Centre interuniversitaire de recherche
en économie quantitative

CIREQ

Cahier 11-2011

Growth through Experimentation

Alain GABLER and
Markus POSCHKE

CIREQ

Cahier 11-2011

Growth through Experimentation

Alain GABLER and Markus POSCHKE

CIREQ, Université de Montréal
C.P. 6128, succursale Centre-ville
Montréal (Québec) H3C 3J7
Canada

téléphone : (514) 343-6557
télécopieur : (514) 343-7221
cireq@umontreal.ca
http://www.cireq.umontreal.ca

Dépôt légal - Bibliothèque nationale du Canada, 2011, ISSN 0821-4441
Dépôt légal - Bibliothèque et Archives nationales du Québec, 2011
ISBN-13 : 978-2-89382-620-2

Growth through Experimentation [a]

Alain Gabler [b]

Swiss National Bank

Markus Poschke [c]

McGill University, CIREQ and IZA

September 21, 2011

Abstract

Recent empirical research has documented the importance of shocks to firm-specific productivity, but has provided only limited evidence on their sources. This paper proposes and analyzes purposeful experimentation by firms as a source of such shocks and models industry dynamics in such a setting. We thereby make three contributions. The first is conceptual and consists in providing a micro-foundation to the stochastic process for firm-level productivity typically specified in the macroeconomic literature with firm heterogeneity. The second consists in quantifying the importance of experimentation for aggregate productivity growth to which experimentation, as a generalized form of R&D, contributes. In a setting that allows for growth through experimentation and through market selection among firms, 36% of aggregate embodied productivity growth can be attributed to experimentation. Finally, we show that size-dependent distortions can strongly reduce growth by reducing firm-level incentives to experiment and improve productivity.

JEL codes: E24, L16, O40.

---

[a]The views expressed in this paper do not necessarily reflect those of the Swiss National Bank.

[b]Address: Swiss National Bank, P.O. Box, CH-8022 Zürich.

[c]Corresponding author. Address: McGill University, Department of Economics (855 Sherbrooke St West, Montreal QC, Canada H2J 3R5; Phone: +1 514 398 4400, extension 09194; Fax: +1 514 398 4938), Centre interuniversitaire de recherche en économie quantitative (Cireq) and Institute for the Study of Labor (IZA). email: markus.poschke@mcgill.ca.

Keywords: endogenous growth, firm dynamics, selection, experimentation, entry and exit.

1 Introduction

Recent empirical research has documented the importance of shocks to firm-specific productivity, but has provided only limited evidence on their sources. The two findings that productivity varies a lot across firms even in narrowly defined industries, in large part due to firm-specific shocks to productivity, and that firm entry and exit rates are substantial, positively correlated across industries, and make a substantial contribution to aggregate productivity growth have had an important impact on the macroeconomic literature.[1] In particular, it has been recognized that distortions of allocative efficiency among heterogeneous firms have potentially large aggregate consequences and thereby can contribute to explaining differences in per capita income across countries.[2]

What are the sources of these firm-specific productivity shocks? This paper proposes and analyzes purposeful experimentation by firms as a source of such shocks and links them to aggregate growth. It is known from firm-level analyses that firms can influence the risk they take (see e.g. Coles, Daniel, and Naveen 2006) and that experimenting with new products and processes is a defining feature of innovation at the firm level. For instance, every year, about 25% of consumer goods for sale are either new or will be discontinued the next year, at least 40% of new goods are sold only for a single year, and plants adopt only between half and a third of the technologies they try (McGuckin, Streitwieser, and Doms 1996; Broda and Weinstein 2010; see also

[1]For some important contributions to this empirical literature, see e.g. Baldwin (1995); Geroski (1995); Sutton (1997); Caves (1998); Foster, Haltiwanger, and Krizan (2001, 2006); Hsieh and Klenow (2009); Gabaix (2010); Syverson (forthcoming).

[2]See e.g. Hopenhayn and Rogerson (1993); Barseghyan (2008); Guner, Ventura, and Xu (2008); Restuccia and Rogerson (2008); Hsieh and Klenow (2009); Poschke (2009, 2010); Moscoso Boedo and Mukoyama (2010); Midrigan and Xu (2010).

Lentz and Mortensen 2008 and Bernard, Redding, and Schott 2010). There is a broad management literature that interprets this process of churning at different levels as "innovation through experimentation" (see e.g. Thomke 2003). The finding that R&D outcomes are very uncertain (Doraszelski and Jaumandreu 2009) points in the same direction. All of this suggests that to some degree, firms deliberately expose themselves to "productivity risk" in order to improve their productivity, but can control the extent of this risk by choosing their experiments.

The first contribution of this paper consists in modelling experimentation at the firm level to provide very simple micro-foundations for a stochastic process for firm-level productivity. Existing heterogeneous-firm work in macroeconomics takes this process as exogenous (see e.g. Hopenhayn 1992; Samaniego 2006; Luttmer 2007; Restuccia and Rogerson 2008). If it is instead endogenous, analyses of the impact of e.g. micro-level distortions on aggregate productivity that presuppose an exogenous process miss the effect of distortions on risk-taking and therefore miss part of the effect on aggregate growth, as we show below. Therefore, modeling experimentation is important.

We model experimentation in a very simple way: in the setting of a heterogeneous-firm model in the tradition of Hopenhayn (1992), we allow firms to experiment with their production process every period. The experiment is modelled as drawing a random innovation to the firm's productivity.[3] (We also allow for additional shocks the firm cannot influence.) Firms can choose how risky or crazy they want their experiment to be; riskier experiments are draws from a distribution with a higher variance. Firms are not forced to stick with the outcomes of failed experiments; they can undo experiments that reduce their productivity, though possibly only incompletely. Because of this option (think of not implementing R&D findings or pulling an unsuccessful new product

[3]When productivity is measured as revenue productivity, as is the case in almost all data sets used in productivity measurement, fluctuations in product quality or consumer tastes are indistinguishable from productivity fluctuations. For this reason, our setting in terms of productivity risk and experimentation with processes can also be interpreted in terms of experimentation with products.

off the market), the expected value of experimenting is positive and increases in the riskiness of the experiment.[4] This is balanced by a higher cost of conducting ambitious experiments compared to more incremental/marginal ones so that in equilibrium, firms choose experiments with limited risk. We first derive some results in a simple analytical model and then integrate the experimentation process into a full quantitative model.[5]

The possibility to reject failed experiments implies that in expectation, an experimenting firm's productivity grows. In promoting productivity, experimentation is akin to R&D and can in fact be interpreted as a generalized form of R&D, which after all essentially is directed experimentation. More formally, compared to the endogenous growth literature with R&D, our modelling strategy implies that a firm's costly innovation activities always yield a result (not only with some probability), but that this result may well be unsatisfactory (e.g. not a productivity improvement). Moreover, in our setting, investing more in innovation activities raises the expected productivity increase because new, better outcomes become attainable, not because the probability of making an innovation of a fixed size increases.[6] Our approach is more general than the typical modeling of R&D for two reasons. Firstly, it allows capturing the activity

[4]Well-known reversed experiments are Coca Cola's New Coke, which served as Coca Cola's flagship product for less than 3 months in 1985 (thanks to Pedro Bento for suggesting this example), and Denver airport's automated baggage handling system, which was turned off in 2005 without ever having been fully used (for this and some further examples, also see Holmes, Levine, and Schmitz, 2008).

[5]Because the way we model experimentation is designed to fit well in a macro model, it is quite distinct from the theoretical literature on experimentation (see e.g. Bolton and Harris, 1999; Keller, Rady, and Cripps, 2005; Acemoglu, Bimpikis, and Ozdaglar, 2011). These papers consider two- or multi-armed bandits. These correspond to the choice between discrete projects with returns following unknown stochastic processes which can be learned by experimentation. Our setting in contrast can be interpreted either as experimenting with changing an existing project or with adding a new project to a portfolio, where the stochastic process for returns is known and can be influenced. The two types of models thus correspond to very different settings, involving different types of uncertainty.

[6]Both formulations are consistent with the finding by Castro, Clementi, and Lee (2009) that the variance of idiosyncratic shocks is larger in industries with higher R&D intensity. The formulation of R&D closest to ours is Kortum (1997), who allows for stochastic R&D outcomes on a continuum when modeling research at the frontier. Yet, predictions for the frontier do not help explain firm dynamics for the bulk of firms. Doraszelski and Jaumandreu (2009) propose a setting where firms engage in R&D with stochastic outcomes resulting in a stochastic process for productivity, which they estimate using the Spanish ESEE firm-level data set. Their paper is empirical and focusses on the firm level; it does not consider aggregate implications.

of the large portion of firms which do not report patenting or R&D spending but still innovate. (These are non-negligible; see also Francois and Lloyd-Ellis (2003); Klette and Kortum (2004); Syverson (forthcoming).) Secondly, in our setting, policies that do not affect the cost of R&D or its rate of return but increase adjustment costs still affect growth by penalizing risky experiments.

In the aggregate, experimentation at the firm level translates into productivity growth, with the idiosyncratic risk present at the firm level smoothed out. Unsuccessful firms slowly fall behind as others' productivity progresses, until they find that profits do not cover fixed costs anymore and they exit. Coupling this with entry and closing the model as in Luttmer (2007) and Gabler and Licandro (2007) results in a balanced growth path. In these papers, the growth rate depends on the (exogenous) variance of idiosyncratic shocks and on the (endogenous) rates of entry and exit. A larger variance of firm-level productivity shocks (crazier experiments) raises the aggregate productivity growth rate. Intuitively, this occurs because firms always have the option to exit after negative shocks, so that an increase in variance is beneficial for productivity growth.[7] This is true even if the average shock does not improve a firm's productivity. In our setting, this effect is even stronger, as failed experiments can be undone while continuing operations. Thus, by allowing for a more realistic response to idiosyncratic shocks and by endogenizing their variance, our setting allows to quantify experimentation and its contribution to growth. Doing so provides information on the contribution of firm-level purposeful improvement activity to aggregate growth that is broader than R&D or patent figures, and thus complements them.

[7] In this spirit, the process of entry and exit, which involves a lot of "churn" and exhibits positive correlation of entry and exit rates across industries, can also be interpreted in terms of experimentation, as proposed by Haltiwanger, Jarmin, and Schank (2003) and by Bartelsman, Perotti, and Scarpetta (2008): entry can result in spectacular success (sometimes) or in failure (often), is associated with large uncertainty about success, and therefore is essentially an experiment. This view fits in naturally with the importance of entry and exit for aggregate productivity growth documented by Foster, Haltiwanger, and Krizan (2001, 2006).

The second contribution of the paper exploits this and consists in quantifying the contribution of experimentation to aggregate productivity growth by calibrating the model to U.S. data on firm dynamics such as firm and job turnover rates, survival rates, and the contribution of entry and exit to aggregate productivity growth. The model suggest that 36% of aggregate productivity growth can be attributed to experimentation, and that firms invest resources worth 7.2% of output in experimentation.

Finally, we use the calibrated model to assess the effect of distortions on growth. This complements the recent literature on the effects of distortions on the level of aggregate productivity (see the references in footnote 2). There are two main results here. Firstly, the ability to reverse failed experiments is key for growth: even small adjustment frictions which make experiments imperfectly reversible have a large impact on growth, both because they reduce incentives to experiment and because they imply a less productive experimentation technology.

Secondly, size-dependent distortions in the form of higher taxes or stricter regulation of more productive firms reduce the incentive to innovate and thus reduce growth. This effect is very powerful; already a 10% tax on the 5% most productive firms can reduce the aggregate productivity growth rate by almost one quarter of a percentage point. This strong effect, which to the best of our knowledge has not been shown before, is only natural: the objective of innovation is profit growth, but size-dependent policies – unlike uniform taxes – do not simply reduce profits but reduce marginal profits even more, discouraging innovation. This effect comes out clearly in an endogenous growth model with firm dynamics like ours but does not necessarily arise in other endogenous growth models with more limited firm heterogeneity. The result is of practical importance because many rules affecting firms are size-dependent; they may be enforced more strictly for larger firms or even apply only to firms above a certain size (for some examples, see Section 4 or Guner, Ventura, and Xu, 2008). Size-dependent policies thus are even more harmful than suggested by Guner, Ventura, and Xu (2008) and Restuccia and Rogerson (2008).

The paper is organized as follows. In Section 2, we present a simple model of experimentation that allows for an analytical solution. Section 3 presents the full quantitative model and its calibration. The effects of distortions are described in Section 4, and Section 5 concludes.

2 A Simple Model of Endogenous Experimentation

In this section we present a simple model in which productivity growth results from experimentation by firms. We model this process of experimentation by assuming that firms are hit by idiosyncratic productivity shocks whose variance they can choose. We interpret these as experiments with random outcomes through which firms try to improve their productivity. Firms can choose how conservative or "crazy" an experiment is; this is reflected in the variance of the shock they receive. If the result of an experiment is not as desired, a firm can revert to its previous level of productivity. We assume that experimentation is costly in terms of current output. The simple specification chosen here, together with some convenient functional form assumptions, allows for an analytical solution for firm behavior and the growth rate. As usual, simplicity has benefits, but also costs. Therefore, we numerically explore a generalized model with some additional realistic features in the Section 3.

The economy is populated by a representative household and by a continuum of measure 1 of goods-producing firms. In this section, we abstract from entry and exit. There also is a sector of perfectly competitive financial intermediaries. Time is discrete.

2.1 Preferences

Household preferences are given by

$$\sum_{t=0}^{\infty} \beta^t U(c_t) = \begin{cases} \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma} & \text{for } \sigma > 0, \sigma \neq 1 \\ \sum_{t=0}^{\infty} \beta^t \log(c_t) & \text{for } \sigma \to 1, \end{cases}$$

where $\beta \in (0,1)$. Households can consume or save by investing in shares of output-producing firms via a sector of perfectly competitive financial intermediaries. In equilibrium, these intermediaries don't make profits, all hold the market portfolio, and, absent aggregate uncertainty, pay a net return r on consumers' investments. A household's budget constraint then is

$$c_t = w_t l + a_t (1 + r_t) - a_{t+1},$$

where a_t denotes assets held at the beginning of period t and household labor supply is constant at l. The Euler equation for the accumulation of assets then is

$$c_t^{-\sigma} = \beta \mathbb{E}\left[(1 + r_{t+1}) c_{t+1}^{-\sigma}\right]. \tag{1}$$

2.2 The Problem of the Firm

Firms produce output with the production function

$$y(z, s) = [z \cdot \theta(s)]^{\alpha} l(z)^{1-\alpha} \tag{2}$$

and sell it in a competitive market. We normalize the price of output to 1. Firms differ in their productivity z and choose their labor input $l(z)$ as a function of it. The term $\theta(s)$ indicates a disruption cost of experimentation, where s denotes the experimentation intensity. This kind of cost is analogous to Holmes, Levine, and Schmitz (2008), who assume the presence of similar "switchover disruption costs" in technology adoption in their analysis of the link between competition and productivity. With the optimal choice of labor input l given productivity z and a choice of s, profits are

$$\Pi(z, w, s) = \alpha z \cdot \theta(s) \left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}}. \tag{3}$$

Each period, the firm chooses how much to experiment. Experimenting implies drawing an innovation ε from a distribution with *cdf* $\Phi_s(\varepsilon)$, resulting in a new level of productivity $z' = \varepsilon z$. The firm can choose the risk it takes in its experiment. This choice is represented by the parameter s, which is related to the variance of ε. In this section, we assume that ε follows a logistic distribution with mean μ and variance $s^2\pi^2/3$. The firm can conduct one experiment per period, implying that the choice of s can be adjusted every period. Choosing riskier experiments is costly as it is disruptive of current production: we assume that $\theta'(s) < 0$ and $\theta''(s) < 0$, so the cost of experimenting is convex.[8]

Firms can discard the results of unsuccessful experiments, i.e. those resulting in draws of $\varepsilon < 1$, which would imply a fall in productivity if the result of the experiment was adopted. In this case, the firm still suffers the disruption cost of the experiment, but can directly revert to its previous technology with productivity z.

The value of a firm is

$$V(z, w) \quad = \quad \max_s \left(\Pi(z, w, s) + \frac{1}{1+r}\left[V(z, w)\,\Phi_s(1) + \int_1^\infty V(z\varepsilon, w)\,d\Phi_s(\varepsilon)\right]\right).$$

[8]Clearly, it should also be the case that $\theta(0) = 1$: it is costless not to experiment. However, this does not imply that experimenting very little should be almost costless. Therefore, it is reasonable to consider cost functions with $\lim_{s\downarrow 0} = \bar{\theta}$, $\bar{\theta} \in (0, 1)$, which essentially feature a fixed cost of experimenting of $1 - \bar{\theta}$ plus a marginal cost of larger experiments.

If we assume that $V(z,w) = \zeta(w)z$, i.e. linear in z just as the profit function, we have

$$\zeta(w) = \max_s \left(\theta(s)\,\alpha \left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}} + \zeta(w)\frac{1}{1+r}\left[\Phi_s(1) + \int_1^\infty \varepsilon d\Phi(\varepsilon)\right]\right) \tag{4}$$

using equation (3).

The first-order condition for s then is

$$\theta'(s^\star)\,\alpha\left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}} + \zeta(w)\frac{\gamma'(s^\star)}{1+r} = 0,$$

where $\gamma(s) = \Phi_s(1)\cdot 1 + \int_1^\infty \varepsilon d\Phi(\varepsilon)$ is the expected gross growth rate of z given a choice of s. Inserting this condition into (4), we obtain

$$-\frac{\theta'(s^\star)}{\theta(s^\star)} = \frac{\gamma'(s^\star)}{1+r-\gamma(s^\star)}, \tag{5}$$

which implies that the marginal cost of experimentation (left-hand side) should equal its expected marginal capitalized benefit (right-hand side).[9]

With a logistic distribution for ε with mean $\mu = 1$ and variance $s^2\pi^2/3$, $\gamma(s) = 1 + s\ln 2$.[10] Assuming that the disruption cost function takes the form

$$\theta(s) = (\bar{s} - s)^q$$

with $s \in [0, \bar{s}]$ and $q \in (0,1)$, this implies

$$\frac{q}{\bar{s}-s^\star} = \frac{\ln 2}{r - s^\star \ln 2}.$$

[9]As usual, this requires restricting parameters such that $\gamma - 1 < r$ for boundedness. A sufficient condition for the second order condition for a maximum to be satisfied is that γ'' is non-positive, which holds for logistic innovations ε. Moreover, under the assumptions made on $\theta(s)$ and with $\gamma'' \le 0$, equation (5) has a unique solution.

[10]The logistic distribution admits a support in $\mathbb{R}$. To be precise and to rule out negative productivity, reallocate the probability mass on $\mathbb{R}^-$ to form an atom at $\varepsilon = 0$.

The optimal choice of experimentation for an individual firm then is

$$s^{\star} = \frac{\bar{s}\ln 2 - qr}{(1-q)\ln 2}$$

and its expected net growth rate of z is

$$\gamma - 1 = s^{\star}\ln 2 = \frac{\bar{s}\ln 2 - qr}{1-q}.$$

The degree of experimentation increases in $\bar{s}$ and falls in r. This is of course typical for an investment decision (given q, raising $\bar{s}$ reduces the cost of the investment).

With all firms choosing the same $s^{\star}$ and conducting independent experiments, there is no aggregate uncertainty, and the growth rate of aggregate productivity equals the expected productivity growth rate of each firm, $\alpha(\gamma-1)$. This implies an aggregate output growth rate g of

$$\alpha(\gamma-1) = \alpha s^{\star}\ln 2 = \alpha\frac{\bar{s}\ln 2 - q\rho}{1+(\alpha\sigma-1)\,q},$$

an interest rate

$$1+r = (1+\alpha s^{\star}\ln 2)^{\sigma}(1+\rho) = 1+\rho+\alpha\sigma s^{\star}\ln 2,$$

and a firm value

$$\zeta(w)z = \alpha\left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}}\theta(s^{\star})\frac{1+r}{r-\gamma}z = \alpha\left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}}\theta(s^{\star})\frac{1+\rho+\alpha\sigma s^{\star}\ln 2}{\rho+(\alpha\sigma-1)\,s^{\star}\ln 2}z,$$

where $\rho \equiv 1/\beta - 1$ and the last equality in each equation uses a log approximation to r. Firm value is indeed linear in z, as guessed.

Given the positive relationship of r with ρ and σ stemming from the Euler equation, experimentation intensity and growth fall with ρ. The aggregate growth rate increases with α because the direct effect of higher α that translates a given γ into a higher output growth rate at the firm level outweighs the general equilibrium effect operating through a higher interest rate. Finally, the growth rate falls in q if $\alpha\sigma > 1$. Smaller q implies that the experimentation cost function is more kinked and the marginal cost of experimenting at low levels of s is flatter, resulting in an optimal choice of s closer to $\bar{s}$.

This simple model allows for a balanced growth path, albeit one where the variance of the firm productivity distribution grows without bound. The endogenously generated aggregate growth rate depends on patience and risk attitudes, on the importance of technology relative to inputs in generating output (α), and on factors related to the cost of experimentation ($\bar{s}, q$).

These growth rate determinants and their comparative statics are similar to those found in other endogenous growth models. What is particular here is that innovation is very "informal" and occurs so to speak on the shop floor, not in the R&D department. R&D can of course be seen as a special case of our perspective on innovation, consisting in directed experimentation using dedicated resources. In fact, our setting essentially is a generalization of common ways of modeling R&D: An experimenting firm's expected productivity improvement increases in s since higher s implies that a productivity improvement is larger, conditional on being an improvement. The probability of the experiment being successful is not affected by s in the simplest specification. This is in contrast to the standard specification of knowledge production through R&D in the endogenous growth literature, where more innovative activity raises the probability of obtaining a given productivity improvement. If only the expected productivity improvement mattered, the two settings would be very close. However, given the variance of firm-level productivity, the uncertainty of R&D success found in the data, and the amount of firms not reporting R&D, the typical R&D specification appears to abstract from many relevant factors that our specification can capture. In particular, if the process of improving productivity entails potentially substantial fluctuations in input use, as can be the case if experimentation is important, then not only the expected value of higher productivity matters for how much innovation activity is undertaken, but adjustment costs encountered in the process also affect the choice of how much to experiment. This factor of course is absent from the typical R&D specification.

2.3 Extensions

2.3.1 Other Sources of Shocks

Of course, a firm's experimentation with products or technology may not be the only source of innovations to its (measured) productivity. There may also be idiosyncratic shocks affecting it coming from changes in customers' tastes or from exogenous shocks to its technology. This section shows that allowing for such additional shocks does not affect the basic conclusions of the previous section.

Suppose that every period, after a firm's productivity is modified by the firm's experiment, its productivity is also hit by an exogenous multiplicative innovation u with *cdf* $F(u)$. Because this is uncorrelated with the level of z, it does not break linearity of firm value in z; it just increases volatility. A firm's problem as given in (4) then becomes

$$\zeta(w) = \max_{s} \left(\theta(s)\,\alpha \left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}} + \zeta(w)\frac{1}{1+r}\left[\Phi_s(1) + \int_1^{\infty} \varepsilon d\Phi(\varepsilon)\right] \int u dF(u) \right). \tag{6}$$

The first order condition for s then is

$$\theta'(s^{\star})\,\alpha\left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}} + \zeta(w)\frac{\gamma'(s^{\star})}{1+r}\int u dF(u) = 0,$$

implying

$$-\frac{\theta'(s^\star)}{\theta(s^\star)} = \frac{\gamma'(s^\star)}{1+r-\gamma(s^\star)} \int u dF(u).$$

If the expected value of u is 1, which is a rather natural assumption, this equation reduces to equation (5) above. Then the solutions for $s^\star$, r and the growth rate are unaffected. We therefore abstract from the u shock in the remainder of this section for notational simplicity. In the full quantitative model explored in Section 3, the exogenous shock gains some more importance, because it leads to growth through selection as in Luttmer (2007) and Gabler and Licandro (2007). In this simple analytical model, this does not occur because there is no entry and exit.

2.3.2 Imperfect reversibility

Up to here, we have assumed that experiments can be reversed freely and fully. In this subsection, we continue to assume free reversibility, but relax full reversibility. For instance, while a firm could take a failed new product off the market, it may well be that the failed introduction affects its reputation and therefore consumers' taste for the firm's other products that it continues to sell. In the data, this would be reflected as reduced productivity.

We implement this assumption by assuming that after an experiment with outcome ε, the firm cannot return to its previous productivity z, but can only return to the value $(1-\lambda)z + \lambda\varepsilon z$, with $\lambda \in (0,1)$. Full reversibility as above corresponds to $\lambda = 0$. So λ can be said to index the degree of irreversibility (between 0 and 1) of experimental outcomes. With $\lambda \in (0,1)$, firms still reject experiments with $\varepsilon < 1$. However, failed experiments here have permanent costs.[11]

[11] A fixed cost of reverting to the old technology and productivity would imply an acceptance threshold for ε strictly below 1. Some experiments would not be worth undoing if this is costly. We abstract from this possibility in the following because it is hard to identify this cost.

Using linearity again, the firm's value function here is

$$\begin{aligned}\zeta(w) &= \max_s \left(\theta(s)\alpha \left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}} + \zeta(w)\frac{1}{1+r}\left[\int_0^1 [(1-\lambda)+\lambda\varepsilon]d\Phi(\varepsilon) + \int_1^\infty \varepsilon d\Phi\left(\varepsilon\right)\right]\right) \\ &= \max_s \left(\theta(s)\alpha \left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}} + \zeta(w)\frac{1}{1+r}\gamma^\lambda(s)\int udF(u)\right). \end{aligned} \tag{7}$$

Note that irreversibility affects the expected growth rate. With $\lambda > 0$, the expected growth gross rate of z after an experiment of intensity s is $\gamma^\lambda(s) = \Phi(1)(1-\lambda) + \lambda\int_0^1 \varepsilon d\Phi(\varepsilon) + \int_1^\infty \varepsilon d\Phi(\varepsilon)$.

The first order condition for s then is

$$\theta'\left(s^\star\right)\alpha\left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}} + \zeta(w)\frac{(\gamma^\lambda)'\left(s^\star\right)}{1+r} = 0,$$

implying

$$-\frac{\theta'\left(s^\star\right)}{\theta\left(s^\star\right)} = \frac{(\gamma^\lambda)'\left(s^\star\right)}{1+r-\gamma^\lambda\left(s^\star\right)}.$$

Using the properties of the lognormal distributions, the gross expected growth rate of z given s then is approximately[12]

$$\gamma^\lambda(s) = 1 + (1-\lambda)s\ln 2. \tag{8}$$

Any given s results in less growth because reduced reversibility ($\lambda > 0$) constitutes a direct impairment of the efficiency of the experimentation technology. As a consequence, firms experiment less: optimal s is given by

$$s^\star = \frac{(1-\lambda)\bar{s}\ln 2 - qr}{(1-\lambda)(1-q)\ln 2}, \tag{9}$$

[12] $\gamma^\lambda(s) = \Phi(1)(1-\lambda) + \lambda\int_0^1 \varepsilon d\Phi(\varepsilon) + \int_1^\infty \varepsilon d\Phi(\varepsilon) = (1-\lambda)(1+s\ln 2) + \lambda s\ln(1+e^{1/s}) \simeq (1-\lambda)(1+s\ln 2) + \lambda = 1 + (1-\lambda)s\ln 2$, with a very good approximation for small s of less than 0.3.

which decreases in λ. Lower reversibility thus directly reduces growth by making the experimentation technology less productive, and reduces it further by discouraging experimentation.

3 Quantitative analysis

We now turn to a quantitative analysis of experimentation in the U.S. economy. In this section, we introduce endogenous entry and exit of firms into the model, describe the calibration strategy and present results on the importance of experimentation in the U.S. economy.

3.1 The full model with firm entry and exit

Firm turnover is substantial in real economies. Since it is robustly related to productivity dynamics and to aggregate growth (see e.g. Foster, Haltiwanger, and Krizan, 2001), it is important to consider it in our analysis. Incorporating this into the model requires going beyond the simple model presented in the previous section.

Firm value and the exit decision. Suppose that there is a probability δ every period that a firm exits exogenously. In addition, surviving firms need to pay a fixed operating cost of κ_f labor units per period to produce. Also suppose that they can costlessly decide to exit endogenously after learning their realizations of ε and u. The value of a firm then is

$$V(z, w) = \max_s \left\{ \Pi(z, w, s) - \kappa_f + \frac{1-\delta}{1+r} \mathbb{E}_s \max\left[V(zu, w), V(\varepsilon zu, w), 0\right] \right\}. \tag{10}$$

The value function embodies three decisions by the firm. The first max operator requires optimal choice of s, the second one optimal acceptance or rejection of the experimental outcome and optimal exit or continuation. The last two decisions are taken knowing the realized values of ε and u. The expectation is taken over these two random variables. The subscript s indicates that the distribution of ε, and thus the expectation, depend on the firm's choice of s.

For the continuation decision, the optimal strategy is to choose a threshold level z_x for the continuation productivity below which the firm exits. The optimal threshold satisfies

$$V(z_x, w) = 0.$$

The optimal acceptance/rejection decision still implies accepting if $\varepsilon > 1$. These decisions combined with the exogenous shocks u and δ determine a firm's probability of exiting the market or of remaining active and transiting from productivity z to productivity z'. Let the productivity transition operator which summarizes the effect of these transitions on the firm productivity distribution be Q.

Entry. There is free entry, and entry requires a sunk investment of κ_e units of labor.[13] Entrants draw their initial productivity from a distribution $\eta(z)$ with mean ϕ times the average z of incumbents and with variance σ_n. Entry is optimal as long as its value exceed its cost. Under free entry it must therefore be that in equilibrium, the expected value of entry equals its cost whenever there is positive entry.

The entry rate e thus is endogenous. Together with the endogenous exit rate, it drives the evolution of the measure of firms in the market. The law of motion of the productivity distribution of active firms, μ, then is

$$\mu' = Q\mu + e\eta.$$

[13]Specifying κ_f and κ_e in units of labor ensures that they grow in line with the economy, allowing for balanced growth.

The dependence of entrants' potential productivity on that of incumbents can be interpreted as imitation by entrants as in Luttmer (2007). The ease of imitation is parameterized by ϕ. This setup implies that the technological advances that are generated within the model are embodied: an individual firm's productivity gains are tied to its production facilities and do not spill over to other active firms, except for entrants who can benefit from them by making a new investment κ_e. Technology is thus embodied in existing or in new firms.

Growth. In the following, we restrict our analysis to the balanced growth path of this economy. This is defined as a situation where all aggregate variables grow at constant rates, except for employment, which is constant.

There are two sources of endogenous growth in this economy: growth through selection and through experimentation. Growth through selection is as in Luttmer (2007): the least productive firms exit and are replaced by more productive entrants, leading to growth through entry and exit as documented by Foster, Haltiwanger, and Krizan (2001). Growth through experimentation is as in the simple model presented in section 2. Both sources of growth lead to growth in embodied technology.

The amount of growth we attribute to experimentation is given by the output-weighted productivity growth rate of incumbents

$$GTE = \int \gamma(s^*(z))\, y(z)\mu(z)\, dz, \tag{11}$$

where $\gamma(s)$ is defined as in section 2 and $y(z)$ is optimal output of a firm with productivity z. This is a rather conservative measure of how much growth is due to experimentation since it excludes entry, which can also be interpreted as containing a component of experimentation given how much churning there is (see also footnote 7).

3.2 Calibration

In this section, we calibrate the model to U.S. data. This allows measuring the importance of experimentation. In the next section, we can then use the calibrated model to measure the potential effect of distortions.

To calibrate the model, we use commonly used values from the literature for some baseline parameters and choose the remaining ones jointly to minimize the mean squared relative deviation between a set of model moments and corresponding data moments on firm dynamics such as firm and job turnover rates, survival hazards, and the contribution of entry and exit to aggregate productivity growth. We use a genetic algorithm to find the global minimum distance.

The length of a time period is set to one year. The parameters which are set based on a priori information are a coefficient of relative risk-aversion σ of 2, a discount factor β of 0.96, and a labor share $1-\alpha$ of 64%. Of the two parameters of the disruption cost function, q and $\bar{s}$, we can only identify one using the model. To deal with this, we assume that the disruption cost function is inverse quadratic, i.e. $q = 0.5$, calibrate $\bar{s}$, and conduct some robustness checks below. We also assume that experiments are fully reversible ($\lambda = 0$) in the benchmark case and explore the consequences of imperfect reversibility in the next section. Finally, we assume that the exogenous idiosyncratic productivity shock u and the productivity of new firms are log-normally distributed with standard deviations σ_e and σ_n and with means that imply that expected profit growth due to exogenous shocks is zero for continuing firms and that entrants' relative productivity is ϕ on average.

The remaining parameters which need to be set then are $\bar{s}$, the upper bound for the risk of experiments; the standard deviation of the productivity distribution of entrants, σ_n; the standard deviation of the exogenous productivity shock, σ_e; the parameter determining the relative productivity of entrants, ϕ; and the fixed costs of production,

κ_f, and of entry, κ_e.

These six parameters are set to minimize the distance to six informative target moments observed in U.S. data. Because of the non-linearity of the model, it is not possible to indicate which individual target identifies which individual parameter, but we will suggest which targets are particularly informative for which parameter.

First of all, growth in embodied productivity – the result of the processes of selection and experimentation – accounts for about 60% of U.S. productivity growth according to Greenwood, Hercowitz, and Krusell (1997). With overall productivity growth of 2% per year, this amounts to 1.2% annual growth in embodied productivity. The model should generate the same growth rate. We do not restrict how much of this is due to experimentation and how much to selection. This target is informative about $\bar{s}$ and σ_e, the main parameters driving growth. The job turnover rate – 32% per year in the U.S. according to Davis, Faberman, Haltiwanger, Jarmin, and Miranda (2008, Table 2) – is also informative about σ_e. Importantly, because of the option to discard failures, experimentation affects job turnover far less than exogenous shocks do. Therefore, the use of these two targets allows to separate growth through selection from growth through experimentation.

The share of aggregate growth that is due to entry and exit is informative about the dispersion of entrants' productivity σ_n. Data reported in Foster, Haltiwanger, and Krizan (2001, 2006) suggests that this is between 26% (in manufacturing) and 100% (in the retail sector). Following Luttmer (2007), we target an intermediate value of 50%. The relative productivity of firms less than five years old of 99% (Foster, Haltiwanger, and Krizan, 2001, Table 9) contains information about ϕ. Finally, the size-weighted firm exit rate (7% per year) and entrants' four-year survival rate of 63% (both as reported in Bartelsman, Haltiwanger, and Scarpetta, 2004, Tables 5 and 7) are informative about δ and the ratio of κ_e and κ_f. The levels of κ_e and κ_f only affect the number of active

Table 1: Calibration: Model statistics, calibration targets

Statistic	Data	Model
calibration targets:		
embodied productivity growth	1.2%	1.2%
job turnover rate	32.2%	32.6%
contribution of entry and exit to growth	50%	49.6%
relative productivity of young firms	99%	99%
employment-weighted exit rate, incumbents	7.0%	7.2%
4-year survival rate, new firms	63.1%	63.0%
not used in calibration:		
productivity dispersion 9th/2nd decile	$2-4$	3.1
fraction of employment accounted for by		
5% largest firms	60%	59.4%
14% largest firms	75%	77.7%
27% largest firms	88%	88.4%
proportion of embodied growth due to experimentation		36.1%

establishment and thus average establishment size, but not entry, exit, growth etc. We set them to normalize average employment to 1.

Since there is no closed-form solution, the model is solved numerically. Table 1 reports the values of the target moments for the data and the model, and Table 2 lists the chosen parameters.

The calibration fits well even in dimensions that were not targeted. For instance, the dispersion in productivity between firms in the 9th and the 2nd decile of the productivity distribution is 3, in the middle of the range from 2 to 4 reported by Dwyer (1998). The firm size distribution also looks similar to the one in the data, with the shares of employment in the 5%, 14% and 27% largest firms close to the data counterparts reported by Henly and Sánchez (2009, Figure 5A) using U.S. Census Bureau County Business Patterns data.

Calibrated parameter values appear reasonable: entrants are on average about a

Table 2: Calibration: Parameter values

Parameter	Value	Parameter	Value
σ	2	α	0.36
β	0.96	q	0.5
σ_n	1.765	ϕ	0.773
σ_e	0.247	κ_e/κ_f	187.3
$\bar{s}$	0.230	δ	0.069

quarter less productive than incumbents, but their productivity distribution features large dispersion. The variance of exogenous productivity shocks is substantial, but a bit smaller than in papers that only allow for these exogenous shocks such as Luttmer (2007), Gabler and Licandro (2007) or Poschke (2009), who all find values above 0.3.[14] The largest feasible experiment (setting $s = \bar{s}$) would involve a variance of the shock about as large as that of the exogenous shock. However, firms choose to invest significantly less in experimentation, choosing s^* of 3.2% on average. Given $\bar{s}$ and q, this implies that they invest an average of 7.2% of output in experimentation. This is substantially larger than reported R&D investment for the US of about 2.7% of GDP in 2007 (World Bank World Development Indicators). This is only natural, since experimentation is much broader than R&D. While it includes R&D, it also includes activities that are aimed at improving productivity that are either done in R&D-conducting firms but not classified as R&D, or even those done in firms not reporting R&D at all.

Although this amount of experimentation may appear small, in particular compared to the variance of the exogenous shock, it makes a substantial contribution to growth: the proportion of growth due to experimentation, computed using equation (11), is a bit larger than a third. The reason is that in our model, the implications of firm-specific productivity shocks for aggregate growth vary greatly depending on the source of the

[14]Note that 0.3 is also the annual standard deviation of individual stock returns for firms listed on NYSE or NASDAQ according to Campbell, Lettau, Malkiel, and Xu (2001). In our model, value is proportional to productivity, so the standard deviation of productivity growth is the same as that of stock returns.

shocks. In particular, since unsuccessful experiments can be discarded, shocks which are due to experimentation lead to much higher growth than exogenous productivity shocks of the same size. As a result, a substantial fraction of productivity growth is due to experimentation, although the variance of experiments is much smaller than the variance of exogenous shocks to productivity.

4 Distortions and growth

Recent work such as Restuccia and Rogerson (2008) or Hsieh and Klenow (2009) has stressed the importance of distortions for understanding differences in levels of productivity across countries. In this section, we conduct an exercise analogous to that of Restuccia and Rogerson (2008), but with the objective of gauging the potential impact of distortions on differences in growth rates. In doing so, we will consider both aggregate and size-dependent distortions. Just as Restuccia and Rogerson (2008), we model distortions in a very general way as taxes on output.

4.1 Aggregate distortions

At optimal input choice, profits of a firm with productivity z and experimentation intensity s that is subject to an output tax at rate τ are

$$\Pi\left(z,w,s\right)=\left(1-\tau\right)^{\frac{1}{\alpha}}\alpha z\cdot\theta\left(s\right)\left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}}.$$

The first order condition for optimal choice of s then is

$$-\left[1-\tau\left(z\right)\right]^{\frac{1}{\alpha}}\alpha z\theta'\left(s\right)\left(\frac{1-\alpha}{w}\right)^{\frac{1-\alpha}{\alpha}}=\frac{1}{1+r}\frac{\partial\mathbb{E}_s\left[\max(V\left(zu,w'\right),V(\varepsilon zu,w'),0)\right]}{\partial s}. \quad (12)$$

If taxes are the same for all firms, that is, $\tau=\bar{\tau}\,\forall\,z$, they affect current and future profits in the same way. As a result, they affect both the cost of experimentation in terms of reduced current profits (left-hand side of equation 12) and the benefits of experimentation in terms of higher expected future profits (right-hand side of 12) in the same way. Because of this, a uniform output tax on all firms does not affect the level of experimentation, and therefore the growth rate, in our model economy.

4.2 Size-dependent distortions

Many taxes are not uniform. Similarly, as pointed out by Guner, Ventura, and Xu (2008), regulation typically has a size-dependent component. For instance, in many countries, rules and regulations are enforced more strictly for larger firms. In some cases, regulations explicitly vary by firm size; a well-known example are stricter firing restrictions on firms with more than 15 employees in Italy (see e.g. Schivardi and Torrini, 2008).

Papers that study such size-dependent policies typically analyze their impact on incentives to adjust factors, e.g. by discouraging firms from crossing a certain employment threshold. However, they have a deeper effect that goes beyond this: they may discourage firms from becoming so productive that it would be desirable to cross the regulation threshold in the first place. Our framework is well-suited to analyzing this effect. It is clear from equation (12) that higher taxes on more productive firms discourage experimentation by reducing its benefits.

In this section, we quantify the effect of size-related taxes. For simplicity, we assume that the tax rate depends only on a firm's current productivity z. In particular, we assume that firms above a certain productivity threshold z_τ are taxed at a rate τ, while firms below the threshold are subsidized at a rate b that balances the government's budget. While such a tax may appear unrealistic (taxes are more likely to depend on observable characteristics, such as employment), it captures the essence of size-

Table 3: Taxing the 5% most productive firms

	Tax rate τ			
	0%	5%	10%	20%
benchmark ($q = 0.5$):				
embodied productivity growth	1.20%	1.09%	0.97%	0.89%
robustness checks:				
$q = 0.25$				
embodied productivity growth	1.20%	1.06%	0.95%	0.85%
$q = 0.75$				
embodied productivity growth	1.20%	1.10%	1.00%	0.92%

Notes: Tax on the 5% most productive firms. The remaining firms are subsidized at a rate that balances the government's budget. In the robustness checks, the economy is recalibrated with the different value of q.

dependent taxes. In particular, it puts the spotlight on the dynamic effect of size-dependent taxes on productivity-promoting activities, while abstracting from strategic effects on input choice around thresholds where taxes increase.[15]

Table 3 shows that already quite low tax rates can substantially affect productivity growth. In the benchmark case with $q = 0.5$, already a small tax of 5% on the 5% most productivity firms reduces embodied technological progress by 0.1 percentage points. A larger tax of 20% reduces embodied technological by a quarter, and total growth by about 0.3 percentage points per year. Results are very similar if the economy is recalibrated with different values of q.

Table 4 shows that taxing more firms initially reduces the growth rate even further. Interestingly, once a very large fraction of firms is taxed, the effect on the growth rate diminishes again. This occurs because the less productive half of firms is very small, and therefore a reduction in its growth rate in response to the tax hardly affects aggregate

[15] Also note that, since uniform taxes do not affect experimentation, effects of a pair of taxes and subsidies (τ, b) on growth will be equivalent to those of a pair $(\tau + c, b - c)$.

Table 4: Varying the tax base

	Fraction of firms taxed				
	1%	5%	10%	17.5%	50%
embodied productivity growth	1.11%	0.97%	0.88%	0.78%	0.85%
employment of smallest taxed firm relative to average	12.7	3.3	1.8	1	0.2

Notes: Tax of 10% on the top 1, 5, 10, 17.5% or 50% of firms. The remaining firms are subsidized at a rate that balances the government's budget.

growth. (In the model, the size of the median firm is only 20% of average size, and the smaller half of firms accounts for only about 5% of employment. This is similar in U.S. data.) In contrast, the effect of taxes on growth is particularly high when taxes kick in and reduce experimentation in a part of the firm size distribution that accounts for an important portion of output.

4.3 Distorting the experimentation technology

Another type of distortion that firms may face are limits to their ability to adjust their technology. While the taxes considered above have similar effects to limits to experimenting, another margin also matters: the ability to undo the outcomes of failed experiments. Whereas in the case of failed R&D projects from the lab it is sufficient not to pursue the project further, different problems can arise when productivity is promoted through experimentation. For instance, unions can make it difficult to change work arrangements, government regulation may make it more costly to adjust factors, or contract incompleteness may limit to which degree it is feasible or worthwhile to phase out failed experiments. In addition, technology or reputation may make it hard to fully undo failed experiments. In general, any factor making adjustments difficult or costly may negatively affect the ability to undo failed experiments. We quantify the effect of such frictions modelled in a very general way by considering limits to the reversibility of

Table 5: Irreversibility

	Irreversibility (λ)	
	0%	5%
optimal choice of s:		
embodied productivity growth	1.20%	0.97%
s fixed at the benchmark level:		
embodied productivity growth	1.20%	1.10%

Notes: No taxes. λ is the degree of irreversibility of an experiment's outcome as defined in Section 2.3.2.

experiments, i.e. $\lambda > 0$, where λ is as defined in Section 2.3.2 on imperfect reversibility.

Table 5 shows that the implications of irreversibility are quite large: if only 95% of the effect of an experiment on productivity can be reversed ($\lambda = 0.05$), annual productivity growth decreases by 0.2 percentage points. The second line of the table shows results keeping firms' choice of s at the benchmark level (with full reversibility). This shows that about half the reduction in growth brought about by less than full reversibility is due to less experimentation, and half due to the fact that imperfect irreversibility makes the experimentation technology less productive for any choice of s. Reversibility thus plays an important role in generating growth in our framework.

5 Conclusion

Firms experiment purposefully. Their experiments can induce fluctuations in observed productivity. They also are an engine of growth, since results of failed experiments can be discarded.

We have proposed a simple model of experimentation by firms, which provides a very simple micro-foundation to (part of) the stochastic process for firm-level productivity typically specified in the macroeconomic literature with firm heterogeneity. The implied process for firm-level productivity is theoretically appealing: positive shocks, which may be due to the firm's own purposeful experimentation behavior, may be permanent, whereas negative shocks may be reversed.

Integrated in a realistic model of firm dynamics with productivity fluctuations, entry and exit, our model of experimentation suggests that 36% of aggregate embodied productivity growth in the U.S. can be attributed to experimentation. Firms' investment in experimentation is somewhat larger than recorded investment in R&D, which we consider a specific, more targeted form of experimentation.

Our setting is ideal for analyzing the impact of distortions on growth. First, the ability to reverse failed experiments is crucial for growth. Aggregate growth declines substantially even if institutions or adjustment frictions prevent reversing only a small part of failed experiments. Secondly, we show that size-dependent distortions in the form of higher taxes or stricter regulation for more productive firms can strongly reduce the incentive to innovate and thereby reduce growth, with already a 10% tax on the 5% most productive firms reducing aggregate growth rate by almost one quarter of a percentage point.

References

Acemoglu, D., K. Bimpikis, and A. Ozdaglar (2011): "Experimentation, Patents, and Innovation," *American Economic Journal: Microeconomics*, 3(1).

Baldwin, J. R. (1995): *The Dynamics of Industrial Competition.* Cambridge University Press, Cambridge.

Barseghyan, L. (2008): "Entry Costs and Cross-Country Differences in Productivity and Output," *Journal of Economic Growth*, 13(2), 145–167.

Bartelsman, E., J. Haltiwanger, and S. Scarpetta (2004): "Microeconomic Evidence of Creative Destruction in Industrial and Developing Countries," *World Bank Policy Research Working Paper*, 3464.

Bartelsman, E. J., E. Perotti, and S. Scarpetta (2008): "Barriers to Exit, Experimentation and Comparative Advantage," *mimeo, Tinbergen Institute.*

Bernard, A. B., S. J. Redding, and P. K. Schott (2010): "Multiple-Product Firms and Product Switching," *American Economic Review*, 100(1), 70–97.

Bolton, P., and C. Harris (1999): "Strategic experimentation," *Econometrica*, 67(2), 349–374.

Broda, C., and D. E. Weinstein (2010): "Product creation and destruction: Evidence and price implications," *American Economic Review*, 100(3), 691–723.

Campbell, J. Y., M. Lettau, B. G. Malkiel, and Y. Xu (2001): "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk," *Journal of Finance*, 56(1), 1–43.

Castro, R., G. L. Clementi, and Y. Lee (2009): "Cross-sectoral variation in firm-level idiosyncratic risk," *Federal Reserve Bank of Cleveland Working Paper*, 08-12.

Caves, R. E. (1998): "Industrial Organization and New Findings on the Turnover and Mobility of Firms," *Journal of Economic Literature*, 36(4), 1947–1982.

Coles, J. L., N. D. Daniel, and L. Naveen (2006): "Managerial incentives and risk-taking," *Journal of Financial Economics*, 79(2), 431–468.

Davis, S. J., R. J. Faberman, J. Haltiwanger, R. Jarmin, and J. Miranda (2008): "Business Volatility, Job Destruction and Unemployment," *NBER Working Paper*, 14300.

Doraszelski, U., and J. Jaumandreu (2009): "R&D and productivity: Estimating endogenous productivity," *mimeo, Harvard.*

Dwyer, D. W. (1998): "Technology Locks, Creative Destruction, and Nonconvergence in Productivity Levels," *Review of Economic Dynamics*, 1(2), 430–473.

Foster, L., J. Haltiwanger, and C. J. Krizan (2001): "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," in *New Developments in Productivity Analysis*, ed. by C. R. Hulten, E. R. Dean, and M. J. Harper, NBER Studies in Income and Wealth. University of Chicago Press, Chicago.

——— (2006): "Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s," *Review of Economics and Statistics*, 88(4).

Francois, P., and H. Lloyd-Ellis (2003): "Animal Spirits through Creative Destruction," *American Economic Review*, 93(3), 530–550.

Gabaix, X. (2010): "The Granular Origins of Aggregate Fluctuations," *mimeo, NYU.*

Gabler, A., and O. Licandro (2007): "Endogenous Growth through Selection and Imitation," *European University Institute Economics Working Paper*, 2007/26.

Geroski, P. A. (1995): "What do we know about entry?," *International Journal of Industrial Organization*, 13(4), 421–440.

Greenwood, J., Z. Hercowitz, and P. Krusell (1997): "Long-run implications of investment-specific technological change," *American Economic Review*, 87(3), 342–362.

Guner, N., G. Ventura, and Y. Xu (2008): "Macroeconomic implications of size-dependent policies," *Review of Economic Dynamics*, 11(4), 721 – 744.

Haltiwanger, J., R. Jarmin, and T. Schank (2003): “Productivity, Investment in ICT and Market Experimentation: Micro Evidence from Germany and the U.S.,” *U.S. Bureau of the Census Center for Economic Studies (CES) Research Paper*, 03-06.

Henly, S. E., and J. M. Sánchez (2009): “The U.S. Establishment-Size Distribution: Secular Changes and Sectoral Decomposition,” *Federal Reserve Bank of Richmond Economic Quarterly*, 95(4), 419–454.

Holmes, T. J., D. K. Levine, and J. A. Schmitz (2008): “Monopoly and the Incentive to innovate when adoption involves switchover disruptions,” *NBER Working Paper*, 13864.

Hopenhayn, H. (1992): “Entry, Exit, and Firm Dynamics in Long Run Equilibrium,” *Econometrica*, 60(5), 1127–1150.

Hopenhayn, H., and R. Rogerson (1993): “Job Turnover and Policy Evaluation: A General Equilibrium Analysis,” *Journal of Political Economy*, 101(5), 915–938.

Hsieh, C.-T., and P. J. Klenow (2009): “Misallocation and Manufacturing TFP in China and India,” *Quarterly Journal of Economics*, 124, 1403–1448.

Keller, G., S. Rady, and M. Cripps (2005): “Strategic experimentation with exponential bandits,” *Econometrica*, 73(1), 39–68.

Klette, T. J., and S. Kortum (2004): “Innovating Firms and Aggregate Innovation,” *Journal of Political Economy*, 62, 986–1018.

Kortum, S. (1997): “Research, patenting, and technological change,” *Econometrica*, 65(6), 1389–1419.

Lentz, R., and D. T. Mortensen (2008): "An empirical model of growth through product innovation," *Econometrica*, 76(6), 1317–1373.

Luttmer, E. (2007): "Selection, Growth and the Size Distribution of Firms," *Quarterly Journal of Economics*, 122(3), 1103–1144.

McGuckin, R. H., M. L. Streitwieser, and M. E. Doms (1996): "The Effect of Technology Use on Productivity Growth," *U.S. Bureau of the Census Center for Economic Studies (CES) Working Paper*, CES 96-2.

Midrigan, V., and D. Xu (2010): "Finance and Misallocation: Evidence from Plant-level Data," *mimeo, NYU*.

Moscoso Boedo, H. J., and T. Mukoyama (2010): "Evaluating the Effects of Entry Regulations and Firing Costs on International Income Differences," *mimeo, U Virginia*.

Poschke, M. (2009): "Employment protection, firm selection, and growth," *Journal of Monetary Economics*, 56(8), 1074–1085.

——— (2010): "The regulation of entry and aggregate productivity," *The Economic Journal*, 120(549), 1175–1200.

Restuccia, D., and R. Rogerson (2008): "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments," *Review of Economic Dynamics*, 11(4), 707–720.

Samaniego, R. M. (2006): "Employment protection and high-tech aversion," *Review of Economic Dynamics*, 9(2), 224–241.

Schivardi, F., and R. Torrini (2008): "Identifying the effects of firing restrictions through size-contingent differences in regulation," *Labour Economics*, 15(3), 482 – 511.

Sutton, J. (1997): "Gibrat's Legacy," *Journal of Economic Literature*, 35(1), 40–50.

Syverson, C. (forthcoming): "What determines Productivity?," *Journal of Economic Literature*.

Thomke, S. H. (2003): *Experimentation Matters: Unlocking the Potential of New Technologies for Innovation*. Harvard Business Press, Cambridge, MA.