Vertical Differentiation in Frictional Product Markets*

James Albrecht
Georgetown University and IZA

Guido Menzio
NYU and NBER

Susan Vroman
Georgetown University and IZA

March 2023

Abstract

We study a search-theoretic model of imperfect competition in product markets where sellers make an ex ante investment in the quality of their variety of the product. Equilibrium exists and is unique. 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. Equilibrium is efficient. As search frictions decline, the market becomes more and more unbalanced, as a vanishing 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 at constant rates.

JEL Codes: D43, L13, D83,O40 .

Keywords: Vertical differentiation, Search frictions, Imperfect competition, Stiglerian growth.

*Albrecht: Department of Economics, Georgetown University, 37th and O Streets, ICC Building, Washington DC 20007 (email: albrecht@georgetown.edu). Menzio: Department of Economics, New York University, 19 West 4th Street, New York, NY 10012 (email: gm1310@nyu.edu). Vroman: Department of Economics, Georgetown University, 37th and O Streets, ICC Building, Washington DC 20007 (email: susan.vroman@georgetown.edu). We are especially grateful to Manuel Amador, Andy Atkeson, Anmol Bhandari, V.V. Chari, Greg Kaplan, Virgiliu Midrigan and Rob Shimer for suggestions that greatly improved the paper. We also wish to acknowledge comments by Marco Bassetto, Simon Gilchrist, Tom Holmes, Boyan Jovanovic, Erzo Luttmer, Ezra Oberfield, Christina Patterson, Chris Phelan, Thomas Philippon, Harald Uhlig, Pierre-Olivier Weill and audiences at the Federal Reserve Bank of Minneapolis, the University of Chicago, Boston University, Rutgers University, Georgetown University and NYU.

1 Introduction

The search-theoretic model of Butters (1977), Varian (1980), and Burdett and Judd (1983) has become one of the leading frameworks to study imperfect competition in the product market. The model is simple. Sellers set prices for the product. Buyers comes into contact with a finite number of randomly selected sellers and decide whether and where to purchase the product. As the fraction of buyers who are in contact with multiple sellers relative to the fraction of buyers who are in contact with a single seller goes from zero to infinity, equilibrium outcomes span the spectrum from pure monopoly to perfect competition. This framework has been usefully applied to the analysis of price dispersion (e.g., Baye, Morgan and Scholten 2006, Hong and Shum 2006, Kaplan et al. 2019, Bethune, Choi and Wright 2020), sales (Varian 1980, Menzio and Trachter 2018), price stickiness (Head, Liu, Menzio and Wright 2012, Burdett, Trejos and Wright 2017, Burdett and Menzio 2018), and macroeconomic phenomena (Kaplan and Menzio 2016, Nord 2022).

Despite its popularity, there is room to further develop this search-theoretic model of imperfect competition. In particular, it is typically assumed that different sellers either carry varieties of the product of the same quality or that they carry varieties with different exogenous qualities. We contribute to the development of the framework by endogenizing the quality of the varieties carried by different sellers. We consider a product market that is populated by ex-ante identical sellers and ex-ante identical buyers. Each seller chooses the quality y of its variety, where y represents the buyers' utility from consuming that variety. A seller's cost from choosing quality y is an increasing, convex, and isoelastic function of y. After a seller chooses y, it decides how much to charge for its variety. Each buyer demands one unit of the product. A buyer comes into contact with a finite number n of randomly selected sellers, where n is drawn from a Poisson distribution with mean λ. A buyer observes the quality and the price of the variety offered by each of the n sellers and then decides whether and where to purchase the good. Sellers produce their variety of the product to meet demand and do so at a zero marginal cost.

In Section 2, we establish the existence, uniqueness and properties of equilibrium. The marginal distribution $H(y)$ of the quality of the varieties designed by sellers is atomless, and its support is some non-degenerate interval $[y_\ell, y_h]$ with $0 < y_\ell < y_h$. The marginal distribution $F(s)$ of the surplus offered by sellers, where surplus s is defined as the difference between quality y and price p, is atomless, and its support is some non-degenerate interval $[s_\ell, s_h]$, with $0 = s_\ell < s_h$. The marginal distributions of quality and surplus are linked by a strictly increasing function $s(y)$ that maps the quality y of the variety designed by the seller to the surplus s offered by the seller to its customers. The equilibrium is efficient, in the sense that the equilibrium allocation solves the problem of a utilitarian social planner. In Section 3, we show that these properties of equilibrium extend to versions of the model in which: (i) the design cost function $c(y)$ is a generic function that is strictly increasing and strictly convex and such that $c(0) = 0$, $c'(0) = 0$ and $c'(\infty) = \infty$; (ii) the meeting process is a generic probability distribution over the number n of sellers contacted by a buyer, where the probability of $n = 1$ and the probability of $n > 1$ are both strictly positive; and (iii) the seller's constant marginal cost of production depends on the quality of the seller's variety. The third version of the model is isomorphic to the baseline model with net quality, defined as the difference between quality and marginal cost of production, taking the place of quality.

From the descriptive point of view, the main finding is that search frictions in the product market require sellers to vertically differentiate themselves. There is a simple intuition for this finding. Imagine an equilibrium in which all sellers design varieties of the same quality y. The sellers must offer different surpluses to their customers, for the same reasons as in Butters (1977), Varian (1980) and Burdett and Judd (1983). Namely, if all the sellers offered the same surplus to their customers, an individual seller could strictly increase its volume of sales by offering ϵ more surplus and stealing all the contested buyers. Now, consider the seller that offers the highest surplus. This seller's marginal benefit from increasing the quality of its variety is equal to its volume, which is relatively high because the seller offers the highest surplus. Next, consider the seller that offers the lowest surplus. That seller's marginal benefit from increasing the quality of its variety is also equal to its volume, but this is relatively low because the seller offers the lowest surplus. While the marginal benefit of increasing quality is different for the two sellers, their marginal cost is the same. Therefore, at least one of them can increase its profit by either increasing or decreasing the quality of its variety.

Since sellers with a variety of strictly higher quality find it optimal to offer strictly more surplus to their customers, the quality distribution H that eliminates all profitable deviations must have the following features: (i) the lowest quality y_ℓ on the support of H 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 y_h on the support of H 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 H is such that the marginal cost of increasing quality is equal to the quantity of output sold by the seller at the x-th quantile of the surplus distribution. These features allow us to solve for the equilibrium quality distribution $H(y)$ in closed form. The equilibrium surplus function $s(y)$ is then given by the solution to a differential equation that equates, for every y, the marginal cost and the marginal benefit of offering more surplus. The equilibrium surplus function $s(y)$ also admits a closed-form solution.

From the normative point of view, the main finding is that the equilibrium is efficient. Let us provide some intuition for this finding. For the seller and the planner, the marginal cost of increasing y is the same. For the seller, the marginal benefit of increasing y is equal to the seller's volume, which is determined by the seller's rank in the surplus distribution. For the planner, the marginal benefit of increasing y is also equal to the seller's volume, but volume depends on the seller's rank in the quality distribution. In equilibrium, the seller's rank in the surplus distribution is the same as the seller's rank in the quality distribution. Therefore, the marginal benefit of increasing y is the same for the seller and for the planner.

After establishing the robustness of our results in Section 3, in Section 4 we examine the effect of declining search frictions on the structure of the product market. We interpret declining search frictions as improvements in information, communication, or transportation technologies that make it easier for buyers to locate and access more sellers. We model declining search frictions as an increase in λ, the expected number of sellers contacted by each buyer. In equilibrium, as λ increases, high-quality sellers make more sales (they are larger) and enjoy higher revenues. Low-quality sellers make fewer sales (they are smaller) and they enjoy lower revenues. A decline in search frictions leads to an increase in sale concentration, in the sense that the share of sales made by large sellers increases, while the share of sales by small sellers decreases. A decline in search frictions leads to an increase in revenue concentration, in the sense that the share of sales made by high-revenue sellers increases, while the share of sales by low-revenue sellers decreases. A decline in search frictions leads to quality polarization—in the sense that high-quality sellers further improve their varieties, while low-quality sellers reduce the quality of their varieties. As search frictions become smaller and smaller, a vanishing measure of sellers takes over the market with varieties of increasing quality, while a growing measure of sellers designs varieties of lower and lower quality and experiences vanishing sales and revenues. These phenomena are the natural and efficient response of the market to declining search frictions. As search frictions become smaller, buyers come into contact with more sellers. Hence, high-quality sellers can serve more buyers and find it optimal to invest more in the quality of their variety. Low-quality sellers serve fewer buyers and find it optimal to invest less in the quality of their variety.

We then examine the effect of declining search frictions on gains from trade, i.e., the sum of the buyers' and the sellers' payoffs after the investment costs are made, and on welfare, i.e., the sum of the buyers' and sellers' payoffs including the sellers' cost of investment. A decline in search frictions leads to an increase in the gains from trade for both buyers and sellers. A decline in search frictions leads to an increase in welfare, to an increase in the buyers' welfare and to a decline in the sellers' welfare. These findings are intuitive. The decline in search frictions has positive effects on welfare and gains from trade since it allows buyers to locate higher quality sellers. The decline in search frictions increases the gains from trade captured by buyers because it has two countervailing effects on ex post competition. On the one hand, the decline in search frictions increases ex post competition by increasing the number of sellers from which a buyer can purchase. On the other hand, the decline in search frictions reduces ex post competition because it leads to a fanning out of the quality distribution. The decline in search frictions decreases the welfare captured by sellers because it unambiguously increases ex-ante competition.

If search frictions decline at a constant rate, in the sense that λ increases at a constant rate g_λ, the structure of the market does not follow a Balanced Growth Path (BGP).

Indeed, the shape of the quality distribution H and the shape of the surplus distribution F change and become more and more skewed. The market payoffs, however, do eventually converge to a BGP. In particular, gains from trade grow at a rate that converges to $g_\lambda/(\gamma-1)$, the buyers' shares of the gains from trade converges to $(\gamma-1)/\gamma$ and the sellers' share of the gains from trade converges to $1/\gamma$, where γ denotes the elasticity of the product design cost with respect to quality. Welfare grows at a rate that converges to $g_\lambda/(\gamma-1)$, the buyers' share of welfare converges to 1, and the sellers' share of welfare converges to 0. Declining search frictions lead to economic growth. The growth rate is $g_\lambda/(\gamma-1)$. This 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.

We believe that the model studied in this paper applies quite broadly. The model applies to many markets for durable consumer products. For example, in the market for espresso machines, sellers design machines that have different quality or have a different marginal cost of production, and buyers may not be aware of all of varieties available on the market because of information frictions or they may not be able to trade with some of the sellers because of physical costs. In this example, a decline in search frictions, broadly construed, would predict the rise of varieties with the largest difference between quality and marginal cost (e.g., the Nespresso machines). The model may also apply to the market for intermediate goods. For example, in the market for accounting software, developers design software that differs in quality and businesses may have limited information about all of the varieties of software that are available on the market. The model may also apply to trading platforms. E-commerce developers design platforms of different quality (as measured by, say, their visibility) and the sellers that want to trade their goods through a platform may have limited information about all the available alternatives. It is also useful to point out the limits to the applicability of the model. The model does not apply to markets in which buyers have very heterogeneous preferences, as this would induce sellers to differentiate horizontally rather than vertically (e.g., clothing). Nor does the model apply to markets in which buyers need (and search for) a particular variety of the product rather than any variety in a product category (e.g. the textbook for their Econ class). The model also does not apply to markets in which sellers face increasing marginal costs (e.g., the market for dentists).

Related literature. 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 allow for seller heterogeneity with respect to the quality of the product (see, e.g., Bontemps, Robin and Van den Berg 2000), the marginal cost (see, e.g., Menzio and Trachter 2018), or both (see, e.g., Wildenbeest 2011). In all these models, however, heterogeneity is exogenous. We contribute to this literature by studying a model 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 these models and their types 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 also find that, in any equilibrium, firms choose different levels of capital.

The second part of the paper relates to the literature on rising market concentration (see, e.g., Autor et al. 2020, De Loecker, Eeckhout and Unger 2020, De Loecker, Eeckhout and Mongey 2021, Kehrig and Vincent 2020). These papers document a recent increase in sales and revenue concentration, a decline in the labor share, and a rise in productivity dispersion. We show that declining search frictions may be a contributing factor to rising concentration measured both in terms of sales and revenues. Moreover, we show that declining search frictions lead to a fanning out of the productivity distribution. Interestingly, we show that the increase in concentration and the increase in quality polarization are efficient responses to an environment that becomes more and more competitive as search frictions become smaller and smaller. Hence, to the extent that increasing concentration and polarization are driven by declining search frictions, they are benign phenomena. In contrast, Gutierrez and Philippon (2019), De Loecker, Eeckhout and Unger (2020), 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 (2022) 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. 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' design 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.

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 that identical sellers choose different qualities and those who choose a higher quality offer higher surplus to their customers. We derive a 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 $y \geq 0$. In this section, we assume that $c(y)$ has the isoelastic form $c(y) = y^\gamma$ with $\gamma > 1$. In Section 3, we consider a general strictly increasing and strictly convex function $c(y)$ such that $c(0) = 0$, $c'(0) = 0$ and $c'(\infty) = \infty$. After paying the design cost $c(y)$, the seller can produce any quantity of its variety of the good at a constant marginal cost. In this section, we assume that the marginal cost of production is 0. In Section 3, we consider a general marginal cost of production that is allowed to depend on the quality of the seller's variety y. 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

[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.

seller that the buyer has contacted. A buyer contacts n randomly selected sellers. In this section, we assume that n is drawn from a Poisson distribution with mean $\lambda > 0$.[2] In Section 3, we consider the case in which n is drawn from an arbitrary distribution such that $n = 1$ with strictly positive probability and $n > 1$ with strictly positive probability. 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 highest surplus $s \geq 0$ by multiple sellers, he randomizes over these sellers.[3]

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. We then 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

[2] For search aficionados, this is the same meeting process as in Butters (1977) and, more recently, Menzio (2022). The meeting process is well-suited for studying the effect of vanishing search frictions, as it allows the expected number of meetings per buyer to grow to infinity. The alternative meeting process of Burdett and Judd (1983) assumes that a buyer either contacts one or two sellers. Clearly, this meeting process is less suitable for studying declining search frictions, as the average number of meetings per buyer is bounded. In the appendix, we analyze the properties of equilibrium and the effect of vanishing search frictions under the Burdett-Judd meeting process.

[3] In Lester et al. (2019) and Menzio and Trachter (2018), sellers post 2-dimensional offers. In our model, sellers post a quality y and a price p. In Lester et al. (2019) and Menzio and Trachter (2018), the sellers' volume depends separately on both dimensions of the offer. Therefore, solving for the equilibrium offer distribution requires solving for a multidimensional cumulative distribution function, which, in turn, requires finding a mapping between the different dimensions of a seller's offer. In our model, the sellers' volume only depends on the difference s between quality y and price p. Therefore, solving for the equilibrium offer distribution amounts to solving for a one-dimensional cumulative distribution function just as in the basic Burdett and Judd (1983) model. Our model is closer to models with exogenously heterogeneous sellers, where single-crossing is used to show that the ranking of sellers by productivity is the same as the ranking of sellers by surplus (see, e.g., Bontemps, Robin and Van den Berg 2000, Wildenbeest 2011).

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 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 offering a surplus $s \geq 0$ trades with a buyer who is in contact with k other sellers with probability

$$\pi_k(s) = F(s-)^k + \sum_{j=1}^{k} \binom{k}{j} \frac{\mu(s)^j F(s-)^{k-j}}{j+1}, \text{ for } k = 0, 1, 2, ... \tag{2.3}$$

where $F(s-)$ denotes the fraction of sellers that offer a surplus strictly smaller than s, and $\mu(s)$ denotes the fraction of sellers that offer a surplus equal to s. Obviously, a seller offering a surplus $s < 0$ never trades with a buyer who is in contact with k other sellers, for $k = 0, 1, 2, ...$ The expression in (2.1) is the probability that a buyer contacts n sellers, given that n is drawn from a Poisson distribution with mean λ. The expression in (2.2) is the measure of contacts by buyers who are in touch with exactly $k+1$ sellers. Since the measure of sellers is normalized to 1, θ_k is the expected number of buyers who contact a particular seller and are in touch with k other sellers. The expression in (2.3) is the probability that the seller offering s to its customers trades with a buyer who is in touch with k other sellers. The first term is the probability that the buyer is in touch with k other sellers who offer a surplus strictly smaller than s. In this case, the buyer purchases from the seller with probability 1. The second term is the probability that the buyer is in touch with j other sellers who offer a surplus equal to s and with $k - j$ other sellers who offer a surplus strictly smaller than s. In this case, the buyer purchases from the seller with probability $1/(j + 1)$.

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

$$\tilde{V}(y, 0) = -y^\gamma + \theta\lambda e^{-\lambda} y. \tag{2.4}$$

The first term on the right-hand side of (2.4) 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).

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 contacted buyers who are only in touch with other sellers that offer $s < s_0$. Moreover, the seller trades with a fraction of the contacted buyers who are in touch with other sellers that offer s_0 and with no other sellers that offer $s > s_0$. If the seller were to offer the surplus $s_0 + \epsilon$, it could trade with all of the contacted buyers who are only in touch with other sellers offering $s \leq s_0$. Since there is a mass of sellers offering exactly s_0, the seller discretely increases the quantity sold by offering $s_0 + \epsilon$ rather than s_0. 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,

$$\begin{aligned} V(y,s) &= -y^\gamma + \left[\sum_{k=0}^{\infty} \theta(k+1)\lambda^{k+1}\frac{e^{-\lambda}}{(k+1)!}F(s)^k\right](y-s) \\ &= -y^\gamma + \theta\lambda e^{-\lambda}\left[\sum_{k=0}^{\infty}\frac{\lambda^k F(s)^k}{k!}\right](y-s) \\ &= -y^\gamma + \theta\lambda e^{-\lambda(1-F(s))}(y-s), \end{aligned} \tag{2.5}$$

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

$$\begin{aligned} V(y_0, s_{\ell,0}) &= -y_0^{\gamma} + \theta\lambda e^{-\lambda(1-F(s_{\ell,0}))}(y_0 - s_{\ell,0}), && (2.6)\\ V(y_0, s_{h,0}) &= -y_0^{\gamma} + \theta\lambda e^{-\lambda(1-F(s_{h,0}))}(y_0 - s_{h,0}), && (2.7)\end{aligned}$$

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 (2.6) and (2.7) with respect to y are

$$\begin{aligned} V_y(y_0, s_{\ell,0}) &= -\gamma y_0^{\gamma-1} + \theta\lambda e^{-\lambda(1-F(s_{\ell,0}))}, && (2.8)\\ V_y(y_0, s_{h,0}) &= -\gamma y_0^{\gamma-1} + \theta\lambda e^{-\lambda(1-F(s_{h,0}))}. && (2.9)\end{aligned}$$

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

$$\begin{aligned} \theta\lambda e^{-\lambda(1-F(s_1))}(y_1 - s_1) &\geq \theta\lambda e^{-\lambda(1-F(s_2))}(y_1 - s_2), && (2.10)\\ \theta\lambda e^{-\lambda(1-F(s_2))}(y_2 - s_2) &\geq \theta\lambda e^{-\lambda(1-F(s_1))}(y_2 - s_1). && (2.11)\end{aligned}$$

Combining (2.10) and (2.11) yields

$$\left[e^{-\lambda(1-F(s_2))} - e^{-\lambda(1-F(s_1))}\right](y_2 - y_1) \geq 0. \qquad (2.12)$$

The inequality (2.12) 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

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

$$V(y_2, s) = -y_2^\gamma + \theta\lambda e^{-\lambda(1-F(s))}(y_2 - s), \tag{2.14}$$

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

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

$$V(y_0, s_{h,0}) = -y_0^\gamma + \theta\lambda e^{-\lambda(1-F(s_{h,0}))}(y_0 - s_{h,0}). \tag{2.16}$$

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 $(y_0, s_{h,0})$ 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 F is strictly increasing, $s(y_1) = s(y_2)$. 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

$$V^* = -y^\gamma + \theta\lambda e^{-\lambda(1-F(s(y)))}(y - s(y)). \tag{2.17}$$

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 (2.17) 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

$$\lambda F'(s(y))(y - s(y)) - 1 = 0. \tag{2.18}$$

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 (2.17) with respect to y must be equal to zero. That is, for all $y \in [y_\ell, y_h]$, we must have

$$-\gamma y^{\gamma-1} + \theta\lambda e^{-\lambda(1-F(s(y)))} = 0. \tag{2.19}$$

Using the fact that $F(s(y)) = H(y)$, we can rewrite (2.19) as

$$\gamma y^{\gamma-1} = \theta\lambda e^{-\lambda(1-H(y))}. \tag{2.20}$$

Equation (2.20) 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 (2.20) with respect to the equilibrium quality distribution H yields

$$H(y) = 1 + \frac{\gamma - 1}{\lambda} \log y + \frac{1}{\lambda} \log\left(\frac{\gamma}{\theta\lambda}\right). \tag{2.21}$$

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

$$y_\ell = \left(\frac{\theta\lambda e^{-\lambda}}{\gamma}\right)^{\frac{1}{\gamma-1}}, \tag{2.22}$$

$$y_h = \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}. \tag{2.23}$$

Using (2.22), we can write the equilibrium quality distribution H and its density as

$$H(y) = \frac{\log y - \log y_\ell}{\log y_h - \log y_\ell}, \tag{2.24}$$

$$H'(y) = \frac{1}{\log y_h - \log y_\ell}\frac{1}{y}. \tag{2.25}$$

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 (2.18) as

$$(\gamma - 1)\frac{y - s(y)}{s'(y)y} = 1. \tag{2.26}$$

The expression above states that the marginal benefit of offering the surplus $s(y)$, which is given by the left-hand side of (2.26), must equal the marginal cost of offering the surplus $s(y)$, which is given by the right-hand side of (2.26). 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 (2.26) that satisfies $s(y_\ell) = 0$ is

$$s(y) = \frac{\gamma - 1}{\gamma}\left(y - y_\ell^\gamma y^{1-\gamma}\right). \tag{2.27}$$

The unique candidate equilibrium is given by the quality distribution H in (2.24) and the surplus function $s(y)$ in (2.27). 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

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

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

$$V(y, s(y_h)) = -y^\gamma + \theta\lambda(y - s(y_h)). \tag{2.29}$$

The expression in (2.28) 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 (2.29) 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 (2.24) and is log-uniform over the support $[y_\ell, y_h]$ with $0 < y_\ell < y_h$, where y_ℓ* is *given by (2.22) and y_h is given by (2.23).*
2. *The surplus function $s(y)$ is given by (2.27), and is such that $s(y_\ell) = 0$ and $s'(y_\ell) > 0$. The surplus distribution $F(s)$ is continuous over the support $[s(y_\ell), s(y_h)]$.*

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. The equilibrium function $s(y)$ mapping the quality of the seller's variety into the surplus that it offers is strictly increasing, as sellers with a better variety have a higher rate of substitution between surplus and quantity sold. The shape of the equilibrium surplus function is such that the marginal cost and marginal benefit from offering more surplus are equalized for every seller's variety y.[4]

[4] The focus of our paper is not price dispersion. Yet, it is useful to briefly remark on the properties of

The proof of Theorem 1 is quite lengthy. There is, however, a simple intuition for the properties of the equilibrium quality distribution H. 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 trades 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 trades 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. This is accomplished by an atomless quality distribution H. Any other profitable deviation is eliminated by a quality distribution H with the following properties: (i) The lower bound y_ℓ of H is where the marginal cost of increasing quality is equal to $\theta\lambda \exp(-\lambda)$, which is the marginal benefit for the lowest-quality seller (as this seller finds it optimal to offer the lowest surplus); (ii) The upper bound y_h of H is where the marginal cost of increasing quality is equal to $\theta\lambda$, which is the marginal benefit for the highest-quality seller (as this seller finds it optimal to offer the highest surplus); (iii) For any $y \in (y_\ell, y_h)$, the marginal cost of increasing quality is equal to $\theta\lambda \exp(-\lambda(1 - H(y)))$, which is the marginal benefit for a seller at the $H(y)$ quantile

the equilibrium price distribution. The equilibrium features dispersion in the price of different varieties. It is easy to verify that the price function $p(y) = y - s(y)$ is strictly convex and, hence, there are sellers that post different prices. The equilibrium does not feature dispersion in the price of the same variety. Trivially, this is the case in our model where every seller is a producer/retailer that carries its own variety of the product. Less trivially, this would also be the case in a version of our model where there are multiple retailers that carry the variety of the same producer. Intuitively, the retailers' mixed strategy over prices would be purified by the producers' mixed strategy over quality. Does this finding imply that Butters (1977), Varian (1980) and Burdett and Judd (1983) are not robust theories of price dispersion for the same variety of the product? We think that one ought to be careful before jumping to this conclusion. First, in a version of the model in which different retailers carry the variety of the same producer and the number of producers is discrete, the equilibrium would feature price dispersion for the same variety (see, e.g., Burdett and Mortensen 1998). Second, in a version of the model in which different retailers carry the variety of the same producers and retailers are heterogeneous, the equilibrium would feature price dispersion for the same variety (see, e.g., Menzio 2022). Third, our model assumes that buyers are searching for a product category. In some cases, though, it may be more sensible to assume that buyers are searching for a specific variety. In these cases, the market would be better described by a different Burdett-Judd submarket for each variety, with each submarket featuring price dispersion.

of the quality distribution (as this seller finds it optimal to be at the $H(y)$ quantile of the surplus distribution). 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 benefit 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. Theorem 1 completes this intuition by showing that only a distribution H with properties (i), (ii) and (iii) eliminates all profitable deviations.

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

$$\begin{aligned} W \;=\; & \max_{y(x)} \left(-\int_0^1 y(x)^\gamma dx + \int_0^1 \theta\lambda e^{-\lambda(1-x)} y(x) dx \right) \\ & \text{s.t. } y(x) \text{ non-decreasing in } x. \end{aligned} \tag{2.30}$$

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 $y(x)$.

Let us consider a relaxed version of (2.30) in which we abstract from the monotonicity constraint on $y(x)$.[5] In the relaxed version of (2.30), the optimality condition for $y(x)$ is given by

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

[5] We are grateful to Manuel Amador for suggesting this approach.

The left-hand side of (2.31) is the marginal cost of increasing the quality $y(x)$ of the variety produced by a seller at the x-th quantile of the distribution. The right-hand side is the marginal benefit of increasing the quality $y(x)$ of the variety produced by a seller at the x-th quantile of the distribution, which is equal to the quantity sold by the seller. Isolating $y(x)$ in (2.31) yields

$$y(x) = \left(\frac{\theta\lambda}{\gamma} \exp(-\lambda(1-x)) \right)^{\frac{1}{\gamma-1}}. \tag{2.32}$$

Notice that $y(x)$ in (2.32) is strictly increasing in x. Therefore, $y(x)$ in (2.32) is not only the solution to the relaxed version of the social planner's problem, but it is also the solution to the original social planner's problem in (2.30).

In equilibrium, the quality distribution H is given by (2.24). 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

$$\begin{aligned} y(x) &= y(0)\exp(x\lambda/(\gamma-1)) \\ &= \left(\frac{\theta\lambda}{\gamma} \exp(-\lambda(1-x)) \right)^{\frac{1}{\gamma-1}}, \end{aligned} \tag{2.33}$$

where the second line makes use of the fact that $y(0) = y_\ell$, where y_ℓ is given by (2.22). Comparing (2.32) and (2.33), it is immediate to 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 of the quality distribution is the same in equilibrium and in the solution to the social planner problem. In equilibrium, the quantity sold by a seller at the x-th quantile is given by the number of buyers who are in contact with the seller and are not in touch with any other seller that offers them more surplus than the seller. In the planner's solution, the quantity sold by a seller at the x-th quantile is given by the number of buyers who are in contact with the seller and are not in touch with any other seller that carries a variety of higher quality. In equilibrium, the sellers that offer more surplus are the sellers that carry a variety of higher quality. Therefore, the quantity sold by a seller at the x-th quantile is the same in equilibrium as in the planner's solution.

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 cost by spreading the quality distribution around y^*. Yet, since a 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 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.

Theorem 2 states that the socially efficient quality distribution is indeed the same as in equilibrium. To understand this, let us contrast the marginal cost and marginal benefit of increasing $y(x)$ for the social planner and for a private seller. The marginal cost of increasing $y(x)$ for the planner is $\gamma y(x)^{\gamma-1}$. The marginal cost of increasing $y(x)$ for the seller is also $\gamma y(x)^{\gamma-1}$. The marginal benefit of increasing $y(x)$ for the planner is $\theta\lambda \exp\left(-\lambda(1-x)\right)$, which is equal to the quantity of output sold by a seller at the x-th quantile of the quality distribution. The marginal benefit of increasing $y(x)$ for the seller is equal to the increase in its revenues $\theta\lambda\exp(-\lambda(1-F(s(y(x))))) \cdot (y(x)-s(y(x)))$. Since the seller's choice of surplus is optimal, the envelope condition guarantees that we can compute the increase in the seller's revenues while keeping the surplus offered by the seller unchanged. Thus, the increase in the seller's revenues is equal to $\theta\lambda\exp(-\lambda(1-F(s(y(x)))))$, the quantity of output sold by the seller. Since the seller finds it optimal to choose a surplus such that its rank in the surplus distribution is the same as in the quality distribution, $F(s(y(x))) = x$. Thus, the increase in the seller's revenues is $\theta\lambda\exp(-\lambda(1-x))$, which is the same as the marginal benefit of increasing $y(x)$ to the planner. Since the marginal cost and the marginal benefit of increasing $y(x)$ are the same for the planner and for a private seller, it follows that the socially efficient quality distribution is the same as in equilibrium. Moreover, since the quantity sold by a seller at a given point of the quality distribution is the same in the solution to the planner's problem as in equilibrium, it also follows that the solution to the social planner's problem prescribes the same allocation that emerges in equilibrium.

3 Generalizations

In this section, we prove that Theorem 1 (existence, uniqueness and qualitative properties of equilibrium) and Theorem 2 (efficiency of equilibrium) extend to: (i) a version of the model in which the product design cost function is not isoelastic, but a generic strictly increasing and strictly convex function $c(y)$ of seller's quality y, such that $c(0) = 0$,

$c'(0) = 0$ and $c'(\infty) = \infty$; (ii) a version of the model in which the meeting process is not Poisson, but a generic probability distribution over the number n of sellers contacted by the buyer, such that the probability of $n = 1$ is strictly positive and the probability of $n > 1$ is strictly positive; (iii) a version of the model in which the marginal cost of producing a variety is not equal to zero, but to some constant that may depend on the seller's quality.

3.1 General design cost function

We first relax the assumption of an isoelastic design cost function $c(y) = y^{-\gamma}$, with $\gamma > 1$. Specifically, we consider a general design cost function $c(y)$ that is strictly increasing, strictly convex, and such that $c(0) = 0$, $c'(0) = 0$ and $c'(\infty) = \infty$. Lemmas 1-6 hold under this general cost function. Therefore, in any equilibrium, the profit of a seller must attain its maximum for all $(y, s(y))$ with $y \in [y_\ell, y_h]$. That is,

$$V^* = -c(y) + \theta\lambda e^{-\lambda(1-F(s(y)))} \left(y - s(y)\right). \tag{3.1}$$

Since $(y, s(y))$ maximizes the seller's profit, the derivative of the right-hand side of (3.1) with respect to s must be equal to 0 when evaluated at any $(y, s(y))$ with $y \in [y_\ell, y_h]$. That is,

$$\lambda F'(s(y))(y - s(y)) - 1 = 0. \tag{3.2}$$

For the same reason, the derivative of the right-hand side of (3.1) with respect to y must be equal to 0 when evaluated at any $(y, s(y))$ with $[y_\ell, y_h]$. That is,

$$-c'(y) + \theta\lambda e^{-\lambda(1-F(s(y)))} = 0. \tag{3.3}$$

Using the optimality conditions (3.2) and (3.3) and the strict monotonicity of $s(y)$, we can solve for the quality distribution $H(y)$ and the surplus function $s(y)$. Specifically, the strict monotonicity of $s(y)$ implies that $F(s(y))$ equals $H(y)$. Using $F(s(y)) = H(y)$, we can solve the optimality condition (3.3) with respect to $H(y)$ and obtain

$$H(y) = 1 - \frac{1}{\lambda}\log\frac{\theta\lambda}{c'(y)}. \tag{3.4}$$

Using $H(y_\ell) = 0$ and $H(y_h) = 1$, we can solve for y_ℓ and y_h and obtain

$$y_\ell = \psi\left(\theta\lambda e^{-\lambda}\right), \quad y_h = \psi\left(\theta\lambda\right), \tag{3.5}$$

where $\psi(\cdot)$ is defined as the inverse of the marginal design cost, $c'^{-1}(\cdot)$. Similarly, using $F'(s(y))s'(y) = H'(y)$ and the density of the quality distribution given by the derivative of (3.4), we can rewrite the optimality condition (3.2) as

$$s'(y) = \frac{c''(y)}{c'(y)}(y - s(y)). \tag{3.6}$$

The quality distribution $H(y)$ given by (3.4) and the surplus function given by the solution to the differential equation (3.6) together with the boundary condition $s(y_\ell) = 0$ constitute the unique candidate equilibrium. Following the same arguments as in Section 2.2, it is easy to verify that the unique candidate equilibrium is, in fact, an equilibrium. The equilibrium has the same properties as in the baseline model, as it features quality dispersion and a strictly positive relation between quality and surplus.

It is easy to show that the equilibrium remains efficient under the general cost function. The problem of a utilitarian social planner is

$$W = \max_{y(x)} \left(- \int_0^1 c(y(x))dx + \int_0^1 \theta\lambda e^{-\lambda(1-x)} y(x)dx \right) \quad \text{s.t. } y(x) \text{ non-decreasing in } x. \tag{3.7}$$

As in Section 2.3, we consider a relaxed version of the social planner's problem (3.7) in which we remove the constraint on the monotonicity of the quantile function $y(x)$. We then verify that the solution to the relaxed problem satisfies the constraint and, hence, it is also a solution to (3.7).

In the relaxed version of the social planner's problem, the optimality condition for $y(x)$ is

$$-c'(y(x)) + \theta\lambda e^{-\lambda(1-x)} = 0. \tag{3.8}$$

Isolating $y(x)$ yields

$$y(x) = \psi\left(\theta\lambda e^{-\lambda(1-x)}\right). \tag{3.9}$$

Since ψ is a strictly increasing function, it follows that $y(x)$ in (3.9) is strictly increasing in x and, hence, it is not only the solution to the relaxed social planner's problem, but also the solution of the original problem (3.7).

In equilibrium, the x-th quantile of the quality distribution is such that

$$x = 1 - \frac{1}{\lambda} \log \frac{\theta\lambda}{c'(y(x))}. \tag{3.10}$$

Isolating $y(x)$ yields

$$y(x) = \psi\left(\theta\lambda e^{-\lambda(1-x)}\right). \tag{3.11}$$

Since (3.11) is identical to (3.9), it follows that the quality distribution in equilibrium coincides with the quality distribution in the solution to the planner's problem. In equilibrium, buyers purchase from the seller that offers them the highest surplus. In the solution to the social planner's problem, buyers are instructed to purchase from the seller with the highest quality. Since, in equilibrium, the seller with the highest surplus is the seller with the highest quality, the quantity sold by a seller at the x-th quantile of the quality distribution is the same in equilibrium as in the solution to the social planner's problem. Thus, the equilibrium is efficient.

3.2 General meeting process

We now relax the assumption that a buyer contacts a number n of randomly selected sellers, where n is drawn from a Poisson distribution with coefficient λ. Specifically, we assume that a buyer contacts n randomly-selected sellers with some generic probability $\alpha_n \in [0,1]$. We assume that $\alpha_1 > 0$, which means that a buyer contacts a single seller with strictly positive probability, that $\alpha_n > 0$ for some $n > 1$, which means that a buyer contacts multiple sellers with strictly positive probability, and that $\sum_{n=0}^{\infty} \alpha_n = \lambda$, which means that a buyer contacts an average of λ sellers. Clearly, the Poisson meeting process that we assumed in Section 2 is just a special case of this general meeting process. The same is true for the Burdett-Judd meeting process in which a buyer either meets 1 or 2 sellers. Given the popularity of the Burdett-Judd meeting process in the search-theoretic literature, we shall discuss it in additional detail in the appendix.

Lemmas 1-6 hold under the general meeting process described above. Therefore, in any equilibrium, the profit of a seller must attain its maximum for all $(y, s(y))$ with $y \in [y_\ell, y_h]$. That is,

$$V^* = -y^\gamma + v(F(s(y)))\,(y - s(y))\,, \tag{3.12}$$

where $v(F(s(y)))$ denotes the quantity sold by a seller at the $F(s(y))$-quantile of the surplus distribution and is such that

$$v(x) = \theta \sum_{n=1}^{\infty} n\alpha_n x^{n-1}. \tag{3.13}$$

Since $(y, s(y))$ maximizes the seller's profit, the derivative of the right-hand side of (3.12) with respect to s must be equal to 0 when evaluated at any $(y, s(y))$ with $y \in [y_\ell, y_h]$. That is,

$$v'(F(s(y)))F'(s(y))(y - s(y)) - v(F(s(y))) = 0. \tag{3.14}$$

For the same reason, the derivative of the right-hand side of (3.12) with respect to y must be equal to 0 when evaluated at any $(y, s(y))$ with $y \in [y_\ell, y_h]$. That is,

$$-\gamma y^{\gamma-1} + v(F(s(y))) = 0. \tag{3.15}$$

Using the optimality conditions (3.14) and (3.15) and the strict monotonicity of $s(y)$, we can solve for the quality distribution $H(y)$ and the surplus function $s(y)$. Specifically, the strict monotonicity of $s(y)$ implies that $F(s(y))$ equals $H(y)$. Using $F(s(y)) = H(y)$, we can solve the optimality condition (3.15) with respect to $H(y)$ and obtain

$$H(y) = v^{-1}\left(\gamma y^{\gamma-1}\right). \tag{3.16}$$

Using $H(y_\ell) = 0$ and $H(y_h) = 1$, we can solve for y_ℓ and y_h and obtain

$$\begin{aligned} \gamma y_\ell^{\gamma-1} &= v(0) = \theta\alpha_1 \quad \Longrightarrow y_\ell = \left(\frac{\theta\alpha_1}{\gamma}\right)^{\frac{1}{\gamma-1}}, \\ \gamma y_h^{\gamma-1} &= v(1) = \theta\lambda \quad \Longrightarrow y_h = \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}. \end{aligned} \tag{3.17}$$

Similarly, using $F'(s(y))s'(y) = H'(y)$, $F(s(y)) = H(y)$ and $v(H(y)) = \gamma y^{\gamma-1}$, we can rewrite the optimality condition (3.14) as

$$s'(y) = (\gamma - 1)\left(\frac{y - s(y)}{y}\right). \tag{3.18}$$

Together with the boundary condition $s(y_\ell) = 0$, the solution to the differential equation (3.18) is given by the surplus function

$$s(y) = \frac{\gamma - 1}{\gamma}\left(y - y_\ell^\gamma y^{1-\gamma}\right). \tag{3.19}$$

It is easy to verify that the unique candidate equilibrium given by the quality distribution $H(y)$ in (3.16) and by the surplus function $s(y)$ in (3.19) is, in fact, an equilibrium. Interestingly, the details of the meeting process affect the equilibrium quality distribution $H(y)$ but, conditional on y_ℓ, they have no impact on the surplus function $s(y)$, which in fact is given by the same expression as in Section 2. The equilibrium has the same properties as in the baseline model, as it features quality dispersion and a strictly positive relation between quality and surplus.

The equilibrium remains efficient under the general meeting process. The problem of a utilitarian social planner is

$$\begin{aligned} W \;=\; &\max_{y(x)} \left(-\int_0^1 y(x)^\gamma dx + \int_0^1 v(x)y(x)dx\right) \\ &\text{s.t. } y(x) \text{ non-decreasing in } x. \end{aligned} \tag{3.20}$$

As usual, we consider a relaxed version of the social planner's problem (3.20) in which we remove the constraint on the monotonicity of the quantile function $y(x)$. In the relaxed version of the social planner's problem, the optimality condition for $y(x)$ is

$$-\gamma y(x)^{\gamma-1} + v(x) = 0. \tag{3.21}$$

Isolating $y(x)$ yields

$$y(x) = \left(\frac{v(x)}{\gamma}\right)^{\frac{1}{\gamma-1}}. \tag{3.22}$$

Since v is a strictly increasing function, it follows that $y(x)$ in (3.22) is strictly increasing in x and, hence, it is not only the solution to the relaxed social planner's problem, but also the solution of the original problem (3.20).

In equilibrium, the x-th quantile of the quality distribution is such that

$$v(x) = \gamma y(x)^{\gamma-1}. \tag{3.23}$$

Isolating $y(x)$ yields the same expression as in (3.22) and hence the quality distribution in equilibrium coincides with the quality distribution in the solution to the planner's problem. In equilibrium, buyers purchase from the seller that offers them the highest surplus. In the solution to the social planner's problem, buyers are instructed to purchase from the seller with the highest quality. Since, in equilibrium, the seller with the highest surplus is the seller with the highest quality, the quantity sold by a seller at the x-th quantile of the quality distribution is the same in equilibrium as in the solution to the social planner's problem. Therefore, the equilibrium is efficient.

3.3 General marginal cost of production

In Section 2, we assumed that the seller could produce a unit of its variety at a constant marginal cost of 0. We now relax this assumption. We assume that a seller faces a marginal cost of production that need not be 0 and that may depend on the quality of its variety. We maintain the assumption that the seller faces a marginal cost of production that is constant. Let $y \in Y$ denote the quality of the seller's variety, where $Y \subset \mathbb{R}_+$. Let $m \in M(y)$ denote the marginal cost of production of the seller's variety, where $M(y) \subset \mathbb{R}_+$ denotes the set of feasible marginal costs given quality y. Let $\hat{y} = y - m$ denote the net quality of the seller's variety, defined as the difference between the quality of the variety and the marginal cost of producing it. Let $\hat{Y}$ denote the set of net qualities that can be achieved with some combination of $y \in Y$ and $m \in M(y)$ and assume that $\mathbb{R}_+ \subset \hat{Y}$. We assume that the seller's cost of designing a variety depends on its net quality $\hat{y} = y - m$. As in Section 2, we assume that the cost of designing a variety of net quality $\hat{y}$ is $c(\hat{y}) = \hat{y}^\gamma$, with $\gamma > 1$. Since $\hat{y} = y - m$ may be negative, we also assume that $c(\hat{y}) = 0$ for all $\hat{y} \leq 0$. This extension of the model encompasses several sensible cases. It encompasses the case in which the marginal cost of production is increasing in quality, i.e. $M(y) = m(y)$ and $m(y)$ is a strictly increasing function. It encompasses the case in which the marginal cost of production is strictly decreasing in quality, i.e. $M(y) = m(y)$ and $m(y)$ is a strictly decreasing function. It also encompasses cases in which $M(y)$ is some interval that depends on y.

As in Section 2, the buyer purchases the good from the seller that offers the highest surplus s, where s is defined as the difference between the quality y of the seller's variety and the price p of the seller's variety. Therefore, conditional on the surplus distribution F, the quantity sold by a seller that offers the surplus s is the same as in Section 2. If the seller chooses a variety with quality y and with marginal cost m and offers the surplus s, its profit is given by the quantity it sells times the profit per unit sold, net of the cost of designing the variety. That is, the profit of the seller $V(y, m, s)$ is given by

$$\begin{aligned} V(y,m,s) &= -(y-m)^{\gamma} + \sum\nolimits_{k=0}^{\infty} \theta_k \pi_k(s)(p-m), \\ &= -\hat{y}^{\gamma} + \sum\nolimits_{k=0}^{\infty} \theta_k \pi_k(s)(\hat{y}-s) \\ &\equiv \hat{V}(\hat{y}, s), \end{aligned} \tag{3.24}$$

where $p - m$ equals $\hat{y} - s$ because the surplus s is defined as $y - p$. The second line in (3.24) shows that seller's profit function $V(y, m, s)$ depends only on the net quality $\hat{y} = y - m$ and on the surplus s. Hence, the seller's profit function $V(y, m, s)$ can be written as $\hat{V}(\hat{y}, s)$.

The problem of the seller is to choose a net quality $\hat{y} \in \hat{Y}$ and a surplus s so as to maximize the profit function $\hat{V}(\hat{y}, s)$. Since θ_k and $\pi_k(s)$ are the same as in Section 2, the seller's profit function $\hat{V}(\hat{y}, s)$ is the same as in Section 2, with net quality $\hat{y}$ taking the place of quality y. Since $\mathbb{R}_+ \subset \hat{Y}$ and negative values of $\hat{y}$ are never optimal, the seller's relevant choice set is $\hat{y} \geq 0$ and $s \geq 0$, the same as the choice set for y and s in Section 2. Therefore, the seller's problem is the same as in Section 2, and we can immediately apply all the results derived there. Theorem 1 tells us that the equilibrium exists and is unique. The equilibrium net quality distribution $H(\hat{y})$ is given by

$$H(\hat{y}) = \frac{\gamma - 1}{\gamma} \left(\log \hat{y} - \log \hat{y}_{\ell}\right), \text{ for } \hat{y} \in [\hat{y}_{\ell}, \hat{y}_h], \tag{3.25}$$

where $\hat{y}_{\ell}$ and $\hat{y}_h$ are respectively given by

$$\hat{y}_{\ell} = \left(\frac{\theta \lambda e^{-\lambda}}{\gamma}\right)^{\frac{1}{\gamma-1}}, \quad \hat{y}_h = \left(\frac{\theta \lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}. \tag{3.26}$$

The equilibrium surplus function $s(\hat{y})$ is given by

$$s(\hat{y}) = \frac{\gamma - 1}{\gamma} \left(\hat{y} - \hat{y}_{\ell}^{\gamma} \hat{y}^{1-\gamma}\right). \tag{3.27}$$

Theorem 2 tells us that the equilibrium is efficient, in the sense that the quantile function $\hat{y}(x)$ is the same in equilibrium as in the solution to the social planner's problem, and the quantity sold by a seller at the x-th quantile of the net quality distribution is the same in equilibrium as in the solution to the social planner's problem.

Even though the expressions for the equilibrium objects are exactly the same as in the baseline model, the interpretation is different. In the baseline model, it is sellers with higher quality y that offer more surplus to their customers and sell more units (i.e., are larger). In the version of the model with marginal costs of production, net quality $\hat{y}$ takes the place of quality y. In general, it need not be the case that the sellers with higher net quality are those with higher quality. Hence, it need not be the case that sellers with higher quality offer more surplus and are larger. It may be the case, for instance, that all sellers design varieties of the same quality, but some sellers invest more than others in lowering their marginal cost. In this case, it is the sellers with lower marginal cost that offer more surplus and are larger. It may also be the the case that the sellers with higher net quality are those that design inferior varieties that can be produced at a significantly lower cost. In this case, it is the sellers of inferior varieties that offer more surplus and are larger.

4 Declining search frictions

We now use the baseline model 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 and revenues concentration – in the sense that the share of output sold by the biggest sellers becomes larger and the share of revenues enjoyed by the biggest sellers becomes larger – and increase quality polarization – in the sense that the varieties designed by the biggest sellers become better, while the varieties designed by the smallest sellers become worse. In terms of payoffs, we find that declining search frictions increase total gains from trade unboundedly. As search frictions vanish, both the buyers' and the sellers' share of the gains from trade converge to strictly positive constants. Similarly, declining search frictions increase total welfare unboundedly. As search frictions vanish, the buyers' share of welfare converges to 1 and the sellers' share of welfare converges to 0. The difference between the behavior of ex-post payoffs (gains from trade) and ex-ante payoffs (welfare) is due to the endogenous response of the quality distribution.

4.1 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

$$q(x) = \theta\lambda e^{-\lambda(1-x)}. \tag{4.1}$$

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

$$\frac{dq(x)}{d\lambda} = q(x)\frac{1-\lambda(1-x)}{\lambda}. \tag{4.2}$$

The derivative in (4.2) 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 (4.1), 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 (4.2), 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 quality of its variety, $Q(z)$ is given by

$$Q(z) = \frac{\int_{1-z}^{1} q(x)dx}{\int_{0}^{1} q(x)dx} = \frac{1-e^{-\lambda z}}{1-e^{-\lambda}}. \tag{4.3}$$

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

$$\frac{dQ(z)}{d\lambda} = \frac{e^{-\lambda(1+z)}\left(1-z+ze^{\lambda}-e^{\lambda z}\right)}{\left(1-e^{-\lambda}\right)^2} > 0. \tag{4.4}$$

Next, we consider the effect of declining search frictions on the quality distribution. As shown earlier in (2.33), the quality $y(x)$ of the variety produced by a seller at the x-th quantile of the H distribution is given by

$$y(x) = \left(\frac{\theta\lambda}{\gamma}e^{-\lambda(1-x)}\right)^{\frac{1}{\gamma-1}}. \tag{4.5}$$

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

$$\frac{dy(x)}{d\lambda} = \frac{y(x)}{\gamma-1}\cdot\frac{1-\lambda(1-x)}{\lambda}. \tag{4.6}$$

The derivative (4.6) 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 (2.27), (2.22), and (4.5), the surplus $s(x) \equiv s(y(x))$ offered by a seller at the x-th quantile of the distribution can be written as

$$s(x) = \frac{\gamma - 1}{\gamma} \left(\frac{\theta \lambda}{\gamma} \right)^{\frac{1}{\gamma - 1}} e^{-\frac{\lambda(1-x)}{\gamma - 1}} \left(1 - e^{-\frac{\lambda \gamma x}{\gamma - 1}} \right). \tag{4.7}$$

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

$$\frac{ds(x)}{d\lambda} = \frac{y(x)}{\gamma} \left[\left(1 - e^{-\frac{\lambda \gamma x}{\gamma - 1}} \right) \frac{1 - \lambda(1-x)}{\lambda} + \gamma x e^{-\frac{\lambda \gamma x}{\gamma - 1}} \right]. \tag{4.8}$$

The derivative in (4.8) 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

$$\begin{aligned} r(x) &= \theta \lambda e^{-\lambda(1-x)} (y(x) - s(x)) \\ &= \left(\frac{\theta \lambda e^{-\lambda}}{\gamma} \right)^{\frac{\gamma}{\gamma - 1}} \left(e^{\frac{\lambda \gamma x}{\gamma - 1}} + (\gamma - 1) \right). \end{aligned} \tag{4.9}$$

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

$$\frac{dr(x)}{d\lambda} = y(x)^{\gamma} \left(\frac{\gamma}{\gamma - 1} \right) \left[\frac{1 - \lambda(1-x)}{\lambda} - \left(1 - \frac{1}{\lambda} \right) (\gamma - 1) e^{-\frac{\lambda \gamma x}{\gamma - 1}} \right]. \tag{4.10}$$

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 (4.9). 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. The effect is unambiguously positive for all $x > \overline{x}(\lambda)$, where $\overline{x}(\lambda) > 1/\lambda$ and $\overline{x}(\lambda) < 1$ for λ

large enough.

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

$$R(z) = \frac{\int_{1-z}^{1} r(x)dx}{\int_{0}^{1} r(x)dx} = \frac{1 - e^{-\frac{\lambda\gamma z}{\gamma-1}} + \lambda\gamma z e^{-\frac{\lambda\gamma}{\gamma-1}}}{1 - e^{-\frac{\lambda\gamma}{\gamma-1}} + \lambda\gamma e^{-\frac{\lambda\gamma}{\gamma-1}}}. \tag{4.11}$$

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

$$\begin{aligned}\frac{dR(z)}{d\lambda} = \Phi \quad & \left\{ z \left[(\lambda\gamma - 1)\, e^{-\frac{\lambda\gamma z}{\gamma-1}} - (\gamma-1)\, e^{-\frac{\lambda\gamma}{\gamma-1}} + e^{\frac{\lambda\gamma(1-z)}{\gamma-1}} - \gamma(\lambda-1) - 1 \right] \right. \\ & \left. + \gamma\,(\lambda-1) \left(1 - e^{-\frac{\lambda\gamma z}{\gamma-1}} \right) \right\},\end{aligned} \tag{4.12}$$

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

$$\begin{aligned}\frac{dR(z)}{d\lambda} = \Phi \quad & \left\{ z \left[e^{\frac{\lambda\gamma(1-z)}{\gamma-1}} - 1 + (\gamma-1)e^{-\frac{\lambda\gamma z}{\gamma-1}} \left(1 - e^{-\frac{\lambda\gamma(1-z)}{\gamma-1}} \right) \right] \right. \\ & \left. + \gamma(\lambda-1)(1-z) \left(1 - e^{-\frac{\lambda\gamma z}{\gamma-1}} \right) \right\},\end{aligned} \tag{4.13}$$

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 highest-sale sellers is increasing in λ for all $z \in [0,1]$; the share $R(z)$ of revenues of the z highest-revenue 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 quality distribution is increasing in λ for $x > 1 - 1/\lambda$, and decreasing in λ for $x < 1 - 1/\lambda$.*

3. *Higher sales and revenues at the top: The sales $q(x)$ of a seller at the x-th quantile of the sales distribution are increasing in λ for all $x > 1 - 1/\lambda$; the revenues $r(x)$ of a seller at the x-th quantile of the revenue distribution are increasing in λ for all*

$x > \overline{x}(\lambda)$, where $\overline{x}(\lambda) > 1 - 1/\lambda$ and $\overline{x}(\lambda) < 1$ for λ large enough.

4. *Lower sales and revenues at the bottom: The sales $q(x)$ and revenues $r(x)$ of a seller at the x-th quantile of the sales and revenue distributions are decreasing in λ for all $x < 1 - 1/\lambda$.*

Theorem 3 contains a description of the structural transformation that a product market undergoes as search frictions become smaller.[6] When search frictions are 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. When search frictions become smaller, the product market becomes more unbalanced. High-quality sellers improve the quality of their varieties, while low-quality sellers lower the quality of their varieties. High-quality sellers, which already have higher sales, further increase their sales and their sales share, while low-quality sellers enjoy fewer sales and capture a smaller share of total sales. High-quality sellers, which already have higher revenues, further increase their revenues and their revenues share, while low-quality sellers experience declining revenues in both absolute and relative terms. Moreover, as search frictions become smaller and smaller, the fraction of high-quality sellers that improve the quality of their variety and that enjoy increasing sales becomes vanishingly small, since the quantile cutoff $1 - 1/\lambda$ converges to 1. The fraction of low-quality sellers that reduce the quality of their variety and that experiences lower sales becomes larger and larger. Eventually, as one can see from (4.1), (4.5) and (4.9), these low-quality sellers offer varieties of negligible quality, they make approximately no sales, and they enjoy negligible revenues. The transformation that the product market undergoes as search frictions decline is efficient and, as we shall see momentarily, it is associated with growing welfare and buyers' surplus.

The structural transformation generated by declining search frictions is a close relative of both the superstar theory by Rosen (1981) and the span of control theory by Lucas (1978). In Rosen (1981) and Lucas (1978), producers of different quality face some technological or organizational constraint that prevents the best ones from taking over the market. As these constraints are relaxed, the higher quality producers become larger, the lower quality producers become smaller and the market becomes more concentrated. The increase in concentration is efficient and welfare unambiguously increases. Like technological or organizational constraints, search frictions act as a limit to the expansion of the highest quality sellers. However, unlike technological or organizational constraints, search frictions also affect the extent of competition in the market. In Rosen (1981) and

[6] The effect of declining search frictions on the structure of the product market does not depend substantively on the specification of the design cost function and of the meeting process. As long as declining search frictions lead to an increase in the average number of buyers contacted by a seller and to a decline in the number of buyers that contact only one seller, high-quality sellers will sell more and low-quality sellers will sell less. As long as the design cost function is increasing and convex, high-quality sellers will improve the quality of their variety and low-quality sellers will lower the quality of their variety. 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 and, hence, to an increase in revenue concentration.

Lucas (1978), competition is always assumed to be perfect. Moreover, search frictions themselves generate and affect the extent of heterogeneity across sellers, while the extent and shape of heterogeneity is exogenously posited in Rosen (1981) and Lucas (1978).

Declining search frictions lead to a structural transformation of the market that matches well several broad empirical trends. Autor et al. (2020) and De Loecker, Eeckhout and Unger (2020) document an upward trend in sales concentration. De Loecker, Eeckhout and Mongey (2021) document an increase in the dispersion of productivity across firms. Autor et al. (2020) document a positive cross-sectoral relationship between the increase in sales concentration and productivity growth. Increasing sales concentration and higher productivity dispersion can be explained by declining search frictions. A positive relationship between sales concentration and productivity growth can be explained by search frictions declining at different rates in different markets. Declining search frictions provide a benign explanation for these empirical trends, in the sense that these trends are efficient, and they are associated with increasing total welfare and buyers' surplus.[7] The standard explanation of these trends is, however, much bleaker (see, e.g., Gutierrez and Philippon 2019, De Loecker, Eeckhout and Unger 2020, De Loecker, Eeckhout and Mongey 2021). The standard explanation views these trends as the consequence of a small number of firms exerting increasing market power due to, say, higher barriers to entry, and it associates these trends with lower welfare and buyers' surplus. The relative contribution of declining search frictions and higher barriers to entry to these trends as yet to be assessed and it may very well vary from market to market.

4.2 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

$$\begin{aligned} G &= \int_{y_\ell}^{y_h} \theta\lambda e^{-\lambda(1-H(y))} y dH(y) \\ &= \int_{y_\ell}^{y_h} \theta(\gamma-1)\left(\frac{y}{y_h}\right)^{\gamma-1} dy \\ &= \theta\frac{\gamma-1}{\gamma}\left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}\left[1-e^{-\frac{\lambda\gamma}{\gamma-1}}\right], \end{aligned} \tag{4.14}$$

where the second line makes use of the fact that $H(y)$ is given by (2.24), the third line solves the integral in the second line and makes use of the fact that y_ℓ and y_h are given

[7] Other related phenomena that have been documented include rising mark-ups (De Loecker, Eeckhout and Unger 2020, Kerhig and Vincent 2020) and declining labor shares (Autor et al. 2020). Our baseline model has nothing to say about these phenomena, as it abstracts from production costs. We suspect, though, that a version of our model where labor is needed to produce the sellers' varieties could explain also these facts as a consequence of declining search frictions.

by (2.22) and (2.23). The derivative of the total gains from trade with respect to λ is

$$\frac{dG}{d\lambda} = \left(\frac{\gamma\lambda}{\theta}\right)^{-\frac{\gamma}{\gamma-1}} \left[1 - (1-\lambda\gamma)e^{-\frac{\lambda\gamma}{\gamma-1}}\right] > 0. \tag{4.15}$$

The elasticity of G with respect to λ is

$$\frac{dG}{d\lambda}\frac{\lambda}{G} = \frac{1}{\gamma-1} \cdot \frac{1-(1-\lambda\gamma)e^{-\frac{\lambda\gamma}{\gamma-1}}}{1-e^{-\frac{\lambda\gamma}{\gamma-1}}} > 0. \tag{4.16}$$

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 (4.16), it converges to $1/(\gamma-1)$ as $\lambda \to \infty$.

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

$$\begin{aligned} G_b &= \int_{y_\ell}^{y_h} \theta\lambda e^{-\lambda(1-H(y))} s(y) dH(y) \\ &= \theta \left(\frac{\gamma-1}{\gamma}\right)^2 \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}} \left[1 - \left(1+\frac{\lambda\gamma}{\gamma-1}\right)e^{-\frac{\lambda\gamma}{\gamma-1}}\right], \end{aligned} \tag{4.17}$$

where the second line makes use of the fact that y_ℓ, y_h, $H(y)$ and $s(y)$ are given by (2.22), (2.23), (2.24) and (2.27). The derivative of the buyers' gains from trade with respect to λ is

$$\frac{dG_b}{d\lambda} = \frac{\theta}{\lambda(\gamma-1)} \left(\frac{\gamma-1}{\gamma}\right)^2 \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}} \left\{1 - \left[1+\frac{\lambda\gamma}{\gamma-1}(1-\lambda\gamma)\right]e^{-\frac{\lambda\gamma}{\gamma-1}}\right\} > 0. \tag{4.18}$$

The elasticity of G_b with respect to λ is

$$\frac{dG_b}{d\lambda}\frac{\lambda}{G_b} = \frac{1}{\gamma-1} \cdot \frac{1-\left[1+\frac{\lambda\gamma}{\gamma-1}(1-\lambda\gamma)\right]e^{-\frac{\lambda\gamma}{\gamma-1}}}{1-\left[1+\frac{\lambda\gamma}{\gamma-1}\right]e^{-\frac{\lambda\gamma}{\gamma-1}}} > 0. \tag{4.19}$$

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 (4.19), it converges to $1/(\gamma-1)$ as $\lambda \to \infty$.

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

$$\begin{aligned} G_s &= \int_{y_\ell}^{y_h} \theta\lambda e^{-\lambda(1-H(y))}(y - s(y))dH(y) \\ &= \theta\frac{\gamma-1}{\gamma^2}\left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}\left[1-(1-\lambda\gamma)\,e^{-\frac{\lambda\gamma}{\gamma-1}}\right]. \end{aligned} \tag{4.20}$$

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

$$\frac{dG_s}{d\lambda} = \frac{\theta}{\lambda\gamma^2}\left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}}\left[1-\left(1-\lambda\gamma-\lambda\gamma^2+\lambda^2\gamma^2\right)e^{-\frac{\lambda\gamma}{\gamma-1}}\right] \tag{4.21}$$

The elasticity of G_s with respect to λ is

$$\frac{dG_s}{d\lambda}\frac{\lambda}{G_s} = \frac{1}{\gamma-1}\cdot\frac{1-\left(1-\lambda\gamma-\lambda\gamma^2+\lambda^2\gamma^2\right)e^{-\frac{\lambda\gamma}{\gamma-1}}}{1-(1-\lambda\gamma)\,e^{-\frac{\lambda\gamma}{\gamma-1}}}. \tag{4.22}$$

The sellers' gains from trade are increasing in λ for λ sufficiently large. The elasticity of the sellers' gains from trade with respect to λ converges to $1/(\gamma-1)$ as $\lambda \to \infty$.

The elasticities with respect to λ of the total gains from trade, the buyers' gains from trade, and the sellers' gains from trade all converge to $1/(\gamma-1)$ as $\lambda \to \infty$ and, hence, the share of the gains from trade accruing to buyers and sellers converge to some constants. Specifically, as one can see from (4.14), (4.17) and (4.20), the buyers' share of the gains from trade converges to the constant $(\gamma-1)/\gamma \in (0,1)$ and the sellers' share of the gains from trade converges to the constant $1/\gamma \in (0,1)$. The finding that the sellers' share of the gains from trade converges to a strictly positive constant may seem surprising. In fact, declining search frictions increase ex post competition among sellers and hence induce sellers to offer as surplus a share of the value of their product that converges to 1. This argument would be complete if the distribution of qualities across varieties was constant. When sellers choose the quality of their variety, however, declining search frictions induce low-quality seller to reduce quality and induce high-quality seller to increase quality. The fanning out of the quality distribution reduces the extent of ex post competition faced by high-quality sellers and, hence, it allows high-quality sellers to lower the share of the value of their product that is offered to buyers. The force increasing and the force decreasing ex post competition exactly offset each other as $\lambda \to \infty$ and the market stabilizes at some imperfect level of ex post competition. As a result, the sellers' share of the gains from trade converges to a strictly positive constant. The finding is not only surprising, but also critical. Indeed, if the share of the sellers' gains from trade converged to zero, sellers would not be able to recoup their investments, the quality distribution would converge to a mass point at zero, and, effectively, the market would shut down.

Total welfare W is given by the sum of buyers' welfare W_b and sellers' welfare W_s. The buyers' welfare W_b is equal to the buyers' gains from trade G_b. The sellers' welfare

W_s is equal to the sum across sellers of the difference between the gains from trade and the design cost. Since the difference between gains from trade and design cost is V^* for every seller and the measure of sellers is normalized to 1, it follows that W_s is equal to V^*. Total welfare W is the sum of buyers' and sellers' welfare and it is such that

$$\begin{aligned} W &= W_b + W_s \\ &= \theta \left(\frac{\gamma-1}{\gamma}\right)^2 \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}} \left(1 - e^{-\frac{\lambda\gamma}{\gamma-1}}\right). \end{aligned} \tag{4.23}$$

The derivative of W with respect to λ is

$$\frac{dW}{d\lambda} = \theta \left(\frac{\gamma-1}{\gamma}\right)^2 \left(\frac{\theta\lambda}{\gamma}\right)^{\frac{1}{\gamma-1}} \frac{1}{\lambda(\gamma-1)} \left[1 - (1-\lambda\gamma)e^{-\frac{\lambda\gamma}{\gamma-1}}\right] > 0. \tag{4.24}$$

The elasticity of W with respect to λ is

$$\frac{dW}{d\lambda}\frac{\lambda}{W} = \frac{1}{\gamma-1} \cdot \frac{1-(1-\lambda\gamma)e^{-\frac{\lambda\gamma}{\gamma-1}}}{1-e^{-\frac{\lambda\gamma}{\gamma-1}}} > 0. \tag{4.25}$$

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 (4.25), it converges to $1/(\gamma-1)$ as $\lambda \to \infty$.

Sellers' welfare W_s is equal to V^*. In turn, V^* is equal to the profit for a seller designing a variety of quality y_ℓ, offering a surplus s of 0, and trading only with captive buyers. Therefore, W_s is equal to

$$W_s = (\gamma-1)\left(\frac{\theta\lambda e^{-\lambda}}{\gamma}\right)^{\frac{\gamma}{\gamma-1}}. \tag{4.26}$$

The derivative of W_s with respect to λ is

$$\frac{dW_s}{d\lambda} = -\left(\frac{\theta\lambda e^{-\lambda}}{\gamma}\right)^{\frac{\gamma}{\gamma-1}} \frac{\gamma(\lambda-1)}{\lambda}. \tag{4.27}$$

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. Buyers' welfare W_b is equal to total welfare W net of sellers' welfare W_s. Since W increases with λ and W_s decreases with λ, it follows that W_b is strictly increasing with λ. Since W_s converges to zero as $\lambda \to \infty$, it follows that W_b converges to W.

As $\lambda \to \infty$, the sellers' share of welfare converges to 0 and the buyers' share of welfare converges to 1. These findings are intuitive. As search frictions vanish, the fraction of captive buyers converges to zero and so does the quantity sold by the lowest-quality seller. In response to vanishing sales, the lowest-quality seller progressively reduces the quality of its variety y_ℓ towards zero. Therefore, the profit of the lowest-quality seller converges toward zero. Since all sellers make the same profit, sellers' welfare converges toward zero as well. Since total welfare keeps growing as search frictions become smaller and smaller, it follows that the sellers' share of welfare converges to 0. In this sense, vanishing search frictions lead to perfect ex ante competition among sellers.

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

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 grow without bound as search frictions become smaller, the share of the gains from trade accruing to sellers converges to $1/\gamma$, and the share of the gains from trade accruing to buyers converges to $(\gamma - 1)/\gamma$. Total welfare also grows without bound as search frictions become smaller, but the share of welfare accruing to sellers converges to 0, and the share of welfare accruing to buyers converges to 1. The fact that the sellers' share of the gains from trade converges to a positive constant implies that ex post competition stabilizes at some imperfect level. The fact that the sellers' share of welfare converges to zero implies that ex ante competition eventually 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 differentiate themselves more and more as search frictions decline.[8]

Theorems 3 and 4 can be used together to understand the path taken by the market as search frictions decline over time 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

[8] Some of the effects of declining search frictions on gains from trade and welfare do not depend on the specification of the design cost function and the meeting process. As long as declining search frictions lead to an increase (in the first-order stochastic dominance sense) in the distribution of the number of sellers contacted by a buyer, aggregate welfare and aggregate gains from trade must increase. As long as declining search frictions drive the number of buyers who contact a single seller to zero, sellers' welfare must go to zero and so does the sellers' share of welfare. The sellers' share of the gains from trade can never converge to 0, as this would not be consistent with increasing investment by the highest-quality sellers. The finding that the declining search frictions lead to a constant growth in aggregate gains from trade and aggregate welfare does, on the other hand, depend on the isoelastic specification of the design cost function and on the Poisson specification of the meeting process.

these distributions grows at the same, constant 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 gains from trade eventually grow at the constant rate $g_\lambda/(\gamma-1)$, and the shares going to buyers and sellers eventually become constant. As λ grows at the rate g_λ, total welfare eventually 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 aggregate gains from trade and welfare to eventually grow at the rate $g_\lambda/(\gamma-1)$. This is what Menzio (2022) and Martellini and Menzio (2020, 2021) refer to as Stiglerian growth. The term $1/(\gamma-1)$ is the rate of return to declining search frictions, and it 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 fraction of sellers to design varieties with higher quality and to serve a larger fraction of 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 aggregate 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 aggregate gains from trade and welfare is lower.

Lastly, let us use Theorems 3 and 4 to remark on the limiting behavior of the market as search frictions disappear. As time goes by and search frictions become smaller and smaller, a vanishing measure of sellers who produce varieties of increasing quality take over the entire market. Even though the market becomes dominated by a vanishing measure of sellers, these sellers do not price monopolistically (i.e., at the buyers' reservation quality). Even though search frictions in the market converge to zero, these sellers do not price competitively (i.e., at their marginal cost). Instead, these sellers set prices equal to their average cost, since their ex-ante profits converge to zero. This property of equilibrium relates to the analysis of natural monopolies in frictionless markets, and to the question of whether natural monopolists should be regulated. Baumol, Panzar and Willig (1982) and Baumol (1982) argued, in the spirit of Demsetz (1968), that a natural monopolist would have to price at average cost if the market is contestable. Intuitively, if the monopolist priced above average cost and, thus, made profits, a competitor could come in and steal the whole market from the monopolist. For this reason, Baumol, Panzar and Willig (1982) argue that a natural monopolist of a contestable market should not be regulated. As pointed out by Brock (1982), however, the argument of Baumol, Panzar and Willig (1982) rests on the assumption that the monopolist cannot adjust its price after the entry of a competitor. If the monopolist could, competition between the monopolist and the competitor would drive prices to marginal cost. Therefore, irrespective of the monopolist's current price, the competitor would have no incentive to enter and the monopolist could safely set its price above average cost. For this reason, a natural monopolist would have to be regulated even if the market is contestable.[9] Theorems 3 and 4 show that, in a world with small but positive search frictions, a natural monopolist prices at average cost (as in Baumol, Panzar and Willig 1982) even though prices are fully flexible (as in Brock 1982).

5 Conclusions

In this paper, we contribute to the development of the search-theoretic framework of imperfect competition of Butters (1977), Varian (1980) and Burdett and Judd (1983) by studying a version of the framework in which sellers make an ex ante investment in the quality of their variety of the product. We show that the equilibrium exists and it is unique. The equilibrium distribution $H(y)$ of quality across sellers is atomless, and its support is an interval. The equilibrium map $s(y)$ between the quality of the variety designed by a seller and the surplus that offered by a seller to its customers is a strictly increasing function. The equilibrium quality distribution $H(y)$ and the surplus function $s(y)$ have simple closed-form solutions. We show that the equilibrium is efficient. That is, the equilibrium quality distribution and the equilibrium allocation of varieties to buyers is exactly the same as in the solution to problem of a utilitarian social planner.

We also contribute to the recent literature on rising market concentration (Autor et al. 2020, De Loeker, Eeckhout and Unger 2020, De Loeker, Eeckhout and Mongey 2021). We show that, as search frictions become smaller and smaller, the structure of the market becomes increasingly unbalanced. Sales and revenues become more concentrated in the hands of fewer and fewer sellers, and the distribution of quality across sellers becomes more and more dispersed. As search frictions become smaller and smaller, welfare and gains from trade keep growing, the sellers' share of welfare converges to 0, and the sellers' share of the gains from trade converges to a positive constant. The recent upward trends in market concentration therefore can be viewed as the efficient and welfare-improving response of the market to the decline in search frictions brought about by improvements in communication, information and transportation technology that allow buyers to locate and trade with an increasing number of sellers. This is quite different from the standard view of increasing concentration as the consequence of increasing barriers to entry.

[9] We are grateful to V.V. Chari for alerting us to this classic controversy in industrial organization.

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] Baumol, W. 1982. "Contestable Markets: An Uprising in the Theory of Industry Structure." *American Economic Review*, 72: 1-15.

[6] Baumol, W., J. Panzar, and R. Willig. 1982. *Contestable Markets and the Theory of Industry Structure*, San Diego: Harcourt Brace Jovanovich.

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

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

[9] 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.

[10] Brock, W. 1983. "Contestable Markets and the Theory of Industry Structure: A Review Article." *Journal of Political Economy*, 91: 1055-1066.

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

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

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

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

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

[16] Demsetz, H. 1968. "Why Regulate Utilities?" *Journal of Law and Economics*, 11: 55-65.

[17] 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.

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

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

[20] 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

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

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

[23] Kaplan, G., G. Menzio, L. Rudanko, and N. Trachter. 2019. "Relative Price Dispersion: Evidence and Theory," *American Economic Journal: Microeconomics*, 11: 68-124.

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

[25] Lester, B., A. Shourideh, V. Venkateswaran, and A. Zetlin-Jones. 2019. "Screening and Adverse Selection in Frictional Markets." *Journal of Political Economy*, 127: 338-377.

[26] Lucas, R. 1978. "On the Size Distribution of Business Firms." *Bell Journal of Economics*, 9: 508-523.

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

[28] Martellini, P., and G. Menzio. 2021. "Jacks of All Trades and Masters of One: Declining Search Frictions and Unequal Growth." *American Economic Review: Insights*, 3: 336-352.

[29] Menzio, G. 2022. "Optimal Product Design: Implications for Competition and Growth under Declining Search Frictions." *Econometrica*. Forthcoming.

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

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

[32] Nord, L. 2022. "Shopping, Demand Composition and Equilibrium Prices." Manuscript, European University Institute.

[33] 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.

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

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

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

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

[38] Wildenbeest, M. 2011. "An Empirical Model of Search with Vertically Differentiated Products," *RAND Journal of Economics*, 42: 729-757.

Appendix

A Burdett-Judd meeting process

In this section, we show that all of the substantive results of the paper hold when we replace the Poisson meeting process with the Burdett-Judd meeting process that is most common in the literature. The Burdett-Judd meeting process is such that a buyer contacts one randomly-selected seller with probability $1-\alpha$ and contacts two randomly-selected sellers with probability α, with $\alpha \in (0,1)$. The parameter α is an inverse measure of search frictions, as a buyer contacts an average λ of $1+\alpha$ sellers.

Since the Burdett-Judd meeting process is a special case of the general meeting process described in Section 3.2, we can directly apply the results derived there. Therefore, the equilibrium exists, is unique, and is efficient. In equilibrium, the quality distribution $H(y)$ is

$$H(y) = \frac{\gamma}{2\theta\alpha}\left(y^{\gamma-1} - y_\ell^{\gamma-1}\right), \text{ for } y \in [y_\ell, y_h], \tag{A.1}$$

where the lower bound y_ℓ and the upper bound y_h of the support of the distribution are

$$y_\ell = \left(\frac{\theta(1-\alpha)}{\gamma}\right)^{\frac{1}{\gamma-1}}, \quad y_h = \left(\frac{\theta(1+\alpha)}{\gamma}\right)^{\frac{1}{\gamma-1}}, \tag{A.2}$$

and the surplus function $s(y)$ is

$$s(y) = \frac{\gamma-1}{\gamma}\left(y - y_\ell^\gamma y^{1-\gamma}\right). \tag{A.3}$$

The quantity $q(x)$ sold by a seller at the x-th quantile of the quality distribution is

$$q(x) = \theta\left(1-\alpha+2\alpha x\right). \tag{A.4}$$

It is clear that the quantity sold by a seller at the x-th quantile of the quality distribution is strictly increasing in x. In other words, sellers of varieties of higher quality are larger. The derivative of $q(x)$ with respect to α is

$$\frac{dq(x)}{d\alpha} = \theta\left(2x-1\right). \tag{A.5}$$

The quantity sold by a seller at the x-th quantile of the quality distribution is strictly increasing in α if $x > 1/2$ and strictly decreasing in α if $x < 1/2$. In other words, as search frictions decline, the quantity sold by a seller increases if the seller's quality is above the median and decreases if the seller's quality is below the median.

The quality $y(x)$ of the variety produced by a seller at the x-th quantile of the quality distribution is

$$y(x) = \left(\frac{\theta}{\gamma}\right)^{\frac{1}{\gamma-1}}\left(1-\alpha+2\alpha x\right)^{\frac{1}{\gamma-1}}. \tag{A.6}$$

The derivative of $y(x)$ with respect to α is given by

$$\frac{dy(x)}{d\alpha} = \frac{1}{\gamma - 1}\left(\frac{\theta}{\gamma}\right)^{\frac{1}{\gamma-1}} (1-\alpha+2\alpha x)^{\frac{1}{\gamma-1}-1} (2x-1). \tag{A.7}$$

The quality of the variety produced by a seller at the x-th quantile of the quality distribution is strictly increasing in α if $x > 1/2$ and strictly decreasing in α if $x < 1/2$. In other words, as search frictions decline, the quality produced by a seller increases if the seller's quality is above the median and it decreases if the seller's quality is below the median. Quality becomes more polarized as search frictions decline.

The surplus $s(x)$ offered by a seller at the x-th quantile of the quality distribution is

$$s(x) = \frac{\gamma - 1}{\gamma}\left(\frac{\theta}{\gamma}\right)^{\frac{1}{\gamma-1}} (1-\alpha+2\alpha x)^{\frac{1}{\gamma-1}} \left[1 - \left(\frac{1-\alpha}{1-\alpha+2\alpha x}\right)^{\frac{\gamma}{\gamma-1}}\right]. \tag{A.8}$$

It is immediate to see from (A.8) that the surplus offered by a seller at the x-th quantile of the quality distribution is strictly increasing in x. From (A.8), it is also easy to see that the derivative of $s(x)$ with respect to α is strictly positive if $x > 1/2$ and strictly negative if $x < 1/2$. Hence, as search frictions decline, the surplus offered by the highest quality sellers increases.

The revenues $r(x)$ for a seller at the x-th quantile of the quality distribution are

$$r(x) = \left(\frac{\theta}{\gamma}\right)^{\frac{\gamma}{\gamma-1}} \left[(1-\alpha+2\alpha x)^{\frac{\gamma}{\gamma-1}} + (\gamma-1)(1-\alpha)^{\frac{\gamma}{\gamma-1}}\right]. \tag{A.9}$$

The revenues enjoyed by a seller at the x-th quantile of the quality distribution are strictly increasing in x. The derivative of $r(x)$ with respect to α is strictly increasing in x and, hence, it is positive if and only if $x > \overline{x}(\alpha)$, where $\overline{x}(\alpha) > 1/2$ and $\overline{x}(\alpha) < 1$ for α large enough. Hence, as search frictions decline, the revenues of the highest quality sellers increase.

The share of output $Q(z)$ sold by the z largest sellers is given by

$$Q(z) = \frac{\int_{1-z}^{1} q(x)dx}{\int_{0}^{1} q(x)dx} = 1 - \alpha + \alpha\left(1 - z^2\right). \tag{A.10}$$

It immediate to see from (A.10) that the share of output sold by the z largest sellers is strictly increasing in α for any $z \in (0,1)$. Therefore, sales become more concentrated at the top as search frictions decline. Similarly, we can show that the share of revenues $R(z)$ enjoyed by the z largest sellers (who are also those with the highest revenues) is strictly increasing in α for any $z \in (0,1)$. Therefore, revenues become more concentrated at the top as search frictions decline.

The total gains from trade G are given by

$$\begin{aligned} G &= \int_{y_\ell}^{y_h} \theta\left(1-\alpha+2\alpha H(y)\right) y dH(y) \\ &= \int_{y_\ell}^{y_h} \frac{\gamma^2(\gamma-1)}{2\alpha\theta} y^{2\gamma-2} dy \\ &= \frac{\gamma^2(\gamma-1)}{2\alpha\theta}\frac{1}{2\gamma-1}\left(\frac{\theta}{\gamma}\right)^{\frac{2\gamma-1}{\gamma-1}}\left[(1+\alpha)^{\frac{2\gamma-1}{\gamma-1}}-(1-\alpha)^{\frac{2\gamma-1}{\gamma-1}}\right], \end{aligned} \tag{A.11}$$

where the first line is the definition of the total gains from trade, the second line makes use of the expression for the quality distribution H, and the third line makes use of the expressions for y_ℓ and y_h. It is easy to verify that the total gains from trade are strictly increasing in α. In other words, as search frictions decline, the sum of the gains from trade accruing to buyers and sellers increases.

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

$$\begin{aligned} G_b &= \int_{y_\ell}^{y_h} \theta(1-\alpha+2\alpha H(y))s(y)dH(y) \\ &= \int_{y_\ell}^{y_h} \frac{\gamma(\gamma-1)^2}{2\alpha\theta} y^{2\gamma-2}\left(1-y_\ell^\gamma y^{-\gamma}\right)dy \\ &= \frac{\gamma(\gamma-1)^2}{2\alpha\theta}\frac{1}{2\gamma-1}\left(\frac{\theta}{\gamma}\right)^{\frac{2\gamma-1}{\gamma-1}}\left[(1+\alpha)^{\frac{2\gamma-1}{\gamma-1}}-(1-\alpha)^{\frac{2\gamma-1}{\gamma-1}}\right] \\ &-\frac{\gamma(\gamma-1)^2}{2\alpha\theta}\frac{1}{\gamma-1}\left(\frac{\theta}{\gamma}\right)^{\frac{2\gamma-1}{\gamma-1}}\left[\left(\frac{1-\alpha}{1+\alpha}\right)^{\frac{\gamma}{\gamma-1}}(1+\alpha)^{\frac{2\gamma-1}{\gamma-1}}-(1-\alpha)^{\frac{2\gamma-1}{\gamma-1}}\right], \end{aligned} \tag{A.12}$$

where the first line is the definition of the buyers' gains from trade, the second line makes use of the expression for the quality distribution H and for the surplus s, and the third line makes use of the expressions for y_ℓ and y_h. A fair amount of algebra allows us to prove that the buyers' gains from trade are increasing in α. The gains from trade G_s accruing to the sellers, which are given by $G-G_b$, need not be monotonic in α.

Total welfare W is given by

$$\begin{aligned} W &= G-\int_{y_\ell}^{y_h} y^\gamma dH(y) \\ &= \int_{y_\ell}^{y_h} \frac{\gamma^2(\gamma-1)}{2\alpha\theta} y^{2\gamma-2} dy - \int_{y_\ell}^{y_h} \frac{\gamma(\gamma-1)}{2\alpha\theta} y^{2\gamma-2} dy \\ &= \left(\frac{\gamma-1}{\gamma}\right)\int_{y_\ell}^{y_h} \frac{\gamma^2(\gamma-1)}{2\alpha\theta} y^{2\gamma-2} dy \\ &= \left(\frac{\gamma-1}{\gamma}\right) G, \end{aligned} \tag{A.13}$$

where the first line is the definition of welfare as the total gains from trade net of the total product design costs, the second line makes use of the expressions for G and for H, the third line collects terms, and the last line uses the definition of G again. Total welfare W is just a fraction of the total gains from trade G. Since G is strictly increasing in α, it follows that W is strictly increasing in α as well. In other words, as search frictions decline, welfare increases.

Sellers' welfare W_s is given by

$$W_s = (\gamma - 1)\left(\frac{\theta(1-\alpha)}{\gamma}\right)^{\frac{\gamma}{\gamma-1}} \tag{A.14}$$

Sellers' welfare is equal to the sum of the profits of all the sellers in the market. Since every seller makes a profit of V^* and the measure of sellers in the market is normalized to 1, it follows that sellers' welfare is equal to V^*, which is given by to the expression on the right-hand side of (A.14). From (A.14), it follows immediately that sellers' welfare is strictly decreasing in α. The buyers' welfare W_b is given by $W - W_s$. Since total welfare W is strictly increasing in α and sellers' welfare is strictly decreasing in α, buyers' welfare is strictly increasing in α. As search frictions decline, sellers' welfare decreases and buyers' welfare increases.

It is useful to examine the properties of equilibrium in the limits for $\alpha \to 0$, where every buyer is in contact with a single seller, and for $\alpha \to 1$, where every buyer is in contact with two sellers. As $\alpha \to 0$, the equilibrium quality distribution $H(y)$ converges to a mass point at $y = (\theta/\gamma)^{1/(\gamma-1)}$ and every seller offers a surplus that converges to 0. These findings are intuitive. As $\alpha \to 0$, a seller understands that he will meet θ captive buyers and 0 non-captive buyers. In turn, the seller understands that he will be able to sell θ units of output by offering a surplus of 0 to its customers, and that offering more surplus would not increase its volume. Therefore, the seller will choose the quality of its variety so as to maximize $-y^\gamma + \theta y$. In aggregate, sellers capture all of the gains from trade and buyers capture none of them. That is, $G_s/G \to 1$ and $G_b/G \to 0$. Similarly, sellers capture all of the welfare and buyers capture none of it. That is, $W_s/W \to 1$ and $W_b/W \to 0$. It is immediate to see that the limit of the equilibrium as $\alpha \to 0$ coincides with the equilibrium of a model in which $\alpha = 0$, i.e. a model with a continuum of ex-ante identical monopolists.

As $\alpha \to 1$, almost every buyer is in contact with two sellers. As $\alpha \to 1$, the equilibrium quality distribution $H(y)$ is such that the lower bound of its support y_ℓ converges to 0, the upper bound of its support converges to $(2\theta/\gamma)^{1/(\gamma-1)}$, and the density converges to some strictly positive value everywhere on the support. When every buyer is in contact with two seller, the worst seller cannot make any sales. For this reason, the lower bound on the support of the quality distribution converges to 0. Not every seller, though, can choose a quality of 0, since an individual seller could then make a strictly positive profit by choosing $y = (\theta/\gamma)^{1/(\gamma-1)} > 0$ and offering some small but strictly positive surplus.

Therefore, the upper bound of the support of the quality distribution cannot converge to 0, and the distribution cannot become degenerate. The surplus function converges to a simple linear function, in which the surplus offered by the seller equals a fraction $(\gamma-1)/\gamma$ of the seller's quality y. Therefore, in aggregate, sellers capture a fraction $1/\gamma$ of the gains from trade and buyers capture a fraction $(\gamma-1)/\gamma$ of the gains from trade. That is, $G_s/G \to 1/\gamma$ and $G_b/G \to (\gamma-1)/\gamma$. Since the lower bound of the quality distribution converges to 0, it follows that the seller's profit V^* converges to 0 as well. Therefore, sellers capture none of the welfare and buyers capture all of it. That is, $W_s/W \to 0$ and $W_b/W \to 1$. It is easy to verify that the limit of the equilibrium coincides with the equilibrium of a model in which $\alpha = 1$, a model in which every buyer can choose from two sellers.

It is useful to compare these findings with the limits for $\alpha \to 0$ and for $\alpha \to 1$ in Burdett and Judd (1983), i.e. in a model where all sellers carry a variety of the good with the same, exogenous quality. In Burdett and Judd (1983), as $\alpha \to 0$, the equilibrium outcomes converge to monopoly outcomes, in the sense that every seller charges a price equal to the buyer's valuation of the good. As $\alpha \to 1$, the equilibrium outcomes converge to competitive outcomes, in the sense that every seller charges a price equal to the seller's marginal cost of producing the good. In our model with endogenous quality choice, the limit for $\alpha \to 0$ is monopolistic, exactly as it is in Burdett and Judd (1983). The limit for $\alpha \to 1$ is different than in Burdett and Judd (1983), as it is not perfectly competitive. In fact, it is not the case that every seller charges a price equal to the marginal cost. Intuitively, sellers have an incentive to design vertically differentiated varieties, and, when different sellers carry varieties of different quality, the fact that every buyer is in contact with two sellers is not enough to drive prices down to marginal cost. The fact that every buyer is in contact with two sellers is, however, sufficient to drive the seller's ex-ante profits to zero. Therefore, prices are not driven down to marginal cost of production but to the average cost of design and production.