Optimal Product Design: Implications for Competition and Growth under Declining Search Frictions*

GUIDO MENZIO
NYU and NBER

December 2022

Abstract

As search frictions in the market for a consumer product decline, buyers are able to locate and access more and more sellers. In response, sellers choose to design varieties of the product that are more and more specialized in order to take advantage of the heterogeneity in buyers' preferences. I find conditions on the fundamentals of the market under which the increase in specialization exactly offsets the decline in search frictions. Under these conditions, the extent of competition and the extent of price dispersion remain constant over time even though search frictions are vanishing. Buyer's surplus and seller's profit, however, grow over time at a constant endogenous rate, as the increase in specialization allows sellers to cater better and better to the preferences of individual buyers.

JEL Codes: O40, D43, L13, E23.

Keywords: Price Dispersion, Markups, Imperfect Competition, Search Frictions, Growth.

*Menzio: Department of Economics, New York University, 19 West 4th Street, New York, NY 10012 (email: gm1310@nyu.edu). I am extremely grateful to the editor, Chad Jones, and three anonyous referees for comments that have greatly improved the paper. I also with to thank Jim Albrecht, Costas Arkolakis, Andy Atkeson, Jess Benhabib, Jaroslav Borovicka, Tal Correm, Manolis Galenianos, Boyan Jovanovic, Julian Kozlowski, Ben Lester, Sephora Mangin, Paolo Martellini, Guillaume Nevo, Ezra Oberfield, Jesse Perla, Ran Spiegler, Bruno Sultanum, Ronald Wolthoff, Pierre-Olivier Weill, Randy Wright and participants at several seminars and conferences for their feedback. A major source of inspiration for this paper is Section III of Nobu Kyiotaki and Randy Wright's 1993 classic "A Search Theoretic Approach to Monetary Economics."

1 Introduction

The standard theory of price dispersion for consumer products is based on search frictions (e.g., Butters 1977, Varian 1980, and Burdett and Judd 1983). The theory is simple. Buyers cannot purchase a particular product from just any seller in the market, but only from the sample of sellers that they managed to locate and access as a result of their search. As long as the outcome of the buyers' search is such that some buyers can purchase from only one seller and some buyers can purchase from multiple sellers, sellers find it optimal to post different prices. The theory has sharp predictions. As search frictions become smaller—in the sense that the fraction of buyers who can purchase from multiple sellers increases—the market becomes more competitive, prices fall towards the marginal cost, and price dispersion eventually disappears. From the perspective of the theory, it is then puzzling that improvements in information and communication technology (e.g., phones, mobile phones, the internet, smartphones, etc...), which ostensibly made it easier for buyers to locate and access sellers, have not led to a sizeable reduction in markups or price dispersion. Indeed, the extent of price dispersion in the 1970s (e.g. Pratt, Wise and Zeckhouser 1979), in the1990s (Lach 2002) and in the 2000s (e.g. Kaplan and Menzio 2015) is very similar. Similarly, competition does not appear to be stronger and price dispersion does not appear to be lower in online than offline markets (e.g., Brynjolfsson and Smith 2000, Baye, Morgan and Sholten 1994, Ellison and Ellison 2014).

In this paper, I argue that—once the firms' incentives to design different types of product varieties are taken into account—the predictions of the theory can be reconciled with the observation that competition and price dispersion appear to be constant in the face of declining search frictions. The basic insight is simple and inspired by Kyiotaki and Wright (1993). Declining search frictions not only allow buyers to locate and access more sellers, but they also allow sellers to reach a larger pool of buyers. For this reason, firms have the incentive to design more and more specialized varieties of the product—varieties that cater better and better to the preferences of a smaller and smaller fraction of buyers. Under some conditions on the fundamentals of the product market, the endogenous increase in specialization exactly offsets the exogenous decline in search frictions, and the product market follows a Balanced Growth Path (BGP), along which competition and price dispersion remain stable. Consistently with the argument advanced in this paper, Anderson (2008) and Nieman and Vavra (2021) document a significant rise in product differentiation.

In the first part of the paper, I develop a simple, partial-equilibrium model in order to formalize my argument. I consider the market for some consumer good. The market is populated by firms, retailers, and buyers. Each firm designs its own variety of the good and, in particular, it chooses whether to design a variety that is more generic—in the sense that it appeals to a larger fraction of buyers but gives them lower utility—or more specialized—in the sense that it appeals to a smaller fraction of buyers but gives them higher utility. Each firm sells its variety of the good through a chain of retailers—whose features have a small, random impact on the buyers' valuation of the firm's variety. Each buyer samples a Poisson number of retailers, inspects the variety of the product that they carry, the price that they charge, and decides whether and where to purchase one unit of the good. The firm's profit, the buyer's utility, and all prices are measured in units of a numeraire good. I capture declining search frictions as a positive growth rate in the average number of retailers contacted by an individual buyer.

I seek conditions for the existence of a BGP of the product market—a path along which all endogenous variables grow at some constant rate. I show that a BGP may exist only if the buyer's utility function is iso-elastic in the degree of specialization of the variety that he consumes (conditional on finding it appealing), the quantity of input used by a firm to design a new variety depends on the desired level of specialization of the variety relative to the aggregate level of specialization of the varieties designed in the past, and the price of the input used by the firm to design a variety grows at the same rate as the buyers' surplus. Moreover, I show that under some additional, generic conditions on the fundamentals of the product market, a BGP does exist.

I then characterize the properties of the product market along a BGP. I show that, along a BPG, the degree of specialization of the varieties designed by firms increases at exactly the same rate at which search frictions decline. Hence, even though buyers locate an increasing number of sellers, they locate a constant number of sellers that carry a variety that they find appealing. For this reason, the extent of competition remains constant—in the sense that the share of surplus accruing to buyers and the share of surplus accruing to retailers do not vary—and the extent of price dispersion remains constant—in the sense that the distribution of prices posted by retailers normalized by the average posted price does not vary. Declining search frictions do not make the product market more competitive, nor they make price dispersion disappear. Declining search frictions, however, lead to economic grow—in the sense that the surplus accruing to buyers, the surplus accruing to retailers, and the profits accruing to firms grow over time. The growth rate of the economy is equal to the rate at which search frictions decline multiplied by the elasticity of the buyers' utility function with respect to the degree of specialization of the varieties that they consume—a parameter that captures the extent of heterogeneity in the preferences of different buyers.

In the second part of the paper, I develop a general equilibrium version of the baseline model—which allows me to close the model by endogenizing the price of the input used by firms to design new varieties. In the general equilibrium model, households purchase a continuum of goods in frictional markets, and they sell labor in a competitive market. Firms design their varieties of the goods by hiring labor, and they sell their varieties through chains of retailers. I show that, if households have King-Plosser-Rebelo (1988) preferences over the numeraire good and leisure, the wage endogenously grows at the same rate as the buyers' surplus—a property that I showed to be necessary for the existence of a BGP, and that I had to exogenously impose in the baseline model. Given that households have King-Plosser-Rebelo (1988) preferences and given the two other necessary conditions for a BGP,[1] I show that a BGP exists if the firm's marginal cost of increasing specialization at the same rate at which search frictions decline belongs to some non-degenerate interval. Along a BGP, the increase in specialization exactly offsets the decline in search frictions, competition remains constant, and so does price dispersion. Moreover, the economy converges to the BGP from any initial conditions. If, on the other hand, the firm's marginal cost of increasing specialization at the same rate at which search frictions decline is too low or too high, a BGP does not exist. In this case, the equilibrium dynamics are either such that specialization outpaces the decline in search frictions—in which case the market becomes more and more monopolistic—or such that specialization lags the decline in search frictions—in which case the market becomes more and more competitive. Either way, price dispersion vanishes.

Armed with a general equilibrium model, I can also formulate the problem of a social planner and evaluate the efficiency properties of equilibrium. I show that the equilibrium is statically efficient but dynamically inefficient—in the sense that the equilibrium level of specialization is generally inefficient, but it would be efficient if the social planner did not place any weight on the households' future utility. Intuitively, the equilibrium is dynamically inefficient because, when a firm decides how much to specialize its variety of the product, it fails to internalize the effect of its decision on the future production cost of other firms. The equilibrium is statically efficient because, even though the product market is not competitive, the firm's static return from designing a marginally more specialized variety coincides with the social return.

In the last part of the paper, I develop a version of the baseline model with firm dynamics. In this version of the model, firms have the option of selling the same variety of the product for an extended period of time. The firms' product design problem turns into an optimal stopping problem—in which firms decide when to scrap their existing varieties and what type of new varieties to design—and the equilibrium of the product market becomes an (S,s) equilibrium—in which the distribution of varieties evolves according to

[1] Namely, the buyer's utility function–measured in units of the numeraire good–is iso-elastic in the degree of specialization of the variety that he consumes, and the quantity of labor used by a firm to design a new variety depends on the degree of specialization of the variety relative to the aggregate level of specialization of the varieties designed in the past.

the firms' optimal stopping strategy. I first find conditions for the existence of a BGP. I then characterize the dynamics of an individual firm along a BGP. I show that a firm goes repeatedly through the same cycle. The cycle starts when the firm designs a new variety of the product. At that moment in time, the firm is the best in the market—it sells the most specialized variety, it offers the highest surplus to buyers, and its value is highest. Over the course of the cycle, the firm loses ground relative to the rest of the market—it sells a variety that becomes relatively less specialized, it offers a surplus that becomes relatively lower, and its value declines relative to its competitors. The cycle ends when the firm scraps its variety of the product. At that moment in time, the firm is the worst in the market. This novel theory of firm dynamics emerges from the interaction between the lumpy nature of product design and the continuous evolution of the product market.[2]

The paper is part of an ongoing research project on "Stiglerian Growth", the idea—which dates back to Stigler's original papers on search frictions (Stigler 1961, 1962)—that information frictions cause misallocation and, hence, technological advances that alleviate information frictions lead to economic growth. In Martellini and Menzio (2020), we consider the effect of declining search frictions in the labor market. We show that, if and only if the quality of a firm-worker match is drawn from a Pareto distribution, declining search frictions are consistent with the stationarity of unemployment and vacancies and with the stability of the Beveridge curve that are observed in the data. However, if the quality of a firm-worker match is drawn from a Pareto distribution, declining search frictions generate economic growth—as firms and workers are able to sort into better and better matches. Using data on the growth rate in the number of applications that firms receive before filling their vacancies, we measure the decline in search frictions and its contribution to productivity growth.[3] In this paper, I study declining search frictions in the consumer product market. Here, the key empirical observation is that declining search frictions do not appear to have increased competition nor to have eliminated price dispersion. The observation suggests that, even though declining search frictions directly tend to increase competition, they must also trigger some adjustment that tends to reduce competition. In this paper, I show that increasing product differentiation may be such an adjustment. If this is the case, however, declining search frictions generate economic growth—as households consume products that are better and better tailored to their

[2] From a mathematical point of view, the version of the model with firm dynamics is very related to models of menu cost pricing in both frictionless markets (e.g., Sheshinski and Weiss 1977, Caplin and Spulber 1986, Alvarez, Lippi and Oskolkov 2022) and fritional markets (e.g. Benabou 1988, Diamond 1993, Burdett and Menzio 2017, 2018). In menu cost models, firms adjust their nominal price infrequently to keep up with a declining value of money. In my model, firms adjust their variety infrequently to keep up with declining search frictions.

[3] In Martellini and Menzio (2021), we show that declining search frictions have a different effect on the growth rate of productivity and wages for routine workers (whose productivity is similar across jobs) and non-routine workers (whose productivity is different across jobs).

preferences.

The paper contributes more generally to the theoretical literature on economic growth. Stiglerian growth is related to "Smithian growth", the idea—which dates back to Adam Smith—that the size of the market limits the extent of specialization in production and, for this reason, any increase in the size of the market leads to economic growth. There are several models that formalize Smithian growth. Kelly (1997) considers an economy with multiple locations. Each location has a technological advantage in the production of a particular intermediate good. When locations are not connected, each location inefficiently produces all intermediate goods. When locations become connected, each location specializes in the intermediate good that they produce more efficiently. Locay (1990) considers an economy in which market production features increasing returns to scale, but is subject to agency problems. As the size of the market increases, production moves from the home to the market and productivity increases. The key difference between Stiglerian growth and Smithian growth is that the former is caused by market deepening, while the latter is caused by market expansion. The growth theory proposed in this paper is also related to Romer (1990). In Romer (1990), an endogenous investment in R&D causes the economy to grow by allowing firms to produce more varieties. In this paper, an exogenous decline in search frictions causes the economy to grow by allowing firms to produce more specialized varieties.

The paper also contributes to the strand of the industrial organization literature that tries to explain why price dispersion in online markets does not appear to be lower than in offline markets, even though search frictions in online markets are ostensibly lower. Baye and Morgan (2001) construct a model of online retailing in which a discrete number of sellers decide whether to participate in an online market or not and what price to post. In equilibrium, sellers randomize over participating in the online market and, conditional on participating, they randomize over their online price. Even though the online market has no search frictions, sellers randomize over the price because they may or may not be the only vendor there. Therefore, the standard logic from Butters (1977), Varian (1980) and Burdett and Judd (1983) applies. Ellison and Wolitsky (2012) argue that the lack of search frictions in online markets does not lead to the disappearance of price dispersion because sellers engage in obfuscation—an activity that artificially increases search costs. Ellison and Ellison (2014) compare an online and an offline market for used books. In both markets, each seller carries a single book and meets heterogeneous buyers sequentially. In the online market, sellers meet buyers at a higher rate and they are less likely to meet captive buyers. In the offline market, sellers meet buyers at a lower rate and they are more likely to meet captive buyers. The first difference between the online and the offline market stretches the right-tail of the online price distribution up. The second difference stretches the left-tail of the online price distribution down. Overall, price dispersion may be higher in the online market.

Lastly, the paper contributes to a strand of the industrial organization literature that studies the effect of declining search frictions in the market for consumer products. Dinlersoz and Yorukoglu (2012) study the stationary equilibrium of a frictional product market in which heterogeneous firms choose an advertisement intensity to reach potential buyers and build their customer base. They show that, when the cost of advertisement falls, more efficient firms become larger, while less efficient ones become smaller or exit. Their paper is different from mine, as it does not consider the effect of declining search frictions on product design. Perla (2019) and Guthmann (2020) are similar to Dinlersoz and Yorukoglu (2012), but focus on the life-cycle of a frictional product market as buyers accumulate more and more information about sellers. They find that, as a market ages, it becomes more and more competitive. Bar-Isaac, Caruana and Cunat (2012) is more closely related to my paper. They study a version of the frictional product market of Wolinsky (1986) in which vertically differentiated firms decide to design niche or generic varieties. Higher quality firms choose to make generic varieties (which deliver similar utility to all buyers), while lower quality firms choose to make niche varieties (which deliver higher utility to some buyers and lower utility to others). A one-time decline in search costs increases the fraction of niche firms, increases the share of sales of the best firms, and may also increase the share of sales of the worst firms. Their paper is different from mine because the cost and effect of specialization are modelled differently and because the search process is different. Owing to these differences, their model does not admit a BGP and only generates price dispersion to the extent that firms are vertically differentiated.

2 Baseline Model

In this section, I develop the baseline version of the model. The market for some consumer product is populated by firms, retailers, and buyers. Each firm designs its own variety of the product and, in particular, decides whether to make a broader variety—one that appeals to a larger fraction of buyers but is less valuable to them—or a narrower variety—one that appeals to a smaller fraction of buyers but is more valuable to them. Each firm sells its own variety of the product through a chain of retailers. Each buyer contacts a Poisson number of retailers and, among these, he chooses whether and where to purchase one unit of the product. Improvements in information and communication technology lead to an increase in the number of retailers contacted by each buyer. In Section 2.2, I characterize the static equilibrium of the product market, and show that the equilibrium features price dispersion. In Section 2.3, I characterize the dynamics of equilibrium and, in particular, focus on the conditions for the existence and on the properties of a Balanced

Growth Path—a path of equilibria along which all equilibrium objects grow at a constant rate. I show that, along a BGP, the extent of competition and the extent of price dispersion remain constant, as firms increase the specialization of their varieties at exactly the same rate at which search frictions decline. Buyers' surplus however grows over time, as buyers consume varieties that are better suited to their preferences. In Section 2.4, I show that the restrictions imposed on the fundamentals of the economy in Sections 2.2 and 2.3 are necessary for the existence of a BGP.

2.1 Environment

Consider the market for some consumer product. On one side of the market, there is a measure $1/m$ of firms each designing their own variety of the product and selling it though a chain of directly owned and operated retailers,[4] with $m \in \{1, 2, ...\}$. In every period $t = 0, 1, 2, \ldots$, a firm designs a new variety of the product and, in particular, chooses whether to make a broader variety—one that appeals to a larger fraction of buyers but gives them less utility—or a more specialized variety—one that appeals to a smaller fraction of buyers but gives them more utility. Formally, the firm chooses the breadth x_t of its variety. The firm's retail outlets are heterogeneous and their features—e.g. location, salesforce, etc...–have a small impact on the final breadth of the firm's variety. Formally, the final breadth of the firm's variety at a particular retail outlet is an x drawn from a log-uniform distribution with support $[x_t, x_t \exp(\eta)]$, where $\eta > 0$ is small and captures the extent of heterogeneity across retail outlets. After observing the realizations of x, the firm sets a (possibly different) price p at each of its retail outlets, where p is in units of a numeraire good.

In order to design a variety with breadth x_t, the firm pays a cost $w_t q_t(x_t)$, where w_t denotes the price (in units of the numeraire good) of the input used in the product design process and $q_t(x_t)$ denotes the quantity of input used by the firm. The quantity of input $q_t(x_t)$ is a strictly decreasing function of x_t. That is, it is cheaper for the firm to design a more generic than a more specialized variety. The quantity of input $q_t(x_t)$ is given by $q(x_t/x_{t-1}^{*\gamma})$, where $q(\cdot)$ is a strictly decreasing function, x_{t-1}^* is the market-wide breadth of varieties designed by firms in period $t-1$, and $\gamma \geq 0$ is a parameter that controls the strength of market-wide externalities in the product design process (aggregate know-how). That is, for any $\gamma > 0$, the firm's marginal cost of designing a more specialized variety in period t is lower if the market-wide breadth x_{t-1}^* of varieties designed by firms in period $t-1$ is lower.

[4] Production and retailing are fully integrated within each firm. The assumption is made for simplicity, even though it is not always realistic. While often producers set retail prices, it is not uncommon for retailers to make their own pricing decision. Yet, as long as there is perfect competition among retailers, the pricing decisions of retailers should end up maximizing the profits of the producers.

Table 1: Baseline Environment

	General	BGP Admissible
Firm cost of designing x_t	$w_t q(x_t/x_{t-1}^{*\gamma})$	
Quantity of input	$q(x_t/x_{t-1}^{*\gamma})$	$q(x_t/x_{t-1}^{*})$
Price of input	$w_t = w(\lambda_t, x_t^*, q_t^*/bm)$	$w(\lambda e^{g_\lambda}, xe^{-g_\lambda}, q) = e^{\alpha g_\lambda} w(\lambda, x, q)$
Firm variety x_t at retail	$x \sim \log U[x_t, x_t e^{\eta}]$	
Buyer utility	$u(x)$	$x^{-\alpha}$
Buyer contacts	$n \sim$Poisson(λ_t)	
Search efficiency	$\lambda_{t+1} = \lambda_t e^{g_\lambda}$, $g_\lambda > 0$	
Retailers per firm	m	
Buyers per retailer	b	

On the other side of the market, there is a measure $b > 0$ of buyers. In every period $t = 0, 1, 2, \ldots$, a buyer demands one unit of the product. A buyer searches the market and comes into contact with $n \in \{0, 1, 2, ...\}$ retailers, where n is a random variable drawn from a Poisson distribution with coefficient λ_t. For each of the n retailers with which he comes into contact, the buyer observes the breadth x and the price p of the variety it carries, and learns whether he likes that variety or not. The buyer likes a variety with breadth x with probability x.[5] The buyer then decides whether to purchase a unit of the good and, if so, from which of the n retailers he came into contact. If the buyer ends up purchasing a variety that he likes with breadth x and price p, he obtains the payoff equivalent to $u(x) - p$ units of the numeraire good, where $u(x)$ is a positive and strictly decreasing function. If the buyer ends up purchasing a variety that he does not like with breadth x and price p, he obtains the payoff $-p$. If the buyer ends up not purchasing the good, he obtains a payoff of 0. That is, a more specialized variety appeals to a smaller fraction of buyers, but gives higher utility to those who like it. A broader variety appeals to a larger fraction of buyers, but gives less utility to those who like it.[6]

The environment is non-stationary. The search parameter λ_t grows over time at the

[5] To be clear, it is the breadth x of the firm's variety at a particular retail outlet that determines the buyer's probability of liking the variety and his utility from consuming it, not the breadth x_t of the variety designed by the firm. Again, the idea is that the features of the retail outlets have a (small) effect on the buyer's valuation of the same physical product.

[6] I assume that a buyer has a probability x of liking a variety of the good with breadth $x \in [0, 1]$. If the buyer likes the variety, he enjoys a utility of $u(x)$ from consuming it. If the buyer does not like the variety, he enjoys a utility of 0 from consuming it. This highly stylized approach to modelling specialization and product differentiation is borrowed from Kyiotaki and Wright (1993) and reminiscent of Diamond (1982).

rate $g_\lambda > 0$, in the sense that $\lambda_{t+1} = \lambda_t \exp(g_\lambda)$. Since λ_t is equal to the average number of retailers contacted by an individual buyer, the growth in λ_t implies that, over time, buyers are able to contact more retailers and, in this sense, search frictions become smaller. The growth in λ_t captures the idea that developments in communication and information technology make it easier for buyers to locate and inspect more retailers. In addition, the input price w_t changes over time. I assume that w_t is given by $w(\lambda_t, x_t^*, q_t^*/bm)$, where w is a function[7] that depends positively on the search parameter λ_t, negatively on the current market-wide breadth of varieties x_t^*, and positively on the per-buyer quantity of input used in the product design process q_t^*/bm. The function w is such that $w(\lambda \exp(g_\lambda), x \exp(-g_\lambda), q)$ equals $w(\lambda, x, q) \exp(g_w)$ for some g_w. Since, as I will show in Section 2.4, x_t^* grows at the rate $-g_\lambda$ along any BGP, this property of the function w implies that the input price grows at the rate g_w.

A few comments about the environment, which is summarized in Table 1, are in order. I assume that, when designing its own variety of the product, the firm chooses some breadth x_t but, when retailing its variety of the product, the breadth of the firm's variety is a random draw from a log-uniform distribution with support $[x_t, x_t \exp(\eta)]$. I assume that there is some noise η so that there can be an equilibrium in which all firms choose x_t^*. Without noise, if a mass of firms were to choose x_t^*, there would be a mass of retailers carrying varieties with breadth x_t^*. As shown in Albrecht, Menzio and Vroman (2021), the mass of retailers at x_t^* would create a discontinuity at x_t^* in an individual firm's marginal benefit from designing a more specialized variety and there would be no continuous marginal cost that makes choosing x_t^* optimal for the firm. Without noise, therefore, there could not be no equilibrium in which all firms choose x_t^*. The noise η eliminates the discontinuity in the individual firm's marginal benefit by eliminating the mass of retailers at x_t^*. I assume that the realization of the noise is independent across different retailers operated by the same firm. I make this assumption to preserve the nature of price dispersion in Butters (1977) and Burdett and Judd (1983). The noise purifies the mixed-strategy equilibrium of Butters (1977) and Burdett and Judd (1983). If the realization of the noise were the same across all the retailers operated by the same firm, the purified equilibrium would only feature price dispersion across different varieties of the good (i.e., varieties made by different firms and targeting different buyers). By assuming that the realization of the noise is independent across retailer, the purified equilibrium

[7] The input price w_t should be naturally related to the opportunity cost of using an additional unit of input in the product design process. In turn, the marginal opportunity cost of the input should depend on the quantity of the input used in the product design process and on the value of its alternative uses—which, if the alternative uses of the input are complementary with the product market outcomes—might depend positively on the gains from trade in the product market and, hence, depend negatively on the extent of search frictions and positively on the degree of specialization of varieties. In Section 3, I will present a general equilibrium model in which the input price satisfies these relationships endogenously.

also features price dispersion for the same variety of the good (i.e. a variety made by a particular firm but sold by heterogeneous retail outlets), as it does in Butters (1977) and Burdett and Judd (1983). I assume that the noise η is small. This assumption makes the model comparable to Butters (1977) and Burdett and Judd (1983) and simplifies some derivations by allowing me to take first-order approximations.[8] The noise is assumed to be drawn from a log-uniform distribution. The assumption is convenient as it simplifies the derivative of the firm's problem with respect to x_t. None of the substantive results hinges on it.

I assume that buyers search the market and contact a number n of retailers distributed as a Poisson with coefficient λ_t. When retailers sell varieties that are liked equally by all buyers (as in Butters 1977, Varian 1980 and Burdett and Judd 1983), the extent of competition is determined by the fraction of buyers with n contacts, $n = 0, 1, 2, \ldots$. When retailers sell products that are liked only by a subset of buyers (as in my model), the extent of competition is determined by the fraction of buyers who are in contact with k retailers selling a good that they like, $k = 0, 1, 2....$ This is where the Poisson assumption comes into play. Indeed, if n is distributed as a Poisson with coefficient λ_t and the average breadth of a variety is X_t, then k is distributed as a Poisson with coefficient $\lambda_t X_t$. Therefore, if $\lambda_t X_t$ remains constant over time, the distribution of k is constant and, for that reason, so does competition. Similarly, when retailers sell products that are liked equally by all buyers, the extent of price dispersion depends on the distribution of n. When retailers sell products that are liked by a subset of buyers, the extent of price dispersion depends on the distribution of k. Hence, if $\lambda_t X_t$ remains constant over time, the distribution of k is constant and so is price dispersion.[9]

Finally, I assume that the decision problem of a firm and the decision problem of a buyer are both static. Indeed, I assume that a firm designs a new variety of the product in every period, and that a buyer demands one unit of the product in every period. Given the static nature of the two decision problems, the evolution of the market equilibrium is determined entirely by the effect of changes in the search parameter λ_t and in aggregate level of know-how x^*_{t-1} on the static equilibrium. While the static nature of the decision problems is not realistic, it greatly simplifies the analysis and allows me to clearly identify the economic forces at work.

[8] Specifically, I use the assumption that η is small to analytically derive some properties of the functions $\Psi(\phi)$ and $\Gamma(\phi)$. Numerically, I can verify that the same properties hold for any η.

[9] The Poisson distribution may not be the only one with the property that the number of relevant contacts for a buyer is a distribution that depends on λ_t and X_t only through their product $\lambda_t X_t$. Lester, Visschers and Wolthoff (2015) list other distributions with the same property. The analysis presented in this paper would naturally extend to these distributions.

2.2 Static equilibrium

I am first going to characterize the static equilibrium of the product market. I am then going to characterize the dynamics of the equilibrium of the product market and, in particular, I am going to focus on the conditions for the existence and on the properties of a Balanced Growth Path. I am going to carry out the static and dynamic analysis of equilibrium under the assumption that the buyer's utility from consuming a variety of the product with breadth x is $x^{-\alpha}$ with $\alpha > 1$, the quantity of input needed to design a variety of the product with breadth x is $q(x/x_{t-1}^*)$, and the growth rate g_w of the input price is αg_λ. These assumptions may appear arbitrary. In Section 2.4, however, I will show that they are necessary for the existence of a Balanced Growth Path.

Let me start by characterizing the equilibrium of the product market in some arbitrary period $t = 0, 1, 2 \ldots$. Consider the shopping problem of a buyer. The probability that a buyer comes into contacts with $n = 0, 1, 2 \ldots$ retailers is

$$\lambda_t^n e^{-\lambda_t}/n! \tag{2.1}$$

The buyer obtains a strictly negative surplus from purchasing any variety of the product that he does not like, as the utility from consuming such variety is 0 and the disutility from paying the price $p > 0$ is strictly positive.[10] For this reason, the buyer will only purchase from a retailer that carries a variety that he does like.

The probability that a buyer contacts $k = 0, 1, 2, \ldots$ retailers that sell a variety that he likes is

$$\begin{aligned} & \sum\nolimits_{n=k}^{\infty} \lambda_t^n \frac{e^{-\lambda_t}}{n!} \frac{n!}{k!(n-k)!} X_t^k (1-X_t)^{n-k} \\ = \quad & \lambda_t^k X_t^k \frac{e^{-\lambda_t}}{k!} \sum\nolimits_{n=0}^{\infty} \frac{\lambda_t^n (1-X_t)^n}{n!} \\ = \quad & \lambda_t^k X_t^k \frac{e^{-\lambda_t X_t}}{k!}, \end{aligned} \tag{2.2}$$

where X_t denotes the average breadth of varieties

$$X_t = \int_{x_t^*}^{x_t^* e^\eta} x \frac{1}{\eta x} dx = \frac{x_t^* e^\eta - x_t^*}{\eta}. \tag{2.3}$$

The first line in (2.2) is easy to understand. The buyer comes into contact with $n \geq k$ retailers with probability (2.1). Conditional on contacting $n \geq k$ retailers, the buyer likes the variety of k of them with probability $X_t^k(1-X_t)^{n-k} n!/k!(n-k)!$. The second line in (2.2) is obtained by collecting $\lambda_t^k X_t^k \exp(-\lambda_t)/k!$. The third line in (2.2) is obtained by noting that the summation in the second line equals $\exp(\lambda_t(1-X_t))$. The last expression

[10] It is easy to verify that a retailer can attain a strictly positive profit in equilibrium. For this reason, a retailer finds it optimal to post a strictly positive price.

in (2.2) shows that the number k of retailers contacted by a buyer that carry a variety of the product that the buyer likes is distributed as a Poisson with coefficient $\lambda_t X_t$. Among the k relevant retailers, the buyer purchases from the one that offers him the highest surplus $s \equiv x^{-\alpha} - p$, as long as such surplus is non-negative. The expression for X_t in (2.3) makes use of the fact that the density of the log-uniform distribution is $1/\eta x$.

Next, consider the pricing problem of a retailer carrying a variety of the product with breadth x. The retailer meets b_k of buyers who are in contact with $k = 0, 1, 2, \ldots$ other relevant retailers, where b_k is given by

$$\begin{aligned} b_k &= b \sum_{n=k+1}^{\infty} n \frac{\lambda_t^n e^{-\lambda_t}}{n!} \frac{(n-1)!}{k!(n-k-1)!} X_t^k (1-X_t)^{n-k-1} \\ &= b\lambda_t e^{-\lambda_t} \frac{\lambda_t^k X_t^k}{k!} \sum_{n=0}^{\infty} \frac{(1-X_t)^n \lambda_t^n}{n!} \\ &= b\lambda_t e^{-\lambda_t X_t} \frac{\lambda_t^k X_t^k}{k!}. \end{aligned} \tag{2.4}$$

The first line in (2.4) is easy to understand. The retailer meets $bn\lambda_t^n \exp(-\lambda_t)/n!$ buyers with n contacts, with $n = k+1, k+2, \ldots$ One of the buyer's contacts is the retailer. The probability that exactly k of the other $n-1$ contacts of the buyer carries a variety that he likes is $X_t^k (1-X_t)^{n-k-1}(n-1)!/k!(n-k-1)!$. The second line in (2.4) is obtained by collecting $\lambda_t^{k+1} X_t^k \exp(-\lambda_t)/k!$ in the first line. The third line in (2.4) is obtained by noting that the summation in the second line is equal to $\exp(\lambda_t(1-X_t))$. The last expression has a simple interpretation. The retailer meets $b\lambda_t$ buyers and the probability that one of these buyers is in contact with k other relevant retailers is given by (2.2).

The retailer chooses a price p so as to maximize its profit.[11] Since buyers make their purchasing decision based on the surplus that they are offered by retailers whose variety they like, it is easier to formulate the retailer's problem as choosing the surplus $s = x^{-\alpha} - p$ offered to buyers who like the variety it carries. Formally, the problem of the retailer is

$$R_t(x) = \max_{s \geq 0} \left(\sum_{k=0}^{\infty} b_k x F_t(s)^k \right) (x^{-\alpha} - s), \tag{2.5}$$

where $F_t(s)$ is defined as

$$F_t(s) = \frac{1}{X_t} \int_{x: s_t(x) \leq s} x \frac{1}{\eta x} dx. \tag{2.6}$$

Let me explain (2.6). The cumulative distribution function $F_t(s)$ denotes the fraction of

[11] While retailers are owned by firms, retailers maximize the profits of their parent firms only if they maximize their own profit. It is easy to understand why this is the case. Since the m retailers owned by a parent firm compete for the same buyers with zero probability, their prices do not affect each other's profits. Hence, the parent firm finds it optimal to set prices that maximize the individual profit of each of its m retailers.

retailers that offer a surplus smaller than s, where retailers are weighted by the probability that a buyer likes their variety. In (2.6), $s_t(x)$ denotes the surplus offered by a retailer with a variety x. Next, let me explain (2.5). The retailer meets b_k buyers who are in contact with k other relevant retailers. The probability that one of these buyers chooses to purchase from the retailer rather than from one of the k relevant competitors is $xF_t(s)^k$, where x is the probability that the buyer likes the retailer's variety, and $F_t(s)^k$ is the probability that all of the k competitors offer him a surplus smaller than s. The profit that the retailer enjoys for every unit sold is $x^{-\alpha} - s$.[12]

Substituting (2.4) in (2.5) yields

$$R_t(x) = \max_{s \geq 0} \left(b\lambda_t x \exp^{-\lambda_t X_t (1-F_t(s))} \right) \left(x^{-\alpha} - s \right). \tag{2.7}$$

The search process is such that buyers contact a number n of retailers distributed as a Poisson with coefficient λ_t. Then, buyers choose where to purchase the good based on which of the n retailers' varieties they like, how specialized these varieties are, and at what price they are sold. This seemingly complicated search process leads to the remarkably simple expression for the retailer's profit in (2.7).

Lemma 1 below characterizes the solution of the retailer's pricing problem (2.7). It characterizes the surplus $s_t(x)$ offered by a retailer carrying a variety with breadth x, the maximized profit $R_t(x)$ for a retailer carrying a variety with breadth x, and the surplus distribution $F_t(s)$. To characterize $s_t(x)$, I first use (2.7) to show that $s_t(x)$ is strictly decreasing for all $x \in [x_t^*, x_t^* \exp(\eta)]$. I then show that, in any equilibrium, the lowest surplus offered by a retailer must be 0 and, hence, $s(x_t^* \exp(\eta)) = 0$. Lastly, I use the first-order condition associated with (2.7) to solve for $s_t(x)$. Using the fact that $s_t(x)$ is strictly decreasing, I derive an expression for $F_t(s_t(x))$. Using the expression for $F_t(s_t(x))$, I characterize $R_t(x)$.

Lemma 1. *(Equilibrium Surplus and Profit).*

1. *The surplus $s_t(x)$ offered by a retailer carrying a variety x is*

$$s_t(x) = \frac{\lambda_t}{\eta} \int_x^{x_t^* e^\eta} \hat{x}^{-\alpha} e^{-\frac{\lambda_t}{\eta}(\hat{x}-x)} d\hat{x}, \text{ for } x \in [x_t^*, x_t^* e^\eta]. \tag{2.8}$$

For $x < x_t^*$, $s_t(x) = s_t(x_t^*)$. For $x > x_t^* \exp(\eta)$, $s_t(x) = s_t(x_t^* \exp(\eta))$.

[12] The notation in (2.5) implicitly rules out the existence of mass points in the surplus distribution F_t. If F_t had some mass points, (2.5) would have to be amended to allow for the possibility that some of the buyers who contact the retailer are in touch with another retailer offering the surplus s, in which case the buyers would randomize between the two. As shown in Butters (1977) and Burdett and Judd (1983), however, there exist no equilibria in which the surplus distribution F_t has a mass point.

2. The surplus distribution F_t is such that

$$F_t(s_t(x)) = \frac{x_t^* e^\eta - x}{\eta X_t}, \text{ for } x \in [x_t^*, x_t^* e^\eta]. \tag{2.9}$$

3. The maximized profit $R_t(x)$ for a retailer carrying a variety x is

$$R_t(x) = b\lambda_t x e^{-\frac{\lambda_t}{\eta}(x - x_t^*)} \left(x^{-\alpha} - s_t(x)\right), \text{ for } x \in [x_t^*, x_t^* e^\eta]. \tag{2.10}$$

For $x < x_t^*$, $R_t(x)$ equals $b\lambda_t x \left(x^{-\alpha} - s_t(x_t^*)\right)$. For $x > x_t^* \exp(\eta)$, $R_t(x)$ equals $b\lambda_t x \exp\left(-(x_t^* e^\eta - x_t^*)\lambda_t/\eta\right) x^{-\alpha}$.

Let me now turn to the firm's product design problem. The expected profit for a firm that designs a variety with breadth x_t is

$$V_t(x_t) = m \int_{x_t}^{x_t e^\eta} \frac{R_t(x)}{\eta x} dx - w_t q\left(x_t / x_{t-1}^*\right). \tag{2.11}$$

The expression above is easy to understand. If the firm designs a variety with breadth x_t, one of its retail outlets will carry a variety with actual breadth x, where x is a random variable with density $1/\eta x$ over the support $[x_t, x_t \exp(\eta)]$. Conditional on x, the retail outlet will make a profit of $R_t(x)$. If the firm designs a variety with breadth x_t, it pays a cost $w_t q\left(x_t / x_{t-1}^*\right)$. The first-order condition for the optimal breadth of the firm's variety is

$$-w_t q'\left(x_t / x_{t-1}^*\right) \frac{1}{x_{t-1}^*} = m \left[\frac{R_t(x_t)}{\eta x_t} - \frac{R_t(x_t e^\eta)}{\eta x_t}\right]. \tag{2.12}$$

The left-hand side of (2.12) is the firm's marginal cost of designing a more specialized variety of the product. The right-hand side of (2.12) is the firm's marginal benefit of designing a more specialized variety of the product, which is given by the probability of drawing an x in $[x_t - dx, x_t]$ multiplied by the value of retailing such an x net of the probability of drawing an x in $[(x_t - dx)\exp(\eta), x_t \exp(\eta)]$ multiplied by the value of retailing such an x.

A static equilibrium of the product market is an x_t^* such that a firm finds it optimal to choose x_t^* given that other firms choose x_t^* as well. That is, a static equilibrium is an x_t^* such that

$$\begin{aligned} & -w_t q'\left(\frac{x_t^*}{x_{t-1}^*}\right) \frac{1}{x_{t-1}^*} \\ = \; & bm \frac{\lambda_t x_t^{*-\alpha}}{\eta} \left[1 - \frac{\lambda_t x_t^*}{\eta} \int_1^{e^\eta} z^{-\alpha} e^{-\frac{\lambda_t x_t^*}{\eta}(z-1)} dz - e^{-\frac{\lambda_t x_t^*}{\eta}(e^\eta - 1)} e^{-(\alpha-1)\eta}\right], \end{aligned} \tag{2.13}$$

where above expression is derived from (2.12), using (2.10) to replace $R_t(x_t^*)$ and $R_t(x_t^* \exp(\eta))$, using (2.8) to replace $s_t(x_t^*)$, and changing the variable of integration in $s_t(x_t^*)$ from x to

$z = x/x_t^*$. Given x_t^*, the distribution of varieties across retailers is given by a log-uniform distribution with support $[x_t^*, x_t^* \exp(\eta)]$. A retailer with a variety x offers the surplus $s_t(x)$ to its customers, with $s_t(x)$ given by (2.8). The distribution of surplus offered by different retailers is non-degenerate, as the distribution of x is non-degenerate and $s_t(x)$ is strictly decreasing. A retailer with a variety x posts the price $p_t(x) = x^{-\alpha} - s_t(x)$. The distribution of prices posted by different retailers is non-degenerate, as the distribution of x is non-degenerate and $p_t(x)$ has the property that $p_t'(x) = 0$ implies $p_t''(x) > 0$. For $\eta \to 0$, the distribution of prices posted by different retailers remains non-degenerate, as it converges to the equilibrium price distribution in a version of Burdett and Judd (1983) where buyers contact retailers at the Poisson rate $\lambda_t x_t^*$. Since x is drawn independently for each of the firm's retailers, the distribution of prices posted by different retailers of the same variety is also non-degenerate. For m large, the distribution of prices posted by different retailers of the same variety converges to the distribution of prices posted by all retailers.[13]

In the above characterization of a static equilibrium, I have implicitly assumed that the x_t^* that solves (2.13) maximizes the firm's profit (2.11). However, (2.13) only states that the firm's marginal benefit from designing a more specialized variety of the product is equal to the marginal cost from designing a more specialized variety when the firm chooses x_t^* and all other firms choose x_t^*. Lemma 2 below provides a condition on the cost function q under which the firm's marginal cost is steeper than the firm's marginal benefit and, hence, the x_t^* that solves (2.13) maximizes (2.11). In the remainder of the paper, I will assume that q is always such that the x_t^* that solves (2.13) also maximizes the firm's profit.[14]

Lemma 2. *The x_t^* that solves (2.13) maximizes the firm's profit (2.11) if the cost function q is such that*

$$-q'(x_t/x_{t-1}^*) = -q'(x_t^*/x_{t-1}^*) + q_0 \left[\left(x_t/x_{t-1}^*\right)^{-\beta} - \left(x_t^*/x_{t-1}^*\right)^{-\beta} \right], \tag{2.14}$$

[13]The equilibrium of the model, thus, features dispersion in the prices posted by retailers for any variety of the product, and dispersion in the price posted by retailers for the particular variety of the product produced by a given firm. Kaplan and Menzio (2015) document that both types of price dispersion are pervasive in the data. They measure the first type of price dispersion using the distribution of prices for any item that belongs to some narrowly defined product category (e.g., ketchup) across retailers in a given market (e.g., Manhattan) and in a given point in time (e.g. September 2022). They measure the second type of price dispersion using the distribution of prices for a particular item (e.g., the UPC of a particular bottle of ketchup) in a given market and at a given point in time.

[14]As discussed in Section 2.1, the noise η eliminates the discontinuity at x_t^* in the firm's marginal benefit by eliminating the mass point at x_t^* in the distribution of varieties across retailers. Naturally, though, the smaller is η, the steeper is the firm's marginal benefit at x_t^* and, hence, the steeper the marginal cost function needs to be at x_t^* in order to make the choice optimal for the firm. This is why the noise parameter η appears in the denominator of the lower bound for q_0. Yet, for any fixed $\eta > 0$, there always exist a q_0 high enough to make x_t^* optimal for the firm.

and β and q_0 are parameters such that

$$\beta > \alpha, \text{ and } q_0 > \frac{\lambda_t x_t^*}{\eta} \frac{x_{t-1}^{*-\alpha}}{w_t}. \tag{2.15}$$

2.3 Balanced Growth Path

I now turn to the analysis of the equilibrium dynamics of the product market. I am especially interested in the conditions for the existence and in the properties of a Balanced Growth Path—a path of static equilibria in which all equilibrium objects grow at a constant rate. Specifically, along a BGP, the breadth of varieties designed by firms grows at some constant rate g_x, the distribution of surplus offered by retailers grows at some constant rate[15] g_s, and the distribution of prices posted by retailers grows at some constant rate g_p.

In order to analyze the equilibrium dynamics, it is useful to denote as ϕ_t the product $\lambda_t x_t^*$, i.e. the product between the average number of retailers contacted by an individual buyer—the inverse of search frictions—and the breadth of varieties designed by firms—the inverse of product specialization. It is also useful to remember that the input price w_t is given by $w(\lambda_t, x_t^*, q_t/bm)$, where $w(\lambda, x, q)$ is a function such that that $w(\lambda \exp(g_\lambda), x \exp(-g_\lambda), q)$ equals $w(\lambda, x, q)\exp(g_w)$ for $g_w = \alpha g_\lambda$. Using the definition of ϕ_t and the properties of $w(\lambda, x, q)$, I can rewrite the static equilibrium condition (2.13) as

$$\begin{aligned} -q'\left(\frac{\phi_t e^{-g_\lambda}}{\phi_{t-1}}\right) \frac{\phi_t e^{-g_\lambda}}{\phi_{t-1}} w\left(\lambda_0, \frac{\phi_t}{\lambda_0}, q\left(\frac{\phi_t e^{-g_\lambda}}{\phi_{t-1}}\right)\frac{1}{bm}\right) \frac{\phi_t^\alpha}{\lambda_0^\alpha} \\ = bm\frac{\phi_t}{\eta}\left[1 - \frac{\phi_t}{\eta}\int_1^{e^\eta} z^{-\alpha} e^{-\frac{\phi_t}{\eta}(z-1)} dz - e^{-\frac{\phi_t}{\eta}(e^\eta - 1)} e^{-(\alpha-1)\eta}\right]. \end{aligned} \tag{2.16}$$

The static equilibrium condition (2.16) is satisfied for $\phi_{t-1} = \phi_t = \phi^*$ if

$$\begin{aligned} -q'\left(e^{-g_\lambda}\right) e^{-g_\lambda} w\left(\lambda_0, \frac{\phi^*}{\lambda_0}, \frac{q\left(e^{-g_\lambda}\right)}{bm}\right) \frac{\phi^{*\alpha}}{\lambda_0^\alpha} \\ = bm\frac{\phi^*}{\eta}\left[1 - \frac{\phi^*}{\eta}\int_1^{e^\eta} z^{-\alpha} e^{-\frac{\phi^*}{\eta}(z-1)} dz - e^{-\frac{\phi^*}{\eta}(e^\eta - 1)} e^{-(\alpha-1)\eta}\right]. \end{aligned} \tag{2.17}$$

The right-hand side of (2.17) is a function $\Psi(\phi)$ that is strictly increasing in ϕ, and such that $\Psi(0) = 0$, $\Psi(\infty) = bm\alpha$ and $\Psi'(0) = bm(1 - \exp(-\eta(\alpha - 1)))/\eta$. In Appendix A, I formally establish these properties of $\Psi(\phi)$. Intuitively, $\Psi(\phi)$ is the firm's marginal benefit—scaled by the factor $x_t^{*-\alpha}$—of designing a more specialized variety of the product. The firm's scaled marginal benefit is increasing in ϕ because, when ϕ is higher, the firm

[15] Following Lucas and Moll (2014) and Perla and Tonetti (2014), I say that a distribution grows at a constant rate g when every quantile of the distribution grows at the constant rate g.

meets more buyers and has a stronger incentive to make a more specialized variety. The left-hand side of (2.17) is a function $\Gamma(\phi)$ that inherits the shape of the input price function w, and may be increasing, decreasing, or non-monotonic in ϕ. Intuitively, $\Gamma(\phi)$ is the firm's marginal cost—scaled by the factor $x_t^{*-\alpha}$—of designing a variety of the product that is more specialized than the previous version by the factor $\exp(g_\lambda)$. The firm's scaled marginal cost depends on ϕ through the input price function w. As long as $\Gamma(\phi)$ is such that either $\Gamma(0) > 0$ or $\Gamma'(0) > \Psi'(0)$ and $\Gamma(\infty) < \Psi(\infty)$, there exists a $\phi^* > 0$ that solves (2.17). In turn, if there exists a $\phi^* > 0$ that solves (2.17), there exists a sequence of static equilibria of the product market[16] such that the breadth of the varieties designed by firms grows at the constant rate $g_x = -g_\lambda$.

I now want to characterize the features of the product market along a path of static equilibria such that $\phi_t = \phi^*$. Along such a path, the distribution of varieties across retailers grows at the constant rate $g_x = -g_\lambda$. In period t, firms design varieties of the product with breadth x_t^* and the distribution of varieties across retailers is log-uniform with support $[x_t^*, x_t^* \exp(\eta)]$. Therefore, in period t, the y-th quantile $x_t(y)$ of the distribution of varieties is

$$x_t(y) = x_t^* \exp(\eta y). \tag{2.18}$$

Since x_t^* grows at the rate $-g_\lambda$, it follows from (2.18) that $x_t(y)$ grows at the rate $-g_\lambda$, and the distribution of varieties across retailers grows at the rate $-g_\lambda$.

Along a path of static equilibria such that $\phi_t = \phi^*$, the distribution of surplus offered by retailers grows at the constant rate $g_s = \alpha g_\lambda$. In period t, a retailer carrying a variety at the y-th quantile of the x distribution offers the surplus

$$s_t(x_t(y)) = \frac{\lambda_t}{\eta} \int_{x_t(y)}^{x_t^* e^\eta} x^{-\alpha} e^{-\frac{\lambda_t}{\eta}(x - x_t(y))} dx = x_t^{*-\alpha} \frac{\phi^*}{\eta} \int_{e^{\eta y}}^{e^\eta} z^{-a} e^{-\frac{\phi^*}{\eta}(z - e^{\eta y})} dz, \tag{2.19}$$

where the second equality is obtained by using (2.18) to substitute $x_t(y)$ and by changing the variable of integration from x to $z = x/x_t^*$. Since x_t^* grows at the rate $-g_\lambda$, it follows from (2.19) that the surplus offered by a retailer at a y-th quantile of the variety distribution grows at the rate αg_λ. Since the surplus offered by a retailer at any given quantile of the variety distribution grows at the rate αg_λ, the distribution of surplus across retailers grows at the rate αg_λ.

Similarly, along a path of static equilibria such that $\phi = \phi^*$, the distribution of prices posted by retailers grows at the constant rate $g_p = \alpha g_\lambda$. In period t, a retailer carrying a

[16] Note that, along such a sequence of equilibria, the condition (2.15) is stationary. It is therefore easy to find parameters β and q_0 such that the firm's first-order condition (2.17) is both necessary and sufficient for optimality.

variety at the y-th quantile of the x distribution posts the price

$$p_t(x_t(y)) = x_t^{*-\alpha}\left[e^{-\alpha\eta y} - \frac{\phi^*}{\eta}\int_{e^{\eta y}}^{e^{\eta}} \hat{z}^{-a} e^{-\frac{\phi^*}{\eta}(\hat{z}-e^{\eta y})} d\hat{z}\right]. \tag{2.20}$$

Since x_t^* grows at the rate $-g_\lambda$, it follows from (2.20) that the price posted by a retailer at the y-th quantile of the variety distribution grows at the rate αg_λ. Since the price posted by a retailer at any given quantile of the variety distribution grows at the rate αg_λ, the distribution of prices across retailers grows at the rate αg_λ. Since the price distribution grows over time at the rate αg_λ, so does the variance. A proper comparison of the extent of price dispersion across time, however, requires normalizing prices by their average. And since all quantiles of the price distribution grow at the common rate αg_λ, the distribution of normalized prices remains constant over time, and so does its variance.[17]

Along a path of static equilibria such that $\phi_t = \phi^*$, the total surplus obtained by buyers is

$$S_t^b = \int_{x_t^*}^{x_t^* e^{\eta}} b\lambda_t x s_t(x)\frac{1}{\eta x}dx = x_t^{*-\alpha} b\int_1^{e^{\eta}}\left[\frac{\phi^*}{\eta}\int_z^{e^{\eta}} \hat{z}^{-\alpha} e^{-\frac{\phi^*}{\eta}(\hat{z}-z)} d\hat{z}\right] dz. \tag{2.21}$$

The total surplus obtained by retailers is

$$S_t^r = \int_{x_t^*}^{x_t^* e^{\eta}} b\lambda_t x\left(x^{-\alpha} - s_t(x)\right)\frac{1}{\eta x}dx = x_t^{*-\alpha} b\int_1^{e^{\eta}}\left[z^{-\alpha} - \frac{\phi^*}{\eta}\int_z^{e^{\eta}} \hat{z}^{-\alpha} e^{-\frac{\phi^*}{\eta}(\hat{z}-z)} d\hat{z}\right] dz \tag{2.22}$$

The first equation in (2.21) states that the total surplus obtained by buyers is the sum across all retailers of the number of units sold by each retailer times the surplus offered by each retailer to its customers. The second equation in (2.21) is obtained by using (2.8) to substitute $s_t(x)$ and by changing the variable of integration from x to $z = x/x_t^*$. Similarly, the first equation in (2.22) states that the total surplus obtained by retailers is the sum across all retailers of the number of units sold by each retailer times the price posted by each retailer. The second equation in (2.22) is derived like the second equation in (2.21). It is apparent from (2.21) and (2.22) that the total surplus obtained by buyers and the total surplus obtained by sellers grow at the rate αg_λ. Since both S_t^b and S_t^r grow at the same rate, the fraction of surplus captured by buyers and retailers remains constant over time.

From the above observations, it follows that all the equilibrium objects grow at a constant rate along a path of static equilibria such that $\phi_t = \phi^*$. Hence, such a path of static equilibria constitutes a Balanced Growth Path. In Proposition 3 below, I summarize the conditions for the existence and the properties of such a BGP. In Proposition 6, I rule

[17] In the model, the varieties of the product change over time. The same is true in the data when comparing price dispersion for some narrowly defined product (e.g., ketchup) across time.

out the existence of any other type of BGP.

Proposition 3. *(Existence and Properties of a BGP)*

1. *A BGP in which* $\phi_t = \phi^* > 0$ *for* $t = 0, 1, 2 \ldots.$ *exists if the function* $\Gamma(\phi)$ *is such that either* $\Gamma(0) > 0$ *or* $\Gamma'(0) > bm(1 - \exp(-\eta(\alpha - 1)))/\eta$ *and* $\Gamma(\infty) < bm\alpha$.
2. *Along such a BGP, the breadth of varieties designed by firms grows at the rate* $g_x = -g_\lambda$*; the distribution of surplus offered by retailers grows at the rate* $g_s = \alpha g_\lambda$*; the distribution of prices posted by retailers grows at the rate* $g_p = \alpha g_\lambda$ *and, hence, the distribution of normalized prices remains constant over time and so does its variance.*
3. *Along such a BGP, the total surplus captured by buyers grows at the rate* αg_λ*, the total surplus captured by retailers grows at the rate* αg_λ*, and the ratio between buyer and retailer surplus is constant over time.*

Let me provide some intuition for the results in Proposition 3. As search frictions decline, buyers come into contact with more and more retailers. Taken in isolation, the increase in the number of retailers contacted by a buyer would increase the extent of competition in the market and, eventually, would eliminate all price dispersion. However, as search frictions decline, firms also come into contact with more and more buyers. For this reason, firms find it optimal to design varieties of the product that are tailored to the needs of a smaller and smaller segment of the population. Taken in isolation, the increase in the degree of specialization of varieties would lead to a decline in the extent of competition in the market. Under some conditions on the buyer's utility and the firm's product design cost, the endogenous response in the specificity of varieties exactly offsets the exogenous decline in search frictions. That is, buyers come into contact with a constant number of retailers that carry a variety of the product that they like. As buyers contact a constant number of relevant retailers, the extent of competition in the market remains constant over time and so does price dispersion.

Even though the decline in search frictions does not increase competition nor eliminate price dispersion, it leads to economic growth. In fact, as search frictions decline, firms design more and more specialized varieties of the product and, since buyers get more utility from consuming varieties that are more closely tailored to their needs, the welfare of buyers grows larger and larger and so do the profits of retailers and firms. The rate of economic growth generated by the decline in search frictions is αg_λ, where g_λ is the rate at which search frictions decline and α is the elasticity of the buyers' utility with respect to the degree to which a variety is specialized to their needs. Intuitively, α represents the extent of heterogeneity in buyers' preferences and, as such, the rate of return from firms deciding to make more and more differentiated varieties of the product in response to declining search frictions.

It is interesting to examine how changes in the fundamentals of the product market affect the degree of specialization along a BGP. To this aim, assume that there exists a unique $\phi^* > 0$ that solves (2.17), which, in turn, implies that $\Gamma'(\phi^*) > \Psi'(\phi^*)$. Under this assumption, if the number b of buyers per retailer is higher or if the number m of retailers per firm is higher, ϕ^* is lower—meaning that firms design more specialized varieties along a BGP. This is intuitive. If there are more buyers per retailer or if there are more retailers per firm, the firm's benefit to design a more specialized variety of the product is higher and the firm's cost to design a more specialized variety is lower. Hence, firms find it optimal to design more specialized varieties. If the price of the product design input $w(x, \lambda, q)$ is higher, ϕ^* is higher—meaning that firm design less specialized varieties along a BGP. Similarly, if the marginal quantity of input $q'(\exp(-g_\lambda))$ needed to design a variety that is $\exp(g_\lambda)$ more specialized than the previous one is higher, ϕ^* is higher. These findings are also intuitive.

The following proposition summarizes the comparative statics results.

Proposition 4. *(Comparative Statics) Let there be a unique BGP with $\phi_t = \phi^* > 0$ for $t = 0, 1, 2, \ldots$*

1. *The breadth of varieties along such a BGP is strictly decreasing in the number b of buyers per retailer and in the number of m of retailers per firm.*
2. *The breadth of varieties along such a BGP is strictly increasing in the input price function $w(\lambda, x, q)$ and in the marginal quantity of input $q'(\exp(-g_\lambda))$ needed to design a variety that is more specialized than the previous one by the factor $\exp(g_\lambda)$.*

2.4 Necessary conditions for a BGP

In this section, I show that the functional forms chosen for the utility function u of the buyers, for the quantity q of the input used by the firms to design a new variety of the product, and for the price w of the input are all necessary for the existence of a Balanced Growth Path. That is, for a BGP to exist, it has to be the case that u, q and w have precisely the functional forms adopted in Section 2.2 and 2.3. Additionally, I show that, in any BGP, it has to be the case that the degree of specialization of the varieties grows at the same rate at which search frictions fall and, hence, $\phi_t = \phi^*$ for $t = 0, 1, 2, \ldots$

Let me first provide a formal definition of a BGP.

Definition 5. *A Balanced Growth Path (BGP) is a sequence of equilibria of the product market with the property that:*

1. *The breadth x_t^* of the new varieties grows at some rate g_x;*

2. *The surplus $s_t(x_t^* \exp(\eta y))$ offered by a retailer at the y-th quantile $x_t^* \exp(\eta y)$ of the variety distribution grows at some rate g_s, and the price posted by a seller at the y-th quantile of the variety distribution grows at some rate g_p;*
3. *The price of the input used by firms grows at some rate g_w.*

Given the definition of a BGP, I can now derive a number of necessary implications of a BGP that end up placing restrictions on the form of the functions u, q, and w. Consider a retailer carrying a variety of the product with breadth $x_t^* \exp(\eta)$. Since the retailer is carrying the broadest variety of the product in the market, it finds it optimal to offer the lowest surplus of 0. Hence, the first-order condition of the retailer's problem is satisfied at $s = 0$, i.e.

$$1 = \lambda_t X_t F_t'(0) u(x_t^* e^\eta). \tag{2.23}$$

Notice that the average number λ_t of retailers contacted by a buyer grows at the rate g_λ. The breadth x_t^* of new varieties designed by firms grows at the rate g_x, which implies that the average breadth of varieties X_t across retailers grows at the rate g_x. The surplus offered by a retailer at the y-th percentile of the variety distribution grows at the rate g_s, which implies that $F_t(s \exp(g_s t))$ is equal to $F_0(s)$ and, in turn, that $F_t'(s)$ is equal to $F_0'(s) \exp(-g_s t)$.

In light of the above observations, I can rewrite (2.23) as

$$1 = e^{(g_\lambda + g_x - g_s)t} \lambda_0 X_0 F_0'(0) u(x_0^* e^\eta e^{g_x t}). \tag{2.24}$$

Differentiating (2.24) with respect to t yields

$$0 = (g_\lambda + g_x - g_s) u(x_t^* e^\eta) + g_x u'(x_t^* e^\eta) x_t^* e^\eta. \tag{2.25}$$

The expression in (2.25) is a differential equation for the buyer's utility function $u(x)$. The unique solution is

$$u(x) = u_0 x^{-\alpha}, \tag{2.26}$$

for some $u_0 > 0$ and $\alpha = (g_\lambda + g_x - g_s)/g_x$. The expression in (2.26) states that a BGP may exist only if the buyer's utility function is an iso-elastic function of the breadth x of variety of the product. The coefficient u_0 is a scaling parameter of the utility function, and it can be set to 1 essentially without loss in generality. The parameter α is the elasticity parameter of the utility function. The condition $\alpha = (g_\lambda + g_x - g_s)/g_x$ is not a restriction on the elasticity α. Instead, the condition is a restriction on the endogenous growth rate g_s for an arbitrary elasticity $\alpha > 0$.

Next, consider a retailer carrying a variety of the product with breadth $x_t^* \exp(\eta y)$.

The first-order condition of the retailer's pricing problem is

$$\begin{aligned} 1 &= \lambda_t X_t F_t'(s_t(x_t^* e^{\eta y})) \left(x_t^{*-\alpha} e^{-\alpha \eta y} - s_t(x_t^* e^{\eta y})\right) \\ &= e^{(g_\lambda + g_x - g_s)t} \lambda_0 X_0 F_0'(s_0(x_0^* e^{\eta y})) \left(x_0^{*-\alpha} e^{-\alpha \eta y} e^{-\alpha g_x t} - s_0(x_0^* e^{\eta y}) e^{g_s t}\right). \end{aligned} \tag{2.27}$$

The left-hand side of (2.27) is constant over time. Therefore, the right-hand side of (2.27) must also remain constant over time. From the second line in (2.27), it follows that first term on the right-hand side is constant only if $g_\lambda + g_x - g_s - \alpha g_x = 0$. The second term on the right-hand side is constant only if $g_\lambda + g_x = 0$. These two conditions imply that, in any BGP, $g_x = -g_\lambda$ and $g_s = \alpha g_x$. That is, in any BGP, the degree of specialization of the varieties designed by firms must fall at the same rate at which search frictions decline. Moreover, in any BGP, the surplus offered by retailers at a given quantile of the variety distribution must grow at the rate αg_x.

Lastly, consider the first-order condition of the firm's design problem

$$\begin{aligned} & -w_t q'\left(\frac{x_t^*}{x_{t-1}^{*\gamma}}\right) \frac{x_t^*}{x_{t-1}^{*\gamma}} \\ = \quad & x_t^{*-\alpha} bm \frac{\lambda_t x_t^*}{\eta} \left[1 - \frac{\lambda_t x_t^*}{\eta} \int_1^{e^\eta} z^{-\alpha} e^{-\frac{\lambda_t x_t^*}{\eta}(z-1)} dz - e^{-\frac{\lambda_t x_t^*}{\eta}(e^\eta - 1)} e^{-(\alpha-1)\eta}\right]. \end{aligned} \tag{2.28}$$

Since λ_t grows at the rate g_λ and x_t^* grows at the rate $g_x = -g_\lambda$, the right-hand side of (2.28) grows at the rate αg_λ. Hence, the left-hand side must also grow at the rate αg_λ. That is,

$$w_0 e^{g_w t} q'\left(x_{t-1}^{*1-\gamma} e^{-g_\lambda}\right) x_{t-1}^{*1-\gamma} e^{-g_\lambda} = w_0 q'\left(x_{-1}^{*1-\gamma} e^{-g_\lambda}\right) x_{-1}^{*1-\gamma} e^{-g_\lambda} e^{\alpha g_\lambda t}, \tag{2.29}$$

where g_w denotes the growth rate of the price of the input used by firms in the design process along a Balanced Growth Path.

Differentiating (2.29) with respect to t yields

$$g_w + \frac{q''\left(x_{t-1}^{*1-\gamma} e^{-g_\lambda}\right)}{-q'\left(x_{t-1}^{*1-\gamma} e^{-g_\lambda}\right)}(1-\gamma) g_\lambda - (1-\gamma) g_\lambda = \alpha g_\lambda. \tag{2.30}$$

For $\gamma \neq 1$, (2.30) is a differential equation for $q'(\cdot)$, the derivative of the quantity of input used by firms in the design process. The unique solution to the differential equation is an iso-elastic $q'(\cdot)$, with a constant elasticity that depends on α, the elasticity of the buyer's utility function. I discard this case, since there is no reason why a technology parameter should be related to a preference parameter. For $\gamma = 1$, (2.30) boils down to $g_w = \alpha g_\lambda$. That is, the growth rate of the price of the input must grow at the rate αg_λ along a Balanced Growth Path. In turn, $g_w = \alpha g_\lambda$ if and only if the input price function is such that

$$w(\lambda e^{g_\lambda}, x e^{-g_\lambda}, q) = w(\lambda, x, q) e^{\alpha g_\lambda}. \tag{2.31}$$

The following proposition summarizes the findings from this subsection.

Proposition 6. *(Necessary Conditions for a BGP)*

1. *A BGP may exist only if the utility of the buyer is* $u(x) = x^{-\alpha}$*; the quantity of input used by the firms is* $q(x/x_{t-1}^{*\gamma})$ *with* $\gamma = 1$*; the input price function* $w(\lambda, x, q)$ *is such that* $w(\lambda \exp(g_\lambda), x \exp(-g_\lambda), q)$ *equals* $w(\lambda, x, q) \exp(\alpha g_\lambda)$*.*
2. *Along any BGP, the breadth of varieties designed by the firms grows at the rate* $g_x = -g_\lambda$ *and, hence,* $\phi_t = \phi^*$ *for* $t = 0, 1, 2, \ldots$

3 A General Equilibrium Model

In this section, I develop a general equilibrium version of the baseline model to address some macroeconomic questions. In this version of the model, households purchase differentiated goods in a continuum of frictional markets that operate as in the baseline model, and they sell labor in a competitive labor market. Firms design varieties of the goods by hiring labor from households, and they sell their varieties through a chain of retailers. I show that, if the household's preferences over the numeraire good and leisure have the King-Plosser-Rebelo form, the equilibrium wage endogenously grows at the rate αg_λ along a Balanced Growth Path—a property that I showed to be necessary for the existence of a BGP and that I had to impose exogenously in the baseline model. In Section 3.2, I find conditions under which a BGP exists and show that the economy converges to it from any initial conditions. Along a BGP, the extent of competition and the extent of price dispersion are constant because the increase in the specialization of product varieties exactly offsets the decline in search frictions. If a BGP does not exist, then either the increase in specialization outpaces the decline in search frictions, or the increase in specialization lags the decline in search frictions. In the first case, the product markets become monopolistic. In the second case, the product market becomes competitive. In both cases, price dispersion vanishes. In Section 3.3, I study the welfare properties of equilibrium. I find that the equilibrium is statically efficient, but dynamically inefficient.

3.1 Environment

The economy is populated by a measure b of infinitely-lived households and by a measure $1/m$ of infinitely-lived firms, each operating a chain of m retailers. In every period t, households supply labor in a competitive labor market, they purchase a continuum of different goods indexed by $i \in [0, 1]$ in frictional product markets, and they earn profits. In every period t, firms hire labor to design their varieties of the goods, they sell their varieties of the goods through their retailers, and they pay out profits. A numeraire good is used to execute transactions in both the labor market and the product market.

The household's periodical utility function is

$$U\left(y_t + \int_0^1 x_{i,t}^{-\alpha} di, \ell_t\right), \tag{3.1}$$

where y_t denotes the quantity of the numeraire good consumed by the household, $x_{i,t}^{-\alpha}$ denotes the value to the household of consuming a variety of good i with breadth $x_{i,t}$ expressed in units of the numeraire good,[18] ℓ_t denotes the quantity of leisure enjoyed by the household, and U is a function that is strictly increasing in both of its arguments. The household is endowed with 1 unit of leisure per period. Each household owns an equal share of the firms and, hence, receives a proportional share of their profits.

The market for good i is frictional and operates as in the baseline model. The firm designs a variety of good i with breadth $x_{i,t}$. At each of the firm's retailers, the actual breadth of the variety is a random variable x_i drawn from a log-uniform distribution with support $[x_{i,t}, x_{i,t}\exp(\eta)]$. At each retailer, the firm observes the realization of x_i and then chooses a price p_i, where p_i is measured in units of the numeraire good. A household comes into contact with n retailers, where n is drawn from a Poisson distribution with coefficient λ_t. The household inspects the variety carried by each of the n retailers, it observes the price posted by each of the n retailers, and then decides whether and where to purchase a unit of the good.

The market for labor is competitive. The firm demands $q(x_t/x_{t-1}^*)$ units of labor in order to design varieties of good i with breadth $x_{i,t} = x_t$. The household supplies $1 - \ell_t$ units of labor in order to finance its consumption of goods. Both the firm and the household take the wage w_t as given, where w_t is measured in units of the numeraire good.

The equilibrium conditions for the product market are the same as in the baseline model. As in the baseline model, the household finds it optimal to purchase good i from the retailer that carries a variety that the household likes and that offers to the household the highest surplus $x_{i,t}^{-\alpha} - p_{i,t}$ measured in units of the numeraire good. Since firms own the retailers and firms are owned by households, each firm chooses the breadth of its variety of the good i so as to maximize the total revenue generated by its retailers net of the design cost, where both revenue and cost are measured in units of the numeraire good. Thus, the firm's problem is the same as in the baseline model. Since the firm's retailers do not compete with each other, the firm instructs each of them to offer the

[18] To be clear, $x_{i,t}^{-\alpha}$ denotes the households value of consuming a variety of good i with breadth $x_{i,t}$ conditional on the household liking that variety. If the household consumes a variety of good i that it does not like, its value is 0. If the household does not consume good i, its value is also 0.

Table 2: General Equilibrium Environment

	Expression
Firm production function x_t	$q(x_t/x^*_{t-1})$
Firm variety x_t at retail	$x \sim \log U[x_t, x_t e^{\eta}]$
Household periodical utility	$\frac{1}{1-\sigma}\left(y_t + \int_0^1 x_{i,t}^{-\alpha} di\right)^{1-\sigma} v(\ell_t)$
Household discount rate	ρ
Household contacts	$n \sim$Poisson(λ_t), $\lambda_{t+1} = \lambda_t e^{g_\lambda}$
Labor resource constraint	$b(1-\ell_t) = q(x^*_t/x^*_{t-1})/m$
Retailers per firm	m
Households per retailer	b

surplus $x_{i,t}^{-\alpha} - p_{i,t}$ that maximizes their own revenue, measured in units of the numeraire good. Thus, also the retailer's problem is the same as in the baseline model. Overall, the equilibrium conditions for the product market are given by the surplus function $s_t(x)$ in (2.8) and by the optimality condition for the firm's design problem in (2.13).

The equilibrium conditions for the labor market are novel. The labor supply problem of the household is

$$\begin{aligned} &\max_{\ell_t \in [0,1]} U\left(y_t + \int_0^1 x_{i,t}^{-\alpha} di, \ell_t\right) \\ &\text{s.t. } y_t = w_t(1-\ell_t) + V_t/bm - \int_0^1 p_{i,t} di. \end{aligned} \tag{3.2}$$

The household chooses how much labor $1-\ell_t$ to supply so as to maximize utility subject to a budget constraint. The budget constraint states that the household's consumption of the numeraire good y_t is equal to the sum of the household's labor income, $w_t(1-\ell)$, and the household's share of profits, V_t/bm, net of the household's total expenditures in the non-numeraire goods.

The first-order condition of the household's labor supply problem is

$$w_t U_1\left(y_t + \int_0^1 x_{i,t}^{-\alpha} di, \ell_t\right) = U_2\left(y_t + \int_0^1 x_{i,t}^{-\alpha} di, \ell_t\right), \tag{3.3}$$

where U_1 denotes the partial derivative of the utility function with respect to its first argument, i.e. the marginal utility of consumption of the numeraire good, and U_2 denotes the partial derivative of the utility function with respect to its second argument, i.e. the marginal utility of leisure. Condition (3.3) is standard and states that the households supply labor so as to equate the marginal utility of leisure to the wage multiplied by the marginal utility of consumption.

The first-order condition in (3.3) can be transformed into an equilibrium condition for the wage. To this aim, note that the clearing condition for the labor market is such that the total amount of labor supplied by households is equal to the total amount of labor hired by firms:

$$b(1-\ell_t) = q(x_t^*/x_{t-1}^*)/m. \tag{3.4}$$

The average profit of a firm is equal to the sum of the expected profit of its m net of the design cost:

$$V_t = bm\textstyle\int_0^1 p_{i,t}di - w_t q(x_t^*/x_{t-1}^*). \tag{3.5}$$

Finally, the total utility of a household from consuming the non-numeraire goods is

$$\begin{aligned}\int_0^1 x_{i,t}^{-\alpha}di &= \int_{x_t^*}^{x_t^* e^\eta} \frac{\lambda_t}{\eta} x^{-\alpha} e^{-\frac{\lambda_t}{\eta}(x-x_t^*)}dx \\ &= x_t^{*-\alpha}\frac{\lambda_t x_t^*}{\eta}\int_1^{e^\eta} z^{-\alpha}e^{-\frac{\lambda_t x_t^*}{\eta}(z-1)}dz.\end{aligned} \tag{3.6}$$

Since the household samples retailers independently across the different product markets, the household's total utility from consuming the non-numeraire goods is equal to the household's expected utility from consuming a particular good i. In turn, the household's expected utility from consuming a particular good i is equal to the per-household average utility provided by a retailer of good i. This is the first line in (3.6). The second line in (3.6) is obtained by changing the variable of integration from x to $z = x/x_t^*$.

Using (3.4), (3.5) and (3.6) to substitute out y_t and ℓ_t in (3.3) yields

$$w_t = \frac{U_2\left(x_t^{*-\alpha}\dfrac{\lambda_t x_t^*}{\eta}\displaystyle\int_1^{e^\eta} z^{-\alpha}e^{-\frac{\lambda_t x_t^*}{\eta}(z-1)}dz, 1-\dfrac{q(x_t^*/x_{t-1}^*)}{bm}\right)}{U_1\left(x_t^{*-\alpha}\dfrac{\lambda_t x_t^*}{\eta}\displaystyle\int_1^{e^\eta} z^{-\alpha}e^{-\frac{\lambda_t x_t^*}{\eta}(z-1)}dz, 1-\dfrac{q(x_t^*/x_{t-1}^*)}{bm}\right)}. \tag{3.7}$$

The expression in (3.7) is the equilibrium condition for the wage w_t—the price of the input used by firms in the design of their product varieties. The wage w_t is a function w of the breadth of the firms' varieties x_t^*, the search parameter λ_t, and the quantity of labor $q(x_t^*/x_{t-1}^*)/bm$ supplied by households to firms.

I am going to specialize the household's periodical utility function to be

$$U\left(y_t + \textstyle\int_0^1 x_{i,t}^{-\alpha}di, \ell_t\right) = \frac{1}{1-\sigma}\left(y_t + \textstyle\int_0^1 x_{i,t}^{-\alpha}di\right)^{1-\sigma} v(\ell_t), \tag{3.8}$$

where $\sigma > 0$ and $v(\ell)$ is strictly increasing if $\sigma < 1$ and strictly decreasing if $\sigma > 1$. Given

(3.8), I can rewrite the wage function in (3.7) as

$$w_t = x_t^{*-\alpha} k \left(1 - \frac{q(x_t^*/x_{t-1}^*)}{bm}\right) \frac{\lambda_t x_t^*}{\eta} \int_1^{e^\eta} z^{-\alpha} e^{-\frac{\lambda_t x_t^*}{\eta}(z-1)} dz, \tag{3.9}$$

where $k(\ell) = (v'(\ell)/v(\ell))/(1-\sigma)$. Note that, given the functional form for the periodical utility function in (3.8)—which happens to be the same utility function that is necessary for the existence of a BGP in the neoclassical growth model (King, Plosser and Rebelo 1988)—the wage function in (3.9) endogenously satisfies the condition $w(\lambda \exp(g_\lambda), x \exp(-g_\lambda), q)$ equals $w(\lambda, x, q) \exp(\alpha g_\lambda)$, a condition that was assumed exogenously in Section 2 and that was shown to be necessary for the existence of a BGP.[19]

I can now combine the equilibrium condition for the wage in (3.9) with the equilibrium condition for the breadth of varieties in (2.13) to obtain a single equilibrium condition

$$\begin{aligned} & -q'\left(\frac{x_t^*}{x_{t-1}^*}\right) \frac{x_t^*}{x_{t-1}^*} k \left(1 - \frac{q(x_t^*/x_{t-1}^*)}{bm}\right) \frac{\lambda_t x_t^*}{\eta} \int_1^{e^\eta} z^{-\alpha} e^{-\frac{\lambda_t x_t^*}{\eta}(z-1)} dz \\ = \; & bm \frac{\lambda_t x_t^*}{\eta} \left[1 - \frac{\lambda_t x_t^*}{\eta} \int_1^{e^\eta} \hat{z}^{-\alpha} e^{-\frac{\lambda_t x_t^*}{\eta}(\hat{z}-1)} d\hat{z} - e^{-\frac{\lambda_t x_t^*}{\eta}(e^\eta - 1)} e^{-(\alpha-1)\eta}\right], \end{aligned} \tag{3.10}$$

3.2 BGP and transitional dynamics

I now want to analyze the dynamics of the general equilibrium model. To this aim, it is useful to rewrite the equilibrium condition (3.10) by substituting x_{t-1}^* and x_t^* with ϕ_{t-1}/λ_{t-1} and ϕ_t/λ_t. The substitution yields

$$f\left(\frac{\phi_t}{\phi_{t-1}}\right) \Gamma(\phi_t) = bm \Psi(\phi_t), \tag{3.11}$$

where $f(\delta)$, $\Gamma(\phi)$ and $\Psi(\phi)$ are defined as

$$f(\delta) = -q'\left(\delta e^{-g_\lambda}\right) \delta e^{-g_\lambda} k \left(1 - q(\delta e^{-g_\lambda})/bm\right), \tag{3.12}$$

$$\Gamma(\phi) = \frac{\phi}{\eta} \int_1^{e^\eta} z^{-\alpha} e^{-\frac{\phi}{\eta}(z-1)} dz, \tag{3.13}$$

$$\Psi(\phi) = \frac{\phi}{\eta} \left[1 - \frac{\phi}{\eta} \int_1^{e^\eta} z^{-\alpha} e^{-\frac{\phi}{\eta}(z-1)} dz - e^{-\frac{\phi}{\eta}(e^\eta - 1)} e^{-(\alpha-1)\eta}\right] \tag{3.14}$$

[19] The existence of a BGP requires the wage function w to be such that $w(x \exp(-g_\lambda), \lambda \exp(g_\lambda), q)$ equals $w(x, \lambda, q) \exp(\alpha g_\lambda)$. In turn, the property of w requires that the periodical utility function is such that the marginal rate of substitution between leisure and consumption grows at the rate αg_λ when consumption grows at the rate αg_λ. As we know from King, Plosser and Rebelo (1988), the function (3.8) satisfies this property and so do any of its monotonic transformations. The particular monotonic transformation is irrelevant for the characterization of the BGP in Proposition 7, since agents only make static decisions. The transformation does matter for the solution to the planner's problem in Section 3.4, but it does not affect the comparison between the equilibrium and the solution to the planner's problem in Proposition 8.

The function $\Psi(\phi_t)$ is the firm's marginal benefit—scaled by $x_t^{*-\alpha}$—of designing a more specialized variety of the product. The product $f(\phi_t/\phi_{t-1})\Gamma(\phi_t)$ is the firm's marginal cost—scaled by $x_t^{*-\alpha}$—of designing a variety of the product. The function $\Gamma(\phi_t)$ is the households' consumption measured in units of the numeraire good and scaled by $x_t^{*-\alpha}$. The function $f(\phi_t/\phi_{t-1})$ is the marginal quantity of output needed to make a more specialized variety multiplied by the semi-elasticity of $v(\ell_t)$. The function $\Psi(\phi)$ is strictly increasing in ϕ, and such that $\Psi(0) = 0$ and $\Psi(\infty) = \alpha$. In Appendix B, I show that the function $\Gamma(\phi)$ is strictly increasing in ϕ, and such that $\Gamma(0) = 0$ and $\Gamma(\infty) = 1$. Moreover, $\Gamma(\phi)$ is such that $\Psi(\phi) = \alpha\Gamma(\phi) - \Gamma'(\phi)\phi$, a property that can be used to show that the ratio $\Psi(\phi)/\Gamma(\phi)$ is strictly increasing in ϕ and takes the value $\alpha - 1$ for $\phi \to 0$ and the value α for $\phi \to \infty$. The function $f(\delta)$ depends on the shape of the functions q and k. I assume that $f(\delta)$ is strictly decreasing in δ, and such that $f(\underline{\delta}) = \infty$ and $f(\bar{\delta}) = 0$.[20]

A BGP is a $\phi^* > 0$ that solves (3.11) for $\phi_{t-1} = \phi_t = \phi^*$, i.e.

$$f(1)\Gamma(\phi^*) = bm\Psi(\phi^*). \tag{3.15}$$

From the properties of the ratio $\Psi(\phi)/\Gamma(\phi)$, it follows that there is a $\phi^* > 0$ that solves (3.15) if and only if $f(1)$ is greater than $bm(\alpha - 1)$ and smaller than $bm\alpha$. Moreover, there exists at most one $\phi^* > 0$ that solves (3.15) and, at such ϕ^*, the derivative of the left-hand side of (3.15) is greater than the derivative of the right-hand side. For this reason, all of the comparative statics derived in Section 2 for the baseline model also apply to the general equilibrium model. That is, ϕ^* is decreasing in the number of buyers per retailer, b, it is decreasing in the number of retailers per firm, m, and it is increasing in the marginal quantity of input needed to design a variety of the product that is more specialized than the previous variety by the factor $\exp(g_\lambda)$.

The general equilibrium model provides an equilibrium condition for the price of the input used by firms in the design of product varieties that applies on and off a BGP. For this reason, I can analyze the global dynamics of the general equilibrium model. To this aim, let me rewrite the difference equation (3.12) as

$$f\left(\frac{\phi_t}{\phi_{t-1}}\right) = bm\frac{\Psi(\phi_t)}{\Gamma(\phi_t)}. \tag{3.16}$$

[20] The properties of $f(\delta)$ are natural. The function $f(\delta)$ is given by the product between the additional quantity of labor required by the firm to design a more specialized variety and the semi-elasticity of the household's utility of leisure. If firms design more specialized varieties, i.e. $\delta = \phi_t/\phi_{t-1}$ decreases, it is natural to assume that both the additional quantity of labor required by the firm to design a more specialized variety and the semi-elasticity of the household's utility of leisure increase. If firms design very specialized varieties, the household supplies all of its labor and it is natural to assume that the semi-elasticity of the household's utility of leisure goes to infinity. If, on the other hand, firms design very brad varieties, it is natural to assume that the additional quantity of labor required by the firm to design a more specialized variety goes to zero.

The left-hand side of (3.16) is a positive function that is strictly decreasing in ϕ_t/ϕ_{t-1} and takes the value ∞ for $\phi_t/\phi_{t-1} = \underline{\delta}$, and 0 for $\phi_t/\phi_{t-1} = \bar{\delta}$. The right-hand side of (3.16) is a positive function that is strictly increasing in ϕ_t. From these observations, it follows that the solution of (3.16) with respect to ϕ_t, which I shall denote as $\phi_+(\phi_{t-1})$, exists, is unique and it is strictly increasing in ϕ_{t-1}. Suppose that $f(1) \in (bm(\alpha-1), bm\alpha)$ so that $\phi_+(\phi^*) = \phi^*$ for some $\phi^* > 0$. For $\phi_{t-1} < \phi^*$, $\phi_+(\phi_{t-1}) < \phi^*$. Moreover, since the right-hand side of (3.16) is strictly increasing in ϕ_t and equals $f(1)$ for $\phi_t = \phi^*$, $bm\Psi(\phi_+(\phi_{t-1}))/\Gamma(\phi_+(\phi_{t-1})) < f(1)$ for $\phi_{t-1} < \phi^*$. Since the left-hand side of (3.16) is strictly decreasing in ϕ_t/ϕ_{t-1}, $bm\Psi(\phi_+(\phi_{t-1}))/\Gamma(\phi_+(\phi_{t-1})) < f(1)$ implies $\phi_+(\phi_{t-1})/\phi_{t-1} > 1$ or, equivalently, $\phi_+(\phi_{t-1}) > \phi_{t-1}$. Conversely, for $\phi_{t-1} < \phi^*$, $\phi_+(\phi_{t-1}) > \phi^*$ and $\phi_+(\phi_{t-1}) < \phi_{t-1}$. Overall, $\phi_+(\phi_{t-1})$ is strictly increasing, greater than ϕ_{t-1} for $\phi_{t-1} < \phi^*$, equal to ϕ^* for $\phi_{t-1} = \phi^*$, and smaller than ϕ_{t-1} for $\phi_{t-1} > \phi^*$. Therefore, for any arbitrary initial condition, the economy moves towards the Balanced Growth Path.

The following proposition summarizes the analysis of the dynamics of equilibrium.

Proposition 7. *(General Equilibrium Dynamics)*

1. *A BGP exists if $f(1) \in (bm(\alpha-1), bm\alpha)$. If a BGP exists, it is unique.*
2. *If a BGP exists, the economy converges to it from any initial condition.*
3. *Along a BGP, the breadth of varieties designed by firms grows at the rate $g_x = -g_\lambda$, the distribution of surplus offered by retailers grows at the rate $g_s = \alpha g_\lambda$, the distribution of prices posted by retailers grows at the rate $g_p = \alpha g_\lambda$, and the total surplus captured by households and the total surplus captured by retailers grow at the rate αg_λ.*

It is also interesting to examine the dynamics of the economy when a BGP does not exist. First, consider the case in which a BGP does not exist because $f(1) > bm\alpha$. In this case, $\phi_+(\phi_{t-1})$ is a strictly increasing function of ϕ_{t-1}, but such that $\phi_+(\phi_{t-1}) > \phi_{t-1}$ for all $\phi_{t-1} > 0$ and, hence, ϕ_t diverges to infinity. This is intuitive. When $f(1) > bm\alpha$, the marginal cost of increasing the degree of specialization of varieties by the factor $\exp(g_\lambda)$ exceeds the marginal benefit for all ϕ and, hence, the increase in the specialization of varieties lags the decline in search frictions. As a result, the effective number of retailers ϕ_t contacted by each household increases over time. As ϕ_t diverges to infinity, the product market becomes more and more competitive and price dispersion essentially vanishes.[21] In fact, as ϕ_t diverges to infinity, the share $s_t(x_t(y))/x_t(y)^{-\alpha}$ of the total gains from

[21] These properties are intuitive. For η small, the price distribution is approximately the same as in Burdett and Judd (1983) in which buyers contact sellers at the Poisson rate ϕ. If $\phi \to \infty$, the price distribution in Burdett and Judd (1983) collapses to the competitive price. If $\phi \to 0$, the price distribution in Burdett and Judd (1983) collapses to the monopoly price.

trade offered to households by a retailer at any quantile $y > 0$ of the variety distribution converges to 1. For the same reason, the price posted by a retailer at any quantile $y > 0$ of the variety distribution converges to 0 and price dispersion disappears.

Next, consider the case in which a BGP does not exist because $f(1) < bm(\alpha - 1)$. In this case $\phi_+(\phi_{t-1})$ is a strictly increasing function of ϕ_{t-1} such that $\phi_+(\phi_{t-1}) < \phi_{t-1}$ for all $\phi_{t-1} > 0$ and, hence, ϕ_t converges to zero. This is also intuitive. When $f(1) < bm(\alpha - 1)$, the marginal cost of increasing the degree of specialization of varieties by the factor $\exp(g_\lambda)$ is smaller than the marginal benefit for all ϕ and, hence, the increase in the specialization of varieties outpaces the decline in search frictions. As a result, the effective number of retailers ϕ_t contacted by each household falls over time. As ϕ_t converges to zero, the product market becomes more and more monopolistic and, also in this case, price dispersion vanishes. In fact, as ϕ_t converges to zero, the share $s_t(x_t(y))/x_t(y)^{-\alpha}$ of the total gains from trade offered to households by a retailer at any quantile y of the variety distribution converges to 0. For the same reason, the price posted by a retailer at any quantile y of the variety distribution converges to $x_t(y)$ and, as long as η is small, price dispersion essentially disappears.

3.3 Efficiency

Another advantage of using a general equilibrium version of the model is that I can formulate the problem of a social planner and study the welfare properties of equilibrium. The social planner's problem is

$$\begin{aligned} W_t(x^*_{t-1}) &= \max_{x_t, \ell_t} U\left(\mathbb{E}(x_i^{-\alpha}|x_t), \ell_t\right) + e^{-\rho} W_{t+1}(x_t) \\ \text{s.t. } \mathbb{E}(x_i^{-\alpha}|x_t) &= x_t^{-\alpha} \frac{\lambda_t x_t}{\eta} \int_1^{e^\eta} z^{-\alpha} e^{-\frac{\lambda_t x_t}{\eta}(z-1)} dz, \quad \ell_t = 1 - \frac{q(x_t/x^*_{t-1})}{bm}. \end{aligned} \tag{3.17}$$

The planner maximizes the sum of the current and future household's periodical utility discounted at the rate $\rho > 0$, which is the household's discount rate. The planner instructs firms to design varieties of the goods with breadth x_t and instructs households to supply $1 - \ell_t = q(x_t/x^*_{t-1})/bm$ units of labor. Given x_t, the breadth of the varieties carried by retailers are distributed according to a log-uniform with support $[x_t, x_t \exp(\eta)]$. The planner instructs a household in the market for good $i \in [0, 1]$ to inspect the varieties carried by the n retailers it has contacted, and then purchase the most specialized variety among those that it likes. That is, the planner instructs the household to purchase the variety that delivers the highest $x_{i,t}^{-\alpha}$. In equilibrium, the household inspects the varieties carried by the n retailers it has contacted, and then it purchases the one that offers the highest surplus $s = x_{i,t}^{-\alpha} - p_{i,t}$. Since in equilibrium retailers with a more specialized variety offer a higher surplus, the purchasing instructions of the planner coincide with the purchasing choices of the household. As a result, conditional on x_t, the expected value of $x_i^{-\alpha}$ is the same for the planner as it is in equilibrium.

The first-order condition of the social planner's problem is

$$\begin{aligned} &-U_2\left(\mathbb{E}(x_i^{-\alpha}|x_t^*),\ell_t^*\right)\frac{q'(x_t^*/x_{t-1}^*)}{bmx_{t-1}^*} \\ &= -U_1\left(\mathbb{E}(x_i^{-\alpha}|x_t^*),\ell_t^*\right)\frac{d\mathbb{E}(x_{i,t}^{-\alpha}|x_t)}{dx_t} - e^{-\rho}W_{t+1}'(x_t^*), \end{aligned} \tag{3.18}$$

where

$$\frac{d\mathbb{E}(x_{i,t}^{-\alpha}|x_t)}{dx_t} = -\frac{\lambda_t x_t^{*-\alpha}}{\eta}\left[1-\frac{\lambda_t x_t^*}{\eta}\int_1^{e^\eta} z^{-\alpha}e^{-\frac{\lambda_t x_t^*}{\eta}(z-1)}dz - e^{-\frac{\lambda_t x_t^*}{\eta}(e^\eta-1)}e^{-(\alpha-1)\eta}\right]. \tag{3.19}$$

The derivative of the social planner's value function is

$$W_{t+1}'(x_t^*) = \frac{q'(x_{t+1}^*/x_t^*)}{bm}\frac{x_{t+1}^*}{x_t^{*2}}U_2\left(\mathbb{E}(x_{i,t+1}^{-\alpha}|x_{t+1}^*),\ell_{t+1}^*\right), \tag{3.20}$$

The left-hand side of (3.18) is the cost to the planner of instructing the firm to design marginally more specialized varieties of the goods. The right-hand side is the benefit. The first term is the benefit in the current period—which is the increase in the buyers' expected utility from consuming marginally more specialized varieties. The second term is the benefit in the next period—which, as one can see from (3.20)—is the decline in the future cost of production from producing more specialized varieties in the current period.

Given the specification of the utility function in (3.8) and the definitions of f, Γ, Ψ and ϕ, the first-order condition in (3.18) can be rewritten as

$$\begin{aligned} f\left(\frac{\phi_t^*}{\phi_{t-1}^*}\right) &= bm\frac{\Psi(\phi_t^*)}{\Gamma(\phi_t^*)} \\ &-e^{-\rho}q'\left(\frac{\phi_{t+1}^*e^{-g_\lambda}}{\phi_t^*}\right)\left(\frac{\Gamma(\phi_{t+1}^*)}{\Gamma(\phi_t^*)}\right)^{1-\sigma}\frac{v'\left(\ell_{t+1}^*\right)}{v'\left(\ell_t^*\right)}\left(\frac{\phi_{t+1}^*e^{-g_\lambda}}{\phi_t^*}\right)^{1-\alpha(1-\sigma)}k(\ell_t), \end{aligned} \tag{3.21}$$

Let me compare (3.21) and (3.16)—the difference equations that describes the evolution of ϕ_t in the solution to the planner's problem and in the equilibrium. The left-hand side of (3.21) is the same as the left-hand side of (3.16), and it is decreasing in ϕ_t and increasing in ϕ_{t-1}. The right-hand side of (3.21) is the sum of two terms. The first term is the same as the right-hand side of (3.16), and it is increasing in ϕ_t. The second term is positive. These observations imply that, given the same ϕ_{t-1}, ϕ_t is smaller in the planner's solution than in equilibrium. Moreover, if ϕ_{t-1} is smaller in the planner's solution than in equilibrium, ϕ_t is smaller in the planner's solution than in equilibrium. Hence, given the same initial condition ϕ_0, the path of ϕ_t in the planner's solution is smaller than in equilibrium for all $t = 1, 2, \ldots$ In words, the equilibrium path for the breadth of varieties is inefficiently high.

The equilibrium is inefficient because firms do not internalize the effect of their product design decisions on the accumulation of aggregate know-how—an effect that the social planner takes into consideration and that has a marginal value given by the second term on the right-hand side of (3.21). Absent the knowledge externality (because, e.g., $\rho = \infty$), the equilibrium would be efficient as (3.21) would be identical to (3.16). This finding is somewhat surprising considering that the product market is not competitive. To understand this result, first notice that, conditional on the level of specialization of its variety, a firm's retailer sells an efficient quantity of output. This is because, in equilibrium, a retailer with a more specialized variety offers more surplus to its customers. Second, notice that a firm understands that the effect of designing a marginally more specialized on profit is the same whether its retailers change the amount of surplus that they offer or not. This is because of the envelope condition. These two observations together imply that, when a firm considers designing a marginally more specialized variety, the extra profits the firm earns are equal to the extra social value that the firm generates.

The following proposition summarizes the analysis of the social planner's problem.

Proposition 8. *(Welfare Properties of Equilibrium). For any initial condition x_0, the breadth of varieties x_t designed by firms in equilibrium is inefficiently high for $t = 1, 2, \ldots$*

4 A Firm Dynamics Model

In this section, I consider a version of the baseline model designed to study firm level dynamics. In this version of the model, firms have the option of selling the same variety of the product for an extended period of time. The firms' product design problem turns into an optimal stopping problem—in which firms decide when to scrap their existing varieties and what type of new varieties to design—and the equilibrium of the product market becomes an (S, s) equilibrium—in which the distribution of varieties evolves according to the firms' optimal stopping strategy. In Section 4.2, I find conditions for the existence and characterize the properties of a BGP. In Section 4.3, I characterize the dynamics of an individual firm along a BGP. I show that a firm goes repeatedly through the same cycle. The cycle starts when the firm designs a new variety of the product. At that moment in time, the firm is the best in the market. Over the course of the cycle, the firm loses ground relative to the rest of the market. The cycle ends when the firm scraps its variety of the product. At that moment in time, the firm is the worst in the market.

4.1 Environment

I consider a version of the baseline model in which time is continuous. The product market is populated by a measure $1/m$ of firms. Each firm produces its own variety of the product and sells it through m retail outlets, where $m = 1$ for simplicity.[22] In contrast to the baseline model, each firm can sell the same variety for an extended period of time. When the firm decides to scrap its variety and design a new one with breadth x, it pays a lumpy cost $w_t q(x/X_t)$, where $q(x/X_t)$ is the quantity of input used by the firm to design a new variety with breadth x given that the average breadth of varieties in the market is X_t, and w_t is the price of the input used by the firm to design a new variety. Each firm maximizes the present value of its profits discounted at the rate $\rho > 0$. The firm's product design strategy can be described by a breadth cutoff $x_{h,t}$, above which the firm scraps its current variety of the product, and by the breadth $x_{\ell,t}$ of the new variety of the product designed by the firm.

The product market is also populated by short-lived buyers with unit demand.[23] During each interval of time of length dt, a measure bdt of buyers enters the market. Each buyer contacts n retailers, where n is drawn from a Poisson distribution with coefficient λ_t, and λ_t grows over time at the constant exogenous rate $g_\lambda > 0$. The buyer inspects the variety of the product carried by each retailer, observes the price posted by each retailer, and decides whether and where to purchase the good. If the buyer purchases a variety that he likes with breadth x at the price p, he attains a payoff of $x^{-\alpha} - p$. If the buyer purchases a variety that he does not like at the price p, he attains a payoff of $-p$. If the buyer does not purchase the good, he attains a payoff of 0. The buyer exits the market whether or not he purchases the good.

I assume that the quantity of input $q(x/X_t)$ used by a firm to design a new variety of the product is a strictly decreasing function of x_n, as in the baseline model. However, since in this version of the model time is continuous, I assume that the knowledge spillover is captured by the average breadth of varieties in the market rather than by the average breadth of varieties designed in the previous period. As in the baseline model, I assume that the price w_t of the input is given by $w(\lambda_t, x_{\ell,t}, q_t/b)$, where w is a function with the property that $w(\lambda \exp(-g_\lambda t), x \exp(-g_\lambda t), q)$ is equal to $w(\lambda, x, q) \exp(\alpha g_\lambda t)$.

Before turning to analysis of a BGP, let me highlight a technical point. In the baseline model, I had to introduce some noise in the firm's product design problem so that an

[22] The downside of the simplifying assumption $m = 1$ is that the model will no longer generate dispersion in the price of the same variety of the product—since a particular variety will be carried by one retailer only.

[23] It would be conceptually easy to make the buyer's problem dynamic as well. If buyers had the option of searching for an extended period of time, they would have an endogenous reservation payoff $\underline{u}_t > 0$, which would determine the lowest surplus that they are willing to accept from a retailer. Along a BGP, the endogenous reservation payoff $\underline{u}_t$ would grow at the rate αg_λ, just like all the other values.

equilibrium in which all firms designed varieties with the same breadth could exist. I modelled the noise as a small discrepancy between the breadth of the variety designed by the firm and the breadth of the variety sold by the firm's retail outlets. In this version of the model, I do not need to introduce noise in the firm's product design problem. Indeed, as I will show soon, in any BGP, the distribution of varieties is non-degenerate. Moreover, in any BGP, the distribution of varieties is log-uniform, exactly as posited in the baseline model.

4.2 Existence and properties of a BGP

A BGP is an equilibrium in which all endogenous objects grow at a constant rate. Specifically, the distribution of firms across varieties grows at some constant rate g_x, the surplus offered by a firm at the y-th quantile of the variety distribution grows at some constant rate g_s, the price posted by a firm at the y-th quantile of the variety distribution grows at some constant rate g_p. Also, the breadth $x_{\ell,t}$ of new varieties designed by firms and the breadth $x_{h,t}$ of varieties that are scrapped by firms grow at the rate g_x. Following the same steps as in Section 2.4, it is easy to show that, in any BGP, g_x must be equal to $-g_\lambda$.

I now want to understand whether a BGP does exist. Suppose that, along a BGP, the distribution H_τ of varieties is log-uniform over $[x_{\ell,\tau}, x_{h,\tau}]$ at some arbitrary date $\tau \geq 0$. Over the next interval of time of length dt, all the firms with a variety between $x_{h,\tau}\exp(-g_\lambda dt)$ and $x_{h,\tau}$ scrap their variety of the product and design a new one with breadth between $x_{\ell,\tau}\exp(-g_\lambda dt)$ and $x_{\ell,\tau}$. As a result, the distribution $H_{\tau+dt}$ is such that

$$\begin{aligned} H_{\tau+dt}(xe^{-g_\lambda dt}) &= H_\tau(xe^{-g_\lambda dt}) + 1 - H_\tau(x_{h,\tau}e^{-g_\lambda dt}) \\ &= \frac{\log(xe^{-g_\lambda dt}) - \log(x_{\ell,\tau})}{\log(x_{h,\tau}) - \log(x_{\ell,\tau})} + \frac{\log(x_{h,\tau}) - \log(x_{h,\tau}e^{-g_\lambda dt})}{\log(x_{h,\tau}) - \log(x_{\ell,\tau})} = H_\tau(x) \end{aligned} \tag{4.1}$$

Equation (4.1) means that, if the distribution of varieties is log-uniform over $[x_{\ell,\tau}, x_{h,\tau}]$ at τ, the distribution of varieties grows at the rate $-g_\lambda$ over the next dt units of time. Moreover, since $H_{\tau+dt}(x\exp(-g_\lambda dt))$ is equal to $H_\tau(x)$, it follows that

$$\begin{aligned} H_{\tau+dt}(x) &= H_\tau(xe^{g_\lambda dt}) \\ &= \frac{\log xe^{g_\lambda dt} - \log x_{\ell,\tau}}{\log x_{h,\tau} - \log x} = \frac{\log x - \log x_{\ell,\tau+dt}}{\log x_{h,\tau+dt} - \log x_{\ell,\tau+dt}}. \end{aligned} \tag{4.2}$$

Equation (4.2) means that, if the distribution of varieties is log-uniform over $[x_{\ell,\tau}, x_{h,\tau}]$ at τ, the distribution of varieties at date $\tau + dt$ is log-uniform over $[x_{\ell,\tau+dt}, x_{h,\tau+dt}]$. Since $H_{\tau+dt}$ has the same properties as H_τ, I can iterate the argument forward and conclude that, H_t is log-uniform over $[x_{\ell,t}, x_{h,t}]$ for all $t \geq \tau$. Since τ is arbitrary, this is true for

$\tau = 0$. Hence, given an H_0 that is log-uniform over $[x_{\ell,0}, x_{h,0}]$, H_t is log-uniform over $[x_{\ell,t}, x_{h,t}]$ and grows at the constant rate $-g_\lambda$. It is also easy to show that such an H_0 is the only distribution that is consistent with constant growth given the behavior of individual firms.

The pricing problem of the firm is the same as in the baseline model because, as in the baseline model, the distribution of varieties is log-uniform. Specifically, the surplus $s_t(x)$ offered by a firm carrying a variety of breadth x is given by (2.8), with η now being defined as $\log(x_{h,t}/x_{\ell,t})$. Similarly, the retail profit $R_t(x)$ for a firm carrying a variety of breadth x is given by (2.10). Note that, for a firm carrying a variety $\hat{x} \cdot x_{\ell,t}$ that is $\hat{x}$ times broader than the latest variety $x_{\ell,t}$, the surplus $s_t(\hat{x} \cdot x_{\ell,t})$ is equal to $x_{\ell,t}^{-\alpha}\hat{s}(\hat{x})$, where

$$\hat{s}(\hat{x}) = \frac{\lambda_t x_{\ell,t}}{\eta} \int_{\hat{x}}^{e^\eta} z^{-\alpha} e^{-\frac{\lambda_t x_{\ell,t}}{\eta}(z-\hat{x})} dz. \tag{4.3}$$

Similarly, the profit $R_t(\hat{x} \cdot x_{\ell,t})$ is equal to $x_{\ell,t}^{-\alpha} r(\hat{x})$, where

$$r(\hat{x}) = b\lambda_t x_{\ell,t} \hat{x} e^{-\frac{\lambda_t x_{\ell,t}}{\eta}(\hat{x}-1)}(1 - \hat{s}(\hat{x})). \tag{4.4}$$

In words, (4.3) states that the surplus offered by a firm carrying a variety with relative breadth $\hat{x}$ at time t is equal to the product between $x_{\ell,t}^{-\alpha}$ and a normalized surplus function $\hat{s}(\hat{x})$ that depends only on $\hat{x}$ and not on t. Similarly, (4.4) states that the retail profit of a firm carrying a variety with relative breadth $\hat{x}$ at time t is equal to the the product between $x_{\ell,t}^{-\alpha}$ and a normalized profit function $r(\hat{x})$ that only depend on $\hat{x}$ and not on t. Note that $\hat{s}(\hat{x})$ and $r(\hat{x})$ do not depend on t because, along a BGP, $\lambda_t x_{\ell,t}$ is constant.

The product design problem of the firm is different than in the baseline model, as the firm has the option of selling the same variety for an extended period of time. The product design problem of the firm is

$$V_t(x) = \max_{T, x_n} \int_0^T e^{-\rho\tau} R_{t+\tau}(x) d\tau + e^{-\rho T} \left[V_{t+T}(x_n) - w_{t+T} q\left(x_n / X_{t+T}\right)\right]. \tag{4.5}$$

The value $V_t(x)$ for a firm carrying a variety of the product with breadth x at date t is given by the sum of two terms. The first term is the integral of the firm's retail profits between time t between and $t+T$. During this interval of time, the firm retails its existing variety of the product. The second term is the firm's continuation value of designing a new variety of the product with breadth x_n. The firm discounts profits at the rate ρ. The firm chooses how much to wait before scrapping its existing variety of the product, T, and the breadth of its new variety of the product, x_n, to maximize its value.

I guess that the firm's value function $V_t(x)$ is such that

$$V_t(\hat{x} \cdot x_{\ell,t}) = x_{\ell,t}^{-\alpha} \cdot v\left(\hat{x}\right). \tag{4.6}$$

In words, I guess that the value of a firm with a variety with relative breadth $\hat{x}$ at time t is equal to the product between $x_{\ell,t}^{-\alpha}$ and a normalized value function $v(\hat{x})$ that depends only on $\hat{x}$ and not on t. In order to verify the guess in (4.6), I substitute (4.6) on the right-hand side of (4.5) to obtain

$$\begin{aligned} V_t(\hat{x} \cdot x_{\ell,t}) &= \max_{T,x_n} \int_0^T e^{-\rho\tau} x_{\ell,t+\tau}^{-\alpha} r(\hat{x} e^{g_\lambda \tau}) d\tau + e^{-\rho T} x_{\ell,t+T}^{-\alpha} \left[v\left(\frac{x_n}{x_{\ell,t+T}}\right) - \frac{w_{t+T}}{x_{\ell,t+T}^{-\alpha}} q\left(\frac{x_n}{X_{t+T}}\right)\right] \\ &= x_{\ell,t}^{-\alpha} \cdot \max_{T,\hat{x}_n} \int_0^T e^{-\hat{\rho}\tau} r\left(\hat{x} e^{g_\lambda \tau}\right) d\tau + e^{-\hat{\rho} T} \left[v\left(\hat{x}_n\right) - \hat{w} q\left(\hat{x}_n / \hat{X}_{t+T}\right)\right] \\ &= x_{\ell,t}^{-\alpha} \cdot v\left(\hat{x}\right). \end{aligned} \tag{4.7}$$

In the first line of (4.7), I use the fact that $R_{t+\tau}(\hat{x} \cdot x_{\ell,t})$ is equal to $x_{\ell,t+\tau}^{-\alpha} r(\hat{x}\exp(g_\lambda \tau))$, and I use the guess that $V_{t+\tau}(\hat{x}_n \cdot x_{\ell,t+\tau})$ is equal to $x_{\ell,t+\tau}^{-\alpha} v(\hat{x}_n)$. In the second line, I first change the choice variable from x_n to $\hat{x}_n = x_n / x_{\ell,t+T}$. I then make use of the fact that $x_{\ell,t+\tau} = x_{\ell,\tau}\exp(-g_\lambda \tau)$ and I let $\hat{\rho}$ denote $\rho - \alpha g_\lambda$. Lastly, I let $\hat{X}$ denote $X_{t+T}/x_{\ell,t+T}$ and $\hat{w}$ denote $w_{t+T}/x_{\ell,t+T}^{-\alpha}$, where $\hat{X}$ and $\hat{w}$ are both constant along a BGP. In the third line, I verify the guess (4.6).

From (4.7), it follows that the firm's normalized value function $v\left(\hat{x}\right)$ is

$$v\left(\hat{x}\right) = \max_{T,\hat{x}_n} \int_0^T e^{-\hat{\rho}\tau} r\left(\hat{x} e^{g_\lambda \tau}\right) d\tau + e^{-\hat{\rho} T} \left[v\left(\hat{x}_n\right) - \hat{w} q\left(\hat{x}_n / \hat{X}\right)\right]. \tag{4.8}$$

The firm's normalized value of carrying a variety with relative breadth $\hat{x}$ is given by the sum of two terms. The first term is the integral of the firm's normalized retail profits over the next T units of time. Over this interval of time, the firm's retail profit depends on the relative breadth of the firm's variety, which is constant in absolute terms, but grows at the rate g_λ in relative terms. The second term is the normalized firm's value of designing a new variety with relative breadth $\hat{x}_n$. The effective rate at which the firm discounts its profits is $\hat{\rho} = \rho - \alpha g_\lambda$, i.e. the actual discount rate net of the growth rate of the scaling factor $x_{\ell,t}^{-\alpha}$.

In equilibrium, the firm must find it optimal to design a new variety with relative breadth $\hat{x}_\ell = x_{\ell,t}/x_{\ell,t}$. The first-order condition for the firm's choice of $\hat{x}_n$ in (4.8) is satisfied by $\hat{x}_\ell$ if

$$-\hat{w} q'\left(\hat{x}_\ell / \hat{X}\right) / \hat{X} = -v'(\hat{x}_\ell). \tag{4.9}$$

The left-hand side of (4.9) is the firm's normalized marginal cost from designing a more specialized variety of the product. The right-hand side of (4.9) is the firm's normalized marginal benefit from designing a more specialized version of the product. Under an appropriate choice of q, the first-order condition (4.9) is not only necessary but also sufficient for the optimality of $\hat{x}_\ell$.

In equilibrium, the firm must find it optimal to scrap its old variety of the product when its relative breadth reaches the cutoff $\hat{x}_h = x_{h,t}/x_{\ell,t}$. The first-order condition for the firm's choice of T in (4.8) is effectively a condition on the relative breadth of its variety at time T and it is satisfied by $\hat{x}_h$ if

$$r(\hat{x}_h) = \hat{\rho}\left[v(\hat{x}_\ell) - \hat{w}q(\hat{x}_\ell/\hat{X})\right]. \tag{4.10}$$

The left-hand side of (4.10) is the firm's normalized benefit from scrapping the old variety an instant later. This benefit is given by the firm's normalized profit from retailing $\hat{x}_h$. The right-hand side of (4.10) is the firm's normalized opportunity cost from scrapping the old variety an instant later. This cost is equal to $v(\hat{x}_\ell) - \hat{w}q(\hat{x}_\ell/\hat{X})$ annuitized by the effective discount factor $\hat{\rho}$. Since $r(\hat{x})$ is strictly decreasing, the first-order condition (4.10) is not only necessary but also sufficient for the optimality of $\hat{x}_h$.

Lastly, note that, for all $\hat{x} < \hat{x}_h$ and $dt > 0$ small, (4.8) implies

$$v(\hat{x}) = \int_0^{dt} e^{-\hat{\rho}\tau} r(\hat{x}e^{g_\lambda \tau})d\tau + e^{-\hat{\rho}dt}v(\hat{x}e^{g_\lambda dt}). \tag{4.11}$$

Subtracting $\exp(-\hat{\rho}dt)v(\hat{x})$ from both sides of (4.11), dividing by dt, and taking the limit for $dt \to 0$ yields

$$\hat{\rho}v(\hat{x}) = r(\hat{x}) + v'(\hat{x})\hat{x}g_\lambda. \tag{4.12}$$

In words, (4.12) states that $\hat{\rho}v(\hat{x})$—which is the normalized value of a firm carrying a variety of the product with relative breadth $\hat{x}$ annuitized at the discount rate $\hat{\rho}$—is equal to the firm's instantaneous retail profit, $r(\hat{x})$, plus the time-derivative of the firm's normalized value, $v'(\hat{x})\hat{x}g_\lambda$. The equation in (4.12) is the continuous-time Bellman equation for $v(\hat{x})$.

From the above analysis, it follows that a Balanced Growth Path is a solution to a system of three equations in three unknowns. The three unknowns are the normalized value of a firm that has just designed a new variety of the product, $v^* = v(\hat{x}_\ell)$, the relative breadth of a variety of the product scrapped by a firm, $\hat{x}_h$, and the rate at which buyers contact firms carrying a relevant variety of the product, ϕ^*. Note that, since all breadths are expressed relative to $x_{\ell,t}$, $\hat{x}_\ell$ is equal to 1 and, hence, known.

The unknowns v^*, ϕ^*, and $\hat{x}_h$ must satisfy the following equations.[24] The first equation,

[24]The system of equations is rather complicated and I was not able to prove that it admits a solution, nor to find conditions on parameters under which it does. It is, however, easy to analyze the system of equations numerically and, hence, it is easy to check whether it admits a solution for particular parameter values.

which pins down v^* as a function of ϕ^* and $\hat{x}_h$, is

$$v^* = \int_0^{\frac{1}{g_\lambda}\log\left(\frac{\hat{x}_h}{\hat{x}_\ell}\right)} e^{-\hat{\rho}\tau} r\left(\hat{x}e^{g_\lambda \tau}\right) d\tau + \left(\frac{\hat{x}_h}{\hat{x}_\ell}\right)^{-\frac{\hat{\rho}}{g_\lambda}} \left[v^* - q\left(\frac{\hat{x}_\ell}{\hat{X}}\right) w\left(\lambda_0, \frac{\phi^*}{\lambda_0}, \frac{q(\hat{x}_\ell/\hat{X})}{b}\right) \frac{\phi^{*\alpha}}{\lambda_0^\alpha}\right]. \tag{4.13}$$

Equation (4.13) is obtained from the firm's value function (4.8) evaluated at $\hat{x} = \hat{x}_\ell$, after making use of the fact that the firm finds it optimal to scrap a variety with relative breadth $\hat{x}_\ell$ after $\log(\hat{x}_h/\hat{x}_\ell)/g_\lambda$ units of time and that the firm finds it optimal to design a new variety with relative breadth $\hat{x}_\ell$. Moreover, equation (4.13) makes use of the fact that the price of the input in product design is given by a function $w(\lambda, x, q)$ such that $w(\lambda \exp(g_\lambda t), x \exp(-g_\lambda t), q)$ equals $w(\lambda, x, q) \exp(\alpha g_\lambda t)$. This property of w implies that the normalized input price $\hat{w}$ is equal to $w(\lambda_0, \phi^*/\lambda_0, q(\hat{x}_\ell/\hat{X})/b)\phi^{*\alpha}/\lambda_0^\alpha$.

The second equation, which pins down ϕ^* as a function of v^*, is

$$r(\hat{x}_\ell) = \hat{\rho} v^* - g_\lambda q'\left(\frac{\hat{x}_\ell}{\hat{X}}\right) \frac{\hat{x}_\ell}{\hat{X}} w\left(\lambda_0, \frac{\phi^*}{\lambda_0}, \frac{q(\hat{x}_\ell/\hat{X})}{b}\right) \frac{\phi^{*\alpha}}{\lambda_0^\alpha} \tag{4.14}$$

Equation (4.14) is obtained from the first-order condition (4.9) for the breadth of a newly designed variety, after the continuous-time Bellman equation (4.12) is used to replace $v'(\hat{x}_\ell)$ with $(\hat{\rho} v^* - r(\hat{x}_\ell))/g_\lambda \hat{x}_\ell$ and the properties of the input price function w are used to replace $\hat{w}$ with $w(\lambda_0, \phi^*/\lambda_0, q(\hat{x}_\ell/\hat{X})/b)\phi^{*\alpha}/\lambda_0^\alpha$.

The third equation, which pins down $\hat{x}_h$ as a function of v^* and ϕ^*, is

$$r(\hat{x}_h) = \hat{\rho}\left[v^* - q\left(\frac{\hat{x}_\ell}{\hat{X}}\right) w\left(\lambda_0, \frac{\phi^*}{\lambda_0}, \frac{q(\hat{x}_\ell/\hat{X})}{b}\right) \frac{\phi^{*\alpha}}{\lambda_0^\alpha}\right]. \tag{4.15}$$

Equation (4.15) is obtained from the first-order condition (4.10) for the relative breadth of a scrapped variety $\hat{x}_h$, after making use of the fact that the normalized input price $\hat{w}$ is equal to $w(\lambda_0, \phi^*/\lambda_0, q(\hat{x}_\ell/\hat{X})/b)\phi^{*\alpha}/\lambda_0^\alpha$.

The following proposition summarizes the analysis of a BGP.

Proposition 9. *(BGP with Firm Dynamics)*

1. *A BGP exists if the system of equations (4.13)-(4.15) has a solution with respect to v^*, ϕ^* and $\hat{x}_h$.*
2. *Along a BGP, the distribution of varieties carried by firms is log-uniform over the interval $[x_{\ell,t}, x_{h,t}]$ and grows at the rate $g_x = -g_\lambda$; the distribution of surplus offered by firms grows at the rate $g_s = \alpha g_\lambda$; the distribution of prices posted by firms grows at the rate $g_p = \alpha g_\lambda$; the total surplus accruing to buyers and the total surplus accruing to firms grow at the rate αg_λ.*

4.3 Firm dynamics

Along a BGP, the product market grows at a constant rate. Underneath this constant growth, however, individual firms go through cycles of length $T = \log(\hat{x}_h/\hat{x}_\ell)/g_\lambda$. A firm's cycle begins when the firm designs a new variety of the product. At that moment in time, the firm carries the most specialized variety in the market. The surplus offered by the firm to its customers is the highest in the market, as the normalized surplus function $\hat{s}(\hat{x})$ is strictly decreasing in $\hat{x}$. Similarly, the flow profit and the value of the firm are the highest in the market, as the normalized flow profit function $r(\hat{x})$ and the normalized value function $v(\hat{x})$ are both strictly decreasing in $\hat{x}$. In other words, at the beginning of its cycle, the firm is the best in the market.

Over the course of its cycle, the firm loses ground relative to the rest of the market. The firm carries a variety of the product that becomes less and less specialized relative to the rest of the market, as more and more of the firm's competitors scrap their old varieties of the product and design new ones that are more specialized than the one carried by the firm. Specifically, after τ units of time since the beginning of its cycle, the variety carried by the firm is at the $H_{t+\tau}(x_{\ell,t}) = \tau/T$ quantile of the breadth distribution. Over the course of its cycle, the surplus offered by the firm to its customers declines relative to the rest of the market. In fact, given that $\hat{s}(\hat{x})$ is strictly decreasing in $\hat{x}$, after τ units of time since the beginning of its cycle, the surplus offered by the firm to its customers is at the $1 - \tau/T$ quantile of the surplus distribution. Similarly, the flow profit and the value of firm decline relative to the competition. In fact, since $r(\hat{x})$ and $v(\hat{x})$ are strictly decreasing in $\hat{x}$, after τ units of time since the beginning of its cycle, the firm is at the $1 - \tau/T$ quantile of both the profit and the value distributions.

After T units of time since the beginning of its cycle, the firm is the worst in the market. The firm carries the least specialized variety of the product. Since $\hat{s}(\hat{x})$ is strictly decreasing in $\hat{x}$, the firm offers to its customers the lowest surplus in the market. Since $r(\hat{x})$ and $v(\hat{x})$ are strictly decreasing in $\hat{x}$, the firm's flow profit and the firm's value are the lowest in the market. It is at this point in time that the firm finds it optimal to scrap its variety, pay the lumpy cost, and design a new, more specialized variety of the product. The firm's cycle comes to an end and a new one begins.

I summarize the results on firm dynamics in the following proposition.

Proposition 10. *(Firm Dynamics). Along a BGP, an individual firm goes repeatedly through the same cycle of length* $T = \log(\hat{x}_h/\hat{x}_\ell)/g_\lambda$*.*

1. *The cycle begins with the firm designing a new variety with relative breadth* $\hat{x}_\ell$*.*
2. *The cycle ends with the firm scrapping its variety with relative breadth* $\hat{x}_h$*.*
3. *After* τ *units of time since the beginning of the cycle, the firm is at the* τ/T *quantile*

of the breadth distribution, at the $1 - \tau/T$ *quantile of the surplus distribution, and at the* $1 - \tau/T$ *quantile of the profit and value distribution.*

5 Conclusions

In the face of technological innovations that have ostensibly reduced search frictions in retail markets, competition does not appear to have increased nor price dispersion appear to have declined. This is puzzling from the perspective of the standard theory of price dispersion of Butters (1977), Varian (1980) and Burdett and Judd (1983), which predicts that declining search frictions ought to increase competition and lead to the disappearance of price dispersion. In this paper, I showed that, once one takes into account the effect of declining search frictions on the firms' incentives to design different types of products, the theory can be reconciled with the observed stability of competition and price dispersion. Intuitively, as search frictions decline, firms have an incentive to design product varieties that are increasingly specialized—in the sense that they appeal more and more to a smaller and smaller fraction of buyers. Under some conditions on the fundamentals of the economy, there exists a Balanced Growth Path, along which firms increase the specialization of their varieties at exactly the same rate at which search frictions decline—in the sense that, even though buyers locate an increasing number of retailers, they locate a constant number of retailers that carry a variety of the product that they like. Consequently, along a BGP, the extent of competition and the extent of price dispersion remain constant. However, along a BGP, there is economic growth—as buyers end up consuming product varieties that match their heterogeneous preferences better and better. I established these results in a variety of settings in order to demonstrate the robustness of the key mechanism and to highlight its implications for the behavior of a single product market, for the behavior of the aggregate economy, and for the behavior of an individual firm.

References

[1] Albrecht, J., G. Menzio, and S. Vroman. 2021. "Vertical Differentiation in Frictional Product Markets." NBER Working Paper 29618.

[2] Alvarez, F., F. Lippi, and A. Oskolkov. 2022. "The Macroeconomics of Sticky Prices with Generalized Hazard Functions." Quarterly Journal of Economics, 137: 989-1038.

[3] Anderson, C. 2008. The Long Tail. Hyperion, New York.

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

[5] Baye, M. and J. Morgan. 2001. "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets." American Economic Review, 91: 454-474.

[6] Baye, M., J., Morgan, and P. Scholten. 2004. "Information, Search, and Price Dispersion." Handbook of Economics and Information Systems. Elsevier, Amsterdam.

[7] Bénabou R. 1988. "Search, Price Setting and Inflation." Review of Economic Studies, 55: 353-376.

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

[9] Burdett, K, and G. Menzio. 2017. "The (Q,S,s) Pricing Rule: A Quantitative Analysis." Research in Economics, 71: 784-797.

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

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

[12] Caplin, A., and D. Spulber. 1986. "Menu Costs and the Neutrality of Money." Quarterly Journal of Economics, 102: 703-725.

[13] Diamond, P. 1982. "Aggregate Demand Management in Search Equilibrium." Journal of Political Economy, 90: 881-894.

[14] Diamond, P. 1993. "Search, Sticky Prices, and Inflation." Review of Economic Studies, 60: 53-68.

[15] Dinlersoz, E., and M. Yorukoglu. 2012. "Information and Industry Dynamics." American Economic Review, 102: 884-913.

[16] Ellison, G., and S. Ellison. 2014. "Match Quality, Search, and the Internet Market for Used Books." Manuscript, MIT.

[17] Ellison, G., and A. Wolitzky. 2012. "A Search Cost Model of Obfuscation." RAND Journal of Economics, 43: 417-441.

[18] Guthmann, R. 2020. "Price Dispersion in Dynamic Competition." Manuscript, PUC-Rio.

[19] Kaplan, G., and G. Menzio. 2015. "The Morphology of Price Dispersion." International Economic Review, 56: 1165-1206.

[20] Kelly, M. 1997. "The Dynamics of Smithian Growth." Quarterly Journal of Economics, 112: 939-964.

[21] King, R., C. Plosser, and S. Rebelo. 1988. "Production, Growth and Business Cycles I. The Basic Neoclassical Model." Journal of Monetary Economics, 21: 195-232.

[22] Kiyotaki, N., and R. Wright. 1993. "A Search-Theoretic Approach to Monetary Economics." American Economic Rewiew, 83: 63-77.

[23] Lach, S. 2002. "The Existence and Persistence of Price Dispersion: An Empirical Analysis." NBER Working Paper 8737.

[24] Lester, B., L. Visschers and R. Wolthoff. 2015. "Meeting technologies and optimal trading mechanisms in competitive search markets." Journal of Economic Theory, 155: 1-15.

[25] Locay, L. 1990. "Economic Development and the Division of Production between Households and Markets." Journal of Political Economy, 98: 965-982.

[26] Lucas, R., and B. Moll. 2014. "Knowledge Growth and the Allocation of Time." Journal of Political Economy, 122: 1-51.

[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. Forthcoming.

[29] Neiman, B., and J. Vavra. 2021. "The Rise of Niche Consumption." Manuscript, University of Chicago.

[30] Perla, J. 2019. "A Model of Product Awareness and Industry Life Cycles." Manuscript, UBC.

[31] Perla, J., and C. Tonetti. 2014. "Equilibrium Imitation and Growth." Journal of Political Economy, 122: 52-76.

[32] Pratt, J., D. Wise, and R. Zeckhauser. 1979. "Price Differences in Almost Competitive Markets." Quarterly Journal of Economics, 93: 189–211.

[33] Romer, P. 1990. "Endogenous Technological Change." Journal of Political Economy, 98: S71-S102.

[34] Sheshinski, E., and Y. Weiss. 1977. "Inflation and Costs of Price Adjustment." Review of Economic Studies, 44: 287-303.

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

[36] Stigler, G. 1962. "Information in the Labor Market." Journal of Political Economy, 70: 94-105.

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