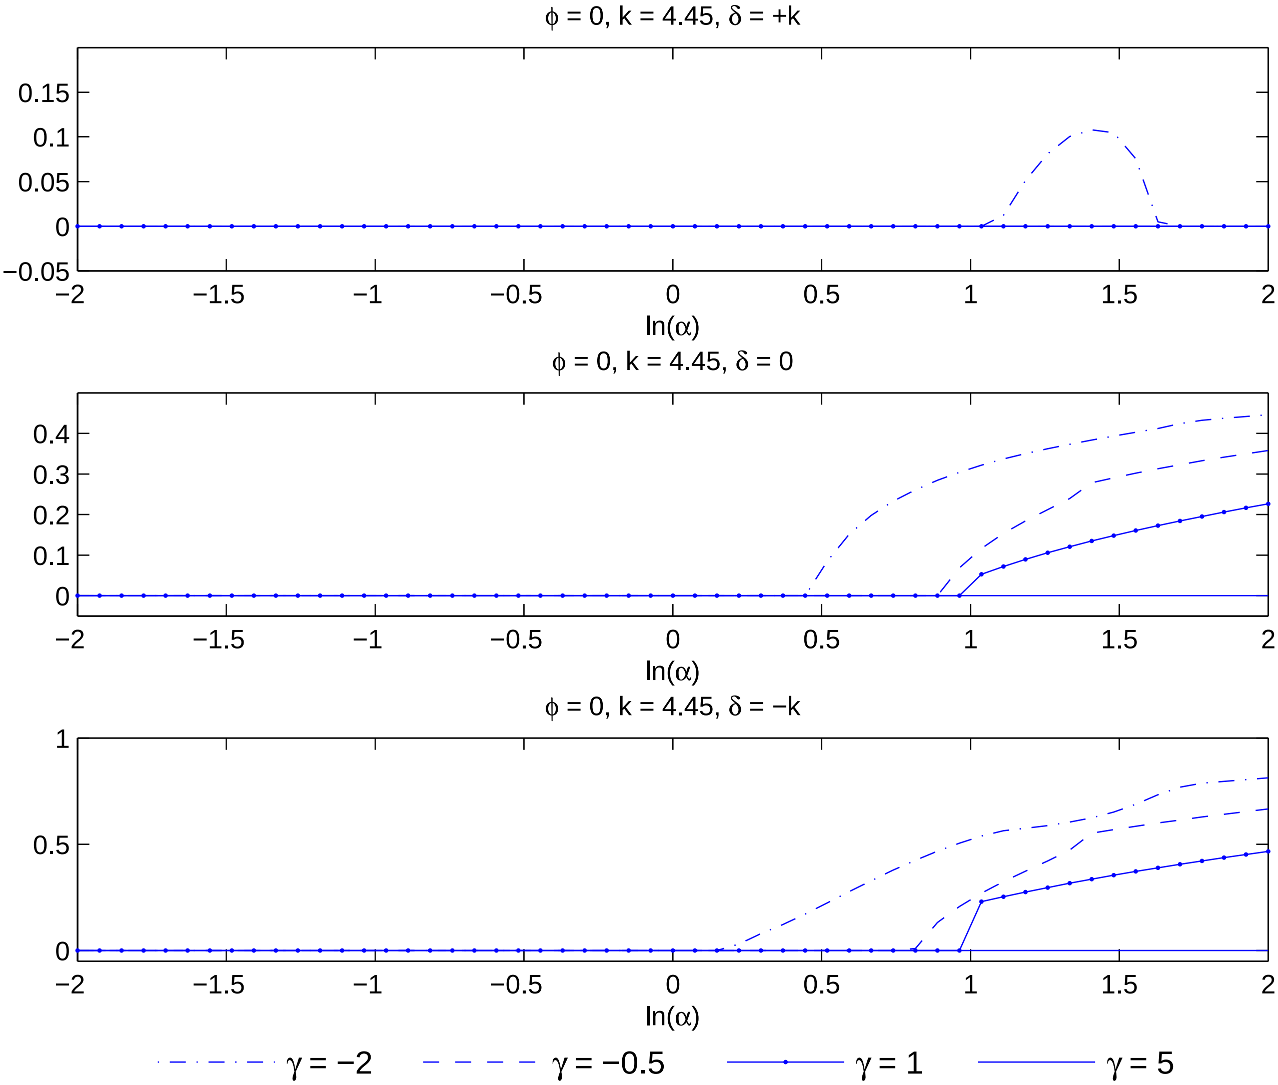

Figure 7: Long-run innovation rate under the equilibrium invariant distribution(s), as a function of α and γ

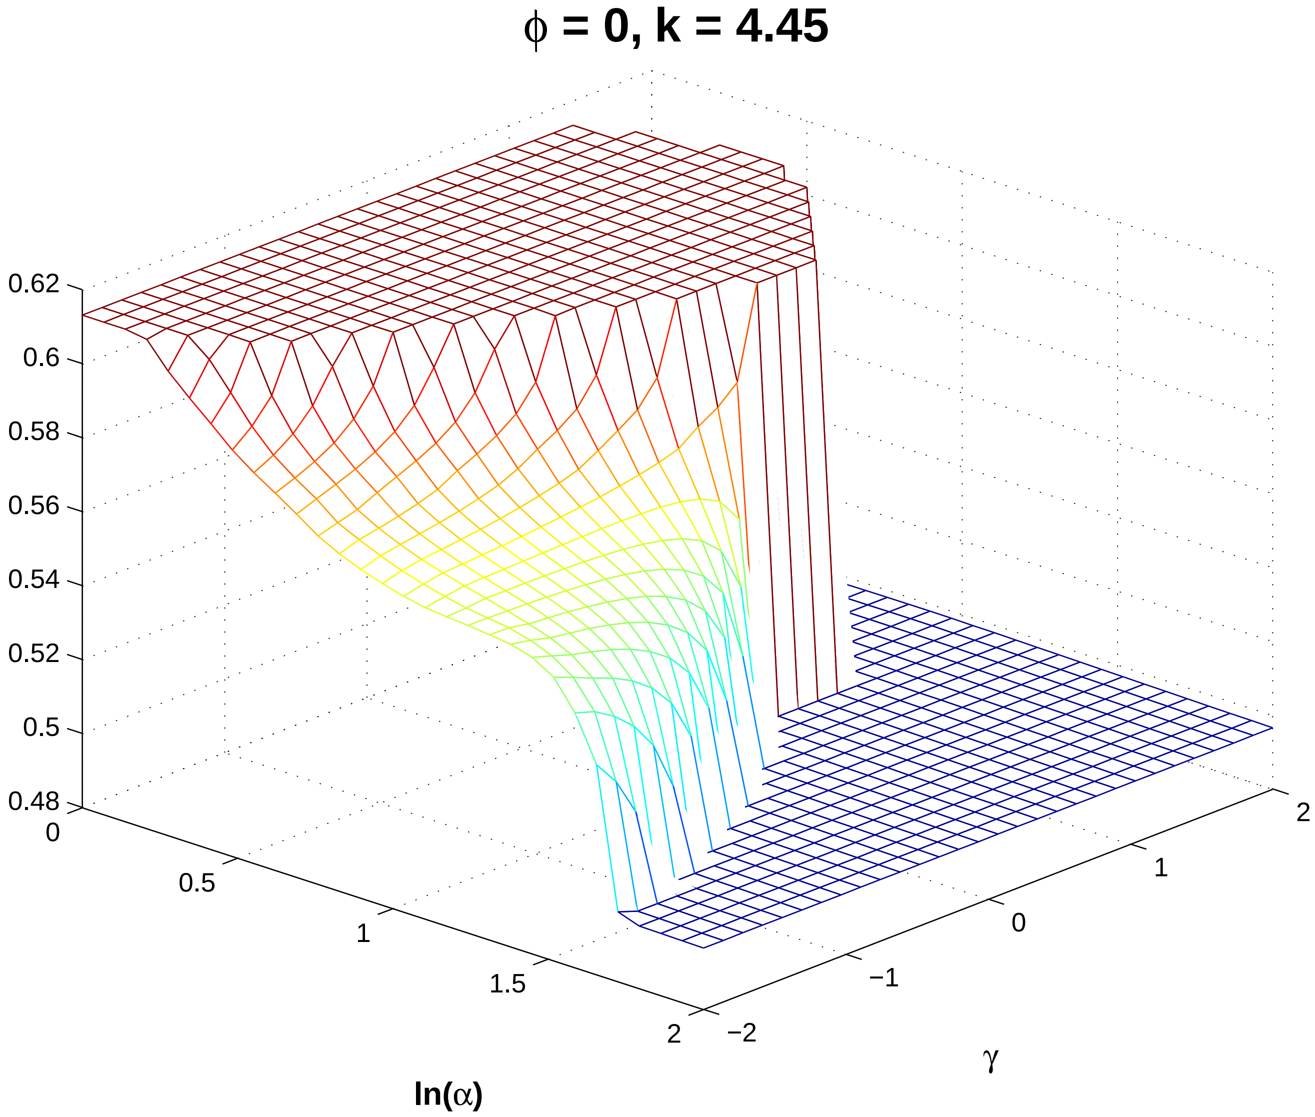

Note: When there are multiple invariant distributions, the case shown is for s=1.

Figure 8: Probability that firms are co-located under the equilibrium invariant distribution(s), as a function of α and γ

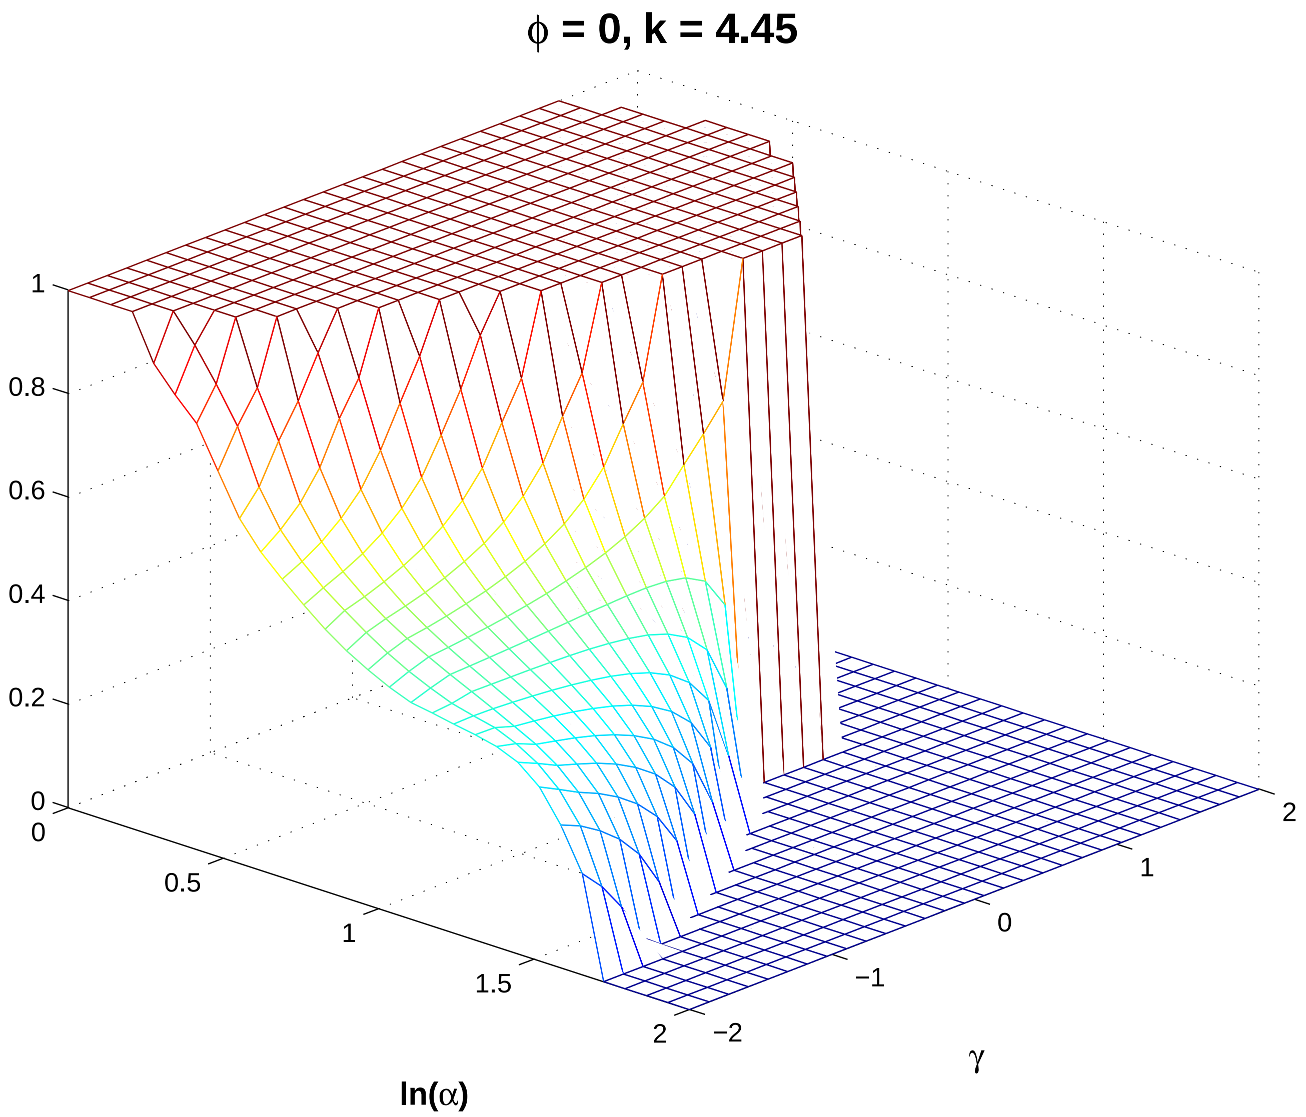

Note: When there are multiple invariant distributions, the case shown is for s=1.

Figure 9: Long-run innovation under the invariant distribution, both unconditional, and conditional on the location state, as functions of α, for $\gamma = -1$