Experimentation by Firms, Distortions, and Aggregate Productivity[a]

Alain Gabler[b] Markus Poschke[c]

Swiss National Bank McGill University

October 4, 2012

Abstract

Recent empirical research has documented that distortions of allocative efficiency among heterogeneous firms can have large aggregate consequences. This paper evaluates the size of these effects when distortions affect not only resource allocation but also the evolution of firm level productivity itself. To this end, we partially endogenize the evolution of firm level productivity in a standard heterogeneous firm model by allowing firms to engage in costly, purposeful experimentation: Firms can engage in risky experiments, which take the form of productivity shocks. Results from failed experiments can be discarded. We then show that endogenous productivity implies up to twice as large effects of productivity-dependent distortions on aggregate consumption.

JEL codes: E24, L16, O40.

Keywords: aggregate productivity, firm dynamics, selection, experimentation, entry and exit.

---

[a]We are grateful for comments from the editor, three anonymous referees, and seminar participants at meetings of the Society of Economic Dynamics and the Canadian Economic Association and at the 8th Meeting of German Economists Abroad. Poschke thanks the Social Sciences and Humanities Research Council (SSHRC) and the Fonds québécois de la recherche sur la société et la culture (FQRSC) for research support. 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. email: markus.poschke@mcgill.ca.

1 Introduction

Recent empirical research has documented 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.

These papers have focussed on the impact of distortions on the allocation of inputs across heterogeneous production units, taking either the distribution of productivity or a stochastic process for productivity as given. This abstracts from a potentially important channel: If firms can affect the evolution of their productivity, then distortions can influence choices along this margin and through this, firm level productivity growth and aggregate productivity. Their effect could thus go beyond resource misallocation. This paper quantitatively assesses the importance of this channel.[1]

To do so, we model firms as using purposeful experimentation to promote their productivity. 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 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

[1]Bhattacharya, Guner, and Ventura (2011) pursue a similar objective and quantify the effect of distortions when managerial skills are endogenous. Further related papers are Bello, Blyde, and Restuccia (2011), Hsieh and Klenow (2012), Ranasinghe (2011; 2012) and Restuccia (2011).

expose themselves to "productivity risk" in order to improve their productivity, but can control the extent of this risk by choosing their experiments.

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.[2] (We also allow for additional shocks the firm cannot influence.) Firms can choose how risky 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.[3] Because of this option (think of not implementing R&D findings or pulling an unsuccessful new product off the market), the expected value of experimenting is positive and increases in the riskiness of the experiment. This is balanced by a higher cost of conducting risky experiments compared to more incremental/marginal ones so that in equilibrium, firms choose experiments with limited risk. We integrate this experimentation process into a full quantitative model of firm dynamics with endogenous entry and exit.[4]

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. Yet, our approach also allows capturing the

[2]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.

[3]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).

[4]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 bandit problems that correspond to the choice between discrete projects.

activity 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 (2011).)

We then calibrate our model using information on firm dynamics in the United States to quantify the effect of distortions on aggregate outcomes. We consider two types of distortions. First, we analyze the effect of productivity dependent distortions. These have been identified in recent research (Guner, Ventura, and Xu, 2008; Restuccia and Rogerson, 2008) as particularly damaging to aggregate productivity because the resulting misallocation of resources is not random, but specifically directs resources from high to low-productivity producers. Secondly, we assess the effect of firing costs, which hamper the efficient reallocation of resources across production units and thereby reduce aggregate productivity (Hopenhayn and Rogerson, 1993).

We find that productivity dependent distortions strongly affect aggregate outcomes. For example, a tax that is linear in productivity with slope and intercept such that the median firm pays no tax, but the 0.1% most productive firms face a 10% tax, reduces aggregate consumption by 2.1%. This large change occurs although on average, firms face a tax rate of 0. More than half of the change is due to the effect of reduced experimentation. Increasing the slope such that top firms face a 20% tax rate reduces experimentation further and reduces consumption by 3.5% compared to an undistorted economy. The productivity dependent specification is key here: Productivity dependent distortions discourage firms from rising to the top, as they would be subject to larger distortions. In contrast, shifting the tax schedule without changing its slope only affects output and consumption by discouraging capital accumulation, but has hardly any effect on experimentation.

We also quantify the effect of firing costs. Compared to productivity dependent distortions, they turn out to have a more limited effect on experimentation. The reason is that, while they discourage some firms from experimenting, they encourage others to experiment more in order to reduce the risk of having to engage in costly layoffs in the future.

Endogenous productivity, here captured through experimentation, thus amplifies the effect of distortions on aggregate productivity substantially. The paper is organized as follows. Section 2 presents a model of firm dynamics with endogenous experimentation. Section 3 presents the calibration. The effects of distortions are described in Section 4, and Section 5 concludes.

2 A Model of Endogenous Experimentation

In this section, we present a model in which firms can influence the evolution of their productivity through experimentation. We model this process of experimentation by assuming that firms are hit by idiosyncratic productivity shocks whose variance they can choose. We interpret these shocks as random outcomes of experiments through which firms try to improve their productivity. Firms can choose how risky an experiment is; this is reflected in the variance of the shock they receive. If the result of a firm's experiment is not as desired, the firm can discard the experiment and revert to the productivity it had before undertaking the experiment. We assume that experimentation is costly in terms of current output.

Because there is free entry and firms have the possibility of exiting, the measure of active goods-producing firms in the economy is endogenous. In addition to these firms, there also is a representative household and 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 and by renting capital to them via a sector of perfectly competitive financial intermediaries. Capital depreciates at a rate δ. In equilibrium, financial 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

$$c_t^{-\sigma} = \beta (1 + r_{t+1}) c_{t+1}^{-\sigma}. \tag{1}$$

2.2 The Problem of the Firm

A given firm i produces output with the production function

$$y = [e^z \cdot \theta(\sigma_\varepsilon)]^{1-\alpha-\gamma} l^\alpha k^\gamma \tag{2}$$

and sells it in a competitive market. We normalize the price of output to 1. Firms differ in their productivity z and choose their labor and capital inputs l and k, which have user costs of w and $R = r + \delta$ per unit, respectively. The term $\theta(\sigma_\varepsilon)$ represents a disruption cost of experimentation, which depends on the experimentation intensity chosen by the firm, σ_ε. 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 and capital inputs given productivity z and a choice of σ_ε, profits are

$$\Pi(z, \sigma_\varepsilon) = (1 - \alpha - \gamma) \left(\frac{\gamma}{R}\right)^{\frac{\gamma}{1-\alpha-\gamma}} \left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha-\gamma}} e^z \theta(\sigma_\varepsilon). \tag{3}$$

A firm's productivity evolves stochastically and is driven both by exogenous shocks and by the outcomes of the firm's experiments. The law of motion of productivity is

$$z' = z + \max(\varepsilon, 0) + u.$$

Here, the innovations u are drawn from a *cdf* $H(u)$ and represent any perturbations to a firm's productivity which are independent from experimentation and over which the firm has no control, such as changes in customers' tastes or exogenous shocks to its technology.

Experimentation is captured in the innovation $\max(\varepsilon, 0)$. Firms can conduct one experiment per period. This consists in drawing an innovation ε from a distribution with *cdf* $\Phi_{\sigma_\varepsilon}(\varepsilon)$. The max operator in the law of motion for productivity reflects that firms can discard the results of unsuccessful experiments. Draws of $\varepsilon < 0$ would imply decreased productivity if the result of the experiment was adopted. Firms can forego results of such failed experiments and instead employ their previously used technology. In this case, they still suffer the disruption cost of the experiment, but can avoid the impact of the experiment on productivity.

The key decision involved in experimentation is that of how much risk to take.

This choice is represented by σ_ε, which determines the variance of ε.[5] Since firms can conduct one experiment per period, they can adjust σ_ε every period. On the one hand, riskier experiments have a larger expected impact on productivity because of the option to reject failures. On the other hand, choosing riskier experiments is costly as it is disruptive of current production: we assume that the disruption cost function $\theta(\sigma_\varepsilon)$ is continuously differentiable for $\sigma_\varepsilon \geq 0$ and that $\theta'(\sigma_\varepsilon) < 0$ and $\theta''(\sigma_\varepsilon) < 0$, so that the cost of experimenting is convex.

We assume that larger experiments are more disruptive because they may be more difficult and costly to implement: In the case of process innovation, fundamentally changing the way a factory works bears a larger potential for productivity changes (improvements if the experiment is successful), but is also more costly to put into practice. In the case of product innovation, introducing a completely new product may also necessitate changes in the production process, and in addition may be more costly in terms of promotional activity. In the case of formal R&D, one can argue that more (potentially) ground-breaking research is more expensive.

While formal R&D can be viewed as a form of experimentation in our model, we interpret experimentation more broadly as encompassing other activities aimed at enhancing productivity. These may be much less formal than R&D. Also, our setup implies that the technological advances that are generated within the model are embodied: an individual firm's productivity gains deriving from its experimentation are tied to its production facilities and do not spill over to other active firms.

Firm value and the exit decision. Continuing firms need to pay a fixed operating cost of κ_f per period to produce. Firms have the option to avoid this by permanently

[5] We assume that experimentation intensity σ_ε has no effect on e^ε, the expectation of gross productivity growth due to experimentation before the accept/reject decision. Since firms can discard failed experiments, riskier experiments do of course have a larger expected benefit if they are successful.

exiting. They may do so after learning their realizations of ε and u. Finally, there is a probability χ every period that a firm has to exit exogenously.

The value of a firm is

$$V(z) = \max_{\sigma_\varepsilon} \left\{ \Pi(z, \sigma_\varepsilon) - \kappa_f + \frac{1-\chi}{1+r} \mathbb{E}_{\sigma_\varepsilon} \max\left[V(z + u + \max(\varepsilon, 0)), 0 \right] \right\}. \qquad (4)$$

This value function embodies three decisions by the firm. The first max operator requires optimal choice of σ_ε, the last one optimal acceptance or rejection of the experimental outcome, and the middle one 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 σ_ε indicates that the distribution of ε depends on the firm's choice of σ_ε.

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

$$V(z_x) = 0.$$

The optimal acceptance/rejection decision for experimental outcomes implies accepting if $\varepsilon > 0$. These decisions combined with the exogenous shocks u and χ 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 the good. New entrants draw their initial level of productivity from a distribution $\eta(z)$. Entry is optimal as long as its value exceeds 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

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

Following Luttmer (2012), a stationary distribution μ exists as long as the exogenous exit rate χ is large enough.

Productivity and optimal experimentation. Figure 2.2 shows optimal experimentation as a function of z, computed for the benchmark equilibrium implied by the calibration described in Section 3 below. Optimal σ_ε is zero for firms that face a probability of exiting at the end of the period that is close to 1, since they hardly benefit from experimenting. As z rises and the exit probability declines, the expected benefit of experimentation rises, and firms experiment more. Once the exit probability becomes negligible, optimal σ_ε becomes flat. Our computations suggest that this occurs for the following reason: Since the profit function is linear in z, the computed value function is also linear in z for firms that are so productive that the option to exit has no value. In this range, higher z then affects the cost and expected benefit of experimentation in the same way for all z. As a consequence, the optimal choice of σ_ε becomes independent of z in this range. Large firms spend a constant fraction of output on experimentation.

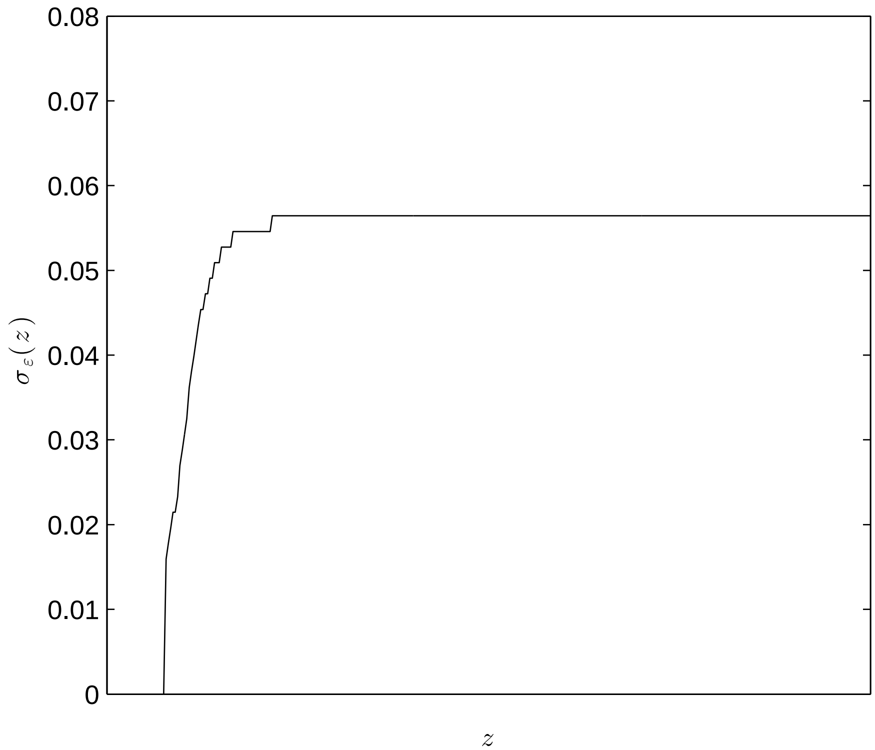

Figure 1: The experimentation policy

Notes: Parameters are from the calibration in Section 3 and are given in Table 2 below.

3 Calibration

In this section, we calibrate the model to U.S. data. The calibrated model then allows us to measure the importance of experimentation and to gauge 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 fit a set of data moments. We pay particular attention to match some static and dynamic features of the U.S. firm size distribution. Matching these is crucial for a realistic assessment of the effect of distortions, since these affect the firm size distribution.

The length of a time period is set to one year. The parameters which are set based on a-priori information are the production function parameters α and γ, the discount factor β, and the depreciation rate δ. The production function parameters α and γ are very important, as they control how much distortions affect a firm's input and output choices. Existing evidence suggests that $1-\alpha-\gamma$, which also is the profit share in output, lies between 0.10 and 0.20. (See e.g. Atkeson, Khan, and Ohanian (1996) or Pavcnik (2002).) We therefore set $1-\alpha-\gamma$ to 0.15.[6] Splitting the remainder into roughly 70% labor income and 30% capital income implies α of 0.6 and γ of 0.25.

In a stationary equilibrium of the model economy, β and δ only matter in so far as they determine the real interest rate and the user cost of capital. We target a real interest rate of 4 percent, implying a β of 0.96. We choose δ of 0.08, which implies a user cost of capital $R = 1/\beta - 1 + \delta$ of 12%. The resulting investment to output ratio of about 16.5% is close to the observed U.S. value for private fixed investment. The implied capital to output ratio is 2.1. As in the data, this number excludes investment in experimentation and firm start-up costs.

We assume that the disruption cost function takes the form

$$\theta\left(\sigma_{\varepsilon}\right) = \begin{cases} \left(\bar{\sigma}_{\varepsilon} - \sigma_{\varepsilon}\right)^{q} & \text{for } \sigma_{\varepsilon} \leq \bar{\sigma}_{\varepsilon} \\ 0 & \text{otherwise,} \end{cases} \tag{5}$$

with parameters $q \in (0,1)$ and $\bar{\sigma}_{\varepsilon} > 0$. Conditional on q, $\bar{\sigma}_{\varepsilon}$ is easy to identify, but we could not find a good target for q. To deal with this, we assume that the disruption cost function is inverse quadratic, i.e. $q = 0.5$, calibrate $\bar{\sigma}_{\varepsilon}$, and report the results of robustness checks in the Appendix. We assume that both the experimentation shock ε and the exogenous shock u are distributed normally, with mean $-\sigma_{\varepsilon}^{2}/2$ and standard deviation σ_{ε} for the experimentation shock, and with mean $-\sigma_{u}^{2}/2$ and standard devi-

[6] In our model, this is the profit share gross of fixed operating costs. Since fixed operating costs are small in the calibrated version (0.5% of aggregate output), this is close to net profits.

ation σ_u for the exogenous shock. This implies that $\mathbb{E}(e^u) = \mathbb{E}(e^\varepsilon) = 1$: without the option to discard results of failed experiments, expected productivity would equal current productivity. Finally, z for new firms is distributed normally with mean $\phi - \sigma_n^2/2$ and standard deviation σ_n. We normalize ϕ to 1.

We set the remaining six parameters $\bar{\sigma}_\varepsilon, \sigma_n, \sigma_u, \chi, \kappa_f$ and κ_e to minimize the distance to six informative target moments observed in U.S. data. We next suggest which targets are particularly informative for which parameter.

How much firms choose to experiment depends crucially on the cost parameter $\bar{\sigma}_\varepsilon$. More success in experimentation implies more firm growth and a more skewed firm size distribution. The share of employment in the 5% largest firms is therefore informative about $\bar{\sigma}_\varepsilon$. (We obtain this and other statistics related to the U.S. firm size distribution from the U.S. Census Bureau's Statistics of U.S. Businesses (SUSB) for 2004.)[7] The job turnover rate – 22.5% per year in the U.S. according to Haltiwanger, Scarpetta, and Schweiger (2008) – is highly informative about σ_u. Importantly, because of the option to discard failures, experimentation affects job turnover far less than exogenous shocks do. Therefore, targeting the size distribution and the job turnover rate allows us to separate exogenous shocks from experimentation.

We obtain information on the parameters driving exit, χ and κ_f, from the exit rate of firms (10% per year according to Bartelsman, Haltiwanger, and Scarpetta, 2009) and the amount of job turnover due to exit, which is about 3% per year (Haltiwanger, Scarpetta, and Schweiger, 2008). The entry cost κ_e affects the number of active firms and thus average firm size, so we also target average firm employment of 20. Finally, we obtain information about σ_n from the share of employment in small firms, which often are young. In the U.S., firms below average size account for 16% of employment.

[7] Results of course depend to some extent on distributional choices. For instance, modelling the exogenous shocks as following a distribution with more mass in the upper tail could imply that experimentation is less important. A more skewed distribution for outcomes of experiments may also be an interesting alternative.

Table 1: Calibration: Model statistics, calibration targets

Statistic	Data	Model
calibration targets:		
share of employment in		
5% largest firms	73.7%	73.6%
firms below average size	16.0%	16.0%
job turnover rate	22.5%	22.6%
job turnover due to exit	3.0%	4.4%
firm exit rate	10.0%	10.2%
average firm size	20	20
not used in calibration:		
share of employment in		
14% largest firms	84.7%	86.4%
39% largest firms	95.0%	95.7%
fraction of firms below average size	87.0%	88.3%

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 reasonably well even in dimensions that were not targeted. In particular, the model fits the distribution of employment across firms very well: different measures of the share of employment in large firms as well as the fraction of employment in small firms are very close to their counterparts in the Census SUSB data. In addition, the model also fits the prevalence of firms below average size very well. The only target that fits less well is the job turnover rate due to exit. Given that some sources report numbers for this target of up to 6.3% p.a. (see e.g. Davis, Faberman, Haltiwanger, Jarmin, and Miranda, 2008, Table 2), the overall fit is acceptable.

Calibrated parameter values are reasonable: the productivity of entrants features large dispersion. The variance of exogenous productivity shocks is substantial, but a

Table 2: Calibration: Parameter values

Parameter	Value	Parameter	Value
$\bar{\sigma}_\varepsilon$	0.124	α	0.6
σ_u	0.213	γ	0.25
σ_n	2.125	β	0.96
χ	0.043	δ	0.08
κ_f	0.101	q	0.5
κ_e	7.990	ϕ	1

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.[8] The largest feasible experiment (setting $\sigma_\varepsilon = \bar{\sigma}_\varepsilon$) would involve a standard deviation of the shock a bit more than half as large as that of the exogenous shock. However, firms choose to invest significantly less in experimentation, choosing σ_ε of 0.046 on average. Highly productive firms are the ones experimenting most, choosing σ_ε of 0.055. Overall, firms invest 3.7% of output in experimentation. In line with our broader interpretation, this number is somewhat larger than reported R&D investment for the US of about 2.7% of GDP in 2007 (World Bank World Development Indicators).[9]

Although innovations due to experimentation may appear small, in particular compared to the variance of the exogenous shock, they make a substantial contribution to productivity: In the calibrated economy, aggregate consumption is 4% larger than in an otherwise identical economy without experimentation. The asymmetry of shocks due to experimentation – unsuccessful experiments can be discarded – implies that ex-

[8]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 close to proportional to productivity, so the standard deviation of productivity growth is close to that of stock returns.

[9]Recalibrating the model with different values of q for the robustness exercises does not yield very different values for the choice of σ_ε. The figure for investment in experimentation changes a bit; it is 3.9% of output for q of 0.25 and 3.1% for q of 0.75. Given the broader interpretation of experimentation compared to R&D, this suggests that q is probably smaller than 0.75.

perimentation raises aggregate productivity. Also, experimentation has a large impact on the expected growth rate of output for an individual firm: a firm that sets σ_ε at the average optimal level of 0.046 expects to grow at a rate of 1.8% due to experimentation. This compares to a zero expected growth rate for non-experimenting firms (since the mean of the exogenous disturbance e^u is one by construction) and a maximum growth rate of about 5% when $\sigma \uparrow \bar{\sigma}_\varepsilon$.[10]

4 Distortions and productivity

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 analyze the impact of a variety of types of distortions on productivity and output. When productivity is endogenous due to experimentation, distortions affect not only the allocation of resources, but also the evolution of firm level productivity itself. We pay particular attention to this effect. To do so, we first analyze aggregate and productivity-dependent distortions, which we model in a very general way as taxes on output, following Restuccia and Rogerson (2008). Then we consider firing costs, a policy that makes it more costly to adjust factors.

4.1 Aggregate distortions

First, consider the effect of aggregate distortions. At optimal input choice, profits of a firm with productivity z and experimentation intensity σ_ε that is subject to an output tax at rate τ are

$$\Pi = (1-\tau)^{\frac{1}{1-\alpha-\gamma}}(1-\alpha-\gamma)\left(\frac{\gamma}{R}\right)^{\frac{\gamma}{1-\alpha-\gamma}}\left(\frac{\alpha}{w}\right)^{\frac{\alpha}{1-\alpha-\gamma}} e^z\theta(\sigma_\varepsilon).$$

[10]Recall though that output goes to zero as $\sigma_\varepsilon \uparrow \bar{\sigma}_\varepsilon$ since $\theta\left(\bar{\sigma}_\varepsilon\right) = 0$.

Denoting the three middle terms by c_0, the first order condition for optimal choice of σ_ε then is

$$-c_0(1-\tau)^{\frac{1}{1-\alpha-\gamma}} e^z \theta'(\sigma_\varepsilon) = \frac{1-\chi}{1+r} \frac{\partial \mathbb{E}_{\sigma_\varepsilon} \max[V(z+u+\max(\varepsilon,0)),0]}{\partial \sigma_\varepsilon}. \tag{6}$$

If the tax rate is constant, it affects current and future profits in the same way. As a result, taxes affect the cost of experimentation in terms of reduced current profits (left-hand side of equation (6)) and the benefits of experimentation in terms of higher expected future profits (right-hand side of (6)) in the same way. Because of this, a uniform output tax on all firms does not affect the level of experimentation in our model economy in a significant way.[11] It will however have a negative impact on capital accumulation, leading to a lower level of output.

4.2 Productivity-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). Another example is the "growth tax" in India applied to revenue beyond a certain level as documented in Little, Mazumdar, and Page (1987).

Work that studies such size-dependent policies typically has analyzed their impact on incentives to adjust factors, e.g. by discouraging firms from crossing a certain employment threshold. However, the distortions have a deeper effect that goes beyond this: they may discourage firms from becoming so productive that it would be desir-

[11] This is strictly true if taxes do not affect the exit decision. In our model economy, the effect of an aggregate tax on exit is small, and experimentation hardly changes.

able to cross the regulation threshold in the first place. Our framework is well-suited to analyzing this effect. It is clear from equation (6) that higher taxes on more productive firms discourage experimentation by reducing its benefits.

In this section, we quantify the effect of productivity-related taxes. For simplicity, we assume that the tax rate depends only on a firm's current productivity z. In particular, we assume that taxes increase linearly in productivity e^z: $\tau(z) = \tau_0 + \tau_1 \exp(z)$. We then examine the effects of different tax regimes by varying both the progressivity of the tax rate and the overall tax burden. Because the parameters τ_0 and τ_1 are hard to interpret, we define tax regimes in terms of taxes on two types of firms: the firms with productivity corresponding to the median and the top 0.1 percentile in the benchmark economy.[12] Varying only the latter changes the slope τ_1 of the tax profile, while changing both rates by the same amount changes the intercept τ_0 but not the slope of the tax profile. We also assume that any net tax revenue is handed back lump-sum to households.

While productivity-dependent taxes may appear unrealistic (taxes are more likely to depend on observable characteristics, such as employment), they capture the essence of size-dependent taxes. In particular, they put the spotlight on the dynamic effect of taxes on productivity-promoting activities, while abstracting from strategic effects on input choice.

Table 3 shows the results for a tax rate for the median firm of zero, and maximum tax rates $\tau_{\max}$ of 10% and 20%, respectively. To assess the importance of experimentation, we consider two cases for each maximum tax rate: one where firms choose their experimentation intensity $\sigma_\varepsilon(z)$ optimally, and one where we restrict firms' experimentation policy to be the same as in the undistorted benchmark economy. Denote that policy by $\tilde{\sigma}_\varepsilon(z)$. In the second case, productivity dependent distortions affect aggregate

[12] To reduce sensitivity to the long tail of the productivity distribution, we assume that all firms in the top 0.1 percentile pay the maximum rate.

Table 3: Productivity-dependent distortions: aggregate outcomes relative to the benchmark economy

	$\tau_{\max} = 10\%$		$\tau_{\max} = 20\%$	
	optimal $\sigma_\varepsilon(z)$	benchmark $\tilde{\sigma}_\varepsilon(z)$	optimal $\sigma_\varepsilon(z)$	benchmark $\tilde{\sigma}_\varepsilon(z)$
Output	0.991	0.992	0.980	0.983
Consumption	0.979	0.995	0.964	0.985
Average firm employment	0.758	0.919	0.703	0.835
Average firm output	0.751	0.911	0.689	0.821
Number of firms	1.320	1.089	1.422	1.198
Average σ_ε	0.904	0.974	0.857	0.957
Firm exit rate	0.965	0.956	0.940	0.929

Notes: Firms are subject to a tax rate that is linear in productivity. The tax rate for the median firm is zero. $\tau_{\max}$ is the tax rate on firms with productivity corresponding to the top 0.1 percentile in the benchmark economy. $\tilde{\sigma}_\varepsilon(z)$ denotes the experimentation policy that is optimal in the benchmark economy. Since we assume that any net tax revenue is handed back lump-sum to the households, the reported values for output and consumption include net tax revenue. Output is reported net of fixed operating costs. All values are relative to outcomes in the undistorted benchmark equilibrium. For results with different values of q, see Table 6 in the Appendix.

productivity in two ways: firstly, they affect firms' input choice and thus the efficiency of resource allocation for a given productivity distribution, and secondly, they affect firms' entry and exit decisions, which contribute to shaping the productivity distribution. This is as in much of the existing literature on misallocation. When firms choose experimentation optimally, distortions also affect aggregate outcomes through their effect on firms' experimentation decisions. The difference between the columns with the optimal and the benchmark experimentation policies thus shows the importance of this new channel.

Results presented in the table show that productivity-dependent distortions lead to lower output and consumption. This occurs even in the case with the fixed experimentation policy $\tilde{\sigma}_\varepsilon(z)$, which combines the effect of distortions on resource allocation analyzed by Restuccia and Rogerson (2008) and their additional effect on the productivity distribution via the exit decision. With such distortions, firms with below-median productivity face a negative tax rate and can thus survive more easily. The consequence is less exit (the exit rate falls by 0.4 to 0.7 percentage points in the different scenarios) and therefore lower average productivity. Lower average productivity also tends to reduce average firm size, both in terms of output and employment. Aggregate output does not fall as much as average output because more firms enter to ensure labor market clearing – the distorted economy features more, smaller firms.[13] This is also why consumption falls more than output: the distorted economies spend relatively more resources on firm entry.[14] Note that even with a fixed experimentation policy, distortions lead to small changes in average experimentation because of their impact on the distribution. Since the "new survivors" don't experiment, average σ_ε is lower in the distorted economies.

The effects of productivity-dependent distortions are even larger when firms choose $\sigma_\varepsilon(z)$ optimally. Increasing tax rates discourage firms from aiming at high productivity, reduce experimentation, and thereby reduce average productivity and average size even more. This is very clear in Figure 4.2, which depicts the experimentation policies in the economies with and without distortions. Distortions strongly discourage firms above median size, which face positive and increasing tax rates, from experimenting. Their effect is slightly reduced for the top firms with $z > z_2$, which face a constant tax rate.

[13] There are many plausible extensions of the model in which the number of firms would respond less strongly. One such case arises if the cost of entry increases in the number of firms. Another possibility is heterogeneity of entrepreneurs. In this case, if the most able entrepreneurs enter first (as in Lucas, 1978), additional firms present in distorted economies are run by less able entrepreneurs. Bhattacharya, Guner, and Ventura (2011) focus on this channel in their analysis of the effect of distortions on aggregate productivity; Poschke (2011) provides a channel through which aggregate technology differences affect which entrepreneurs are active.

[14] The effects of distortions here are a bit larger than in Restuccia and Rogerson (2008) even before the adjustment of σ_ε because of endogenous exit, which implies that distortions do not only imply resource misallocation, but also worsen the productivity distribution by helping low-productivity firms to survive. The evolution of idiosyncratic productivity and endogenous exit are also the reasons why average firm size and the number of firms adjust here, in contrast to Restuccia and Rogerson (2008).

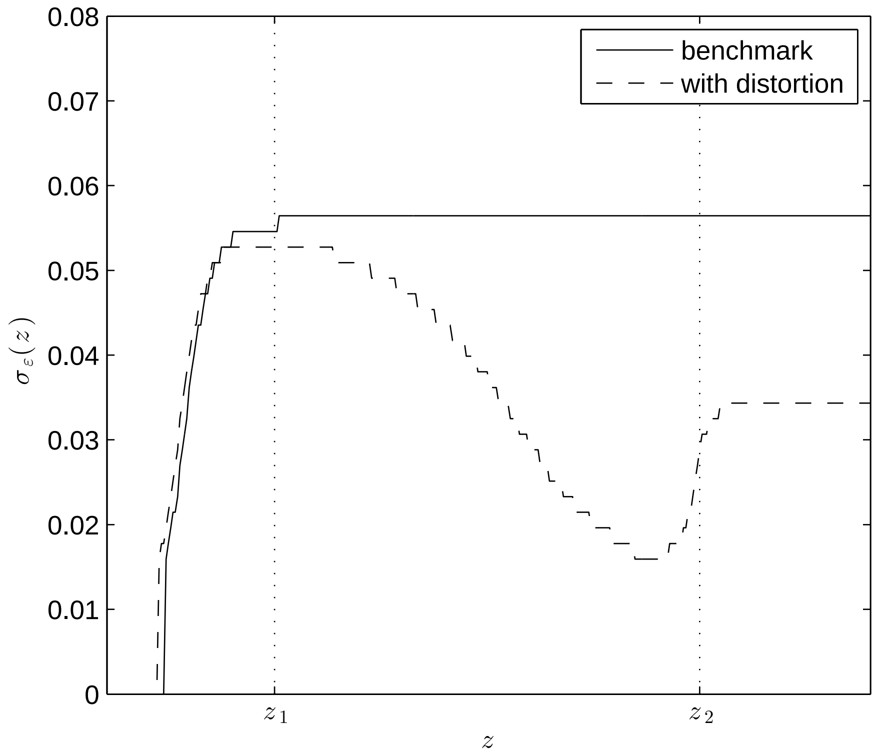

Figure 2: Optimal experimentation with and wihout productivity dependent distortions
Notes: Firms face distortions in the form of taxes on output given by the schedule $\tau(z) = \tau_0 + \tau_1 e^z$. The parameters τ_0 and τ_1 are such that the tax rate is 0 at z_1 and 0.1 at z_2. These productivity levels correspond to the median and the top 0.1 percent firm in the benchmark economy, respectively. The tax rate to the right of z_2 is also 0.1. Model parameters are given in Table 2.

Again, the decline in average productivity is counteracted by an increase in the number of firms – at a cost of more expenditure on establishing firms. As a consequence, reduced experimentation does not imply a much larger output response compared to the case where $\sigma_\varepsilon(z)$ is fixed. Consumption, in contrast, reacts much more strongly, falling by an additional 1.6 to 2.1 percentage points. The distortion of firms' experimentation decisions (the difference between optimal and benchmark $\sigma_\varepsilon(z)$) thus accounts for between 60-75% of the total reduction in aggregate consumption.

Increasing the maximum tax rate and thus the slope of the tax function $\tau(z)$ amplifies the effect of distortions on experimentation. Doubling the slope of the tax function

Table 4: Varying the overall tax burden: aggregate outcomes relative to the benchmark economy

	$\tau_{\text{median}} = 10\%$		$\tau_{\text{median}} = 20\%$	
	optimal $\sigma_\varepsilon(z)$	benchmark $\tilde{\sigma}_\varepsilon(z)$	optimal $\sigma_\varepsilon(z)$	benchmark $\tilde{\sigma}_\varepsilon(z)$
Output	0.923	0.924	0.851	0.855
Consumption	0.937	0.952	0.887	0.902
Average employment	0.892	1.081	1.078	1.296
Average output	0.823	0.999	0.918	1.107
Number of firms	0.925	1.121	0.927	0.772
Average σ_ε	0.896	0.972	0.886	0.970
Firm exit rate	0.952	0.961	0.952	0.948

Notes: Firms are subject to a tax rate that is linear in productivity. The tax rate for the median firm is given in the table. $\tau_{\max}$ is the tax rate on firms with productivity corresponding to the top 0.1 percentile in the benchmark economy and is $\tau_{\text{median}} + 0.1$. The slope of the tax function thus is the same in both columns. $\tilde{\sigma}_\varepsilon(z)$ denotes the experimentation policy that is optimal in the benchmark economy. Since we assume that any net tax revenue is handed back lump-sum to the households, the reported values for output and consumption include net tax revenue. Output is reported net of fixed operating costs. All values are relative to outcomes in the undistorted benchmark equilibrium. For results with different values of q, see Table 7 in the Appendix.

by raising the maximum tax from 10 to 20 percent implies that experimentation declines another 50%, output another 100% and consumption another 70%. The decrease in consumption due to experimentation increases from 1.6 to 2.1 percent points.

Table 4 shows the implications of varying the overall tax burden. This is done by changing the tax rate faced by the median firm, while keeping the slope of the tax function constant by setting the maximum tax rate 10% above the median rate. The additional loss in consumption from a 10% increase in the median tax rate ranges between 4.2 percentage points for an increase from 0 to 10% (results reported in Table 3) and 5 percentage points for an increase from 10% to 20%. As expected, the total effect of the distortion on consumption through the experimentation channel remains practically constant throughout: it is 1.6 percentage points when the median tax rate is zero, and 1.5 percentage points when it is 10 or 20 percent. Average σ_ε hardly changes. As indicated by (6), it is thus the slope of the tax function which matters for experimentation, not the average rate.

Finally, robustness checks shown in Tables 6 and 7 in the Appendix indicate that results are preserved when the model is recalibrated with different values of q. Optimal σ_ε is affected more (less) by distortions if q is larger (smaller) and the cost function less (more) curved. The effect of the distortion through the experimentation channel thus is larger (smaller) for larger (smaller) q. Still, the effect of productivity dependent distortion through endogenous productivity is substantial even when q is only 0.25.

4.3 Firing costs

The first source of distortions to be analyzed in a heterogeneous firm model were firing costs (Hopenhayn and Rogerson, 1993). What is their effect when firms can affect their productivity through experimentation? To answer this question, we analyze the effect of introducing firing costs of one year's wages in our benchmark economy. To focus only on the distortionary effect of firing costs, we assume that they are redistributed lump-sum to the representative household.[15]

As is well known, a firm's productivity and past employment are both state variables in this setting since a firm's past employment determines the cost of adjusting employment downwards (see e.g. Bentolila and Bertola, 1990). A firm's optimal hiring policy exhibits an inaction region. In this region, employment is not adjusted following small shocks, since this may trigger current or future firing cost payments. Employ-

[15] The level of firing costs we impose is close to empirically observed values: The 2012 cross-country average of the amount of severance pay due upon dismissal of a worker with tenure of 5 years is 11 months according to the World Bank's Doing Business project (`www.doingbusiness.org`; see also Botero, Djankov, La Porta, Lopez-De-Silanes, and Shleifer (2004)). In addition, there is a notice period of on average 5 weeks. The mandated severance payment is larger and the notice period longer for workers with longer tenure. In addition to these transfers, there often are substantial administrative costs. All these components contribute to the firing cost incurred by the firm.

Table 5: Firing costs: aggregate outcomes relative to the benchmark economy

	optimal $\sigma_\varepsilon(z)$	benchmark $\tilde{\sigma}_\varepsilon(z)$
Output	0.975	0.976
Consumption	0.960	0.965
Average firm employment	1.061	1.116
Average output	1.034	1.089
Number of firms	0.943	0.896
Average σ_ε	1.121	1.042

Notes: Firing costs of a year's wages. All proceeds are returned lump sum to consumers. Output is reported net of fixed operating costs. All values are relative to outcomes in the undistorted benchmark equilibrium.

ment is only adjusted once the benefit from doing so is sufficiently large compared to the expected penalty. As a consequence, the marginal product of labor is not equalized across firms, and there is resource misallocation of labor across firms.

This can be seen in the rightmost column in Table 5, which shows the consequences of introducing firing costs in an economy where firms use the benchmark experimentation policy. Because of the misallocation of labor they induce, firing costs reduce aggregate output and consumption. Average experimentation increases slightly because of changes in the productivity distribution of firms.

Differently from the previous section, aggregates are very similar with optimal experimentation. The reason for this is that unlike above, firing costs in this setting do not discourage experimentation, but on average slightly encourage it. Figure 3 shows the experimentation policy function in the benchmark economy (solid line) and for a firm of about five times average size in an economy with firing costs (dashed line). Patterns are qualitatively similar for firms of other sizes. z_1 denotes the largest level of productivity at which the probability of optimally exiting at the end of the period is 99% or more.

In the benchmark, optimal experimentation first increases with the survival probability (which raises the expected benefit from experimentation) and then becomes flat. With firing costs, the policy features two increasing and two flat parts. The difference to the benchmark is due to the inaction region in the employment policy, which lies between z_2 and z_3. These thresholds depend on past employment; they lie further to the right the larger past employment is. In this region, firms do not change their employment in response to small changes in productivity. As a consequence, firms in this region benefit less from small productivity increases and are thus discouraged from experimenting. This effect becomes weaker as we move closer to the right edge of the inaction region (z_3), where the probability of optimally moving out of that region, and thus the benefit from experimentation, is larger. Firms to the right (left) of the inaction region grow (shrink) and experiment at constant rates.

Compared to the benchmark economy, growing firms experiment more when there are firing costs. Our computations suggest that this occurs because firing costs make the value function for these firms slightly steeper in z. This happens because productivity losses are more costly when there are firing costs. As a consequence, firms experiment more to reduce the risk of costly layoffs. The increase in experimentation negatively affects aggregate consumption, which falls more with optimal experimentation than when experimentation is kept fixed at the benchmark policy $\bar{\sigma}_{\varepsilon}(z)$.

Still, the difference is small. The effects of firing costs shown in Table 5 are also somewhat smaller than those previously obtained in the literature (e.g. Hopenhayn and Rogerson, 1993). These differences can be attributed to some particular conservative modelling choices that tend to reduce the effect of firing costs. For instance, apart from experimentation, the present model differs from the one in Hopenhayn and Rogerson (1993) in three key respects: firms use both capital and labor in production; labor supply is inelastic; and z is not mean reverting. The presence of capital softens the

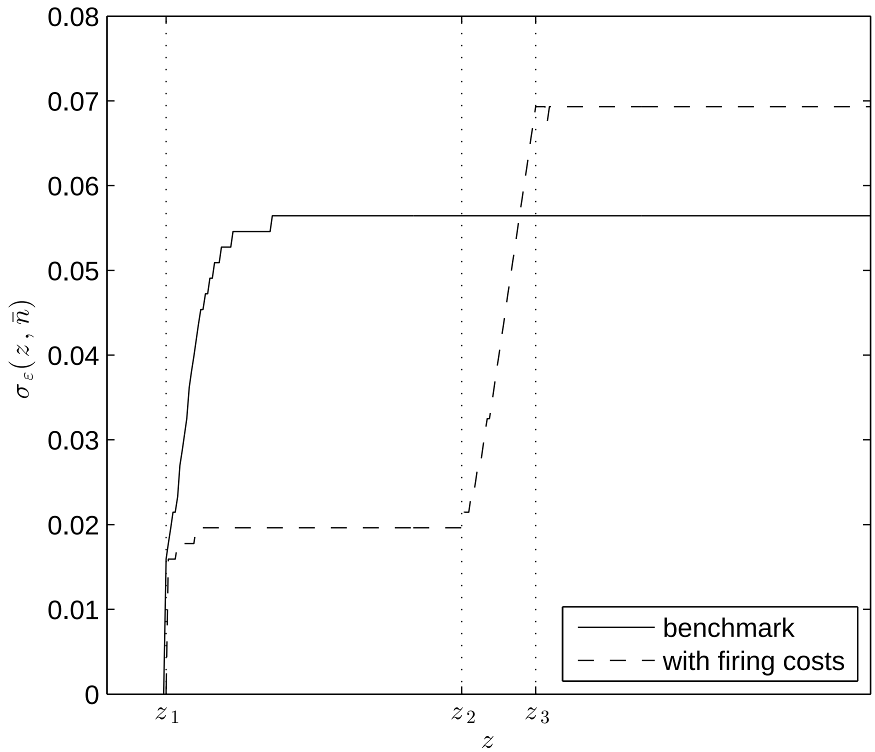

Figure 3: Optimal experimentation with and without firing costs

Notes: Without firing costs, optimal experimentation depends only on z. With firing costs, it also depends on past employment, which is denoted by $\bar{n}$ and here is fixed at a level of about 5 times average employment. z_1 marks the largest level of z at which the probability of choosing to exit at the end of the period exceeds 99%. z_2 and z_3 mark the left and right edge of the inaction region for employment conditional on past employment being $\bar{n}$. Raising (reducing) $\bar{n}$ shifts z_2, z_3 and the increasing part of the experimentation policy to the right (left). Parameters as in Table 2.

blow from firing costs because firms can adjust capital use to partially compensate for reduced adjustment of the labor input. The absence of labor supply reactions to lower wages implied by firing costs further reduces their aggregate consequences. Most importantly for our setting, firing costs would discourage experimentation much more if experimentation outcomes decayed over time, or if z mean reverted. If this was the case, firms would be more reluctant to adjust employment after a successful experiment, as this would involve higher expected future firing costs than with fully persistent z. Since the ability to adjust factors substantially enhances the benefits of successful experiments (the elasticity of firm output with respect to z alone is only 15% in our calibration), firing costs would have much more deleterious effects on experimentation in such a setting.

The results on firing costs illustrate that not all types of distortions reduce experimentation. They thus highlight again the particularly negative effect of productivity-related distortions. Ignoring their effect on firms' productivity-promoting activities may result in significant understatements of the damage they cause.

5 Conclusion

We have proposed a model of experimentation by firms, which provides a 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: successful experiments lead to a permanent increase in productivity, while unsuccessful ones can be reversed. Integrated within a realistic framework of firm dynamics with endogenous entry and exit, the calibrated version of our model replicates a number of moments of the U.S. firm size distribution.

We use this model to quantify the effect of distortions of allocative efficiency in the context of such a partially endogenous productivity process. We find that productivity dependent distortions strongly affect aggregate outcomes. For instance, a distortion that increases in productivity and, while not taxing the median firm, implies a 10% tax rate for the most productive firms implies a 2.1% reduction in aggregate consumption. More than half of this decline is due to reduced experimentation: Taxes that increase in productivity discourage firms from investing in productivity-enhancing activities. As a consequence, distortions do not just cause misallocation of resources across firms of different productivity, but may affect the location of the productivity distribution itself. We also examine the impact of firing costs. While these substantially reduce aggregate output and consumption by inducing misallocation, their effect on experimentation is more limited.

Endogenous productivity thus amplifies the effect of some, but not all distortions. We show this for productivity-dependent distortions – a very stylized way of modelling the heterogeneous incidence of distortions. Analyzing concrete size-dependent distortions (like size thresholds in regulations) may lead to more precise estimates but most likely would not change the basic thrust of our results. Future research on the role of endogenous productivity should, however, also consider the effects of a type of distortion that may have a particularly devastating effect in our setting: distortions that limit firms' flexibility and make it more difficult or costly for firms to undo failed experiments.

Similarly, there are three aspects related to modelling experimentation where further research could allow us to learn more about experimentation as well as to evaluate the effect of distortions more precisely. Firstly, we have assumed that failed experiments are instantaneously and fully reversible. Clearly, this is an extreme assumption that merits being relaxed. However, quantifying barriers to the reversibility of experiments is hard without more informative data. Therefore, secondly, it appears promising to bring to bear more empirical evidence on experimentation. The clearest existing empirical evidence on experimentation concerns the turnover of products (Bernard, Redding, and Schott, 2010). A full investigation should however also consider process innovation, on which there is less information, but which could be very different in terms of both costs and reversibility. Thirdly, both reversibility and other model features, like the variance of exogenous shocks or fixed costs, may well vary across industries. These differences would in turn imply differences in experimentation. Evidence on differences in productivity dynamics and size distributions across industries as documented by Rossi-Hansberg and Wright (2007) and Castro, Clementi, and Lee (2009) could thus also provide further insights into experimentation and into the costs of distortions across industries. We leave these issues for future research.

References

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

Atkeson, A., A. Khan, and L. Ohanian (1996): "Are Data on Industry Evolution and Gross Job Turnover Relevant For Macroeconomics?," *Carnegie-Rochester Conference Series on Public Policy*, 44, 216–250.

Bartelsman, E., J. Haltiwanger, and S. Scarpetta (2009): "Cross-Country Differences in Productivity: the Role of Allocation and Selection," *IZA Discussion Papers*.

Bello, O. D., J. S. Blyde, and D. Restuccia (2011): "Venezuela's Growth Experience," *Latin American Journal of Economics*, 48, 199–226.

Bentolila, S., and G. Bertola (1990): "Firing costs and labor demand: how bad is eurosclerosis?," *Review of Economic Studies*, 57, 381–402.

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

Bhattacharya, D., N. Guner, and G. Ventura (2011): "Distortions, Endogenous Managerial Skills and Productivity Differences," *mimeo*.

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

Botero, J. C., S. Djankov, R. La Porta, F. Lopez-De-Silanes, and A. Shleifer (2004): "The Regulation of Labor," *Quarterly Journal of Economics*, 119(4), 1339–1382.

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.

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.*

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

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

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

Haltiwanger, J., S. Scarpetta, and H. Schweiger (2008): "The Relationship Between Job Flows across Countries: The Role of Industry, Firm Size and Institutions," *NBER Working Paper*, 13920.

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.

——— (2012): "The Life Cycle of Plants in Intia and Mexico," *NBER Working Paper No. 18133*.

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.

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

Little, I. M. D., D. Mazumdar, and J. M. Page, Jr. (1987): *SmallManufacturing Enterprises: A Comparative Analysis of India and Other Economies*. Oxford University Press for the World Bank, New York, NY.

Lucas, R. E. (1978): "On the size distribution of business firms," *Bell Journal of Economics*, 9, 508–523.

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

——— (2012): "Technology Diffusion and Growth," *Journal of Economic Theory*, 147(2), 602–622.

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.

Pavcnik, N. (2002): "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants," *Review of Economic Studies*, 69, 245–276.

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

——— (2011): "The firm size distribution across countries and skill-biased change in entrepreneurial technology," *CIREQ Working Paper*, 08-2011.

Ranasinghe, A. (2011): "Impact of Policy Distortions on Plant-level Innovation, Productivity Dynamics and TFP," *SSRN 1833022*.

——— (2012): "Property Rights, Extortion and the Misallocation of Talent," *mimeo*.

Restuccia, D. (2011): "The Latin American Development Problem," *University of Toronto, Department of Economics Working Paper*.

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

Rossi-Hansberg, E., and M. L. J. Wright (2007): "Establishment Size Dynamics in the Aggregate Economy," *American Economic Review*, 97(5), 1639–1666.

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

Syverson, C. (2011): "What determines Productivity?," *Journal of Economic Literature*, 49(2), 326–365.

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

A Additional Tables and Figures

Table 6: Productivity dependent distortions: aggregate outcomes relative to the benchmark economy – robustness to changes in q

	$\tau_{\max} = 10\%$		$\tau_{\max} = 20\%$	
	optimal $\sigma_\varepsilon(z)$	benchmark $\tilde{\sigma}_\varepsilon(z)$	optimal $\sigma_\varepsilon(z)$	benchmark $\tilde{\sigma}_\varepsilon(z)$
$q = 0.25$				
Output	0.994	0.991	0.980	0.983
Consumption	0.989	0.995	0.970	0.984
Average employment	0.820	0.915	0.746	0.834
Average output	0.815	0.907	0.731	0.820
Number of firms	1.220	1.092	1.340	1.198
Average σ_ε	0.968	0.984	0.953	0.975
Firm exit rate	0.960	0.952	0.934	0.930
$q = 0.75$				
Output	0.992	0.991	0.981	0.983
Consumption	0.978	0.994	0.965	0.984
Average employment	0.747	0.927	0.698	0.833
Average output	0.741	0.919	0.684	0.819
Number of firms	1.339	1.079	1.434	1.200
Average σ_ε	0.846	0.982	0.786	0.964
Firm exit rate	0.974	0.963	0.949	0.927

Notes: Firms are subject to a tax rate that is linear in productivity. The tax rate for the median firm is zero. $\tau_{\max}$ is the tax rate on firms with productivity corresponding to the top 0.1 percentile in the benchmark economy. $\tilde{\sigma}_\varepsilon(z)$ denotes the experimentation policy that is optimal in the benchmark economy. Since we assume that any net tax revenue is handed back lump-sum to the households, the reported values for output and consumption include net tax revenue. Output is reported net of fixed operating costs. All values are relative to outcomes in the undistorted benchmark equilibrium. Parameters are as in Table 2, except for $\bar{\sigma}_\varepsilon$, which is 0.09 when q is 0.25 and 0.1575 when q is 0.75. For results with $q = 0.5$, see Table 3 in Section 4.2.

Table 7: Varying the overall tax burden: aggregate outcomes relative to the benchmark economy – robustness to changes in q

	$\tau_{\text{median}} = 10\%$		$\tau_{\text{median}} = 20\%$	
	optimal $\sigma_\varepsilon(z)$	benchmark $\tilde{\sigma}_\varepsilon(z)$	optimal $\sigma_\varepsilon(z)$	benchmark $\tilde{\sigma}_\varepsilon(z)$
$q = 0.25$				
Output	0.926	0.924	0.855	0.854
Consumption	0.946	0.952	0.896	0.901
Average employment	0.965	1.081	1.147	1.304
Average output	0.894	0.999	0.981	1.113
Number of firms	1.036	0.925	0.872	0.767
Average σ_ε	0.967	0.984	0.962	0.984
Firm exit rate	0.959	0.952	0.948	0.952
$q = 0.75$				
Output	0.924	0.924	0.854	0.854
Consumption	0.937	0.951	0.889	0.902
Average employment	0.873	1.082	1.054	1.302
Average output	0.807	1.000	0.900	1.112
Number of firms	1.145	0.924	0.949	0.768
Average σ_ε	0.833	0.976	0.824	0.976
Firm exit rate	0.962	0.951	0.962	0.951

Notes: Firms are subject to a tax rate that is linear in productivity. The tax rate for the median firm is given in the table. $\tau_{\max}$ is the tax rate on firms with productivity corresponding to the top 0.1 percentile in the benchmark economy and is $\tau_{\text{median}} + 0.1$. The slope of the tax function thus is the same in both columns. $\tilde{\sigma}_\varepsilon(z)$ denotes the experimentation policy that is optimal in the benchmark economy. Since we assume that any net tax revenue is handed back lump-sum to the households, the reported values for output and consumption include net tax revenue. Output is reported net of fixed operating costs. All values are relative to outcomes in the undistorted benchmark equilibrium. Parameters are as in Table 2 except for $\bar{\sigma}_\varepsilon$, which is 0.09 when q is 0.25 and 0.1575 when q is 0.75. For results with $q = 0.5$, see Table 4 in Section 4.2.