Revolutionary Effects of New Information Technologies

Gerard J. van den Berg

July 2005

Forschungsinstitut
zur Zukunft der Arbeit
Institute for the Study
of Labor

Revolutionary Effects of New Information Technologies

Gerard J. van den Berg
*Free University Amsterdam, IFAU Uppsala, CEPR,
Tinbergen Institute, INSEE-CREST and IZA Bonn*

ABSTRACT

Revolutionary Effects of New Information Technologies[*]

In markets with imperfect information and heterogeneity, the information technology affects
the rate at which agents meet, which in turn affects the distribution of production technologies
across firms. We show that in models for such markets there are typically multiple equilibria
because reservation utility levels and the lowest production technology in use affect each
other. The adoption of novel information technologies may then entail a revolution in the
sense of a move from an inefficient to an efficient equilibrium. Inefficient production
technologies are removed even in sectors where the new information technology has only
recently been introduced. The effect is much larger than a marginal comparative-statics effect
on a given equilibrium. The results apply to markets for consumer products, labor,
intermediate goods, and (public) institutional services.

JEL Classification: D43, D83, L11, L15, O33, J42, L86

Keywords: imperfect information, heterogeneity, production technology,
              informational frictions

Corresponding author:

Gerard J. van den Berg
Department of Economics
Free University of Amsterdam
De Boelelaan 1105
1081 HV Amsterdam
The Netherlands
Email: gjvdberg@xs4all.nl

1 Introduction

One may distinguish at least three major effects of the introduction of novel information and communication technologies (IT). First, it may have a direct positive effect on the technology of production at given firms because it improves the technology of information processing within the firm. Secondly, it may affect the search process by way of which agents in the market meet. Thirdly, it may introduce a new type of firms (and a new IT sector) in the economy.

This paper considers implications of the second effect for the composition or distribution of the different production technologies in use in a market. Many models of markets with imperfect information and heterogeneity give rise to multiple equilibria, and in that case the introduction of novel information technologies for agents looking for trading partners may cause a move from a bad to a good equilibrium. In the latter equilibrium, inefficient production technologies are not used.

To shape thoughts, consider a product market in which different firms have different production technologies. In equilibrium, the firms with an inferior production technology set higher prices than the firms with a superior production technology. Nevertheless, due to the presence of informational frictions, the former type of firms have a positive steady-state number of customers and a positive profits level. Thus, these low-productivity firms can coexist with other firms. However, any decrease in the amount of frictions lowers the reservation price of the customers, and this eventually makes the low-productivity firms go bankrupt. This is a comparative-statics composition effect that has been noted in the related theoretical literature on equilibrium search models and that can be related to statements in the early academic and popular literature on the effects of the Internet (Bakos, 1997, The Economist, 1999; see below for a more detailed discussion of the literature).

As we demonstrate, in such markets, multiplicity of equilibria basically means that a market has *either (1)* high-productivity as well as low-productivity firms, with customers using a high reservation price, *or (2)* high-productivity firms only, with customers using a low reservation price. In the second equilibrium, the low reservation price acts as a binding upper bound on the set of production technologies that enable a positive profit per customer. It rules out production at a low productivity level. Now, in general (i.e., in both equilibria), part of the rent of production is distributed to the customers. The resulting price distribution in the first equilibrium dominates the price distribution in the second equilibrium. This in turn justifies the higher reservation price in the first equilibrium.

Suppose that the market is in the first equilibrium, with on average higher prices and less efficient production than in the second equilibrium. Now suppose that a new information technology is introduced to this market. This may have as a result that the first equilibrium cannot exist anymore, so that the market moves to the second equilibrium,

and the low-productivity firms go bankrupt. All equilibrium outcomes make a jump; i.e., they are discontinuous in the amount of frictions. Such effects are fundamentally larger than marginal comparative-statics effects in a given equilibrium.

Now consider an economy where agents are simultaneously active in different markets. If the agents experience that the introduction of the new information technology establishes a change in the equilibrium at some markets, then they may also expect such a change in markets or sectors where the new information technology has only recently been introduced (we call these the more traditional markets). This change in beliefs may generate a shift from the inefficient equilibrium to the efficient equilibrium in those markets, so this effect is transmitted across markets.

The effects described in the previous two paragraphs may entail a revolution in the economy as a whole, meaning that quantitatively small changes in the economic fundamentals may generate enormously different outcomes, where part of this is due to self-fulfilling beliefs. In the paper we show the implications of the results for growth, stock prices, takeovers, and inflation. There is a peak in the growth rate around the moment of the introduction of the new information technology. Stock prices in more traditional sectors of the economy increase even when novel information technologies are barely used there yet. Takeovers and mergers of firms become more common during the high growth years, market concentration increases to a higher level, and inflation remains low.

The results apply to many different types of markets, including consumer product markets and labor markets, as well as to non-market situations where individual agents face heterogeneous institutions, like schools for children of the individual or hospitals or police stations in the area, or even local government. In all these cases individuals need choose among such heterogeneous institutions, and the surplus of the match between individual and institution depends on the efficiency of the institution. The analysis may be even more relevant for markets for intermediate products, where firms trade amongst each other.

One may wonder how important the above-mentioned mechanisms are for explaining the over-all effects of the recent introduction of digital information technologies. After all, the first and third effect mentioned in the first paragraph of the paper may also contribute to this over-all effect. We do not consider the introduction of novel sectors in the economy, notably IT-sectors (and, by implication, we do not consider the "IT bubble" of the late 1990s). Even then, some of the implications of the analysis also follow if the information technology does not lead to a change in frictions but to an over-all productivity increase, or if it leads to entry of a cohort of new firms with higher productivity than the others. We make two observations that support the relevance of our analysis. First, there is abundant empirical evidence that search frictions have decreased during and after the introduction of new digital IT. Kuhn and Skuterud (2000, 2004) and Stevenson (2003) provide evidence that the Internet has recently become a quantitatively important channel to find jobs in the

U.S.. Stevenson (2003) also shows that the Internet has become qualitatively important in this respect. More and more consumer and production inputs are organized and traded by way of the Internet. In particular, this applies to products in many markets that have not seen a recent major entry of new producers with superior production technology, and that have not been subject to recent dramatic productivity improvements (see e.g. Bakos, 2001, and a recent survey in The Economist, 2004). Secondly, the transmittance of equilibrium shifts across markets may be more likely in case of uncertainty about the magnitude of frictions (and the ensuing uncertainty about the distribution of production technology) than in case of uncertainty about just the distribution of production technology.

The theoretical and empirical literature on IT effects, while recognizing the effect on market frictions, has not considered the type of multiplicity of equilibrium that we address. A number of other literatures can also be related to the present paper. First, there is the empirical literature on job and worker flows and the empirical literature on the so-called "productivity puzzle". These show that there are persistent productivity differences across firms and establishments within industries. Within-firm productivity growth among existing firms is typically low. At the same time, there appears to be a lot of growth-enhancing reallocation across firms. Studies on productivity often use panel data of firms that exist throughout the observation period. By discarding observations on firms that enter or exit during the observation period, a large fraction of such reallocation effects is not picked up. It turns out that a low firm-specific productivity has a strong effect on the probability of market exit. Aggregate productivity growth seems to have been to a large extent driven by reallocation of outputs and inputs from less productive dying firms to more productive firms (see Foster, Haltiwanger and Krizan, 2001, for a detailed overview).

Another branch of related literature is the literature on industry "shakeouts". This resorts to production technology innovations to explain sudden dramatic increases in market concentration. These innovations either originate externally or they evolve from a competitive process between heterogeneous firms (see Klepper and Simons, 2001, for an overview). In the second variety, a shakeout is typically seen as a stage in the evolution of an industry following mass entry of firms, and as such this focuses on market dynamics along a different time scale that in the present paper. Note that the analysis in the present paper can contribute to an understanding of such shakeouts by emphasizing the effect of improvements in information technology on the functioning of the product market, that may occur if an industry ages. The first variety of shakeout theories may also cover the first IT effect mentioned in the first paragraph of the paper. However, the IT effects considered in the present paper can also contribute to the explanation of such shakeouts.

Yet another related literature concerns the theoretical analysis of equilibrium search models (see e.g. Mortensen and Pissarides, 1999, and Mortensen, 2003, for overviews). Below we will refer extensively to that literature.

To derive the main insights, we deliberately adopt the simplest possible model framework. For a given information technology, one may adopt a model in which prices are determined by bargaining or a model where they are determined by price setting. Similarly, one may consider sustained matches or continued search after having formed a match, and one may consider fixed arrival rates of trading opportunities or endogenous arrival rates depending on the numbers of active agents in the market. Throughout much of the paper we assume bargaining, sustained matches, and fixed arrival rates, but the results in the paper apply to all these approaches. It should be emphasized, however, that we always assume a sequential arrival of trading opportunities. Because of this, the analysis is not very suitable for markets that are primarily located on the Internet where the whole price distribution can be found within seconds and the desired price can be chosen by one click (for example using shopbots; see e.g. Ellison and Fisher Ellison, 2001; in such cases the market structure is completely different). Apart from that, the analysis is applicable not only to the recent introduction of digital information technologies, but also to older changes, like the introduction of the print press (15th century) and the increase in newspaper circulation (19th century). For expositional reasons we formulate the analysis of this paper in terms of markets for consumer products.

The outline of the paper is as follows. In Section 2 we present the basic model. In Section 3 we examine the introduction of a new information technology. We also extend the model in various ways. Amongst other things, we allow for endogenous arrival rates of trading opportunities depending on the numbers of actively searching agents in the market, and we consider price setting and continued search. We show that the main results are not affected. In Section 3 we also consider an implication for the empirical literature on the effect of novel IT on price dispersion of consumer goods. In this literature the hypothesis of interest is that prices converge. We argue that the measures of price dispersion used are often not informative on the effect of novel IT on market equilibrium outcomes and on the degree of convergence towards an efficient market. Section 4 concludes.

2   A simple model framework

2.1   The behavior of agents

2.1.1   Search

In this subsection we start with a description of the behavior of the participants in a market with a given distribution of production technologies across producers. Subsequently, we consider market entry of producers with production technologies that are superior to those in use. Throughout the whole Section 2 we take the technology of information processing as fixed and given.

Consider a product market consisting of a continuum of consumers and producers. The measure of consumers (or individuals, or customers) is denoted by m and the measure of producers (or firms) by n. From the point of view of the consumer, all products that are traded on the market are equal in terms of the instantaneous utility flow u that they provide. The products may however differ in their price p.

The supply side of the model is equivalent to a standard partial model of product search. New consumers enter the market without information on the location of firms or their production technology or product price. At this stage, the individual's instantaneous utility flow equals b (so $u - b$ is the "willingness to pay"). Information on the existence and location of firms arrives according to a Poisson process with rate λ. In particular, each arrival is characterized by a random drawing from the distribution of firms. Each firm has the same "probability" of being drawn. A contact between the individual and the firm results in a particular value of the product price p. The drawing of a firm therefore entails a random drawing from the distribution F of product prices p. Whenever this information arrives, the decision has to be made whether to start buying the product from this firm or whether to reject the firm and search further for a firm with a lower product price. In equilibrium, F is an endogenous outcome that depends on the behavior of the market participants.

The trading relationship (or match) between an individual and a firm is dissolved at the constant exogenous rate δ. Individuals maximize their expected discounted lifetime utility, where we assume intertemporal separability and where the discount rate ρ includes the rate of death.

The parameter λ represents the technology of information processing. The model framework adopts a number of *initial assumptions* that are made purely for convenience (to keep the exposition clear and to make the derivations manageable) but which can be dropped without affecting the key results. The first of these assumptions is that λ is exogenous at the individual level. It is actually useful to think of λ as being a function of the costs and benefits and therefore of the importance attached to the acquisition of information on this particular product. An individual may not want to spend a great deal of time and money to find the cheapest car dealer, bookseller or flower shop in town. The second of the initial assumptions is that λ does not depend on the number of market participants. The third is that individuals do not search for lower prices once they have started to buy the product from a given firm. In Subsection 3.5 we demonstrate the robustness of the key results with respect to these assumptions.

Concerning the model parameters, we assume that $0 < \lambda, \delta, \rho < \infty$ and $0 \leq b < u < \infty$. By analogy to partial product search models, it follows that the optimal strategy of the individual is to accept a product if and only if its price is smaller than a reservation price ϕ which satisfies

$$u - \phi = b + \frac{\lambda}{\rho + \delta} \int_0^\phi F(p) dp \tag{1}$$

(see Appendix 1 for a derivation). This equation has a unique solution $\phi \in (0, u - b]$ given the other variables and functions.[1] For future reference it is useful to note that ϕ is related to the expected present value of being a searching customer R, by way of $R = (u - \phi)/\rho$. This follows immediately from the fact that at a product price ϕ the customer is indifferent between purchase forever and further search.

2.1.2 Prices

Now let us turn to the equilibrium determination of F. First, there are no equilibrium prices above ϕ, because there would be no trade at such prices. Therefore, in equilibrium, F is necessarily such that $F(\phi) = 1$, so the support of F is in the interval $(0, \phi]$. Now consider a firm that is able to make the product at a cost price q. We assume that q does not depend on the size of the production at the firm. The value of q represents the production technology of the firm. The lower q, the more advanced the technology. At this stage, we take the distribution of q among active firms as given. Note that it is necessary that $0 < q < \phi$. Firms have fixed search costs, and they are able to serve as many customers as show up. Indeed, as long as the price exceeds q, they always engage in trade with additional customers.

Suppose that an individual discovers a firm with production technology q. We assume that the product price in that case is determined by a rent sharing condition (i.e., as the axiomatic Nash bargaining solution). Consider the surplus value $S(q)$ of the trading relationship (as compared to the return when no such relationship is formed after the discovery). This equals the expected discounted utility gain from production minus the value of further search by the individual, over a period of which the length has an exponential distribution with parameter δ. The utility gain flow equals $u - q$ and the flow value of continued search equals ρR, which equals $u - \phi$. As a result,

$$S(q) = \frac{u - q - \rho R}{\rho + \delta} = \frac{\phi - q}{\rho + \delta}$$

The firm receives a fraction $1 - \beta$ of this, where $\beta \in (0, 1)$ is a parameter that represents the relative bargaining power of the customer vis-à-vis the firm. This rent sharing condition $(p - q)/(\rho + \delta) = (1 - \beta)S(q)$ determines the price,

$$p = \beta q + (1 - \beta)\phi \tag{2}$$

---

[1] The derivative of the left-hand side with respect to ϕ is -1, whereas the derivative of the r.h.s. w.r.t. ϕ is nonnegative. At $\phi := 0$ (at $\phi := u - b$) the l.h.s. is larger (smaller) than the r.h.s..

Since individuals share the same value of ϕ, equation (2) implies that p can be thought of as a firm characteristic.

Of course, price determination can be modeled in many different ways. Note that we deal with a market with imperfect competition where all firms and individuals have some market power. If they do not agree to trade after having discovered each other, then they lose the rent (or surplus value) of their match. The present approach is common in the literature (see e.g. Acemoglu, 2001, and Albrecht and Vroman, 2002, for applications in labor markets with heterogeneous firms and homogeneous workers). The above bargaining solution can also be generated as the equilibrium outcome of a decentralized strategic bargaining game (see Bester, 1988, who examines a similar model in which bargaining takes time and is explicitly modeled; the outcome is as above except for the fact that there is some residual randomness in the price depending on who makes the first proposal).

An alternative approach would be to assume that firms set prices before they meet customers. In such models, typically, the optimal price and the profit flow of a firm with technology q are monotone increasing and decreasing functions of q, respectively, like in the model based on equation (2). In fact, this often requires some additional heterogeneity in the model; for example customers differ in b, or customers of other firms search further for lower prices, or customers may buy more than a single product at a firm (Reinganum, 1979, Rob, 1985, Diamond, 1987, Burdett, 1990). However, what is important is that in such models, as in the present model, the division of the rent of the match reflects the relative market power of individuals vis-à-vis firms. We return to this in Subsection 3.5. It should be noted that the price determination rule in equation (2) has been used in the literature as a simple and robust shortcut to capture the true underlying mechanism (see e.g. Mortensen, 1999). In terms of the signs of the relations between q, p, ϕ and profits, it agrees to a wide range of models.

Now let us turn to the implications of the modeling of F. First, the price p is always in the interval $(0, \phi)$, so it is always acceptable to the individual. This means that the individual decides to buy at the first firm he discovers. One may argue that this reduces the model to a trivial search model and that ϕ is irrelevant. However, this is not correct, because the value of ϕ determines the set of production technologies q that can be profitably employed. It demarcates the profitable firms with $q < \phi$ from the unprofitable (and therefore non-existing) firms with $q \geq \phi$. The reservation price drives out inefficient firms with high prices, and therefore there is no need to reject a price and search further for even lower prices. Once we examine the effects of changes in the model determinants, this constraint has to be taken into account. In any case, in Subsection 3.5, we also discuss extensions of this basic model in which individuals do actually reject product prices from time to time (e.g. because they are heterogeneous in b or because they search for even lower prices while buying at a certain price).

We denote the distribution of q among active firms by Γ, so $\Gamma(0) = 0$ and $\Gamma(\phi) = 1$. There holds that $F(p) = \Gamma(q)$, with p and q satisfying equation (2). As a result,

$$F(x) = \Gamma(\frac{x - (1 - \beta)\phi}{\beta}) \tag{3}$$

and the support of F can be narrowed down to $((1 - \beta)\phi, \phi)$. By substituting (3) into (1) we obtain

$$u - \phi = b + \frac{\lambda\beta}{\rho + \delta} \int_0^\phi \Gamma(q)dq \tag{4}$$

2.1.3 Entry and exit of producers

As noted above, the value of ϕ determines the set of profitable q. This means that Γ is not a structural determinant. In order to obtain some foundation of it, we consider market entry and exit of producers. This is regarded as an extremely-long-run process, in the sense that individual customers and active firms do not take this into account in their strategy.[2]

At the beginning, the market was created by firms using the best available technology at the time, say q_0, with $q_0 < u - b$. The investments that lead to a given production technology are sunk costs. After the beginning, production technologies were designed that were superior to those in use at that time. At the same time, a stock of potential new entrants was built up. We assume that only new entrants can adopt the new technologies. A potential new entrant enters the market with the most advanced technology at the time iff the expected discounted profit flow exceeds the cost of entry plus a certain minimum amount of expected profit. The magnitude of the cost of entry and the number of potential new entrants, as well as the way they change over time, are determined outside the model. What is important is that this process gives rise to a sequence of production technologies that are used in production,

$$q_0, q_1, q_2, q_3, \ldots \qquad \text{with} \quad u - b > q_0 > q_1 > q_2 > \cdots > 0.$$

with associated measures of firms $n_0, n_1, n_2, \ldots$. At the beginning, Γ is a degenerate distribution at q_0. Later, mass points $q_1, q_2, \ldots$ are added to Γ. This causes the reservation price ϕ to decrease, such that it may fall below the level of the oldest technology in use at the time. In that case the firms using that technology are driven out of the market.

The result of this process is that at any given point of time the distribution Γ is a discrete distribution with a countable number of mass points at adjacent values of q, say

---

[2]It is not difficult to relax this if we specify a Poisson process for the arrival of cohorts of new entrants. For expositional clarity we do not pursue this.

$q_i, q_{i+1}, \ldots, q_j$, with $0 \le i \le j$ and $i, j \in \mathbb{R}$. From the point of view of the customers, the distribution of Γ is then such a discrete distribution.[3,4]

The above simple model is in agreement with a number of stylized facts from establishment and firm level data. Foster, Haltiwanger and Krizan (2001) show that there are persistent productivity differences across establishments within industries, and that a large part of national GDP growth is due to within-sector reallocation, notably entry and exit of individual establishments. A low establishment-specific productivity is a good predictor of market exit. In the service sectors that are considered, basically all sector-level productivity growth comes from entry (of high-productivity establishments) and exit (of low-productivity establishments).

It is useful to define the distribution Γ_0 of q across all active and latently existing firms at a given point of time. At that point of time, this is the distribution of all production technologies across firms that are at most as efficient as the most efficient technology then in use. Clearly,

$$\Gamma(q) = \min\{\frac{\Gamma_0(q)}{\Gamma_0(\phi)}, 1\} \tag{5}$$

This requires $\Gamma_0(\phi) > 0$, which follows from optimality of ϕ.

2.2   Equilibria

We examine equilibria for given Γ_0. We are concerned with non-cooperative steady-state equilibria (see e.g. Bontemps, Robin and Van den Berg, 2000),

Definition. In a given environment $(\Gamma_0, n, m, \lambda, \beta, \delta, \rho, u, b)$, a market equilibrium is characterized by a reservation price ϕ, a distribution F of product prices faced by searching customers, and a distribution Γ of production technologies q among active firms Γ, such that

(1) ϕ is the customers' optimal behavior given the active firms' aggregate behavior, i.e. it is determined by equation (1), given F.

(2) F summarizes the price determination mechanism, i.e. it is determined by equation (3), given Γ and ϕ.

(3) An active firm's profits are positive, i.e., Γ is determined by equation (5) given ϕ.

---

[3]The distribution of q can also be rationalized along the lines of Acemoglu and Shimer (2000) and Robin and Roux (2002), by assuming that ex ante identical firms choose their capital endowments simultaneously before production starts.

[4]It should be stressed that the results in the paper are not critically dependent on the discrete nature of the distribution of production technologies. In the next subsection we give an example with a continuous distribution.

Note that this definition applies to discrete as well as continuous Γ_0, To make the basic point it is sufficient to restrict attention to a market in which there are two possible production technologies. This entails no loss of generality. Denote the two levels of q by q_0 and q_1, with $q_0 > q_1$, and denote the corresponding fractions of firms by $1 - \gamma$ and γ, respectively, with $0 < \gamma < 1$. We examine equilibria for given values of q_0, q_1 and γ.

First, we examine whether there is an equilibrium in which only type-q_1 firms are active. This is the desirable equilibrium because all products are produced in the most efficient manner. We start by *assuming* that only type-q_1 firm types are active in equilibrium, and then we verify this by checking whether in equilibrium $\phi \leq q_0$. If the latter does not hold then there is an incentive for type-q_0 firms to enter the market.

If only type-q_1 firms are active then we have a model with homogeneous firms. From equation (4),

$$\phi = \frac{(\rho + \delta)(u - b) + \lambda\beta q_1}{\rho + \delta + \lambda\beta} \tag{6}$$

which is a weighted average of $u - b$ and q_1. There is only one equilibrium price p_1, which can be expressed in the parameters by way of $p_1 = \beta q_1 + (1 - \beta)\phi$. The reservation price satisfies $\phi \leq q_0$ if and only if

$$\frac{u - b - q_0}{q_0 - q_1} \leq \frac{\lambda\beta}{\rho + \delta} \tag{7}$$

Note that this condition does not depend on γ, n or m.

Now suppose that firms of both types are active in equilibrium. This must again be verified, by checking whether in equilibrium $\phi > q_0$. If the latter does not hold then type-q_0 firms will disappear from the market. From equation (4),

$$\phi = \frac{(u - b)(\rho + \delta) + \lambda\beta(\gamma q_1 + (1 - \gamma)q_0)}{\rho + \delta + \lambda\beta} \tag{8}$$

which is a weighted average of $u - b, q_0$ and q_1. This expression is similar to expression (6), the only difference being that the productivity level in the homogeneous case is replaced by the mean productivity in the present case. There are equilibrium prices p_0, p_1, which can be expressed in the parameters by way of $p_i = \beta q_i + (1 - \beta)\phi$. The reservation price satisfies $\phi > q_0$ if and only if

$$\frac{u - b - q_0}{q_0 - q_1} > \frac{\gamma\lambda\beta}{\rho + \delta} \tag{9}$$

As a result, both equilibria are supported by exactly the same parameter values if and only if

$$\frac{\lambda\beta}{\rho+\delta} \geq \frac{u-b-q_0}{q_0-q_1} > \gamma \cdot \frac{\lambda\beta}{\rho+\delta} \tag{10}$$

Note that the left-hand term is larger than the right hand term (due to $0 < \gamma < 1$), so that there always exists at least one equilibrium. Moreover, the left- and right-hand sides only depend on rates and fractions, whereas the middle term only depends on monetary flow variables (all three terms are dimensionless). Thus, the model can have multiple equilibria. In one equilibrium, all active firms are type-q_1 firms, whereas in the other, both types are active. Note that there are always values of $u-b, q_1, q_0$ such that there are multiple equilibria.[5]

The multiplicity is an implication of the interaction between the reservation price and the distribution of production technologies that allow for profitable production. The reservation price ϕ affects the distribution of q among active firms by way of the restriction that these q should be smaller than ϕ. Thus, ϕ affects the *upper bound of the support* of the distribution of q among active firms. Conversely, the distribution of q among active firms affects the reservation price, because part of the rent of production is distributed to the customers. This effect works by way of the *mean* production technology (by way of the mean price).[6]

Multiplicity of equilibrium reflects a coordination failure. In the "bad" equilibrium, production is inefficient, but there is no coordinating market mechanism to move the market to the "good" equilibrium.

The next example illustrates that multiplicity can be supported by reasonable parameter values.

Example 1. To fix thoughts, let a month and a U.S. dollar be the units of measurement. Meetings between customers and firms occur at the rate of one per month ($\lambda = 1$), and matches are dissolved at a rate of once every five months ($\delta = 0.2$). The distribution of production technologies is such that half of the firms has $q_0 = 500$ while the other half has $q_1 = 400$ (so $\gamma = 0.5$). The utility gain $u-b$ of the customers equals 650. Customers and firms have the same bargaining power ($\beta = 0.5$), and we take a 12% yearly discount rate ($\rho = 0.01$). The three terms in (10) have values 2.38, 1.5, and 1.19, respectively. In one

---

[5]There is a third type of equilibrium in which only a fraction of the type-q_0 firms is active and $\phi = q_0$. This hybrid equilibrium exists if and only if both other equilibria exist. One can argue that it is less stable than the other two, and in any case it does not play a role in the analysis.

[6]Bester (1988) describes such multiplicity in a similar model. That model includes strategic bargaining behavior, while frictions are represented by a fixed waiting time between offer arrivals. There are no production costs, but product quality may be heterogeneous across firms. This can be translated into our model with heterogeneous positive production costs and identical quality, and as such Bester (1988)'s discussion and example of multiplicity antedate the results in this subsection.

equilibrium $\phi = 509$, both types of firms are active, and the prices are 455 and 505. In the other equilibrium, $\phi = 474$, only type-q_1 firms are active, and $p = 437$.

Suppose that the parameter values $u, b, q_0, q_1, \lambda, \beta, \delta, \rho$ are such that both equilibria are possible. Consider the moment at which the type-q_1 firms enter the market at which so far only type-q_0 firms have been active. The resulting equilibrium is the "bad" equilibrium. The type-q_0 firms still make a profit because of the customers with whom they formed a trade relationship, so there is no reason for them to leave the market. This means that the new customers face a market with both types of firms, which will induce them to set a reservation price higher than q_0. More in general, if multiple equilibria are possible, then the process of entry of firms is such that the resulting equilibrium will be the worst of the possible equilibria.

Consider the bad equilibrium if there are two possible types of firms and two possible equilibria. The expected flow of profit π_0^B out of the current production of a type-q_0 firm equals

$$\pi_0^B = \frac{m\lambda}{n(\lambda + \delta)}(1 - \beta)(\phi^B - q_0)$$

where we ignore the repayment of the loan used to finance the sunk investment costs, and where ϕ^B is the reservation price in the bad equilibrium. The ratio at the right-hand side equals the average number of customers per firm, which is the average number of individuals per firm times the fraction of individuals who are currently in a trading relationship. Similarly, we define

$$\pi_1^B = \frac{m\lambda}{n(\lambda + \delta)}(1 - \beta)(\phi^B - q_1)$$

$$\pi_1^G = \frac{m\lambda}{\gamma n(\lambda + \delta)}(1 - \beta)(\phi^G - q_1)$$

After substitution of (8) and (6) for ϕ^B and ϕ^G, respectively, it follows that

$$\pi_0^B < \pi_1^B < \pi_1^G$$

so the type-q_1 firms are better off in the good equilibrium. The price of their products decreases but this is more than offset by the increase in the size of their clientèle. Note that the total profits of all firms that are active in equilibrium are larger in the good equilibrium. Finally, it is trivial to see that each type-q_0 firm is better off before the entry of the type-q_1 firms.

We finish this section by emphasizing that the multiplicity of equilibria and the finiteness of the number of equilibria do not depend on the assumption that the number of firm

types is finite. If q has a continuous distribution then it is also possible to have multiple equilibria, and, in particular, to have a small number of equilibria. We demonstrate this by way of a numerical example. (Appendix 2 contains background details as well as a proof that there always exists an equilibrium regardless of whether q is discrete or continuous.) Note that a continuous distribution of q cannot be generated by the firm entry process described in the previous subsection.

Example 2. Let Γ_0 have a density equal to $c_0 q \exp(c_1 q^2)$ on the support $(0, \overline{q})$. This distribution has parameters $c_1 > 0$ and $\overline{q} \in (0, u - b)$, while c_0 follows from the constraint that the density integrates to one. The density is upward sloping, so most probability mass is at the oldest production technologies. We take $c_1 = 0.75$ and $\overline{q} = 2.5$, so that the support is $(0, 2.5)$ and the three quartiles equal 2.11, 2.31, and 2.42, respectively. Figure 1 displays this density. Furthermore, $u - b = 8.5$ and $\lambda \beta / (\rho + \delta) = 20$ (e.g. because $\beta = 0.5, \rho = 0.01, \delta = 0.1$ and $\lambda = 4.4$).

There are exactly three equilibria. In the first one, $\phi = 2.51$, so that all possible production technologies are in use. In the second one, $\phi = 2.35$, and only the 57% most productive firms are active. In the third one, $\phi = 1.56$, and only the 5% most productive firms are active.

Intuitively, what is important for multiplicity is not that q is discrete or continuous, but that the reservation price depends on an *expectation* of q among active firms. For the customers it is not relevant whether this expectation is generated by a discrete or by a continuous distribution. The equilibrium reservation price is determined as the solution of an equation in which the upper bound of the support of the (truncated) distribution of q among active firms is related to the expectation of this truncated distribution. In general this can have any number of solutions. If the upper bound of the support of the (truncated) distribution of q increases then the expectation of this truncated distribution can increase more than proportionally.

2.3   Arrival rates depending on the measures of customers and active firms

It is plausible that the arrival rate of trading opportunities for customers λ depends on the measure of active firms $\Gamma_0(\phi)n$, which is γn in the good equilibrium and n in the bad equilibrium. Whether $\lambda(\gamma n) \gtrless \lambda(n)$ depends on the matching technology in the market and the search behavior of customers and firms. In the good equilibrium, on the one hand, there are less firms to be sampled by a customer. On the other hand, the firms are larger so they may be easier to locate by searching customers. If these effects counterbalance then

λ is independent of the measure of active firms, and the results of the previous subsections follow.

In this subsection we adopt a more general approach. From the derivation of equation (10) it follows that there are parameter values for which both equilibria exist if and only if $\lambda(\gamma n) > \gamma \lambda(n)$. At first sight it seems that this is always satisfied if the meeting rate displays decreasing returns to scale in the measure of active firms, which is a fairly innocuous assumption: for example, it follows if the matching rate displays constant returns to scale in the measure of customers and the measure of active firms taken together.

However, this ignores that the fraction of customers who are searching now typically also depends on the fraction of active firms. To demonstrate this, we assume that the flow of contacts between searching customers and active firms is a CRS Cobb-Douglas function M of the measures of searching customers m_a and active firms n_a,

$$M(m_a, n_a) = \eta \cdot (m_a)^{1-\omega} n_a^{\omega} \tag{11}$$

where of course $n_a = \Gamma_0(\phi)n$ and $m_a = m\delta/(\lambda + \delta)$. In line with the literature, M is called the matching function. For ease of exposition we rule out that $\omega = 1$, so $0 \leq \omega < 1$. Note that $\omega = 0$ produces the model of the previous section. Now, in general, $\lambda = M(m_a, n_a)/m_a = \eta(n_a/m_a)^{\omega}$. We thus have two equations in λ and m_a, and we need to solve them to study how λ depends on n_a, and in particular whether $\lambda(\gamma n) > \gamma \lambda(n)$. It is easy to show that the systems of equations for $n_a := n$ and for $n_a := \gamma n$ lead to

$$\gamma = \left[\frac{\lambda(\gamma n)}{\lambda(n)} \right]^{\frac{1}{\omega}} \frac{\delta + \lambda(n)}{\delta + \lambda(\gamma n)}$$

This equation is formally equal to a special case of equation (A.2) in the Appendix of Van den Berg (2003). With minor modifications, we can apply the derivations there to obtain the following result: if

$$\omega \leq \frac{1}{2}$$

then one can always find values of the other model parameters such that multiple equilibria exist. This condition is by no means necessary. For $\omega > 1/2$ it is also often possible to construct multiple equilibria. A necessary and sufficient condition is

$$\omega < \frac{1}{1 - \frac{1}{\log \gamma} \cdot \log \left(\frac{\delta + \lambda(n)}{\delta + \gamma \lambda(n)} \right)}$$

but this is not a condition in terms of model primitives, since it depends on $\lambda(n)$. Note that the right-hand side is always between 0 and 1.

We conclude that the results of the previous subsection follow for a wide range of values of the matching function parameter. With $\omega = 1$ there is no multiplicity, illustrating that in equilibrium search and matching models, multiplicity is due to a coordination failure.

3 The introduction of new information and communication technologies

3.1 The revolution

New information and communication technologies are modeled by way of an increase in λ. Suppose that the parameter values before the introduction of the new information technology are such that both equilibria are possible, that is, (10) holds, and suppose that the actual equilibrium is the bad equilibrium. At some point, an infinitesimally small increase of λ implies that the bad equilibrium is not sustainable anymore, and the only possible equilibrium is the good equilibrium. In particular, we need λ to satisfy

$$\lambda \geq \frac{(\rho + \delta)(u - b - q_0)}{\gamma\beta(q_0 - q_1)}$$

If λ crosses the right-hand side threshold value, the equilibrium outcomes jump from the bad outcomes to the good outcomes. In particular, the reservation price ϕ jumps from marginally above q_0 to the value

$$\phi = q_0 - \frac{(1 - \gamma)(q_0 - q_1)(u - b - q_0)}{u - b - (1 - \gamma)q_0 - \gamma q_1} \tag{12}$$

which is strictly smaller than q_0.[7] This is a discontinuous jump. As a result, the new distribution of production technologies is such that only the efficient technology q_1 is used. The mean productivity therefore makes an upward jump as well, as does social welfare. The mean price makes a downward jump. The discontinuities are what distinguish these effects from ordinary comparative statics effects of parameter changes on equilibrium outcomes. The latter effects are marginal in the sense that infinitesimally small changes have infinitesimally small effects. Here, infinitesimally small changes of an economic fundamental may for some time have infinitesimally small effects, until a point at which suddenly the outcome variables jump discretely. This can be labeled a "revolution".

Again, the discontinuity is not due to the assumption that the distribution Γ of q is discrete. We illustrate this in Figure 2, which displays the possible equilibrium reservation prices as a function of λ, for Example 2 above. The horizontal axis denotes k, which is

---

[7]To avoid confusion, note that after the jump ϕ is functionally independent of q_0, but equation (12) evaluates it at a value of λ that is made to depend on q_0.

defined as $\lambda\beta/(\rho + \delta)$. As k increases up to 20.4, ϕ decreases until it reaches the value 2.5. Then, as k increases marginally more, the equilibrium shifts, and ϕ jumps down to 1.46. Figure 4 displays similar graphs for the example where Γ_0 is a mixture of two normal distributions both truncated at 0 and $u - b$. In the latter example, the other parameter values are taken as in Example 2. Figure 3 plots the density corresponding to Γ_0 used for Figure 4.

Formally, the effects of λ on ϕ and Γ are qualitatively the same as the effects of β on these outcomes and qualitatively opposite to the effects of δ and ρ on these outcomes. This is because these parameters affect these outcomes by way of $\lambda\beta/(\delta + \rho)$ only. As an example, if the firms' market power strongly increases if frictions decrease due to new IT, so if β decreases in λ, then the discrete jump in outcomes may be replaced by a smoother transition.

The good equilibrium was in principle also attainable before the introduction of the new information technology, but due to a lack of coordination it was not attained.[8] Now suppose that new information technologies are successfully introduced in a number of sectors of the economy. Agents may be simultaneously active in different sectors, or they may be able to observe fundamentals and outcomes in other sectors. If the agents experience that the introduction of the new information technology establishes a change in the equilibrium in some sectors, then they may also expect such a change in sectors where the new information technology has only recently been introduced. As a result, they may start to behave as if the new information technologies are already implemented in the latter sectors. This may induce a shift in the equilibrium outcomes.[9] The change in beliefs may generate a shift from the inefficient equilibrium to the efficient equilibrium in those markets, in a situation where both equilibria are attainable. Such reasoning is in line with the approach proposed by Morris and Shin (2000) to single out an equilibrium in the case that multiplicity is possible. Morris and Shin (2000) allow for uncertainty about the fundamentals in a market, and show that public information (such as media coverage of the use of new information technologies) may have a large impact on the actual outcome. A formal analysis would be beyond the scope of the paper.

Taken together, these effects may entail a revolution in the economy as a whole, meaning that quantitatively small changes in economic fundamentals generate enormously different outcomes, where part of this is due to self-fulfilling beliefs.

---

[8]We do not provide a formal analysis of the issues of equilibrium selection and the out-of-equilibrium path. These are notoriously difficult to analyze. It is plausible that very small increases in λ that are not foreseen do not induce an equilibrium switch unless the current type of equilibrium is not sustainable anymore.

[9]This resembles the mechanism behind the successive revolutions in different countries in Eastern Europe in 1989.

3.2 Effects

Upon the introduction of novel information technology, the growth rate displays an upward jump. The average price goes down, so the exit of type-q_0 firms increases the customers' purchasing power, even though at the same time the profits of type-q_1 firms increase.[10] Concerning the recent introduction of digital information technologies, this may explain why there has not been much inflation during the prolonged economic upturn of the 1990s. Inflationary pressures are counterbalanced by the price decreases due to the equilibrium shift. The model prediction on average prices is also consistent with empirical findings that show that prices of items that are often purchased over the Internet are lower (see e.g. Brown and Goolsbee, 2002, and Clay, Krishnan and Wolff, 2001), that the average price of an item is lower on the Internet than in conventional outlets (see e.g. Brynjolfsson and Smith, 2000), and that the price elasticity of items traded on the Internet is very high (Ellison and Fisher Ellison, 2001).

When relating a firm's share price to its profit flow, the above model also explains rising share prices among surviving firms. Moreover, in the so-called traditional sectors, where at first sight the market fundamentals have not changed, share prices may also rise, in particular among the more productive firms in those sectors. This explains the apparent paradox that share prices of large firms in traditional sectors have increased during the late 1990s whereas productivity at those firms has not increased with a similar order of magnitude, and novel information technologies had not been used yet to a large extent to facilitate the contact with potential customers.

The model is amenable to including the possibility of takeovers of less productive firms by more productive firms. To see this, note that the profit rate of high-productivity firms suddenly jumps upward. This creates funding for the takeover of less productive firms. Moreover, the less productive firms have become cheap. Subsequently, the more efficient production technology may be injected into the less productive firms. This enhances the efficiency gains.

3.3 A note on the effect on price variation

Recently, an empirical literature has emerged on the effect of the adoption of new information and communication technologies on price dispersion. The hypothesis of interest is that prices converge over time. Convergence is often operationalized as a decrease of the standard deviation and/or interquartile range of the distribution of observed prices. The underlying limiting argument is that these measures are zero if frictions are absent. The empirical studies examine price dispersion of Internet-traded items over time (see e.g.

---

[10]Recall the results in the previous section on profits in different equilibria.

Brown and Goolsbee, 2002, Baye, Morgan and Scholten, 2004b, Clay, Krishnan and Wolff, 2001, Baylis and Perloff, 2002, and Baye, Morgan and Scholten, 2004a) or the difference between the price dispersion of an item on the Internet and the price dispersion of the same item in conventional outlets (see e.g. Brynjolfsson and Smith, 2000).

Let us consider this in the context of our model, where again q may be discrete and/or continuous. We are interested in the effect of λ on the standard deviation of the distribution F of p. From equation (2) it follows that $\text{var}(p) = \beta^2 \text{var}(q)$, so it suffices to examine the effect of λ on the standard deviation of Γ. This effect solely works by way of the upper bound ϕ of the support of Γ. In words, price variation changes because the reservation price decreases.

From results in e.g. Karlin (1982) and Burdett and Muus (1989) it follows that the standard deviation $\sqrt{\text{var}}(x|x < a)$ of a random variable x with a distribution that does not depend on a does not necessarily increase with a. In other words, truncation of a distribution may actually result in a larger standard deviation.

We illustrate this with the model of Example 2, where q and p are continuous.[11] Consider the distribution $\Gamma_0(q)$ and the associated price distribution, both truncated from above at a certain level. Figure 5 plots the standard deviation of the truncated price distribution, as a function of the truncation point. Obviously, this function is not monotonous. In Example 2 we showed that a certain set of parameter values is compatible both with $\phi = 2.51$ and with $\phi = 2.35$. Suppose that an external event forces ϕ to move from 2.51 to 2.35. This amounts to a truncation of the distributions of q and p at 2.35. However, the standard deviation of p *increases* from 0.1571 to 0.1626.

Similar results hold for the interquartile range.[12] The interquartile range of p equals β times the interquartile range of q, so again the effect of λ works by way of the upper bound ϕ of the support of Γ. But truncation of a distribution may actually result in a larger interquartile range. Consider again Example 2. Figure 6 plots the first and third quartile of the truncated price distribution, as a function of the truncation point. The difference increases on the interval considered. For $\phi = 2.51$ the range equals 0.316, whereas for $\phi = 2.35$ it equals 0.336.

We conclude that the standard deviation and interquartile range are not informative measures of the extent of price convergence. As a consequence, the model is not at odds with the empirical literature that does not find decreases in price variation in response to increased use of the Internet (see the above references; interestingly, Brown and Goolsbee, 2002, find that the price standard deviation first increases and then decreases upon the adoption of Internet search sites).

---

[11]Of course, with a discrete distribution with two mass points, the standard deviation never increases with truncation.

[12]The analysis can also be applied to measures of market concentration like the Herfindahl index.

3.4 Long run effects

The long run growth rate of a market is completely determined by the rate at which new production technologies are developed and the rate at which new cohorts of firms enter the market. In the model, these rates are not affected by the information technology used in the market. The information technology affects the heterogeneity of production technologies in use. A primitive information technology generates a market structure in which firm cohorts with inefficient production technologies remain active for a long time, so that the average production technology lags behind the newest production technology. However, there is no reason for why the growth rate of this average is lower than the growth rate of the newest production technologies in use.

As a result, the model does not predict that the long run growth rate after the introduction of new information and communication technologies is higher than the long run growth rate before it. Indeed, under the assumption that the rates at which new production technologies are developed and new cohorts of firms enter the market are constant over time, the long run growth rate after the introduction of new information technologies is the same as before the introduction. The introduction generates a short run effect on welfare growth, by way of a jump in production.

There are some lasting effects on firm dynamics. First, the average time in which cohorts of firms with a particular production technology (and therefore single firms) remain active decreases once and for all to a lower level. The new reservation price will always closely follow the lowest cost price in the market, so that it takes less time before a given production technology becomes unprofitable. So the average firm lifetime decreases once and for all to a lower level.

Secondly, if the cohorts have the same size then the number of active firms at a given point of time is lower after the introduction of new information and communication technologies. This effect may be further enhanced by takeovers. So there is a single upward jump in the market concentration. The firms serve more customers than before, so they are larger. Note that this prediction depends crucially on the absence of free entry.

3.5 More general models

The results are robust with respect to many model assumptions. First, new information technologies may also enhance the productivity of existing firms, and indeed may lead to the creation of a newer type of efficient firms. This may be particularly relevant in specific sectors in which firms' operating costs depend heavily on computer networks. The effects strengthen the results of the previous subsections.

As noted earlier in the paper, the results in the paper can also be generated in models where customers who are engaged in a match search further for better prices, and where

firms post prices before they meet potential customers. Suppose that customers, after having found an acceptable trading partner, can continue to search for a better match. This creates an incentive for firms to set prices below the willingness to pay of customers who have not found a trading partner yet. A lower price implies that the exit rate of customers at the firm is smaller, and that a larger fraction of all customers is willing to prefer the firm over their current trading partner. Hence, in terms of total profits of a firm, there is a trade-off between the profit per customer and the number of customers at the firm. It turns out that in equilibrium, firms with a better production technology (i.e. lower q) offer a lower price, have more customers, and have a higher total profit flow. The gap between the price and q reflects the market power of the firms. If customers can move easily between firms then firms have little market power, and this gap is large.

By analogy to the analysis in the present paper (see also Van den Berg, 2003) it can be shown that such models give rise to multiple equilibria, provided that the arrival rate λ is larger when the customer is without a trading partner than otherwise. The underlying intuition is the same as in Subsection 2.2. The reservation price ϕ of customers without a trading partner affects the distribution of q among active firms by way of the restriction that these q should be smaller than ϕ. Conversely, the location of the distribution of q among active firms affects ϕ, because part of the rent of production is distributed to the customers. If customers have a larger λ when without a trading partner then ϕ decreases if the price distribution becomes more favorable for the customers. An increase in a multiplicative factor in λ can again induce an equilibrium shift.

Instead of allowing customers to search while engaged in a match, one can also let them be different in other ways ensuring that at any point of time there is an incentive for at least part of the firms to offer prices below the monopoly price. For example, customers may differ in b. We conjecture that multiple equilibria are also possible if customers are allowed to purchase any number of products upon meeting a firm, although a formal analysis seems intractable.

There may also be forces that oppose the mechanisms discussed in this paper. Firms may exploit the decrease in search costs, e.g. by collecting information on customers, by product differentiation, and by using obfuscation strategies. Also, maintenance of information technology facilities may be costly. Daripa and Kapur (2001) provide an extensive overview of these topics.

4 Conclusion

This paper focuses on the effect of novel information technologies on the distribution of production technologies across firms, where the causality runs from IT to the amount of search frictions in the market to the set of profitable production technologies. The results

demonstrate a shift from a bad to a good equilibrium with more efficient production. This effect is much larger than a conventional marginal effect: a quantitatively small change in an economic fundamental generates a large effect on outcomes. Secondly, markets in which the new information technologies have not yet (or only recently) been introduced may experience a similar shift. The analysis is relevant for many different types of markets, including consumer product markets and labor markets, as well as for non-market situations where individual agents face heterogeneous institutions, like schools for children of the individual or hospitals or police stations in the area, or even local government. As a side issue, we have demonstrated that despite their popularity the standard deviation and interquartile range are not informative measures of the extent of price convergence.

A topic for further research would be to formalize the out-of-equilibrium dynamics arising upon the introduction of a novel information technology. This may facilitate the formal analysis of the transmission of equilibrium shifts across markets.

Another topic for further research would be to empirically assess the importance of the different ways in which information technology affects market outcomes. Different mechanisms may have different implications for the way in which the distribution of productivity across firms changes over time. In particular, the framework of the present paper predicts that changes operate by way of the lower bound of the productivity distribution (i.e. at the highest production cost q), whereas explanations that focus on the adoption of IT in the production process predict that the upper bound of this distribution changes in tune.
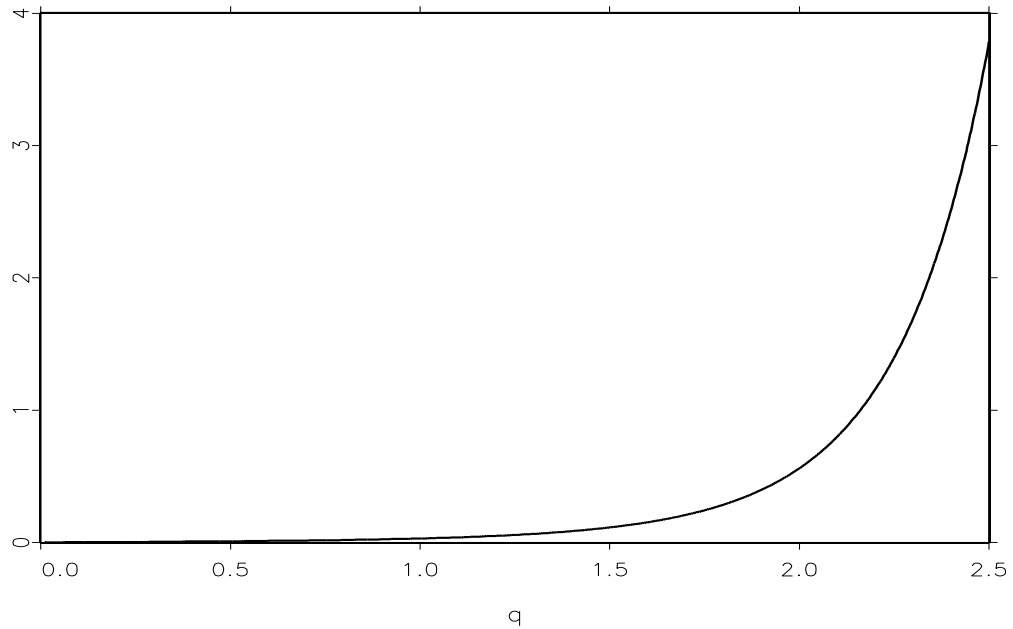
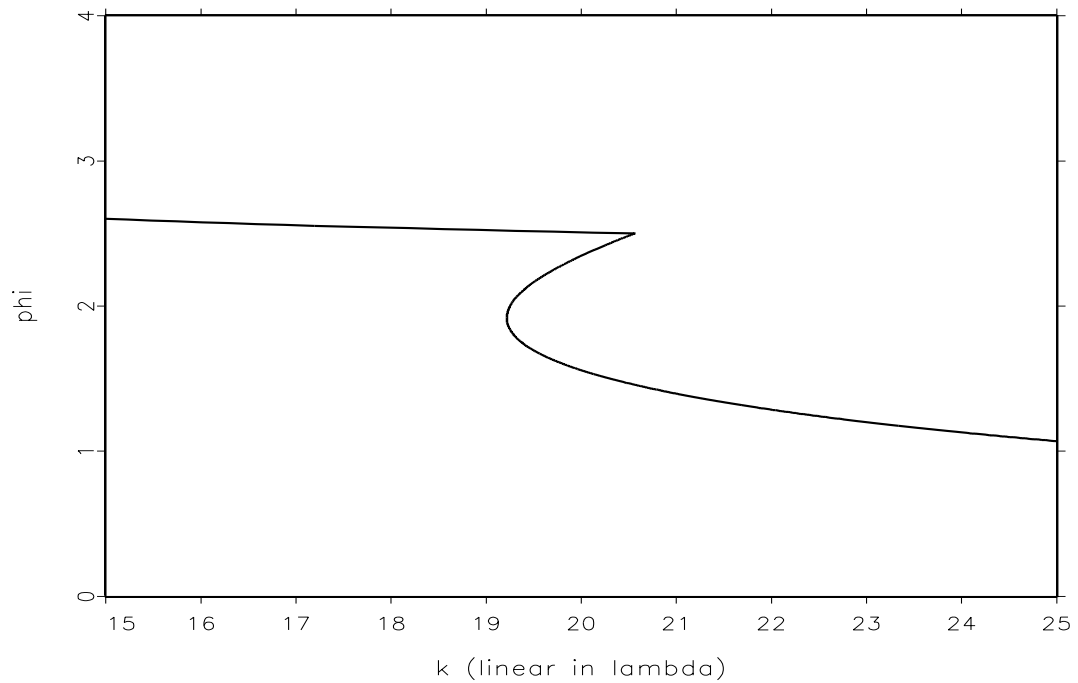Figure 1: Density of $\Gamma_0(q)$ in Example 2. Note: support is $(0, 2.5)$.

Figure 2: Equilibrium reservation prices in Example 2. Note: support of $\Gamma_0(q)$ is $(0, 2.5)$, while $u - b = 8.5$.
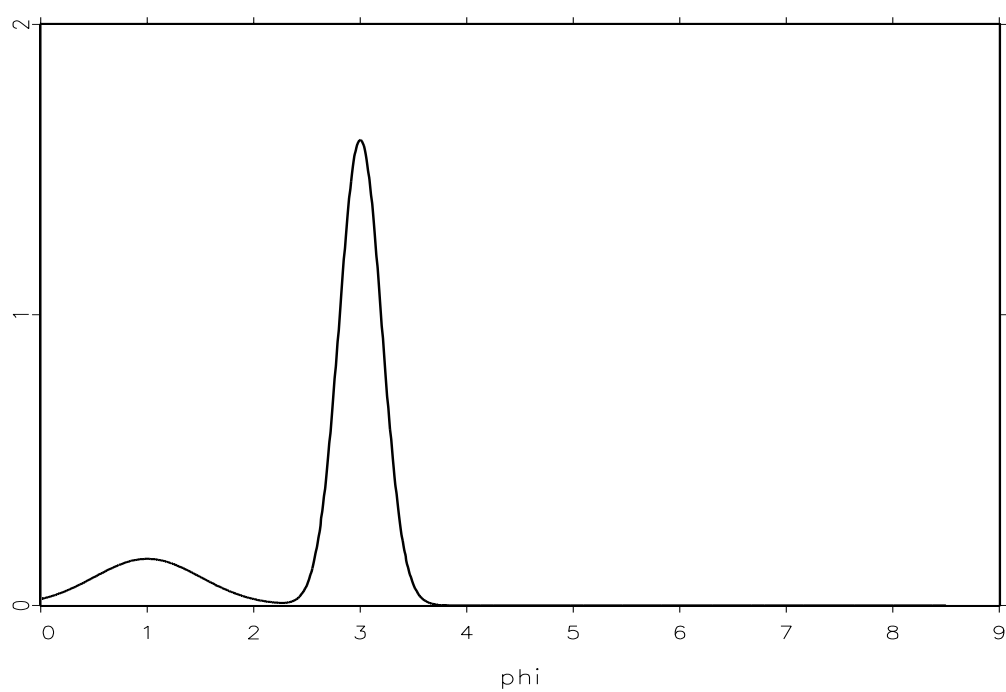
22

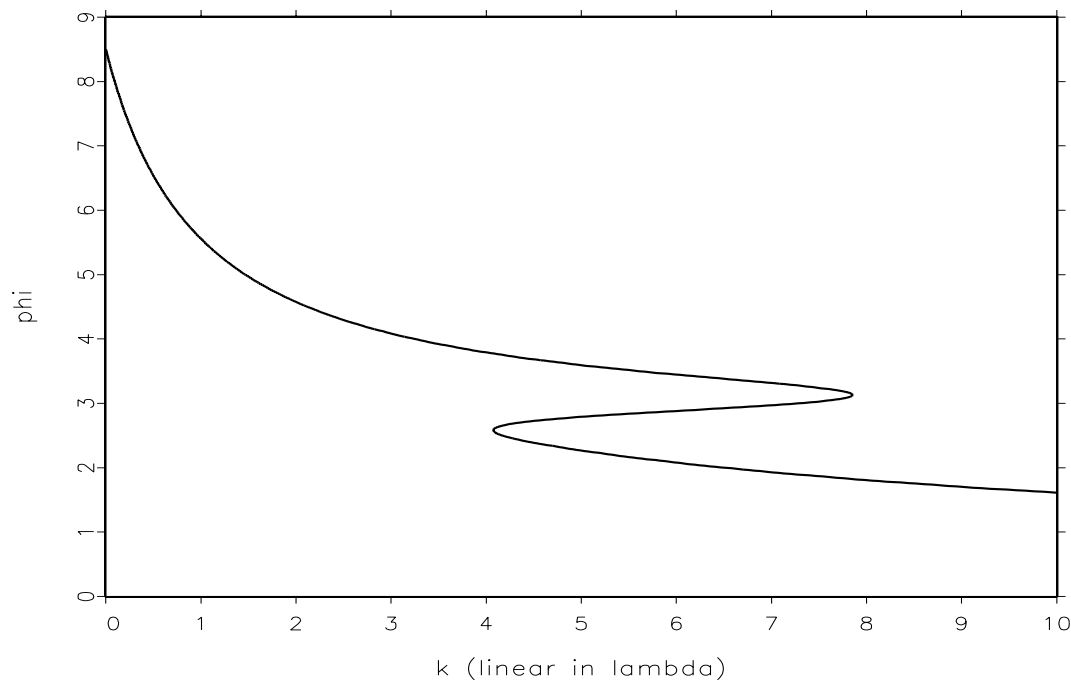Figure 3: Another density of $\Gamma_0(q)$. Note: support is $(0, 8.5)$.

Figure 4: Equilibrium reservation prices if $\Gamma_0(q)$ is as in Figure 3. Note: support of $\Gamma_0(q)$ is $(0, 8.5)$, and $u - b = 8.5$.
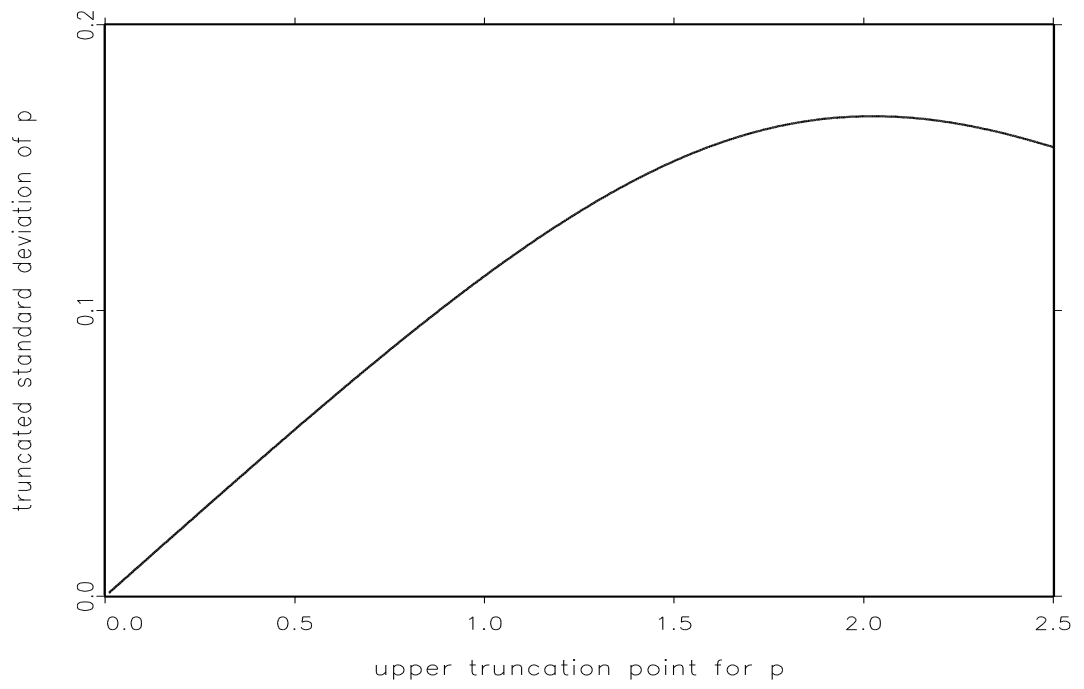
Figure 5: Standard deviation of the price distribution truncated from above, as a function of the truncation point.
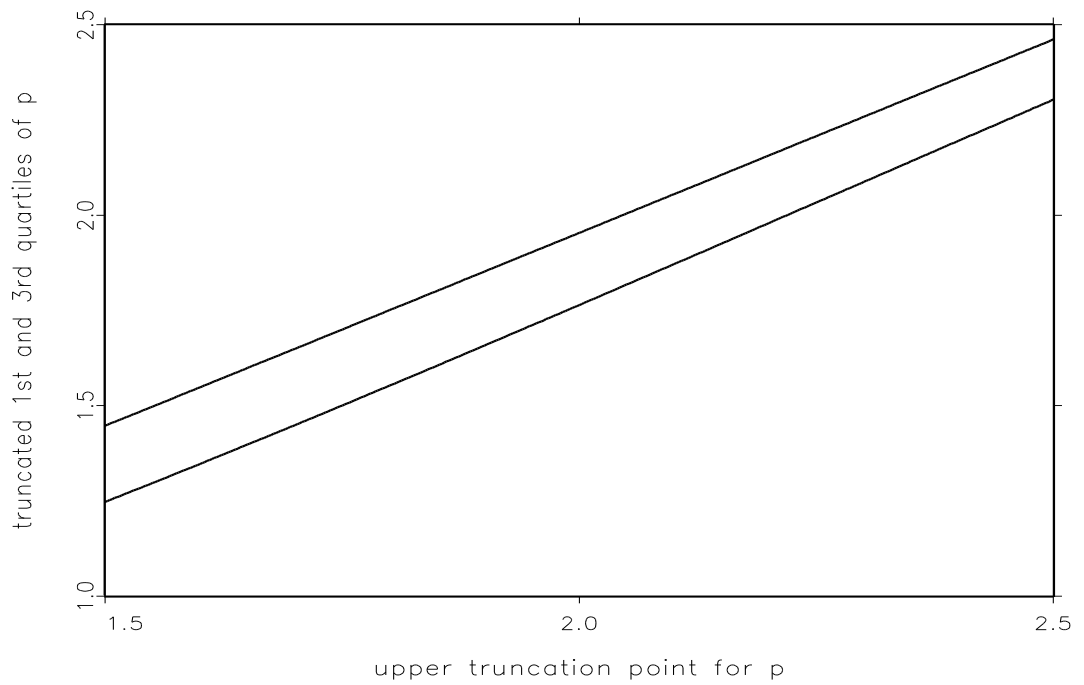

Figure 6: 1st and 3rd quartiles of the price distribution truncated from above, as a function of the truncation point.

References

Acemoglu, D. (2001), "Good jobs versus bad jobs", *Journal of Labor Economics*, 19, 1–22.

Acemoglu, D. and R. Shimer (2000), "Wage and technology dispersion", *Review of Economic Studies*, 67, 585–608.

Albrecht, J.W. and S.B. Vroman (2002), "A matching model with endogenous skill requirements", *International Economic Review*, 43, 283–305.

Bakos, J.Y. (1997), "Reducing buyer search costs: implications for electronic marketplaces", *Management Science*, 43, 1676–1692.

Bakos, J.Y. (2001), "The emerging landscape for retail e-commerce", *Journal of Economic Perspectives*, 15, 69–80.

Baye, M.R., J. Morgan, and P. Scholten (2004a), "Persistent price dispersion in online markets", in *The New Economy*, University of Chicago Press, Chicago.

Baye, M.R., J. Morgan, and P. Scholten (2004b), "Price dispersion in the small and in the large: evidence from an internet price comparison site", *Journal of Industrial Economics*, forthcoming.

Baylis, K. and J.M. Perloff (2002), "Price dispersion on the internet: good firms and bad firms", *Review of Industrial Organization*, 21, 305–324.

Bester, H. (1988), "Bargaining, search costs and equilibrium price distributions", *Review of Economic Studies*, 55, 201–214.

Bontemps, C., J.M. Robin, and G.J. van den Berg (2000), "Equilibrium search with continuous productivity dispersion: theory and non-parametric estimation", *International Economic Review*, 41, 305–358.

Brown, J.R. and A. Goolsbee (2002), "Does the internet make markets more competitive? Evidence from the life insurance industry", *Journal of Political Economy*, 110, 481–507.

Brynjolfsson, E. and M. Smith (2000), "Frictionless commerce? A comparison of internet and conventional retailers", *Management Science*, 46, 563–585.

Burdett, K. (1990), "Search market models: a survey", Working paper, University of Essex, Colchester.

Burdett, K. and L. Muus (1989), "A note on truncated means and variances", Working paper, Cornell University, Ithaca.

Clay, K., R. Krishnan, and E. Wolff (2001), "Prices and price dispersion on the web: evidence from the online book industry", *Journal of Industrial Economics*, 49, 521–539.

Daripa, A. and S. Kapur (2001), "Pricing on the internet", *Oxford Review of Economic Policy*, 17, 202–216.

Diamond, P.A. (1987), "Consumer differences and prices in a search model", *Quarterly Journal of Economics*, 102, 429–436.

Economist (1999), "Frictions in cyberspace", page 112, November 20, The Economist Newspaper Ltd., London.

Economist (2004), "A perfect market; a survey of e-commerce", May 15, The Economist Newspaper Ltd., London.

Ellison, G. and S. Fisher Ellison (2001), "Search, obfuscation, and price elasticities on the internet", Working paper, MIT, Cambridge MA.

Foster, L., J. Haltiwanger, and C.J. Krizan (2001), "Aggregate productivity growth: lessons from microeconomic evidence", in C.R. Hulten et al., editors, *New Developments in Productivity Analysis*, NBER, Cambridge MA.

Karlin, S. (1982), "Some results on optimal partitioning of variance and monotonicity with truncation level", in G. Kallianpur et al., editors, *Statistics and probability: essays in honor of C.R. Rao*, North-Holland, Amsterdam.

Klepper, S. and K.L. Simons (2001), "Industry shakeouts and technological change", Working paper, Carnegie-Mellon University, Pittsburgh.

Kuhn, P. and M. Skuterud (2000), "Job search methods: internet versus traditional", *Monthly Labor Review*, October, 3–11.

Kuhn, P. and M. Skuterud (2004), "Internet job search and unemployment durations", *American Economic Review*, 94, 218–232.

Morris, S. and H.S. Shin (2000), "Rethinking multiple equilibria in macroeconomic modelling", Working paper, Yale University, New Haven.

Mortensen, D.T. (1986), "Job search and labor market analysis", in O. Ashenfelter and R. Layard, editors, *Handbook of Labor Economics*, North-Holland, Amsterdam.

Mortensen, D.T. (1999), "Modeling matched job-worker flows", Working paper, Northwestern University, Evanston.

Mortensen, D.T. (2003), *Wage Dispersion*, MIT Press, Cambridge.

Mortensen, D.T. and C.A. Pissarides (1999), "New developments in models of search in the labor market", in O. Ashenfelter and D. Card, editors, *Handbook of Labor Economics, Volume III*, North-Holland, Amsterdam.

Reinganum, J. (1979), "A simple model of equilibrium price dispersion", *Journal of Political Economy*, 87, 851–858.

Rob, R. (1985), "Equilibrium price distributions", *Review of Economic Studies*, 52, 487–504.

Robin, J.M. and S. Roux (2002), "An equilibrium model of the labor market with endogenous capital and two-sided search", *Annales d'Économie et de Statistique*, 67–68, 257–308.

Stevenson, B. (2003), "The Internet, job search, and worker mobility", Working paper, Harvard University, Cambridge.

Van den Berg, G.J. (2003), "Multiple equilibria and minimum wages in labor markets with informational frictions and heterogeneous production technologies", *International Economic Review*, 44, 1337–1357.

Appendix

Appendix 1   The reservation price equation

Let R denote the expected present value of being a searching customer who follows the optimal strategy. The customer's perception of the future is independent of time or search duration, so the optimal strategy is constant, and R does not depend on the elapsed duration of search. It is well known (see e.g. Mortensen, 1986) that the Bellman equation for R can be written as an asset flow equation,

$$\rho R = b + \lambda \mathrm{E}_p \max\{0, V(p) - R\} \tag{13}$$

In this equation, the expectation is taken over the distribution F, and $V(p)$ is the expected present value of being a customer who buys the product at price p and who follows the optimal strategy after dissolution of the match with the firm. The return of the asset R in a small interval around t equals the sum of the instantaneous utility flow in this interval and the expected excess value of finding a trading partner in this interval.

    Similarly, the Bellman equation for $V(p)$ can be written as

$$\rho V(p) = u - p + \delta(R - V(p)) \tag{14}$$

so $V(p)$ decreases in p. The optimal strategy of the searching customer is to accept a match iff $V(p) > R$, and this can be characterized by a reservation price ϕ satisfying $V(\phi) = R$. Substitution of the latter into (14) gives $\rho R = u - \phi$ and $\rho V(p) = [\rho(u-p)+\delta(u-\phi)]/(\rho+\delta)$. Substitution of this into (13) and subsequent partial integration gives equation (1) in the main text.

Appendix 2   Continuous distribution of firm types

We do not impose discreteness or continuity of Γ_0, and we assume that its support equals $[\underline{q}, \overline{q}]$, with $0 < \underline{q} \le \overline{q} < u - b$. Equation (5) can be substituted into (4),

$$u - b = \phi + \frac{\lambda\beta}{\rho+\delta} \frac{\int_0^\phi \Gamma_0(q)dq}{\Gamma_0(\phi)} \tag{15}$$

The equilibrium ϕ is the solution of this equation. If q is continuous then multiple solutions may arise, but it turns out to be difficult to characterize in a transparent way the set of structural parameters for which multiplicity arises. It can be shown that necessary for multiplicity is that the density of Γ_0 increases on at least part of the support. For discrete q, equations (6) and (8) follow from (15).

There always exists a solution to (15). Consider the right-hand side. If $\phi = u - b$ then it exceeds the left-hand side. In general, the right-hand side can be rewritten as

$$\phi + \frac{\lambda\beta}{\rho + \delta}[\phi - \mathrm{E}_0(q|q < \phi)]$$

where the expectation is taken over Γ_0. If $\phi \downarrow 0$ then this right-hand side converges to zero, so then it is smaller than the left-hand side. On $[0, u - b]$, the right-hand side is continuous in ϕ. By invoking Brouwer's fixed point theorem, it follows that there exists a solution in $(0, u - b)$.