NBER WORKING PAPER SERIES

VERTICAL DIFFERENTIATION IN FRICTIONAL PRODUCT MARKETS

James Albrecht
Guido Menzio
Susan Vroman

Working Paper 29618
http://www.nber.org/papers/w29618

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2021

We are grateful to Boyan Jovanovic, Virgiliu Midrigan, Ezra Oberfield and Pierre-Olivier Weill for useful comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Vertical Differentiation in Frictional Product Markets
James Albrecht, Guido Menzio, and Susan Vroman
NBER Working Paper No. 29618
December 2021
JEL No. D43,D83,L13,O40

ABSTRACT

We consider a version of the imperfect competition model of Butters (1977), Varian (1980) and Burdett and Judd (1983) in which sellers make an ex-ante investment in the quality of their variety of the product. Equilibrium exists, is unique and is efficient. In equilibrium, search frictions not only cause sellers to offer different surpluses to buyers but also cause sellers to choose different qualities for their varieties. That is, equilibrium involves endogenous vertical differentiation. As search frictions decline, the market becomes more and more unequal as a smaller and smaller fraction of sellers produces varieties of increasing quality, offers increasing surplus to their customers, and captures an increasing share of the market, while a growing fraction of sellers produces varieties of decreasing quality. Gains from trade and welfare grow. Under some conditions, the growth rate of gains from trade and welfare is constant.

James Albrecht
Department of Economics
Georgetown University
Washington, DC 20057
albrecht@georgetown.edu

Guido Menzio
Department of Economics
New York University
19 West 4th Street,
New York, NY 10012
and NBER
gm1310@nyu.edu

Susan Vroman
Department of Economics
Georgetown University
Washington, DC 20057
vromans@georgetown.edu

1 Introduction

The search-theoretic model of the product market of Butters (1977), Varian (1980), and Burdett and Judd (1983) has become one of the leading frameworks to study imperfect competition. The model is simple. Sellers set prices. The search process is such that each individual buyer comes into contact with a finite number of randomly selected sellers. The model spans the competitive spectrum going from perfect monopoly to perfect competition as the number of buyers who are in contact with multiple sellers relative to the number of buyers who are in contact with only one seller goes from zero to infinity. The framework has been usefully applied to the analysis of price dispersion (e.g., Baye, Morgan and Scholten 2006, Hong and Shum 2006, Bethune, Choi and Wright 2020), sales (Varian 1980), price stickiness (Head, Liu, Menzio and Wright 2012, Burdett, Trejos and Wright 2017, Burdett and Menzio 2018, Wang, Liu, and Wright 2020), macroeconomic fluctuations (Kaplan and Menzio 2016). Despite its popularity, the framework can still be improved in a fundamental way. In particular, sellers are typically assumed to be either homogeneous or heterogeneous in an arbitrary and exogenous way. In this paper, we contribute to the development of the framework by characterizing the properties of equilibrium in a version of the model in which sellers make an ex-ante investment in the quality of their variety of the product. This opens up the possibility of endogenous heterogeneity across sellers.

We study a frictional product market in the spirit of Butters (1977), Varian (1980) and Burdett and Judd (1983) in which sellers invest in the quality of their variety of the product and, hence, can vertically differentiate themselves. The market is populated by ex-ante identical sellers and ex-ante identical buyers. Each seller chooses how much to invest in the quality y of its variety of the product, where y represents the buyers' utility from consuming that variety. The seller then chooses how much to charge for its variety. Each buyer has a unit demand for the product. The buyer comes into contact with a finite number n of randomly selected sellers, where n is drawn from a Poisson distribution with mean λ. The buyer observes the quality and the price of the variety offered by each of the n sellers, and decides whether and where to buy. Lastly, sellers produce their variety of the product to meet their demand, and do so at a marginal cost normalized to zero.

In the first part of the paper, we establish the existence, uniqueness and properties of equilibrium. Our main finding is that, in equilibrium, the distribution of quality across sellers must be non-degenerate. Imagine an equilibrium in which all sellers choose the same quality for their variety of the product. These sellers offer different surpluses to their buyers for the same reason that they charge different prices in Burdett and Judd (1983). The reason that the distribution of quality across sellers must be non-degenerate in our equilibrium is different. The seller that offers the lowest surplus only trades with the buyers it meets who contact no other seller. The seller that offers the highest surplus trades with all of the buyers it meets. If the seller that offers the lowest surplus deviates and marginally decreases its quality, it would still find it optimal to offer the same surplus.

Hence, this seller's fall in revenues would be proportional to the quantity of output that it sells, which is low. If the seller that offers the highest surplus deviates and marginally increases its quality, it would still find it optimal to offer the same surplus as before. Hence, this seller's increase in revenues would be proportional to the quantity of output that it sells, which is high. The absolute value of the marginal design cost is the same for the two sellers. Since the absolute value of the change in revenues is strictly smaller for the seller offering the lowest surplus than for the seller offering the highest surplus, at least one of them would have an incentive to deviate.

Intuitively, search frictions call for an equilibrium in which there is heterogeneity in the surplus offered by different sellers. In turn, the heterogeneity in surplus generates heterogeneity in the quantity of output sold by different sellers. And the heterogeneity in the quantity of output sold by different sellers generates heterogeneity in the marginal benefit of increasing quality for different sellers. The heterogeneity of marginal benefit is not compatible with all sellers having the same marginal cost of increasing quality and, hence, with an equilibrium in which they all choose the same quality.

The equilibrium must eliminate the discontinuity in the marginal benefit of increasing quality for sellers offering a different surplus. This is accomplished by a non-degenerate distribution of quality across sellers with the following properties: (i) the lowest quality is such that the marginal cost of increasing quality is equal to the quantity of output sold by the seller offering the lowest surplus; (ii) the highest quality is such that the marginal cost of increasing quality is equal to the quantity of output sold by the seller offering the highest surplus; and (iii) the quality at the x-th quantile of the distribution is such that the marginal cost of increasing quality is equal to the quantity of output sold by a seller at the x-th quantile of the surplus distribution. Since sellers with a higher quality find it optimal to offer more surplus to their customers, it follows that the quality distribution does not admit any profitable deviations. If the seller with the lowest quality were to reduce the quality of its variety, the reduction in cost would be no more than the reduction in revenues. If the seller with the highest quality were to increase the quality of its variety, the increase in cost would be no less than the increase in revenues. For any quality between the lowest and the highest, the marginal cost and the marginal benefit of increasing quality are equal.

The equilibrium quality distribution can be derived in closed form using the properties (i), (ii) and (iii) listed above. The equilibrium surplus function, which maps the quality of a seller's variety into the surplus that the seller offers, can be derived in closed form by solving a differential equation obtained from the seller's first-order condition with respect to the surplus. We find that this is the unique equilibrium of the product market. We also find that the unique equilibrium is efficient—in the sense that it decentralizes the solution to the problem of a utilitarian social planner.

In the second part of the paper, we examine the effect of a decline in search frictions due to, say, improvements in information, communication, or transportation technologies that make it easier for buyers to locate and access a larger number of sellers. We model a decline in search frictions as an increase in λ, the average number of sellers contacted by each buyer. We find that a decline in search frictions leads to higher sales concentration—in the sense that it increases the quantity of output sold by larger sellers (sellers who produce a higher-quality variety, offer higher surplus and sell to more buyers) and it lowers the quantity of output sold by smaller sellers (sellers who produce a lower-quality variety, offer less surplus, and sell to fewer buyers). Higher sales concentration affects the sellers' incentives to invest and, in turn, leads to quality polarization—in the sense that it further increases the quality of the variety produced by high-quality sellers, and it further decreases the quality of the variety produced by low-quality sellers. Higher sales concentration and quality polarization together lead to an increase in revenue concentration. These phenomena, which are typically associated with market dysfunction, are here the efficient response to a changing environment. Intuitively, when search frictions become smaller, it is efficient for a small fraction of sellers to invest more in quality and serve a larger fraction of buyers. Moreover, while these phenomena are typically associated with increasing market power, here they are a consequence of the increase in competition brought about by declining search frictions.

We then examine the effect of declining search frictions on the gains from trade—i.e. the sum of the buyers' and sellers' payoffs excluding the sellers' investment in quality—and on welfare—i.e. the sum of the buyers' and sellers' payoffs including the cost of sellers' investment in quality. We find that declining search frictions increase the total gains from trade, the buyers' gains from trade and the sellers' gains from trade. We find that declining search frictions increase total welfare, increase the buyers' welfare and reduce the sellers' welfare. It is not surprising that lower search frictions increase the gains from trade and welfare. Lower search frictions increase welfare because the equilibrium decentralizes the solution of the social planner problem and, from the perspective of the social planner, relaxing constraints on trade is beneficial. Lower search frictions increase the gains from trade because they allow buyers to meet more sellers and, hence, to purchase varieties with higher quality. Moreover, lower search frictions induce the largest sellers to increase the quality of their varieties. It is, however, surprising that lower search frictions have opposite effects on the gains from trade and on the welfare going to the sellers. Intuitively, the welfare of the sellers decreases because lower search frictions increase the extent of ex-ante competition among sellers. The gains from trade of the sellers increase because the increase in ex-ante competition is offset, ex-post, by the growing differences in the quality of the varieties produced by different sellers.

If search frictions decline at a constant rate, in the sense that λ increases at a constant rate g_λ, the market does not follow a Balanced Growth Path (BGP). Indeed, as λ increases, the shapes of the quality distribution and of the surplus distribution change and become more and more skewed. However, if the cost of designing a variety of quality y has a constant elasticity γ with respect to y, the market payoffs do grow at a constant rate.

In particular, the gains from trade grow at a rate that converges to $g_\lambda/(\gamma-1)$, and the shares of the gains from trade accruing to buyers and the sellers converge to two strictly positive constants. Welfare grows at a rate that converges to $g_\lambda/(\gamma-1)$, the share of the welfare accruing to buyers converges to 1, while the share of the welfare accruing to sellers converges to 0. The constant growth rate of market payoffs $g_\lambda/(\gamma-1)$ is equal to the rate at which search frictions decline, g_λ, times the return to declining search frictions, $1/(\gamma-1)$, which captures the elasticity of the quality of the best variety in the market with respect to λ.

The first part of the paper is a contribution to search-theoretic models of imperfect competition in the spirit of Butters (1977), Varian (1980), Burdett and Judd (1983) and Burdett and Mortensen (1998). The original versions of these models assume that sellers are homogeneous with respect to the quality of their product and their marginal cost. Some later versions of these models allow for seller heterogeneity with respect to either the quality of the product (see, e.g., Bontemps, Robin and Van den Berg 2000) or marginal cost (see, e.g., Menzio and Trachter 2018). Heterogeneity is, however, exogenous. We contribute to this literature by studying a version of these models in which sellers choose the quality of their product through an ex-ante costly investment. Our main finding is that search frictions not only cause sellers to offer different surpluses but also cause them to choose different qualities. Our main finding is related to Robin and Roux (2002) and Acemoglu and Shimer (2000), although the models and the type of heterogeneity are quite different from ours. Robin and Roux (2002) consider a version of Burdett and Mortensen (1998) in which firms make an ex-ante investment in capital that affects the marginal productivity of labor. They find that, under some conditions, firms choose different levels of capital. Similarly, Acemoglu and Shimer (2000) consider a directed search model of the labor market in which firms make an ex-ante investment in capital. They find that, in any equilibrium, firms choose different levels of capital.

The second part of our paper relates to the literature on rising market concentration and markups (see, e.g., Autor et al. 2020, De Loecker, Eeckhout and Unger 2020, Kehrig and Vincent 2021). These papers, using a variety of measures, document a rise in market concentration and markups over the last few decades. In our paper, we show that declining search frictions due to improvements in information and communication technology may be a contributing factor to rising concentration of sales and revenues, a finding related to Rosen (1981). Moreover, we show that declining search frictions change sellers' incentives to invest in quality and, in particular, increase the quality at the largest sellers and lower the quality at the smallest sellers. Interestingly, we find that the increase in market concentration and the increase in quality polarization are all efficient responses to an environment that becomes more and more competitive as search frictions become smaller. Hence, to the extent that increasing concentration and polarization are driven by declining search frictions, they are benign phenomena, a view that is close to the one suggested by Autor et al. (2020) and Aghion et al. (2021). In contrast, Gutierrez and Philippon

(2019) and De Loecker, Eeckhout and Mongey (2021) interpret rising concentration as the nefarious consequence of rising barriers to entry.

The second part of the paper also contributes to a recent strand of literature that studies Stiglerian growth (Stigler 1961), i.e., the contribution of declining search frictions to economic growth. Menzio (2021) considers a version of Burdett and Judd (1983) in which sellers choose how much to horizontally differentiate their variety of the product. That is, sellers choose the degree of specificity of their variety of the product, where higher specificity implies that their variety is liked by a smaller fraction of buyers but gives them more utility. Under some conditions, the market follows a Balanced Growth Path as search frictions become smaller. That is, the variety and the price distributions grow at a constant rate, and so do the payoffs to buyers and sellers. The growth rate of payoffs depends on the rate at which search frictions decline and on the elasticity of the buyers' utility function with respect to the degree of specialization of a variety. In this paper, we consider a version of Burdett and Judd (1983) in which sellers choose how much to vertically differentiate their variety of the product. We show that, when differentiation is vertical, the market does not follow a BGP as search frictions decline. Rather, the variety and the surplus distributions change shape and become more and more skewed. The payoffs to buyers and sellers, however, do grow at a constant rate, which depends on the rate at which search frictions decline and on the elasticity of the sellers' cost function.[1] Martellini and Menzio (2020, 2021) consider a version of Mortensen and Pissarides (1994) in which there is horizontal differentiation between workers and firms, in the sense that the productivity of a firm-worker match depends on their distance along a circle. Under some conditions, they find that the labor market follows a BGP as search frictions decline. Specifically, unemployment, vacancies and transition rates remain constant, while labor productivity grows at a constant rate.

In the next section, we set up the model and prove the existence of a unique equilibrium. We solve for the equilibrium in closed form, and we show that the equilibrium is constrained efficient in the sense that the equilibrium distribution of quality across sellers is the same as the distribution that a social planner would choose. In Section 3, we analyze the effects on the market of decreasing search frictions. We show the effect of reducing search frictions on both market structure and welfare. The final section contains concluding remarks.

2 Equilibrium

In this section, we establish the existence, uniqueness and efficiency of equilibrium in a version of Butters (1977), Varian (1980) and Burdett and Judd (1983) in which sellers invest in the quality of their variety of the product. We find that the equilibrium is such

[1] In the context of a different model, Bar-Isaac, Caruana, Cunat (2012) study the effect of declining search frictions on the incentives of sellers to either vertically or horizontally differentiate.

that identical sellers choose different qualities and those who choose a higher quality offer higher surplus to their customers. We derive a simple closed-form expression for both the distribution of quality across sellers and for the mapping between a seller's quality and the surplus the seller offers. We also show that the equilibrium allocation coincides with the allocation that solves the problem of a utilitarian social planner.

2.1 Environment

We consider the market for some consumer good. The market is populated by a positive measure of buyers and by a positive measure of sellers, where $\theta > 0$ denotes the ratio between the measure of buyers and the measure of sellers. Every buyer in the market demands a single unit of the good. A buyer who purchases a unit of the good with quality y at price p obtains a payoff of $y - p$. A buyer who fails to purchase the good receives a payoff of 0.

Every seller in the market designs its own variety of the good. A seller pays the cost $c(y)$ in order to design a variety of the good with quality $y \geq 0$, where $c(y)$ is a strictly increasing and strictly convex function such that $c(0) = 0$, $c'(0) = 0$ and $c'(\infty) = \infty$. For the sake of simplicity, we focus on the case in which the design cost function $c(y)$ has the isoelastic form $c(y) = y^\gamma$ with $\gamma > 1$. After paying the design cost $c(y)$, the seller can produce any quantity of its variety of the good at a constant marginal cost, which we normalize to 0. A seller obtains a payoff of $qp - c(y)$ from designing a variety of the good with quality y and selling q units of it at the price p.

The market is subject to search frictions in the sense that a buyer cannot simply purchase the good from any seller in the market, but rather can only purchase from a seller that the buyer has contacted. A buyer contacts n randomly-selected sellers, where n is drawn from a Poisson distribution with mean $\lambda > 0$. A buyer observes the quality y and the price p of the variety sold by each of the n sellers contacted, and then decides whether to purchase the good and, if so, from which of the n sellers. Clearly, the buyer finds it optimal to purchase the good from the seller that offers the highest surplus $s \equiv y - p$, as long as s is positive. If the buyer is offered the same surplus by multiple sellers, the buyer chooses one at random.

We can now define an equilibrium in this market.

Definition 1. *An equilibrium is a quality distribution $H(y)$ and a surplus distribution $F(s)$ such that:*

1. *Every seller chooses a quality y and a surplus s that maximize its expected payoff, taking as given $H(y)$, $F(s)$, and the buyers' strategies;*
2. *The distributions $H(y)$ and $F(s)$ are consistent with the sellers' strategies.*

2.2 Existence and uniqueness

We now establish the existence and uniqueness of equilibrium and solve for the equilibrium quality distribution H and the equilibrium surplus distribution F in closed form. As in Butters (1977), Varian (1980), and Burdett and Judd (1983), we first derive a number of properties that must hold in any equilibrium and that, when taken together, identify a unique candidate equilibrium. Then, we verify that the unique candidate equilibrium is indeed an equilibrium.

To start, we derive the contact and trade probabilities for buyers and sellers. A buyer contacts n sellers with probability

$$\lambda^n \frac{e^{-\lambda}}{n!} \text{ for } n = 0, 1, 2, ... \tag{2.1}$$

A seller expects to meet $\theta\lambda$ buyers. The probability that any one of those buyers contacts exactly k other sellers is $\lambda^k \exp(-\lambda)/k!$.The seller therefore expects to meet θ_k buyers who are in contact with k other sellers, where

$$\theta_k = \theta(k+1)\lambda^{k+1} \frac{e^{-\lambda}}{(k+1)!} \text{ for } k = 0, 1, 2, ... \tag{2.2}$$

A seller expects to meet θ_0 buyers who are not in contact with any other seller. As long as the seller offers a nonnegative surplus s, the seller trades with all of these buyers. The seller expects to meet θ_1 buyers who are in contact with one other seller. As long as $s \geq 0$, the seller trades with any one of these buyers with probability $F(s-) + \mu(s)/2$, where $F(s-)$ denotes the fraction of sellers offering a surplus strictly smaller than s and $\mu(s)$ denotes the fraction of sellers offering a surplus equal to s. The seller expects to meet θ_2 buyers who are in contact with two other sellers. As long as $s \geq 0$, the seller trades with any one of these buyers with probability $F(s-)^2 + 2\mu(s)/2 + \mu(s)^2/3$. Using the same logic, one can easily derive the probability that the seller trades with a buyer who has a generic number k of other contacts.

Lemma 1 below states that, in any equilibrium, the maximized profit of a seller, which we denote by V^*, is strictly positive. Intuitively, this is because a seller can design a variety with some low quality $y > 0$, offer its customers the surplus $s = 0$, and make a strictly positive profit by trading with the buyers who are not in contact with any other seller.

Lemma 1. *In any equilibrium, the maximized profit of a seller, V^*, is strictly positive.*

Proof. The profit for a seller that designs a variety of quality $y \geq 0$ and offers a surplus of 0 is greater than or equal to $\tilde{V}(y, 0)$, where $\tilde{V}(y, 0)$ is given by

$$\tilde{V}(y, 0) = -y^\gamma + \theta\lambda e^{-\lambda} y. \tag{2.3}$$

The first term on the right-hand side of (2.3) is the seller's cost of designing a variety of quality y. The second term is a lower bound on the seller's revenues when it offers a surplus of 0 to its customers. Given that the seller offers a surplus of 0, only buyers who are not in contact with any other seller will purchase its variety. The seller expects to meet $\theta_0 = \theta\lambda\exp(-\lambda)$ such buyers, and each generates revenue y. Noting that $\tilde{V}(0,0) = 0$ and $\tilde{V}_y(0,0) > 0$, there exists a $y_0 > 0$ such that $\tilde{V}(y_0, 0) > 0$. Therefore, the maximized profit of a seller, V^*, is strictly positive. ■

Lemma 2 below establishes some key properties of the equilibrium surplus distribution F. In the context of a model in which the quality y of the product is the same across all sellers and strictly positive, Butters (1977), Varian (1980) and Burdett and Judd (1983) show that (i) F has no mass points, (ii) the support of F has no gaps, and (iii) the lower bound on the support of F is such that the seller offering the lowest surplus extracts all of the gains from trade. Even though the quality of the product may differ across sellers in our model, the properties of F are the same as in Butters (1977), Varian (1980) and Burdett and Judd (1983) because, as implied by Lemma 1, every seller chooses a quality y that is strictly positive.

Lemma 2. *In any equilibrium, the surplus distribution F has no mass points, and its support is some interval $[s_\ell, s_h]$, with $0 = s_\ell < s_h$.*

Proof. The proof of Lemma 2 is the same as in Butters (1977), Varian (1980) and Burdett and Judd (1983). Let us sketch the proof. The surplus distribution F cannot have a mass point. Suppose there is a mass point at $s_0 \geq 0$. Any seller offering s_0 has a variety of quality $y_0 > s_0$, since Lemma 1 guarantees that $V^* > 0$. A seller that offers s_0 trades with all of the buyers it meets who do not contact any other seller offering a surplus $s \geq s_0$. Moreover, the seller trades with a fraction of the buyers it meets who also meet some other seller(s) offering s_0 but do not contact anyone offering a surplus $s > s_0$. If the seller were instead to offer surplus $s_0 + \epsilon$, it could trade with all of the buyers it meets who do not contact any other seller offering a surplus $s > s_0$. Therefore, by offering $s_0 + \epsilon$ rather than s_0, the seller increases the quantity sold discretely. Since ϵ is an arbitrary positive number and $y_0 > s_0$, the seller's profit is strictly greater at $s_0 + \epsilon$ than at s_0 for some ϵ small enough, which contradicts the definition of equilibrium.

To see that the support of the distribution F has no gaps, suppose there is a gap in the support of F between s_1 and s_2, with $s_1 < s_2$. Since $F(s_1) = F(s_2)$, a seller offering the surplus s_2 would trade with the same number of buyers if it deviated and offered the surplus s_1. Since $s_1 < s_2$, a seller offering the surplus s_2 makes a higher profit per sale by deviating and offering the surplus s_1. Hence, the seller's profit is strictly greater at s_1 than at s_2, which contradicts the definition of equilibrium.

Since F has no mass points and no gaps, its support is some interval $[s_\ell, s_h]$, with $s_\ell < s_h$.

To see that $s_\ell = 0$, suppose that s_ℓ were strictly positive. A seller offering the surplus s_ℓ trades with the same number of buyers by deviating and offering a surplus of 0. The seller's profit would be strictly greater at 0 than at s_ℓ, which contradicts the definition of equilibrium. ■

Lemma 2 implies that the surplus distribution F has no mass points. This allows us to derive a simple expression for the expected profit $V(y, s)$ for a seller that designs a variety of the product with quality $y \geq 0$ and that offers its customers the surplus $s \geq 0$. Specifically,

$$\begin{aligned} V(y,s) &= -y^\gamma + \left[\sum_{k=0}^{\infty} \theta(k+1)\lambda^{k+1}\frac{e^{-\lambda}}{(k+1)!}F(s)^k\right](y-s) \\ &= -y^\gamma + \theta\lambda e^{-\lambda}\left[\sum_{k=0}^{\infty}\frac{\lambda^k F(s)^k}{k!}\right](y-s) \\ &= -y^\gamma + \theta\lambda e^{-\lambda(1-F(s))}(y-s), \end{aligned} \tag{2.4}$$

where the second line is obtained by collecting $\theta\lambda e^{-\lambda}$ in the first line, and the third line is obtained by recognizing that the summation in the second line is equal to $\exp(\lambda F(s))$.

The next lemma states that, in any equilibrium, the quality distribution H has no mass points. This is one of the main findings of our paper. It shows that the existence of search frictions not only calls for a surplus distribution that is non-degenerate, as shown in Butters (1977), Varian (1980) and Burdett and Judd (1983), but also for a quality distribution that is non-degenerate. The reason that the quality distribution must be non-degenerate differs from the reason that the surplus distribution must be non-degenerate. The surplus distribution must be non-degenerate because any mass point in it would create a discontinuity in the level of profit for a seller with respect to the surplus it offers to its customers. Specifically, as shown above, a mass point at some s_0 implies that the profit of a seller would be discretely higher at $s_0 + \epsilon$ than at s_0 because, by offering $s_0 + \epsilon$, the seller can outbid a mass of competitors. The quality distribution must be non-degenerate because any mass point in H would create heterogeneity in the derivative of the profit of a seller with respect to the quality of its variety. Specifically, since the surplus distribution must be atomless, a mass point at some y_0 means that sellers with a variety of quality y_0 must offer different surpluses to their customers and, hence, must have different volumes of trade. The derivative of these sellers' profit functions with respect to their quality, which depends on volume, will be different and, for this reason, cannot be zero for all of them.

Lemma 3. *In any equilibrium, the quality distribution H has no mass points.*

Proof. To show a contradiction, suppose H has a mass point at some quality $y_0 > 0$. Since the surplus distribution F has no mass points, the mass of sellers with a variety of quality y_0 cannot offer a common surplus. Let $s_{\ell,0}$ and $s_{h,0}$ denote the lowest and highest surpluses offered by the mass of sellers with a variety y_0. The profits for a seller choosing $(y_0, s_{\ell,0})$ and for a seller choosing $(y_0, s_{h,0})$ are given by

$$V(y_0, s_{\ell,0}) = -y_0^\gamma + \theta\lambda e^{-\lambda(1-F(s_{\ell,0}))}(y_0 - s_{\ell,0}), \tag{2.5}$$
$$V(y_0, s_{h,0}) = -y_0^\gamma + \theta\lambda e^{-\lambda(1-F(s_{h,0}))}(y_0 - s_{h,0}), \tag{2.6}$$

where $F(s_{h,0}) - F(s_{\ell,0})$ is greater than or equal to the mass of sellers with a variety of quality y_0.

The derivatives of (2.5) and (2.6) with respect to y are

$$V_y(y_0, s_{\ell,0}) = -\gamma y_0^{\gamma-1} + \theta\lambda e^{-\lambda(1-F(s_{\ell,0}))}, \tag{2.7}$$
$$V_y(y_0, s_{h,0}) = -\gamma y_0^{\gamma-1} + \theta\lambda e^{-\lambda(1-F(s_{h,0}))}. \tag{2.8}$$

Since $F(s_{h,0}) > F(s_{\ell,0})$, $V_y(y_0, s_{\ell,0}) < V_y(y_0, s_{h,0})$. If $V_y(y_0, s_{\ell,0}) < 0$, the seller choosing $(y_0, s_{\ell,0})$ can strictly increase its profit by lowering the quality of its variety. If $V_y(y_0, s_{h,0}) > 0$, the seller choosing $(y_0, s_{h,0})$ can strictly increase its profit by increasing the quality of its variety. Since $V_y(y_0, s_{\ell,0}) < V_y(y_0, s_{h,0})$, either $V_y(y_0, s_{\ell,0}) < 0$ or $V_y(y_0, s_{h,0}) > 0$. Therefore, if H has a mass point at y_0, there exists a seller that is not maximizing its profit. ■

Having established that, in any equilibrium, sellers must choose a different quality for their variety of the product, we now want to characterize the relationship between the quality of a seller's variety and the surplus that it offers to its customers. Lemma 4 below shows that sellers that choose a higher quality for their variety find it optimal to offer strictly more surplus to their customers. The intuition behind this property of equilibrium is that the seller's profit function is such that a seller with a higher quality has more to gain from offering a higher surplus to its customers than does a seller with a lower quality.

Lemma 4. *In any equilibrium, the surplus offered by a seller is strictly increasing in the quality of the seller's variety.*

Proof. Let s_1 denote the surplus offered by a seller with a variety of quality y_1, and let s_2 denote the surplus offered by a seller with a variety of quality y_2, where $y_1 < y_2$. We first prove that $s_1 \leq s_2$. Since s_1 is optimal for y_1 and s_2 is optimal for y_2, it follows that

$$\theta\lambda e^{-\lambda(1-F(s_1))}(y_1 - s_1) \geq \theta\lambda e^{-\lambda(1-F(s_2))}(y_1 - s_2), \tag{2.9}$$
$$\theta\lambda e^{-\lambda(1-F(s_2))}(y_2 - s_2) \geq \theta\lambda e^{-\lambda(1-F(s_1))}(y_2 - s_1). \tag{2.10}$$

Combining (2.9) and (2.10) yields

$$\left[e^{-\lambda(1-F(s_2))} - e^{-\lambda(1-F(s_1))}\right](y_2 - y_1) \geq 0. \tag{2.11}$$

The inequality (2.11) implies that $\exp(-\lambda(1-F(s_2)))$ is greater than or equal to $\exp(-\lambda(1-F(s_1)))$ and, hence, s_2 is greater than or equal to s_1.

We now prove that $s_1 < s_2$. To show a contradiction, suppose that $s_1 = s_2 = s$. The profits for the sellers choosing (y_1, s) and (y_2, s) are given by

$$V(y_1, s) = -y_1^\gamma + \theta\lambda e^{-\lambda(1-F(s))}(y_1 - s), \tag{2.12}$$
$$V(y_2, s) = -y_2^\gamma + \theta\lambda e^{-\lambda(1-F(s))}(y_2 - s), \tag{2.13}$$

Both $V(y_1, s)$ and $V(y_2, s)$ must be equal to the maximized profit V^*. Since $V(y, s)$ is strictly concave in y, $V(y_1, s) = V(y_2, s) = V^*$ implies $V(y, s) > V^*$ for all $y \in (y_1, y_2)$, which contradicts the definition of equilibrium. ■

Lemma 4 shows that sellers that choose a different quality for their variety of the product must offer a different surplus to their customers. The next lemma, which follows from Lemma 3, can be interpreted as a converse to Lemma 4, in that it establishes that sellers that choose the same quality for their variety of the product must offer the same surplus to their customers.

Lemma 5. *In any equilibrium, all sellers with a variety of the same quality offer the same surplus.*

Proof. To show a contradiction, suppose that sellers with a variety of quality y_0 offer different surpluses. Let $s_{\ell,0}$ and $s_{h,0}$ denote the lowest and highest surplus offered by these sellers. The profits for a seller choosing $(y_0, s_{\ell,0})$ and for a seller choosing $(y_0, s_{h,0})$ are

$$V(y_0, s_{\ell,0}) = -y_0^\gamma + \theta\lambda e^{-\lambda(1-F(s_{\ell,0}))}(y_0 - s_{\ell,0}), \tag{2.14}$$
$$V(y_0, s_{h,0}) = -y_0^\gamma + \theta\lambda e^{-\lambda(1-F(s_{h,0}))}(y_0 - s_{h,0}). \tag{2.15}$$

Since H cannot have a mass point at y_0, $F(s_{\ell,0}) = F(s_{h,0})$ and, in turn, $V(y_0, s_{\ell,0}) > V(y_0, s_{h,0})$. Therefore, a seller choosing $(y_0, s_{\ell,0})$ does not maximize its profit. ■

Lemma 5 implies that there exists a function $s(y)$ that maps the quality of the variety of a seller into the surplus offered by the seller. Lemma 4 implies that the function $s(y)$ is strictly increasing in the quality of the variety of the seller. When taken together, Lemmas 4 and 5 imply that $F(s(y)) = H(y)$. That is, the fraction of sellers who offer a surplus no greater than $s(y)$ is equal to the fraction of sellers with a variety of quality no greater than y. We use this observation to establish that the support of the quality distribution H has no gaps.

Lemma 6. *In any equilibrium, the support of the quality distribution H has no gaps.*

Proof. To show a contradiction, suppose that the support of H has a gap between y_1 and y_2 with $y_1 < y_2$. Since $s(y)$ is strictly increasing in y, it follows that $s(y_1) < s(y_2)$. Since $F(s(y)) = H(y)$ and $H(y_1) = H(y_2)$, it follows that $F(s(y_1)) = F(s(y_2))$ and, since the support of the F distribution has no gaps, $s(y_1) = s(y_2)$. A contradiction. ■

Lemma 6 states that the support of the quality distribution H has no gaps. Lemma 3 states that the support of the quality distribution H is non-degenerate. When taken together, Lemmas 3 and 6 imply that the support of H is some interval $[y_\ell, y_h]$ with $y_\ell < y_h$. Since Lemma 1 implies that all qualities y on the support of H are strictly positive, it follows that $y_\ell > 0$.

We can now derive the unique candidate equilibrium. In any equilibrium, the profit $V(y, s(y))$ for a seller designing a variety of the product with quality y and offering a surplus $s(y)$ must equal the maximized profit V^* for all y on the support $[y_\ell, y_h]$ of the distribution H. That is, for all $y \in [y_\ell, y_h]$, we must have

$$V^* = -y^\gamma + \theta\lambda e^{-\lambda(1-F(s(y)))}(y - s(y)). \tag{2.16}$$

Since $(y, s(y))$ attains the maximized profit, the surplus $s(y)$ must maximize the seller's profit given y and, hence, the derivative of the right-hand side of (2.16) with respect to s must be equal to zero at $s(y)$. That is, for all $y \in [y_\ell, y_h]$, we must have

$$\lambda F'(s(y))(y - s(y)) - 1 = 0. \tag{2.17}$$

For the same reason, the quality y must maximize the seller's profit given the surplus and, hence, the derivative of the right-hand side of (2.16) with respect to y must be equal to zero. That is, for all $y \in [y_\ell, y_h]$, we must have

$$-\gamma y^{\gamma-1} + \theta\lambda e^{-\lambda(1-F(s(y)))} = 0. \tag{2.18}$$

Using the fact that $F(s(y)) = H(y)$, we can rewrite (2.18) as

$$\gamma y^{\gamma-1} = \theta\lambda e^{-\lambda(1-H(y))}. \tag{2.19}$$

Equation (2.19) states that, for any $y \in [y_\ell, y_h]$, the distribution H is such that the marginal cost of designing a variety of quality y, which is given by $\gamma y^{\gamma-1}$, is equal to the marginal benefit of designing a variety of quality y, which is given by the quantity $\theta\lambda \exp(-\lambda(1 - H(y)))$ of output sold by a seller with a variety of quality y that offers a surplus at the $H(y) = F(s(y))$ quantile of the surplus distribution. Solving (2.19) with respect to the equilibrium quality distribution H yields

$$H(y) = 1 + \frac{\gamma - 1}{\lambda}\log(y) + \frac{1}{\lambda}\log\left(\frac{\gamma}{\theta\lambda}\right). \tag{2.20}$$

The lower bound y_ℓ of the support of H must be such that $H(y_\ell) = 0$ and the upper bound y_h of the support of H must be such that $H(y_h) = 1$. Solving these equations with respect to y_ℓ and y_h yields

$$y_\ell = \left(\frac{\theta\lambda e^{-\lambda}}{\gamma}\right)^{\frac{1}{\gamma-1}}, \tag{2.21}$$

$$y_h = \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}. \tag{2.22}$$

Using (2.21), we can write the equilibrium quality distribution H and its density as

$$H(y) = \frac{\gamma-1}{\lambda}\left[\log y - \log y_\ell\right], \tag{2.23}$$

$$H'(y) = \frac{\gamma-1}{\lambda}\frac{1}{y}. \tag{2.24}$$

The equilibrium quality distribution H is log-uniform over the interval $[y_\ell, y_h]$. The lowest quality on the support of H is such that the marginal cost of designing a product with quality y_ℓ, which is given by $\gamma y_\ell^{\gamma-1}$, is equal to the marginal benefit of designing a product with quality y_ℓ, which is given by the trade volume $\theta\lambda\exp(-\lambda)$ for a seller who offers the lowest surplus in the market and, hence, only trades with captive buyers. The highest quality on the support of H is such that the marginal cost of designing a product with quality y_h, which is given by $\gamma y_h^{\gamma-1}$, is equal to the marginal benefit of designing a product with quality y_h, which is given by the trade volume $\theta\lambda$ for a seller that offers the highest surplus in the market and, hence, trades with all the buyers it meets. For any $H(y) \in (0,1)$, the marginal cost of designing a product with quality y is equal to the marginal benefit, which is given by the trade volume associated with being at the $H(y)$ quantile of the surplus distribution.

Next, we solve for the equilibrium surplus function $s(y)$. Using the fact that $F'(s(y))s'(y) = H'(y)$ and $H'(y) = (\gamma-1)/(\lambda y)$, we can write (2.17) as

$$\frac{\gamma-1}{\lambda}\cdot\frac{y-s(y)}{s'(y)y} = 1. \tag{2.25}$$

The expression above states that the marginal benefit of offering the surplus $s(y)$, which is given by the left-hand side of (2.25), must equal the marginal cost of offering the surplus $s(y)$, which is given by the right-hand side of (2.25). The expression above is a first-order differential equation for $s(y)$. The boundary condition for the solution of the differential equation is $s(y_\ell) = 0$, since Lemma 2, Lemma 6 and Lemma 4 respectively imply that the lower bound on the support of the surplus distribution F is zero, the lower bound on the support of the quality distribution H is y_ℓ, and the surplus function $s(y)$ is strictly increasing in y. The unique solution to (2.25) that satisfies $s(y_\ell) = 0$ is

$$s(y) = \frac{\gamma-1}{\gamma}\left[y - y_\ell^\gamma y^{1-\gamma}\right]. \tag{2.26}$$

The unique candidate equilibrium is given by the quality distribution H in (2.23) and the surplus function $s(y)$ in (2.26). The candidate equilibrium is a legitimate equilibrium. To see why this is the case, consider the optimality of the surplus $s(y)$ given a quality $y \in [y_\ell, y_h]$. The derivative of the seller's profit function with respect to the surplus s is zero at $s = s(y)$, strictly positive for $s < s(y)$, and strictly negative for $s > s(y)$. Hence, for any quality $y \in [y_\ell, y_h]$, a seller maximizes its profit by offering the surplus $s(y)$. Next, consider the optimality of the quality choice y. By construction of H, the seller attains a profit of V^* for any $y \in [y_\ell, y_h]$. For any $y < y_\ell$, the seller finds it optimal to offer the surplus $s(y_\ell) = 0$ and attains profit

$$V(y, s(y_\ell)) = -y^\gamma + \theta\lambda e^{-\lambda} y. \tag{2.27}$$

For any $y > y_h$, the seller finds it optimal to offer the surplus $s(y_h)$ and attains profit

$$V(y, s(y_h)) = -y^\gamma + \theta\lambda(y - s(y_h)). \tag{2.28}$$

The expression in (2.27) takes the value V^* at $y = y_\ell$, and its derivative with respect to y is strictly positive for all $y < y_\ell$. The expression in (2.28) takes the value V^* at $y = y_h$, and its derivative with respect to y is strictly negative for all $y > y_h$. Hence, for any $y \notin [y_\ell, y_h]$, the seller's profit is strictly smaller than V^*.

We have thus established the following theorem.

Theorem 1: (Existence and uniqueness). *There exists a unique equilibrium. In equilibrium:*

1. *The quality distribution $H(y)$ is given by (2.23), and is log-uniform over the support $[y_\ell, y_h]$ with $0 < y_\ell < y_h$, where y_ℓ,* is *given by (2.21) and y_h given by (2.22).*
2. *The surplus function $s(y)$ is given by (2.26), and such that $s(y_\ell) = 0$ and $s'(y_\ell) > 0$.*

The most important result in Theorem 1 is that, if sellers can choose the quality of their variety of the product, the existence of search frictions in the market not only requires dispersion in the payoffs offered by sellers to buyers, but also requires dispersion in the quality chosen by sellers. The equilibrium quality distribution $H(y)$ is such that the marginal cost to a seller from designing a variety of quality y is equal to the marginal benefit, which is given by the quantity of output sold by a seller that is at the $H(y)$ quantile of the surplus distribution. Given an isoelastic design cost function, the equilibrium quality distribution $H(y)$ is log-uniform over the interval $[y_\ell, y_h]$, where y_ℓ is such that the marginal design cost is equal to the volume sold by a seller offering the lowest surplus in the market, and y_h is such that the marginal design cost is equal to the volume sold by a seller offering the highest surplus in the market.

Even though the proof of Theorem 1 is quite lengthy, there is a simpler, albeit incomplete, argument that can explain the fact that the quality distribution is non-degenerate as well as its other properties. Consider an equilibrium in which every seller chooses the same quality y^*. As explained in Butters (1977), Varian (1980) and Burdett and Judd (1983), each seller will offer a different surplus to its buyers but all sellers will enjoy the same revenues. The seller that offers the lowest surplus, $s_\ell = 0$, will make a higher revenue per unit but will only sell to captive buyers. The seller that offers the highest surplus, $s_h > 0$, will make a lower revenue per unit but will sell to all the buyers it meets. If a seller deviates from y^* to $y^* - \epsilon$, its design cost is $(y^* - \epsilon)^\gamma$ and its revenue is $\theta\lambda \exp(-\lambda)(y^* - \epsilon)$ since the seller with the lowest quality in the market finds it optimal to offer a surplus of $s_\ell = 0$ and only trade with captive buyers. If a seller deviates from y^* to $y^* + \epsilon$, its design cost is $(y^* + \epsilon)^\gamma$ and its revenue is $\theta\lambda(y^* + \epsilon - s_h)$ since the seller with the highest quality in the market finds it optimal to offer a surplus of s_h and trade with all the buyers it meets. Therefore, to the left of y^*, the marginal cost of increasing y is $\gamma y^{*\gamma-1}$ and the marginal benefit is $\theta\lambda \exp(-\lambda)$. To the right of y^*, the marginal cost of increasing y is $\gamma y^{*\gamma-1}$ and the marginal benefit is $\theta\lambda$. Since the marginal cost is the same to the left and to the right of y^*, but the marginal benefit is strictly smaller to the left than to the right of y^*, there is always a profitable deviation away from y^*.

In order to eliminate these profitable deviations, the discontinuity in the marginal benefit of increasing y must be eliminated in equilibrium. This is accomplished by generating a quality distribution H with the following properties: (i) At the lower bound y_ℓ of the support of H, the marginal cost of increasing y is equal to the marginal benefit $\theta\lambda \exp(-\lambda)$ associated with being the lowest-quality seller in the market; (ii) At the upper bound y_h of the support of H, the marginal cost of increasing y is equal to the marginal benefit $\theta\lambda$ associated with being the highest-quality seller in the market; (iii) Everywhere on the support of H, the marginal cost of increasing y is equal to the marginal benefit $\theta\lambda \exp(-\lambda(1 - H(y)))$ associated with being a seller at the $H(y)$ quantile. For any $y < y_\ell$, the marginal cost is strictly smaller than the marginal benefit, as the marginal cost is strictly increasing in y and the marginal benefit is constant for all $y \leq y_\ell$. For any $y \geq y_h$, the marginal cost is strictly greater than the marginal benefit, as the marginal cost is strictly increasing in y and the marginal benefit is constant for all $y \geq y_h$. For any $y \in (y_\ell, y_h)$, the marginal cost and the marginal benefits are equal. Therefore, a seller's payoff attains its maximum if and only if $y \in [y_\ell, y_h]$. There are no profitable deviations.

The argument above is incomplete as it only establishes the existence of an equilibrium in which the quality distribution has properties (i), (ii) and (iii). The full proof of Theorem 1 establishes that the unique equilibrium is such that the quality distribution has properties (i), (ii) and (iii). The proof is useful because it shows that search frictions—by generating dispersion in the quantity of output sold by different sellers and, hence, by generating dispersion in the marginal benefit of investing in quality—require the equilibrium to feature dispersion in the quality of the variety designed by different sellers. The proof also shows how to construct the quality distribution. The proof does not rely on the design cost function being isoelastic. Indeed, the argument applies to any design cost function $c(y)$ that is strictly increasing, strictly convex and such that $c'(0) = 0$

and $c'(\infty) = \infty$. In addition, the proof does not rely on the buyers contacting a number of sellers drawn from a Poisson distribution. Indeed, the proof applies to any search process such that a buyer has a strictly positive probability of contacting one seller, and a strictly positive probability of contacting multiple sellers.

2.3 Efficiency

We next formulate and solve the problem of a utilitarian social planner and, in turn, establish the welfare properties of equilibrium. A utilitarian social planner maximizes the sum of the payoffs to buyers and sellers in the market. The planner chooses the quality of the variety for each seller and from which seller each buyer should purchase the product. The first choice is conveniently represented by a non-decreasing function $y(x)$ that maps the quantile x of a seller in the distribution of quality into the quality y of their product. The second choice is trivial. If a buyer contacts n sellers, the planner finds it optimal for the buyer to purchase the product from the seller that has a product with the highest quality since this maximizes the sum of the sellers' and buyers' payoffs.

Formally, the problem of a utilitarian social planner is

$$\begin{aligned} W &= \max_{y(x)} \left(-\int_0^1 y(x)^\gamma dx + \int_0^1 \theta\lambda e^{-\lambda(1-x)} y(x) dx \right) \\ & \quad \text{s.t. } y(x) \text{ non-decreasing in } x. \end{aligned} \tag{2.29}$$

The first term in the parentheses is the total cost to the sellers from designing their varieties of the product. This cost is the integral for the quantile x going from 0 to 1 of the cost $y(x)^\gamma$ of designing a variety of quality $y(x)$. The second term is the total benefit to buyers and sellers from trading. This benefit is the integral for the quantile x going from 0 to 1 of the quantity sold by a seller with quality $y(x)$ times the sum of the payoff to the seller and each of its customers. The quantity sold by a seller with quality $y(x)$ is given by the number of buyers $\theta\lambda \exp(-\lambda(1-x))$ who meet the seller and who are not in contact with any seller at a higher quantile of the distribution. The sum of the payoffs to the seller and each of its customers is simply $y(x)$.

It is useful to express the planner's choice $y(x)$ as $y(0) + \int_0^x y'(z)dz$. That is, it is useful to break down the planner's choice into the choice of the lowest quality designed by a seller, $y(0)$, and the derivative of the quality with respect to the quantile, $y'(x)$. Abstracting momentarily from the constraint $y'(x) \geq 0$, the planner's problem (2.29) can then be written as

$$W = \max_{y(0), y'(x)} -\int_0^1 \left(y(0) + \int_0^x y'(z)dz \right)^\gamma dx + \int_0^1 \theta\lambda e^{-\lambda(1-x)} \left(y(0) + \int_0^x y'(z)dz \right) dx \tag{2.30}$$

The first order condition of (2.30) with respect to $y(0)$ is

$$\int_0^1 \gamma y(x)^{\gamma-1} dx = \int_0^1 \theta\lambda e^{-\lambda(1-x)} dx. \tag{2.31}$$

The first order condition of (2.30) with respect to $y'(x)$ is

$$\int_x^1 \gamma y(z)^{\gamma-1} dz = \int_x^1 \theta\lambda e^{-\lambda(1-z)} dz. \tag{2.32}$$

The expressions (2.31) and (2.32) are intuitive. The left-hand side of (2.31) is the marginal cost of increasing $y(0)$, which is the additional design cost from raising the quality $y(x)$ for all $x \in [0, 1]$. The right-hand side of (2.31) is the marginal benefit of increasing $y(0)$, which is the additional payoff to buyers and sellers from increasing the quality of the product in all trades. The left-hand side of (2.32) is the marginal cost of increasing $y'(x)$, which is the additional design cost from raising the quality of the variety of all sellers at a quantile z above x. The right-hand side of (2.32) is the marginal benefit of increasing $y'(x)$, which is the additional payoff to buyers and sellers from increasing the quality of the variety of all sellers at a quantile z above x.

Differentiating (2.32) with respect to x yields

$$\gamma y(x)^{\gamma-1} = \theta\lambda e^{-\lambda(1-x)}. \tag{2.33}$$

The solution of (2.33) with respect to $y(x)$ is

$$y(x) = \left(\frac{\theta\lambda}{\gamma} \exp(-\lambda(1-x))\right)^{\frac{1}{\gamma-1}}. \tag{2.34}$$

The expression for $y(x)$ in (2.34) is the solution to the relaxed version of the planner's problem, in which the constraint $y'(x) \geq 0$ is removed. Since $y(x)$ in (2.34) is strictly increasing in x, it is also the solution to the actual planner's problem, the problem in which the constraint $y'(x) \geq 0$ is included.

In equilibrium, the quality distribution H is given by (2.23). The x-th quantile, $y(x)$, of the equilibrium quality distribution is given by the solution to the equation $H(y(x)) = x$, which is

$$\begin{aligned} y(x) &= y(0)\exp(x\lambda/(\gamma-1)) \qquad (2.35) \\ &= \left(\frac{\theta\lambda}{\gamma} \exp(-\lambda(1-x))\right)^{\frac{1}{\gamma-1}}, \end{aligned}$$

where the second line makes use of the fact that $y(0) = y_\ell$, which is given by (2.21). Comparing (2.35) and (2.34), one can see that the x-th quantile of the quality distribution is the same in equilibrium and in the solution to the social planner problem. Moreover, the quantity sold by a seller at the x-th quantile is the same in equilibrium and in the solution to the social planner problem since in equilibrium and in the planner's solution, a seller at the x-th quantile trades with all the buyers contacted who do not contact any seller at a higher quantile.

We have thus established the following.

Theorem 2 (Efficiency) *The equilibrium is efficient in the sense that it decentralizes the solution to the problem of a utilitarian social planner.*

Theorem 2 implies that the socially efficient quality distribution is non-degenerate. Why would a social planner ask different sellers to choose different qualities for their varieties of the product? Suppose the social planner asks all sellers to choose the same quality y^*. In this case, a buyer who contacts only one seller generates a social value of y^* and a buyer who contacts multiple sellers also generates a social value of y^*. Suppose now that the social planner asks different sellers to choose a different quality while keeping the average quality fixed at y^*. In this case, a buyer who contacts only one seller generates an expected social value of y^*. A buyer who contacts multiple sellers, however, can choose between different sellers and their different qualities and pick the best. Hence, a buyer who contacts multiple sellers generates an expected social value greater than y^*. By spreading the quality distribution around y^*, the planner increases the expected value of a transaction between a buyer and a seller. On the other hand, since $c(y)$ is convex, the planner increases the total design by spreading the quality distribution around y^*. Yet, since a small spread in the quality distribution around y^* has a first-order effect on the expected value of a transaction between a buyer and a seller and only a second-order effect on the total design cost, the planner finds it optimal to have different sellers choose a different quality for their variety of the product.

More specifically, Theorem 2 states that the socially efficient quality distribution is exactly the same as in equilibrium. To understand this, consider the planner's choice of $y(x)$, the quality for a seller at the x-th quantile of the distribution. The marginal cost of increasing $y(x)$ is the same for the planner and for a seller. The marginal benefit of increasing $y(x)$ for the planner is the value of increasing the sum of the payoffs to the seller and all of its buyers given that the seller trades with a buyer if and only if that buyer is not in contact with a seller at a higher quantile of the distribution. This is the same as the marginal benefit of increasing $y(x)$ for a seller, since a seller at the x-th quantile of the quality distribution finds it optimal to locate at the x-th quantile of the surplus distribution. It then follows that the planner's optimality condition coincides with the one of a private seller, and the solution to the planner's problem coincides with the equilibrium.

The intuition behind Theorem 2 makes it clear that the efficiency of equilibrium does not rely on the assumption that the design cost function $c(y)$ is isoelastic, nor does it depend on the assumption that a buyer contacts a number of sellers drawn from a Poisson distribution. The efficiency of equilibrium extends to any design cost function $c(y)$ that is increasing, convex and such that $c'(0) = 0$ and $c'(\infty) = \infty$. Similarly, the efficiency of equilibrium extends to any search process such that a buyer has a strictly positive probability of contacting only one seller and a strictly positive probability of contacting multiple sellers.

3 Declining search frictions

We now use our model of vertical differentiation in a frictional product market to understand the consequences of declining search frictions due to, for example, improvements in information and communication technologies that make it easier for buyers to come into contact with sellers. We model declining search frictions as an increase in λ, the expected number of sellers that a buyer contacts. In terms of market structure, we find that declining search frictions increase sales concentration – in the sense that the share of output sold by the biggest sellers becomes larger and larger – and increase quality polarization – in the sense that the varieties designed by the biggest sellers become better and better, while the varieties designed by the smallest sellers become worse and worse. In terms of payoffs, we find that declining search frictions increase the total gains from trade while keeping the shares of the gains from trade accruing to buyers and sellers unchanged. Similarly, we find that declining search frictions increase total welfare, but the share of welfare accruing to buyers converges to one, while the share of welfare accruing to sellers converges to zero.

3.1 Declining search frictions and market structure

We begin by considering the effect of declining search frictions on the quantity of output sold by different sellers. A seller at the x-th quantile of the quality distribution H is a seller at the x-th quantile of the surplus distribution F and, hence, the quantity it sells is given by

$$q(x) = \theta\lambda e^{-\lambda(1-x)}. \tag{3.1}$$

The derivative of $q(x)$ with respect to λ is

$$\frac{dq(x)}{d\lambda} = q(x)\frac{1-\lambda(1-x)}{\lambda}. \tag{3.2}$$

The derivative in (3.2) is negative for all $x < 1 - 1/\lambda$, and positive for all $x > 1 - 1/\lambda$. Therefore, as long as $\lambda > 1$, a decline in search frictions causes the quantity sold by sellers at the bottom of the distribution to fall, and the quantity sold by sellers at the top of the distribution to rise. This finding is easy to understand. As search frictions decline, buyers come into contact with more sellers. For this reason, buyers are more likely to purchase the good from a seller that offers a high surplus and are less likely to purchase the good from a seller that offers a low surplus.

Declining search frictions lead to an increase in sales concentration. As one can see from (3.1), sellers at a higher quantile of the distribution are larger in the sense that they sell more units of output than do sellers at a lower quantile of the distribution. As one can see from (3.2), declining search frictions further increase the gap between the quantity sold by sellers at a higher quantile of the distribution and the quantity sold by sellers at a lower quantile of the distribution and, hence, declining search frictions lead to an increase in market concentration. Formally, let $Q(z)$ denote the fraction of output sold by the z largest sellers—i.e., the sellers that are above the $(1-z)$-th quantile of the sales distribution. Since the size of a seller is monotonic in the seller's quality, $Q(z)$ is given by

$$Q(z) = \frac{\int_{1-z}^{1} q(x)dx}{\int_{0}^{1} q(x)dx} = \frac{1-e^{-\lambda z}}{1-e^{-\lambda}}. \tag{3.3}$$

The derivative of $Q(z)$ with respect to λ is

$$\frac{dQ(z)}{d\lambda} = \frac{e^{-\lambda(1+z)}\left(1-z+ze^{\lambda}-e^{\lambda z}\right)}{\left(1-e^{-\lambda}\right)^{2}} > 0. \tag{3.4}$$

Next, we consider the effect of declining search frictions on the quality distribution. As shown earlier in (2.35), the quality $y(x)$ of the variety produced by a seller at the x-th quantile of the H distribution is given by

$$y(x) = \left(\frac{\theta\lambda}{\gamma}e^{-\lambda(1-x)}\right)^{\frac{1}{\gamma-1}}. \tag{3.5}$$

The derivative of $y(x)$ with respect to λ is

$$\frac{dy(x)}{d\lambda} = \frac{y(x)}{\gamma-1}\cdot\frac{1-\lambda(1-x)}{\lambda}. \tag{3.6}$$

The derivative (3.6) is negative for all $x < 1-1/\lambda$, and positive for all $x > 1-1/\lambda$. As long as $\lambda > 1$, declining search frictions lead to a decline in the quality of the varieties designed by sellers at the bottom of the distribution and to an increase in the quality of the varieties designed by sellers at the top of the distribution. In this sense, declining search frictions lead to an increased polarization of the quality distribution. The intuition behind this finding is simple. As search frictions decline, the quantity of output sold by sellers at the bottom of the distribution falls and, as a result, these sellers find it optimal to reduce the quality of their varieties. In contrast, the quantity of output sold by sellers at the top of the distribution increases and, as a result, these sellers find it optimal to increase the quality of their varieties.

Using (2.26), (2.21), and (3.5), the surplus $s(x) \equiv s(y(x))$ offered by a seller at the x-th quantile of the distribution can be written as

$$s(x) = \frac{\gamma-1}{\gamma}\cdot\left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}\cdot e^{-\frac{\lambda(1-x)}{\gamma-1}}\left(1-e^{-\frac{\lambda\gamma x}{\gamma-1}}\right). \tag{3.7}$$

The derivative of the surplus $s(x)$ with respect to λ is

$$\frac{ds(x)}{d\lambda} = \frac{y(x)}{\gamma}\left[\left(1 - e^{-\frac{\lambda\gamma x}{\gamma-1}}\right)\frac{1-\lambda(1-x)}{\lambda} + \gamma x e^{-\frac{\lambda\gamma x}{\gamma-1}}\right]. \tag{3.8}$$

The derivative in (3.8) is strictly positive for all $x > 1 - 1/\lambda$. Declining search frictions affect the surplus offered by sellers through two channels. First, taking as given the quality distribution, declining search frictions increase competition and, hence, drive the surplus offered by a seller towards the quality of the seller's variety. Second, declining search frictions affect the quality distribution. For $x > 1 - 1/\lambda$, sellers increase the quality of their variety and this tends to increase the surplus offered by these sellers. For $x < 1-1/\lambda$, sellers reduce the quality of their variety and this tends to decrease the surplus offered by these sellers. Therefore, for $x > 1-1/\lambda$, the effect of declining search frictions on surplus is positive through both channels. For $x < 1-1/\lambda$, the effect of declining search frictions on surplus is positive through the first channel and negative through the second one, and the surplus may decrease.

The revenues for a seller at the x-th quantile of the distribution are

$$\begin{aligned} r(x) &= \theta\lambda e^{-\lambda(1-x)}(y(x) - s(x)) \\ &= \theta\lambda e^{-\lambda(1-x)}\left(\frac{\theta\lambda e^{-\lambda(1-x)}}{\gamma}\right)^{\frac{1}{\gamma-1}}\left(\frac{1}{\gamma} + \frac{\gamma-1}{\gamma}e^{-\frac{\lambda\gamma x}{\gamma-1}}\right) \\ &= \left(\frac{\theta\lambda e^{-\lambda(1-x)}}{\gamma}\right)^{\frac{\gamma}{\gamma-1}}\left(1 + (\gamma-1)e^{-\frac{\lambda\gamma x}{\gamma-1}}\right) \\ &= \left(\frac{\theta\lambda e^{-\lambda}}{\gamma}\right)^{\frac{\gamma}{\gamma-1}}\left(e^{\frac{\lambda\gamma x}{\gamma-1}} + (\gamma-1)\right). \end{aligned} \tag{3.9}$$

The derivative of $r(x)$ with respect to λ is

$$\frac{dr(x)}{d\lambda} = y(x)^{\gamma}\left(\frac{\gamma}{\gamma-1}\right)\left[\frac{1-\lambda(1-x)}{\lambda} - \left(1 - \frac{1}{\lambda}\right)(\gamma-1)e^{-\frac{\lambda\gamma x}{\gamma-1}}\right]. \tag{3.10}$$

To understand the effect of declining search frictions on the revenues of a seller, it is useful to look at the first line on the right-hand side of (3.9). A decline in search frictions affects the quantity sold by a seller at the x-th quantile of the distribution. This effect is positive for all $x > 1 - 1/\lambda$, and negative otherwise. A decline in search frictions affects the quality of the variety designed by a seller at the x-th quantile of the distribution. This effect is positive for $x > 1 - 1/\lambda$, and negative otherwise. A decline in search frictions affects the surplus offered by a seller at the x-th quantile of the distribution. This effect is positive for all $x > 1 - 1/\lambda$. For $x \leq 1 - 1/\lambda$, the overall effect of declining search frictions on revenues is negative. For $x > 1 - 1/\lambda$, the effect is generally ambiguous. For λ large enough, however, the effect is unambiguously positive for all $x > \overline{x}(\lambda)$, for some $\overline{x}(\lambda) \in (1 - 1/\lambda, 1)$.

Just like decreasing search frictions lead to an increase in sales concentration, they lead to an increase in revenue concentration. Formally, let $R(z)$ denote the fraction of revenues enjoyed by the z highest-revenue sellers—namely the sellers above the $(1-z)$-th quantile in the revenue distribution. Since the revenues of a seller are monotonic in the seller's quality, $R(z)$ is given by

$$R(z) = \frac{\int_{1-z}^{1} r(x)dx}{\int_{0}^{1} r(x)dx} = \frac{1 - e^{-\frac{\lambda\gamma z}{\gamma-1}} + \lambda\gamma z e^{-\frac{\lambda\gamma}{\gamma-1}}}{1 - e^{-\frac{\lambda\gamma}{\gamma-1}} + \lambda\gamma e^{-\frac{\lambda\gamma}{\gamma-1}}}. \tag{3.11}$$

The derivative of $R(z)$ with respect to λ is given by

$$\begin{aligned}\frac{dR(z)}{d\lambda} = \Phi \quad & z\left[(\lambda\gamma - 1)\, e^{-\frac{\lambda\gamma z}{\gamma-1}} - (\gamma-1)\, e^{-\frac{\lambda\gamma}{\gamma-1}} + e^{\frac{\lambda\gamma(1-z)}{\gamma-1}} - \gamma(\lambda-1) - 1\right] \\ & +\gamma\,(\lambda-1)\left(1 - e^{-\frac{\lambda\gamma z}{\gamma-1}}\right),\end{aligned} \tag{3.12}$$

which, with a bit of algebra, can be rewritten as

$$\begin{aligned}\frac{dR(z)}{d\lambda} = \Phi \quad & z\left[e^{\frac{\lambda\gamma(1-z)}{\gamma-1}} - 1 + (\gamma-1)e^{-\frac{\lambda\gamma z}{\gamma-1}}\left(1 - e^{-\frac{\lambda\gamma(1-z)}{\gamma-1}}\right)\right] \\ & +\gamma(\lambda-1)(1-z)\left(1 - e^{-\frac{\lambda\gamma z}{\gamma-1}}\right),\end{aligned} \tag{3.13}$$

where Φ is strictly positive. As long as $\lambda > 1$, the derivative of $R(z)$ with respect to λ is unambiguously positive. This finding is intuitive. Since $r(x)$ is strictly increasing in x, high-quality sellers enjoy higher revenues than low-quality sellers. As search frictions become smaller, the revenues enjoyed by the highest-quality sellers increase, while the revenues enjoyed by the lowest-quality sellers decrease. Hence, as search frictions become smaller, the share of revenues of the highest-quality sellers increases.

The following theorem summarizes our findings.

Theorem 3 (Declining search frictions and market structure). *The effects of declining search frictions on market structure are:*

1. *Higher sales and revenues concentration: The share $Q(z)$ of the output sold by the z largest sellers is increasing in λ for all $z \in [0,1]$; the share $R(z)$ of revenues of the z largest sellers is increasing in λ for all $z \in [0,1]$ and $\lambda > 1$.*
2. *Higher quality polarization: The quality $y(x)$ of the variety sold by a seller at the x-th quantile of the H distribution is increasing in λ for $x > 1-1/\lambda$, and decreasing in λ for $x < 1-1/\lambda$.*
3. *Higher surplus and revenues at the top: The surplus $s(x)$ offered by a seller at the x-th quantile of the H distribution is increasing in λ for all $x > 1-1/\lambda$; the revenues $r(x)$ for a seller at the x-th quantile of the H distribution are increasing in λ for all $x > \overline{x}(\lambda)$ and λ large enough, with $\overline{x}(\lambda) > 1-1/\lambda$.*

Theorem 3 contains a description of the kind of transformation that a product market undergoes as search frictions become smaller. When search frictions are relatively large, the product market is balanced. The difference in the quality of the variety designed by different sellers is small, the differences in the market share of different sellers is small, and all sellers enjoy similar revenues. As search frictions become smaller, the product market becomes more and more unbalanced. A small and decreasing fraction of sellers design a variety of the product of higher and higher quality, while a large and increasing fraction of sellers design a variety with decreasing quality. High-quality sellers capture a larger and larger share of the market and, in this sense, become super-star firms, while low-quality sellers capture a falling share of the market and, in this sense, become marginal. High-quality sellers enjoy larger and larger revenues, while the revenues of low-quality sellers decrease.

Search frictions act as a balancing force in the product market in two ways. First, as noted by Rosen (1981), frictions prevent higher quality sellers from taking over the market even though buyers unanimously prefer their variety of the product. Second, as shown here, frictions limit the difference in the return to quality investment for high- and low-ranked sellers and, hence, they limit the difference in the quality of the products designed by high- and low-ranked sellers. In addition, they limit the difference in the revenues enjoyed by different sellers.

The transformation of the product market that is triggered by declining search frictions leads to a number of phenomena that are typically associated with market dysfunction such as higher sales concentration, higher revenue concentration, and higher quality concentration. Yet, the increase in sales, revenue and quality concentration are an efficient response to a changing environment. Indeed, as search frictions become smaller, it is efficient for quality to become concentrated into a smaller fraction of sellers, and it is efficient for this smaller fraction of sellers to capture a larger share of sales, and it is efficient for this smaller fraction of sellers to capture a larger share of the revenues. Moreover, increasing sales, revenue and quality concentration are typically interpreted as symptoms of increasing market power. Instead, here they are consequences of the increase in competition caused by declining search frictions.

Finally, we note that the effect of declining search frictions described in Theorem 3 does not depend qualitatively on our assumptions about the shape of the design cost function or the nature of the search process. As long as declining search frictions lead to an increase in the number of sellers contacted by each buyer, they will lead to an increase in sales concentration. As long as the design cost function is increasing and convex, the increase in sales concentration will lead to a fanning out of the quality distribution. Since all sellers are ex-ante identical, the fanning out of the quality distribution will lead to an increase in the dispersion of revenues across sellers.

3.2 Declining search frictions, gains from trade and welfare

We next examine the effects of declining search frictions on the total gains from trade realized in the product market and on their division between buyers and sellers. The total gains from trade, G, are given by

$$\begin{aligned} G &= \int_{y_\ell}^{y_h} \theta\lambda e^{-\lambda(1-H(y))} y dH(y) \\ &= \int_{y_\ell}^{y_h} \theta(\gamma-1)\left(\frac{y}{y_h}\right)^{\gamma-1} dy \\ &= \theta\frac{\gamma-1}{\gamma}\left(\frac{1}{y_h}\right)^{\gamma-1}\left[y_h^\gamma - y_\ell^\gamma\right] \\ &= \theta\frac{\gamma-1}{\gamma}\left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}\left[1-e^{-\frac{\lambda\gamma}{\gamma-1}}\right], \end{aligned} \tag{3.14}$$

where the second line makes use of the fact that $H(y)$ is given by (2.23), the third line solves the integral in the second line, and the last line makes use of the fact that y_ℓ and y_h are given by (2.21) and (2.22). The derivative of the total gains from trade with respect to λ is

$$\frac{dG}{d\lambda} = \left(\frac{\gamma\lambda}{\theta}\right)^{-\frac{\gamma}{\gamma-1}}\left[1-(1-\lambda\gamma)e^{-\frac{\lambda\gamma}{\gamma-1}}\right] > 0. \tag{3.15}$$

The elasticity of G with respect to λ is

$$\frac{dG}{d\lambda}\frac{\lambda}{G} = \frac{1}{\gamma-1}\cdot\frac{1-(1-\lambda\gamma)e^{-\frac{\lambda\gamma}{\gamma-1}}}{1-e^{-\frac{\lambda\gamma}{\gamma-1}}} > 0. \tag{3.16}$$

As search frictions decline, the gains from trade increase. This finding is intuitive. As search frictions decline, buyers come into contact with more sellers. For this reason, buyers become more likely to purchase from sellers who offer a high surplus, who happen to be sellers with a high-quality variety and, hence, high gains from trade. Moreover, as search frictions decline, the sellers who offer a high surplus find it optimal to increase the quality of their variety and, hence, the gains from trade. The elasticity of the total gains from trade with respect to λ is positive and, as one can see from (3.16), it converges to $1/(\gamma-1)$ as $\lambda \to \infty$.

The gains from trade accruing to the buyers, G_b, are given by

$$\begin{aligned} G_b &= \int_{y_\ell}^{y_h} \theta\lambda e^{-\lambda(1-H(y))} s(y) dH(y) \\ &= \theta\left(\frac{\gamma-1}{\gamma}\right)^2\left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}\left[1-\left(1+\frac{\lambda\gamma}{\gamma-1}\right)e^{-\frac{\lambda\gamma}{\gamma-1}}\right], \end{aligned} \tag{3.17}$$

where the second line makes use of the fact that y_ℓ, y_h, $H(y)$ and $s(y)$ are given by (2.21), (2.22), (2.23) and (2.26). The derivative of the buyers' gains from trade with respect to

λ is

$$\frac{dG_b}{d\lambda} = \theta \left(\frac{\gamma-1}{\gamma}\right)^2 \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}} \frac{1}{\lambda(\gamma-1)} \left[1-\left(1+\frac{\lambda\gamma}{\gamma-1}(1-\lambda\gamma)\right) e^{-\frac{\lambda\gamma}{\gamma-1}}\right] > 0. \quad (3.18)$$

The elasticity of G_b with respect to λ is

$$\frac{dG_b}{d\lambda}\frac{\lambda}{G_b} = \frac{1}{\gamma-1} \cdot \frac{1-\left(1+\frac{\lambda\gamma}{\gamma-1}(1-\lambda\gamma)\right) e^{-\frac{\lambda\gamma}{\gamma-1}}}{1-\left(1+\frac{\lambda\gamma}{\gamma-1}\right) e^{-\frac{\lambda\gamma}{\gamma-1}}} > 0. \quad (3.19)$$

As search frictions decline, the buyers' gains from trade increase. This finding is also easy to understand. As search frictions decline, buyers come into contact with more sellers and, hence, they are more likely to purchase the good from sellers who offer a high surplus. Moreover, as search frictions decline, the sellers who offer a high surplus find it optimal to increase their offers. The elasticity of the buyers' gains from trade with respect to λ is positive and, as one can see from (3.19), it converges to $1/(\gamma-1)$ as $\lambda \to \infty$.

The gains from trade accruing to the sellers, G_s, are

$$\begin{aligned} G_s &= \int_{y_\ell}^{y_h} \theta\lambda e^{-\lambda(1-H(y))}(y-s(y))dH(y) \\ &= \theta\frac{\gamma-1}{\gamma^2}\left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}} \left[1-(1-\lambda\gamma)\, e^{-\frac{\lambda\gamma}{\gamma-1}}\right]. \end{aligned} \quad (3.20)$$

The derivative of the sellers' gains from trade with respect to λ is

$$\frac{dG_s}{d\lambda} = \theta\frac{1}{\lambda\gamma^2}\left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}} \left[1-\left(1-\lambda\gamma-\lambda\gamma^2+\lambda^2\gamma^2\right) e^{-\frac{\lambda\gamma}{\gamma-1}}\right] \quad (3.21)$$

The elasticity of G_s with respect to λ is

$$\frac{dG_s}{d\lambda}\frac{\lambda}{G_s} = \frac{1}{\gamma-1} \cdot \frac{1-\left(1-\lambda\gamma-\lambda\gamma^2+\lambda^2\gamma^2\right) e^{-\frac{\lambda\gamma}{\gamma-1}}}{1-(1-\lambda\gamma)\, e^{-\frac{\lambda\gamma}{\gamma-1}}}. \quad (3.22)$$

The sellers' gains from trade are increasing in λ for λ sufficiently large. As $\lambda \to \infty$, the elasticity of the sellers' gains from trade with respect to λ converges to $1/(\gamma-1)$, the same as the elasticity of the buyers' gains from trade and of the total gains from trade. This finding is surprising. Indeed, as search frictions decline, ex-post competition increases, and sellers are induced to offer a surplus that is closer and closer to the buyers' valuation. This argument would be complete if the distribution of qualities across varieties were fixed. However, as search frictions decline, low-quality sellers respond by further reducing the quality of their varieties, and high-quality sellers respond by further increasing the quality of their varieties. The fanning out of the quality distribution reduces the extent of ex-post competition faced by high-quality sellers. The force increasing and the force decreasing ex-post competition exactly offset each other, and the sellers' share of the gains from trade remains constant as $\lambda \to \infty$. The finding is not only surprising, but also critical. Indeed, if the sellers' gains from trade were falling as search frictions decline, sellers would not be able to recoup their investments, the quality distribution would fall and, eventually, the market would collapse.

The buyers' welfare, W_b, is equal to G_b. The sellers' welfare, W_s, is equal to V^*. In turn, V^* is equal to the payoff for a seller who designs a variety of the product with quality y_ℓ, offers a surplus of 0, and sells only to captive buyers. That is,

$$W_s = -y_\ell^\gamma + \theta\lambda e^{-\lambda} y_\ell = (\gamma - 1)\left(\frac{\theta\lambda e^{-\lambda}}{\gamma}\right)^{\frac{\gamma}{\gamma-1}}. \tag{3.23}$$

The derivative of W_s with respect to λ is

$$\frac{dW_s}{d\lambda} = -\left(\frac{\theta\lambda e^{-\lambda}}{\gamma}\right)^{\frac{\gamma}{\gamma-1}} \frac{\gamma(\lambda - 1)}{\lambda}. \tag{3.24}$$

As long as $\lambda > 1$, the derivative of the sellers' welfare with respect to λ is negative, and, as $\lambda \to \infty$, the sellers' welfare converges to zero. As search frictions decline, the sellers' gains from trade grow at the constant rate $1/(\gamma - 1)$. But, as search frictions decline, a larger and larger share of the sellers' gains from trade is spent investing in quality. In this sense, while declining search frictions do not increase ex-post competition among sellers, they do increase ex-ante competition and drive the sellers' ex-ante payoffs to zero.

Welfare is given by the sum of buyers' and sellers' payoffs. That is,

$$\begin{aligned} W &= W_b + W_s \\ &= \theta\left(\frac{\gamma-1}{\gamma}\right)^2 \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}} \left[1 - e^{-\frac{\lambda\gamma}{\gamma-1}}\right]. \end{aligned} \tag{3.25}$$

The derivative of W with respect to λ is

$$\frac{dW}{d\lambda} = \theta\left(\frac{\gamma-1}{\gamma}\right)^2 \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}} \frac{1}{\lambda(\gamma-1)} \left[1 - (1 - \lambda\gamma) e^{-\frac{\lambda\gamma}{\gamma-1}}\right] > 0. \tag{3.26}$$

The elasticity of W with respect to λ is

$$\frac{dW}{d\lambda}\frac{\lambda}{W} = \frac{1}{\gamma - 1} \cdot \frac{1 - (1 - \lambda\gamma) e^{-\frac{\lambda\gamma}{\gamma-1}}}{1 - e^{-\frac{\lambda\gamma}{\gamma-1}}} > 0. \tag{3.27}$$

The derivative of the buyers' welfare with respect to λ is positive, and the derivative of the sellers' welfare with respect to λ is negative. The effect of λ on the buyers' payoff dominates and total welfare increases with λ. This finding is not surprising, since the equilibrium allocation decentralizes the solution to the planner's problem, and the maximized value of the planner's problem is obviously increasing in λ. The elasticity of total welfare with respect to λ is positive and, as one can see from (3.27), it converges to $1/(\gamma-1)$ as $\lambda \to \infty$.

The following theorem summarizes our findings.

Theorem 4 (Declining search frictions, gains from trade, and welfare). *The effects of declining search frictions on payoffs are:*

1. *Higher gains from trade: Total gains from trade G increase with λ. As $\lambda \to \infty$, the elasticity of G with respect to λ is $1/(\gamma-1)$, and the buyers' and sellers' shares of the gains from trade remain constant.*
2. *Higher welfare: Welfare W increases with λ. As $\lambda \to \infty$, the elasticity of W with respect to λ is $1/(\gamma-1)$, the buyers' share of welfare converges to 1, and the sellers' share converges to 0.*

Theorem 4 describes the evolution of ex-post and ex-ante payoffs to buyers and sellers as search frictions become smaller. The total gains from trade realized in the market (i.e., the ex-post payoffs) grow unboundedly as search frictions become smaller and smaller, and the shares of the gains from trade going to buyers and to sellers converge to some strictly positive constants. The total welfare (i.e. the ex-ante payoffs) also grows unboundedly as search frictions become smaller and smaller, but the share of welfare going to sellers converges to zero and the share going to buyers converges to one. The fact that the sellers' share of ex-post payoffs converges to a positive constant implies that ex-post competition stabilizes at some imperfect level. The fact that the sellers' share of ex-ante profits converges to zero implies that ex-ante competition becomes perfect. The difference between the behavior of ex-ante and ex-post competition is due to the fact that, while sellers are ex-ante identical, they choose to differentiate themselves more and more as search frictions increase.

Theorems 3 and 4 can be used to understand the path taken by the market as search frictions become smaller. In particular, suppose that search frictions decline at some constant rate, in the sense that λ grows at some constant rate $g_\lambda > 0$. Theorem 3 implies that the quality distribution and the surplus distribution do not follow a Balanced Growth Path, in the sense that they do not follow a travelling wave in which every quantile of these distributions grows at the same rate. Indeed, as λ grows at the rate g_λ, the top quantiles of the quality distribution grow and the bottom quantiles fall. Similarly, as λ grows at the rate g_λ, the top quantiles of the surplus distribution grow and the bottom quantiles fall. Theorem 4, however, shows that the ex-post and ex-ante aggregate payoffs do eventually grow at constant rates. As λ grows at the rate g_λ, the total realized gains from trade grow at the rate $g_\lambda/(\gamma-1)$, and the shares going to buyers and sellers remain constant. As λ grows at the rate g_λ, the total welfare grows at the rate $g_\lambda/(\gamma-1)$, and the share going to buyers settles at 1, while the share going to sellers settles at 0.

Declining search frictions cause ex-ante and ex-post payoffs to grow at the rate $g_\lambda/(\gamma-1)$. The term $1/(\gamma-1)$ is the rate of return to declining search frictions. The rate of return to declining search frictions is inversely related to the elasticity of the design cost function. There is a simple intuition for this finding. As search frictions decline, it is efficient for a smaller and smaller fraction of sellers to design varieties with higher and higher quality and to serve a larger and larger fraction of the buyers. The growth rate of the quality of the top varieties depends inversely on the elasticity of the design cost function. If the design cost function has a lower elasticity, the growth rate of the quality of the top varieties and, in turn, of gains from trade and welfare is higher. If the design cost function has a higher elasticity, the growth rate of the quality of the top varieties and, in turn, of gains from trade and welfare is lower.

Lastly, let us point out that declining search frictions lead to higher gains from trade and higher welfare independently of the specification of the design cost function and of the nature of the search process. However, the fact that the declining search frictions lead to a constant growth in the gains from trade and in welfare does rely on the choice of an isoelastic design cost function and a Poisson search process. With different design cost functions and search processes, declining search frictions would generate a growth rate in gains from trade and welfare that need not be constant over time.

4 Concluding Remarks

In this paper, we contribute to the literature on search-theoretic models of imperfect competition in the spirit of Butters (1977), Varian (1980), Burdett and Judd (1983) by allowing sellers to choose the quality of their product through an ex-ante costly investment, i.e., by allowing for endogenous heterogeneity. Our main finding is that search frictions not only cause sellers to offer different surpluses but also cause them to choose different qualities. We solve for the unique equilibrium in closed form and show that it is efficient in the sense that it coincides with the allocation that would be chosen by a utilitarian social planner.

We also contribute to the literature on market concentration and polarization.. We consider the effect on the market of a reduction in search frictions and show that as search frictions fall, sales and revenues become increasingly concentrated in a smaller and smaller number of sellers and the distribution of quality across sellers becomes more and more skewed. Our model thus offers an explanation for the recent increase in market concentration and polarization that is quite different from the one typically proposed in the literature.

References

[1] Acemoglu, D., and R. Shimer. 2000. "Wage and Technology Dispersion." *Review of Economic Studies*, 67: 585-607.

[2] Aghion, P., A. Bergeaud, T. Boppart, P. Klenow, and H. Li. 2021. "A Theory of Falling Growth and Rising Rents." NBER Working Paper.

[3] Autor, D., D. Dorn, L. Katz, C. Patterson, and J. Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms". *Quarterly Journal of Economics*, 135: 645–709.

[4] Bar-Isaac, H, G. Caruana, and V. Cunat. 2012. "Search, Design, and Market Structure." *American Economic Review*, 102: 1140-1160.

[5] Baye, M., J., Morgan, and P. Scholten. 2006. "Information, Search, and Price Dispersion." *Handbook of Economics and Information Systems*. Elsevier, Amsterdam.

[6] Bethune, Z., M. Choi, and R. Wright. 2020. ""Frictional Goods Markets: Theory and Applications." *Review of Economic Studies*, 87: 691-720.

[7] Bontemps, C., J. Robin and G. Van den Berg. 2000. "Equilibrium Search with Continuous Productivity Dispersion: Theory and Non-Parametric Estimation." *International Economic Review*, 41: 305-358.

[8] Burdett, K., and K. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica*, 51: 955-970.

[9] Burdett, K, and G. Menzio. 2018. "The (Q,S,s) Pricing Rule." *Review of Economic Studies*, 85: 892-928.

[10] Burdett, K., and D. Mortensen. 1998. "Wage Differentials, Employer Size and Unemployment." *International Economic Review*, 39: 257-293.

[11] Burdett, K., A. Trejos, and R. Wright. 2017. "A New Suggestion for Simplifying the Theory of Money," *Journal of Economic Theory*, 172: 423-450.

[12] Butters, G. 1977. "Equilibrium Distributions of Sales and Advertising Prices." *Review of Economic Studies*, 44: 465-491.

[13] De Loecker, J., J. Eeckhout, and G. Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *Quarterly Journal of Economics*, 135: 561-644.

[14] De Loecker, J., J. Eeckhout, and S. Mongey. 2021. "Quantifying Market Power and Business Dynamism in the Macroeconomy." NBER Working Paper.

[15] Gutierrez, G., and T. Philippon. 2019. "The Failure of Free Entry." NBER Working Paper.

[16] Head, A., L. Liu, G. Menzio, and R. Wright. 2012. "Sticky Prices: A New Monetarist Approach." *Journal of the European Economic Association*, 10: 939-973

[17] Hong, H., and M. Shum. 2006. "Using Price Distributions to Estimate Search Costs." *Rand Journal of Economics*, 37: 257-275.

[18] Kaplan, G., and G. Menzio. 2016. "Shopping Externalities and Self-Fulfilling Unemployment Fluctuations." *Journal of Political Economy*, 124: 771-825.

[19] Kehrig, M., and N. Vincent. 2021. "The Micro-Level Anatomy of the Labor Share Decline." *Quarterly Journal of Economics*, 136: 1031-1087.

[20] Martellini, P., and G. Menzio. 2020. "Declining Search Frictions, Unemployment, and Growth." *Journal of Political Economy*, 128: 4387-4437.

[21] Martellini, P., and G. Menzio. 2021. "Jacks of All Trades and Masters of One: Declining Search Frictions and Unequal Growth." *American Economic Review: Insights*. Forthcoming.

[22] Menzio, G. 2021. "Optimal Product Design: Implications for Competition and Growth under Declining Search Frictions." NBER Working Paper.

[23] Menzio, G., and N. Trachter. 2018. "Equilibrium Price Dispersion Across and Within Stores." *Review of Economic Dynamics*, 28: 205-220.

[24] Mortensen, D. and C. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies*, 61: 397-415.

[25] Robin, J., and S. Roux. 2002. "An Equilibrium Model of the Labour Market with Endogenous Capital and Two-Sided Search." *Annales d'Economie et de Statistique*, 67: 257-307.

[26] Rosen, S. 1981. "The Economics of Superstars." *American Economic Review*, 71: 845-858.

[27] Stigler, G. 1961. "The Economics of Information." *Journal of Political Economy*, 69: 213-225.

[28] Varian, H. 1980. "A Model of Sales." *American Economic Review*, 70: 651-659.

[29] Wang, L., L. Liu, and R. Wright. 2020. "Sticky Prices and Costly Credit," *International Economic Review*, 61: 37-70.