

Demand Learning and Firm Dynamics: Evidence from Exporters

Nicolas Berman
Vincent Rebeyrol
Vincent Vicard

Demand learning and firm dynamics: evidence from exporters*

Nicolas BERMAN[†] Vincent REBEYROL[‡] Vincent VICARD[§]

Abstract: This paper provides direct evidence that learning about demand is an important driver of firms' dynamics. We present a model of Bayesian learning in which firms are uncertain about their idiosyncratic demand in each of the markets they serve, and update their beliefs as noisy information arrives. Firms are predicted to update more their beliefs to a given demand shock, the younger they are. We test and empirically confirm this prediction, using the structure of the model together with exporter-level data to identify idiosyncratic demand shocks and the firms' beliefs about future demand. Consistent with the theory, we also find that the learning process is weaker in more uncertain environments.

Keywords: firm growth, belief updating, demand, exports, uncertainty.

JEL classification: D83, F14, L11.

*We thank the editor Emir Kamenica and four referees, as well as Richard Baldwin, Giuseppe Berlingieri, Patrick Blanchenay, Thomas Chaney, Simon Fuchs, Fabien Gensbittel, Joe Francois, Sacha Kapoor, Omar Licandro, Erzo Luttmer, Thierry Mayer, Marti Mestieri, Rahul Mukherjee, Franck Portier, Jim Tybout and seminar participants at the Geneva Trade and Development Workshop, EBRD, Aix-Marseille University, Università Cattolica Milano, Nottingham University, Universidad de Navarra, Tilburg University, Banque de France, the VIth Villars Research Workshop on International Trade, LETC, DEGIT conference, SED conference and Compnet conference at ECB for very useful discussions and comments. The authors gratefully acknowledge financial support from Fondation Banque de France. Nicolas Berman thanks the A*Midex for financial support (grant ANR-11-IDEX-0001-02 funded by the French government "Investissement d'Avenir" program). This paper was partly written when Nicolas Berman was at the Graduate Institute of International and Development Studies, Geneva. The opinions expressed in this paper are those of the authors and do not reflect the views of the Banque de France. This paper features an online appendix containing additional results and available on the authors' webpages.

[†]Aix-Marseille Univ., CNRS, EHESS, Centrale Marseille, AMSE, France, and CEPR. E-mail: nicolas.berman@univ-amu.fr.

[‡]Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France. E-mail: vincent.rebeyrol@ut-capitole.fr

[§]CEPII and Banque de France. E-mail: vincent.vicard@cepii.fr.

1 Introduction

Why do some firms grow faster than others? While some producers rapidly expand after entry, many others do not survive the first few years. After some time however, those surviving firms account for a large share of sales on both domestic or foreign markets (Haltiwanger et al., 2013; Bernard et al., 2009; Eaton et al., 2008). In the case of French firms, those that did not serve foreign markets a decade earlier account for 53.5 percent of total foreign sales, of which 40 percent comes from post-entry growth.¹ Understanding the sources of heterogeneity in post-entry firm dynamics – survival and growth – is therefore crucial to explain the dynamics of aggregate sales and firm size distribution.

Firm dynamics are characterized by a number of systematic patterns, which have been documented by a large body of empirical literature. New firms start small and have higher exit rates. For those that survive, the average growth of their sales declines with their age.² These facts can be rationalized by several theories, relying on different underlying mechanisms such as stochastic productivity growth, endogenous R&D investment, financial constraints, adjustment costs, demand accumulation or demand learning. Yet, empirically, disentangling the role of these specific channels has proven difficult, as it requires identifying separately the contributions of idiosyncratic demand and productivity to the variations of firms sales. For this reason, the literature has followed an indirect approach: it has studied which models are able to replicate the behavior of observables such as sales growth and exit. In contrast, this paper directly tests for the existence of demand learning by identifying firms' beliefs about demand and the signals they receive, and shows that it is an important driver of post-entry firm dynamics.

We first document two novel stylized facts using detailed data from the French customs containing information on firms' sales by destination and 6-digit product between 1994 and 2005. Throughout the paper, we refer to a product-destination pair as a *market*, and define age as the tenure of a firm in a specific *market*. We show that existing results about aggregate firm behavior carry over at the firm-market level. More precisely, sales growth, exit rates and the variance of sales growth within cohort all decrease with the age of the firm in its market. Importantly, these patterns are still present after controlling for firm-market size or conditioning on firm-product-year fixed effects. In addition, we find that the market-specific growth paths after entry are highly heterogeneous across firms: while entrants grow on average in their first years, a significant share of survivors exhibit negative post-entry growth in the markets they serve. For instance, around 40% of the firms that enter a market in 1996 and stay until 2005 sell *less* at the end of the period than in their second year.

We then present a standard model with Bayesian demand learning in the spirit of Jovanovic (1982) that can rationalize these facts.³ Firms operate under monopolistic competition and face CES demand, but at the same time are uncertain about their idiosyncratic demand in each market, and learn as noisy information arrives in each period. These signals determine the firms' posterior beliefs about demand, on which they base their quantity decision. A higher than expected signal

¹These numbers are based on the 1996-2005 period – see online appendix, Section B.

²See Evans (1987), Dunne et al. (1989), Cabral and Mata (2003), Haltiwanger et al. (2013) among many others. Eaton et al. (2008), Buono and Fadinger (2012), Berthou and Vicard (2015), or Ruhl and Willis (2017) show that these dynamics are also observed for exporters.

³In Jovanovic (1982), firms actually learn about their cost parameter. While the learning mechanism is the same, we apply it to demand, as in Timoshenko (2015).

leads younger firms to update more their beliefs than older ones, which implies that the growth rates of young firms are more volatile, even conditional on their size. The model also predicts that market-specific uncertainty limits the extent of belief updating and the impact of age on the updating process. The main contribution of this paper is to test these core predictions, which are specific to the passive learning mechanism.⁴

To do so, we derive from the theory a methodology which allows to separately identify the firms' beliefs and the demand shocks (the signals) they face in each period, in each of the markets they serve. First, we purge market-specific conditions and firm-specific supply side dynamics (e.g. productivity) from quantities and prices. This is made possible by a unique feature of international trade data, in which we can observe the values and quantities sold of a given product by a given firm in different markets. This is key as it enables to cleanly separate productivity from demand variations. In addition, observing different firms selling the same product in the same destination allows to control for aggregate market-specific conditions. Second, we use the fact that, in the model, quantity decisions only depend on the firms' beliefs while prices also depend on the realized demand shocks. This allows to separate out the firms' beliefs from the demand signals. Hence, while requiring few, standard assumptions, our methodology allows to directly test predictions that relate the evolution of firms' beliefs to firm age and demand signals.

We find strong support for the core predictions of the model. Belief updating following demand shocks is stronger for younger firms, with age being defined at the firm-product-destination (i.e. firm-market) level. Further, using a theory-based measure of market-specific uncertainty, we find that the learning process is significantly weakened and less dependent on age in more uncertain environments. We provide several robustness exercises to show that these results are not driven by our main modeling assumptions. Our findings survive after accounting for potential endogenous selection bias, and are extremely stable across alternative samples, specifications and changes in variables' measurement. We also discuss the implications of relaxing several important assumptions of the model related to the timing of price and quantity adjustment, market structure and firms' productivity. We show that even after relaxing these assumptions, our results can still be interpreted as evidence of belief updating. Some of these extensions however require that we control for firm-market size in our estimations, which leaves our results unchanged.

The literature has proposed a number of potential supply or demand side drivers of firm dynamics. But, learning apart, they cannot explain our main result of a smaller quantity adjustment to past demand shocks for older firms. Suppose indeed that firms have full information about demand, except about a stochastic shock each period. If these shocks are *iid*, there is no reason for the firm to adjust quantities the period after, as these shocks do not convey any information and have no relevance beyond the current period. If instead shocks are persistent, there is no reason for older firms to react less to a shock of a given size. Alternative mechanisms are also difficult to reconcile with our stylized facts.

On the supply side, several papers attempt to explain the heterogeneity in firm size with productivity variations only (through stochastic shocks or endogenous decisions).⁵ By construc-

⁴We additionally show in the online appendix that exit behavior is also consistent with the learning model: the exit rate decreases with firms' beliefs and the demand shocks the firms face, and demand shocks trigger more exit in younger cohorts.

⁵See for instance Hopenhayn (1992), Luttmer (2007), Impullitti et al. (2013) for models with stochastic shocks to productivity, Klette and Kortum (2004) or Rossi-Hansberg and Wright (2007) for theories of endogenous pro-

tion, they are not able to generate an age dependence of firm growth, *conditional* on size. In contrast, models introducing additional sources of heterogeneity, such as financial constraints or capital adjustments costs, are able to generate this age dependence.⁶ Yet, since these sources of heterogeneity apply to the firm as a whole, they cannot deliver the heterogeneous firm-market specific dynamics that we find in the data.

Beyond learning, some demand side mechanisms could be affecting firm growth at the market level. Various processes giving rise to demand accumulation have been proposed. Firms could engage in market-specific investments (e.g. Ericson and Pakes, 1995, Luttmer, 2011, Eaton et al., 2014, Fitzgerald et al., 2016), price low in their first years to build a consumer base (Foster et al., 2016, Gourio and Rudanko, 2014, Piveteau, 2016), or face demand that evolves exogenously over time (Ruhl and Willis, 2017). Among the most recent contributions, Ruhl and Willis (2017) use a model with stochastic entry costs and gradual increase in demand to match the average growth and exit rates of Colombian exporters. Arkolakis (2016) shows that a combination of idiosyncratic productivity shocks and market penetration costs is able to reproduce some important patterns of the distribution of US and Brazilian exporters’ growth. Since these models include some mean reversion effects, they can generate an age dependence of firm growth conditional on size; but they fall short at predicting the decline in the variance of growth rates with age, conditional on size.⁷ On the other hand, we show that our estimates of firms’ beliefs reproduce well this observed decline in the variance of sales growth.

The last part of our paper discusses whether alternative demand based theories, possibly on top of a learning effect, could be driving our findings. We show in particular that theories of demand accumulation would have serious difficulties matching the profiles of prices and quantities that we find. Indeed, in our data, once purged from their productivity component, firm-market-specific prices are (slightly) decreasing with age. Such a pattern contradicts models of active demand accumulation through pricing decisions or models featuring learning in which firms set prices rather than quantities. It is however consistent with the passive demand learning model, in which survivors tend to have received relatively more “good news” than exiters, leading them to adjust their prices upwards to take advantage of this unexpectedly high demand. As firms get better informed over time, their prices converge to their optimal pricing rule. However, once composition effects are controlled for, prices – in the model and in the data – are constant as firms have equal probabilities to update upwards or downwards. Similarly, quantities should increase over time, but in the learning model this is mostly due to selection. This prediction is confirmed empirically: when accounting for composition effects triggered by selection, we find that quantities within firms-markets exhibit a very limited positive growth, observed only in the first years. This matches well our second stylized fact: a substantial part of survivors shrinks in size due to their “over-optimistic” beliefs at entry.⁸

ductivity growth.

⁶See for example Cooley and Quadrini (2001) or Clementi and Hopenhayn (2006) for financial constraints, and Clementi and Palazzi (2016) for adjustments costs.

⁷For example, Arkolakis (2016) assumes an exogenous Ornstein-Uhlenbeck process for productivity, which generates an age dependence of firm growth at the market level, conditional on size. But this set-up cannot explain the decline in the variance of growth within cohorts at the market level, as Ornstein-Uhlenbeck processes have a constant variance. What would be needed is a process that implies both smaller shocks over time *and* a smaller variance of these shocks. This is not a standard feature of the most common stochastic processes.

⁸We also perform a test initially proposed by Pakes and Ericson (1998), in which we regress current firm beliefs on immediate past beliefs and initial beliefs. Consistent with a passive learning model we find that initial beliefs

Overall, these results do not preclude alternative mechanisms to be jointly at work, but they clearly suggest that the patterns we identify in our data are unlikely to be driven by demand accumulation processes. Demand learning appears to be an important determinant of the micro-dynamics of firms in narrowly defined markets, which is key as more than half of the variance of sales growth in our sample is due to firm-market factors. This supports the view of several recent works arguing that demand learning models reproduce well some important characteristics of the dynamics of firms and exporters.⁹ Compared to these papers, we follow a different strategy as we propose a direct test of the updating process, which lies at the core of the learning mechanism. Our empirical methodology is close in spirit to Foster et al. (2016, 2008), in that they also separate idiosyncratic demand shocks from firms' productivity, but our paper differs in several ways. In particular, we do not need to measure productivity or other firm-specific determinants of sales to identify demand shocks.

Finally, we assume that the actual sales of a firm in a given product-destination market are the only source of information about demand. In other words, we assume away information spillovers. A firm's belief in a given market might well be affected by its beliefs in other destinations (Albornoz et al., 2012), or about other products in the same destination (Timoshenko, 2015). These effects might be stronger for similar destinations and products (Morales et al., 2014; Defever et al., 2015; Lawless, 2009). The behavior of other firms serving the same market might also play a role (Fernandes and Tang, 2014). Studying the relative importance of these various potential sources of information is an interesting and vast question in itself, that we indeed plan to study in the future, but which is beyond the scope of this paper.

The paper proceeds as follows. In the next section, we describe our data, document new stylized facts about firms' post-entry dynamics, and discuss them in light of existing theories. In section 3 we present the model and our identification strategy. Section 4 contains our main results and section 5 various robustness exercises. Section 6 discusses whether our results could be explained by alternative demand-based mechanisms. The last section concludes.

2 Firm dynamics on foreign markets and export growth

2.1 Data

We use detailed firm-level data by product and destination country provided by the French Customs. The unit of observation is an export flow by a firm i of a product k to a destination j in year t . A product is defined at the 6-digit level (HS6). The data cover the period from 1994 to 2005, and contain information about both the value and quantity exported by firms, which will allow us to compute firm-destination-product specific unit values that we will use as a proxy for prices in the second part of the paper. Section A.1 of the online appendix provides more details on the source data.

Two important notes on the terminology we use throughout the paper. First, what we call a *market* is a product \times destination combination. Second, *age* is defined by market. Our baseline definition of age is the number of years of presence since the last entry of a firm in a product-destination. Age is reset to zero whenever the firm exits for at least a year from a specific market.

are useful to forecast future firms' beliefs throughout their life.

⁹See for instance Albornoz et al. (2012), Timoshenko (2015), Fernandes and Tang (2014) or Eaton et al. (2014).

What we call age is therefore equivalent to market-specific tenure. Section I.2 in the online appendix discusses alternative measures. Note that in all the empirical analysis, to ensure the consistency of our measures of age, we drop firm-product-destination triplets already present in 1994 and 1995, as these years are used to define entry.

Finally, a *cohort* of new exporters in a product-destination market includes all firms starting to export in year t but that were not exporting in year $t - 1$, and we are able to track all firms belonging to a cohort over time.

Our final dataset covers the sales of 3,844 HS6 product categories to 179 destinations by 77,076 firms over the period 1996-2005. All these firms entered at least one market over the period.

2.2 Stylized facts

In this section we provide two novel stylized facts on the post-entry dynamics of firms at the product-destination level. The first is that growth rates and their variance within cohorts decline sharply with age, within *firm-markets and conditional on size*. The second is that among survivors, growth paths are highly heterogeneous, with a large number of firms exhibiting negative growth rates. We will argue that both facts are difficult to reconcile with most theories of firm dynamics apart from the passive learning model.

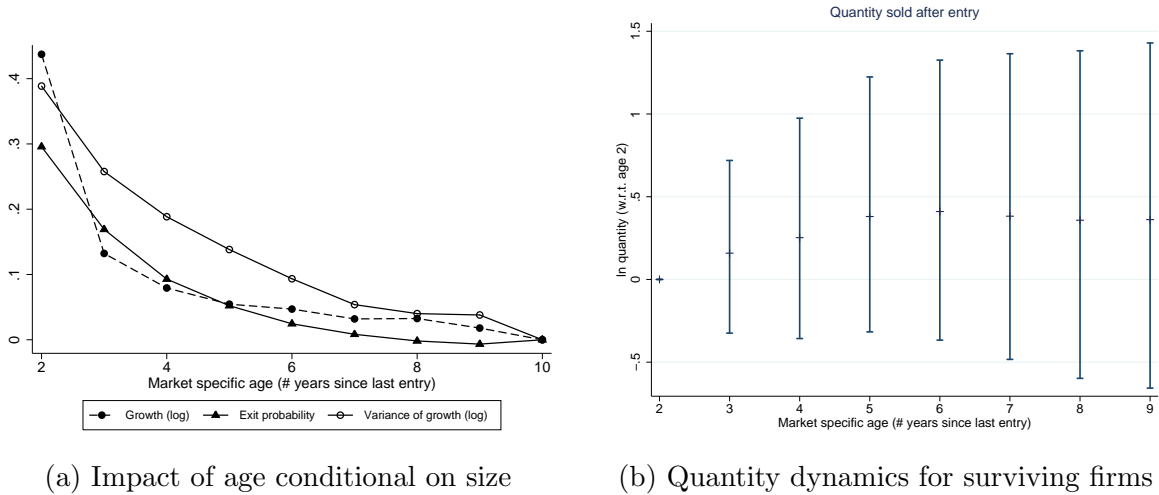
Before explaining these facts in more details, note that our data exhibits patterns that are in line with those found by the literature. Consistent with the results of Eaton et al. (2008) on Colombian data (see also Haltiwanger et al., 2013 and Bernard et al., 2009), we find that new firms-markets contribute disproportionately to aggregate trade growth: new flows account for only 12.3 percent of total export value after a year, but this share reaches 53.5 percent after a decade. Moreover, regressing firm-market sales growth on various sets of fixed effects, we find that market-time and firm-product-time factors only account for 44 percent of the variance of sales growth, a result which echoes the findings of Eaton et al. (2011) or Munch and Nguyen (2014). In other words, firm-market factors are key to explain growth dynamics. The online appendix, section A, provides further discussion of these results.

Fact #1: Firm-market growth and its variance decline with age, conditional on size.

Contrary to most existing papers that have documented facts about the aggregate dynamics of firms or exporters, our data allows to study growth and survival in each market served by the firms. We consider three components of firm-market post-entry dynamics: sales growth, exit rate, and the variance of sales growth within cohort. Figure 1.a plots the coefficients obtained by regressing these different variables on age dummies, controlling for sector and time dummies and, more importantly, for bins of firm size. The full set of results is shown and further discussed in the online appendix, section A.2. All three sets of coefficients sharply decrease with age, with age being defined as firm-product-destination specific tenure. Both the growth rates of firms-markets and the variance of these growth rates within cohort is about 40 percent higher in the second year than after ten years. Importantly, we still find that sales growth declines with age when we include in our regressions firm \times product \times year fixed effects which control for any unobserved supply side factors (like financial constraints) which are common to all markets within a firm (see online appendix, Table A.3, column 2).

Fact #2: Post-entry growth dynamics are heterogeneous across survivors. Our second

Figure 1: Stylized facts



Note: Figure (a) plots the coefficients obtained from a regression of the log change of firm sales (respectively variance of firms' sales and exit) on age bins, firm size and year and sector dummies (see table A.3 in online appendix section A for the full set of results). All coefficients are relative to the omitted category, age of ten years. The variance of firms' sales growth is measured within cohorts of firms on a product-destination market. Similar patterns are obtained when controlling for country-and-sector fixed effects. Figure (b) plots statistics about market-specific firm quantities with respect to age for the cohorts of firms which entered in 1996. Quantities are normalized to 1 in age 2. The upper and lower limits of the lines represent the first and last quartiles of the variable, with the median in between.

stylized fact appears in Figure 1.b, where we plot the log of quantities sold by firms entering a given market in 1996 and staying the entire period (until 2005).¹⁰ Quantities are normalized to one in year two.¹¹ The horizontal lines depict the first quartile, the median and last quartile at each age. Survivors grow after entry consistently with existing evidence (Eaton et al., 2008, 2014; Foster et al., 2016; Ruhl and Willis, 2017; Fitzgerald et al., 2016), a pattern that has motivated theories of demand accumulation. But Figure 1.b makes it clear that growth paths are greatly heterogeneous and that a significant share of firm-markets experience negative growth. More precisely, around 40 percent of the firms shown in this figure sell actually *less* at the end of the period than in their second year.

As mentioned in the introduction, the set of facts shown in Figures 1.a and 1.b is difficult to rationalize using existing theories that do not incorporate learning. Models featuring solely supply side dynamics that are firm or firm-product specific (productivity, financial constraints, capital adjustment costs) cannot help understanding the behavior of firms-products across destinations. Theories introducing both supply and demand mechanisms are better designed to explain a heterogeneity across destinations, but they typically fail to generate the dependence of the variance of growth rate to firm age that we observe in the data. Finally, in models of firm dynamics with demand accumulation, survivors tend to be those that have been able to

¹⁰A similar pattern is obtained with different dates of entry, or using values instead of quantities. See Figures A.1 and A.2 in section A of the online appendix.

¹¹We do not consider the first year because of its potential incompleteness when measured over a calendar year (Berthou and Vicard, 2015). Similarly, we plot the statistics up to 9 years and not 10 because we want to look at flows that will still be present the year after (and 10-year-old flows can only be observed in 2005, which is the last year of our sample). The online appendix section I.3 discusses this point.

accumulate demand. This allows to fit the average growth path of new firms/exporters observed in the data but does not necessarily provide a framework to think about heterogeneous outcomes across firms. On the other hand, the passive learning model naturally generates these patterns. The decline in the variance of growth rates with cohort age is caused by the larger updating of younger firms. The decline in growth rates is mostly driven by selection: firms that decline the most in size exit the market, which implies that the distribution of growth rates is truncated from below. Together with their larger variance, this implies larger growth rates for younger firms, conditional on survival. It should be noted that the passive learning model is also consistent with larger unconditional growth rates.¹² Finally, the high heterogeneity in firms' growth paths after some years comes from the fact that initial prior beliefs may not be accurate, leading some firms to shrink in size over time.

3 A model of firm growth with demand learning

We consider a standard model of international trade with Dixit-Stiglitz monopolistic competition and demand learning in the spirit of Jovanovic (1982). As earlier, we index firms by i , destination markets by j , products by k and time by t .

3.1 Economic environment

Demand. Consumers in country j maximize utility derived from the consumption of goods from K sectors. Each sector is composed of a continuum of differentiated varieties of product k :

$$U_j = \mathbb{E} \sum_{t=0}^{+\infty} \beta^t \ln(C_{jt})$$

$$\text{with } C_{jt} = \prod_{k=0}^K \left(\int_{\Omega_{kt}} (e^{a_{ijkt}})^{\frac{1}{\sigma_k}} c_{kt}(\omega)^{\frac{\sigma_k-1}{\sigma_k}} d\omega \right)^{\frac{\mu_k \sigma_k}{(\sigma_k-1)}}$$

with β the discount factor, Ω_{kt} the set of varieties of product k available at time t , c_{kt} is the consumption level of each variety, and $\sum_k \mu_k = 1$. Demand in market j at time t for a variety of product k supplied by firm i is given by:

$$q_{ijkt} = e^{a_{ijkt}} p_{ijkt}^{-\sigma_k} \frac{\mu_k Y_{jt}}{P_{jkt}^{1-\sigma_k}} \quad \text{where} \quad P_{jkt}^{1-\sigma_k} = \int_{\Omega_{kt}} e^{a_{ijkt}} p_{ijkt}^{1-\sigma_k} d\omega \quad (1)$$

where σ_k is the (sector-specific) elasticity of substitution, Y_{jt} is total expenditure and P_{jkt} is the ideal price index of destination j in sector k , during year t . The demand parameter a_{ijkt} is given by $a_{ijkt} = \bar{a}_{ijk} + \varepsilon_{ijkt}$, with ε_{ijkt} a white noise. \bar{a}_{ijk} is an idiosyncratic constant parameter and is unknown to the firm.

Production. Each period, firms make quantity decisions for their product(s), before observing demand in each market served, i.e. before observing a_{ijkt} . The unit cost function is linear in the marginal cost and there is a per-period fixed cost F_{ijk} to be paid for each product-destination pair. Labor L is the only factor of production. Current input prices are taken as given (firms

¹²This is however generated by functional form assumptions. See the online appendix G for details.

are small) and there is no wedge between the buying and selling price of the input (i.e. perfect reversibility in the hiring decision). Hence, the quantity decision is a static decision.

We do not make any assumption on the evolution of firm productivity. Productivity may also be subject to learning, in which case the firm would base its quantity decision on its beliefs about its costs. As we will not back out learning from firms' productivity, we do not add expectation terms here to save on notations. We only need to assume that unit costs at the firm-product level are *not* destination specific – we come back to this assumption in section 3.3. Per period profits in market j from product k write:

$$\pi_{ijkt} = q_{ijkt}p_{ijkt} - \frac{w_{it}}{\varphi_{ikt}}q_{ijkt} - F_{ijk} \quad (2)$$

where w_{it} is the wage rate in the origin country, φ_{ikt} is the product-time specific productivity of firm i .

Learning. Firm i is uncertain about the parameter \bar{a}_{ijk} . Before observing any signal, its prior beliefs about \bar{a}_{ijk} are normally distributed with mean θ_{ijk0} and variance σ_{jk0}^2 . Different firms may well have different initial beliefs prior to entry (i.e. different θ_{ijk0}). θ_{ijk0} is drawn from a normal distribution with mean \bar{a}_{ijk} and variance σ_{jk0}^2 : prior beliefs may not be accurate, but are unbiased on average.¹³ The firm observes t independent signals about \bar{a}_{ijk} : $a_{ijkt} = \bar{a}_{ijk} + \varepsilon_{ijkt}$, where each ε_{ijkt} is normal with (known) mean 0 and variance σ_ε^2 . According to Bayes' rule, the firm's posterior beliefs about \bar{a}_{ijk} after t signals are normally distributed with mean $\tilde{\theta}_{ijkt}$ and variance $\tilde{\sigma}_{ijkt}^2$, where:

$$\tilde{\theta}_{ijkt} = \theta_{ijk0} \frac{\frac{1}{\sigma_{jk0}^2}}{\frac{1}{\sigma_{jk0}^2} + \frac{t}{\sigma_\varepsilon^2}} + \bar{a}_{ijk} \frac{\frac{t}{\sigma_\varepsilon^2}}{\frac{1}{\sigma_{jk0}^2} + \frac{t}{\sigma_\varepsilon^2}} \quad (3)$$

$$\tilde{\sigma}_{ijkt}^2 = \frac{1}{\frac{1}{\sigma_{jk0}^2} + \frac{t}{\sigma_\varepsilon^2}} \quad (4)$$

and \bar{a}_{ijk} is the average signal value, $\bar{a}_{ijk} = (\frac{1}{t} \sum_t a_{ijkt})$. Note that contrary to $\tilde{\theta}_{ijkt}$, the posterior variance $\tilde{\sigma}_{ijkt}^2$ does not depend on the realizations of the signals and decreases only with the number of signals (i.e. learning reduces uncertainty). Hence, the posterior variance is always smaller than the prior variance, $\tilde{\sigma}_{ijkt}^2 < \tilde{\sigma}_{ijkt-1}^2$.

In the following, it will be useful to formulate the Bayesian updating recursively. Denoting $\Delta\tilde{\theta}_{ijkt} = \tilde{\theta}_{ijkt} - \tilde{\theta}_{ijkt-1}$, we have:

$$\Delta\tilde{\theta}_{ijkt} = g_t (a_{ijkt} - \tilde{\theta}_{ijkt-1}) \quad \text{with } g_t = \frac{1}{\frac{\sigma_\varepsilon^2}{\sigma_{jk0}^2} + t}. \quad (5)$$

Intuitively, observing a higher-than-expected signal, $a_{ijkt} > \tilde{\theta}_{ijkt-1}$ leads the agent to revise the expectation upward, $\tilde{\theta}_{ijkt} > \tilde{\theta}_{ijkt-1}$, and vice versa. This revision is large when g_t is large, which happens when t is small, i.e. when the firm is “young” in market jk .

¹³We could further assume, leaving our results fully unchanged, that the variance of the prior beliefs is firm specific, i.e. σ_{ijk0}^2 . We would need to assume in that case that this firm specific variance is independent from firm characteristics.

3.2 Firm size and belief updating

Firms maximize expected profits, subject to demand. Labelling $G_{t-1}(a_{ijkt})$ the prior distribution of a_{ijkt} at the beginning of period t (i.e. the posterior distribution after having observed $t - 1$ signals), firm i maximizes:

$$\max_q \int \pi_{ijkt} dG_{t-1}(a_{ijkt}) \quad \text{s.t.} \quad p_{ijkt} = \left(\frac{\mu_k Y_{jt} e^{a_{ijkt}}}{q_{ijkt} P_{jkt}^{1-\sigma_k}} \right)^{\frac{1}{\sigma_k}}. \quad (6)$$

Here, we assume for simplicity that aggregate market conditions at time t , i.e. $\mu_k Y_{jt} / P_{jkt}^{1-\sigma_k}$, are observed by firms before making their quantity decision. This leads to the following optimal quantities and prices (see appendix):

$$q_{ijkt}^* = \left(\frac{\sigma_k}{\sigma_k - 1} \frac{w_{it}}{\varphi_{ikt}} \right)^{-\sigma_k} \left(\frac{\mu_k Y_{jt}}{P_{jkt}^{1-\sigma_k}} \right) \left(\mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right] \right)^{\sigma_k} \quad (7)$$

$$p_{ijkt}^* = \left(\frac{\sigma_k}{\sigma_k - 1} \frac{w_{it}}{\varphi_{ikt}} \right) \left(\frac{e^{\frac{a_{ijkt}}{\sigma_k}}}{\mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right]} \right) \quad (8)$$

with $\mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right] = \int e^{\frac{a_{ijkt}}{\sigma_k}} dG_{t-1}(a_{ijkt})$. As firm i makes a quantity decision before observing demand for its product, q_{ijkt}^* depends on expected demand, not on demand realization, contrary to p_{ijkt}^* .

The literature has typically computed correlations between firm age and firm growth rates, and attributed negative ones as potential evidence for a learning mechanism. Indeed the fact that younger firms adjust more their beliefs leads growth rate to decrease with age in absolute value. But of course, as is clear from equations (7) and (8), firm size and therefore firm growth (would it be measured in terms of employment or sales) also depend on the evolution of market-specific conditions and firm productivity, which could be correlated with firm age. Directly testing for the presence of demand learning thus requires either making assumptions about the dynamics of aggregate market conditions and firm productivity or finding a way to account for them. Our methodology follows the second route.

Let us now decompose optimal quantities and prices into three components. They first depend on unit costs, which are a function of wages in country i and firm-product specific productivity φ_{ikt} . This first component is ikt -specific, i.e. is independent of the destination served; we label it C_{ikt} . Second, they depend on aggregate market conditions, which are common to all firms selling product k to destination j . We label this component C_{jkt} . Finally, they depend on the firm i beliefs about expected demand in j for its product k and on the demand shock at time t . This last composite term – labelled Z_{ijkt} – is the only one to be impacted by firm learning about its demand in a specific destination market: it is $ijkt$ -specific. We can now rewrite the above expressions for quantities and prices as:

$$q_{ijkt}^* = C_{ikt}^q C_{jkt}^q Z_{ijkt}^q \quad (9)$$

$$p_{ijkt}^* = C_{ikt}^p Z_{ijkt}^p. \quad (10)$$

The impact of demand learning is fully included in the Z_{ijkt}^q and Z_{ijkt}^p terms. These terms can be understood as the quantity and price of firm i for product k on market j at time t , purged from firm unit costs and aggregate market conditions, and may be very different from the actual firm size and firm price. From a methodological point of view, any prediction about firm demand learning should be based on these Z_{ijkt} terms rather than the actual q_{ijkt}^* and p_{ijkt}^* . This also means that we will not look at the dynamics of firm size (at least per se), but directly at the dynamics of the firms' beliefs about demand. Their growth rate can be expressed as:¹⁴

$$\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right] = \frac{g_t}{\sigma_k} \left(a_{ijkt} - \tilde{\theta}_{ijkt-1} \right) - \frac{g_t}{\sigma_k} \frac{\tilde{\sigma}_{ijkt-1}^2}{2\sigma_k}. \quad (11)$$

At the beginning of period t , firms make quantity decisions based on their beliefs about local demand for their product ($\tilde{\theta}_{ijkt-1}$). Then, demand is realized (a_{ijkt}) and firms update their beliefs. A higher than expected demand leads the firm to update upwards its belief. The opposite is true for a lower than expected demand. Importantly, as is clear from equation (11), this upward or downward updating is larger for younger firms. It follows our main prediction:

Prediction # 1 (updating and age): *A given difference between realized and expected demand leads to a larger updating of the belief, the younger the firm is.*

It is also interesting to note that larger uncertainty (i.e. a higher σ_ϵ^2) reduces the extent of belief updating and the effect of age on belief updating. This is because a signal is less informative when uncertainty is higher. Put differently, the information contained in the realized price will be noisier when σ_ϵ^2 is large, in which case firms will adjust less their beliefs in the next period. This is our second prediction:

Prediction # 2 (updating and uncertainty): *A higher level of market uncertainty reduces the extent of beliefs updating, and the effect of age on belief updating.*

In the next section, we derive our methodology to isolate the Z_{ijkt}^q and Z_{ijkt}^p terms and distinguish the beliefs from the demand shock component.

3.3 Identification and measurement

Identifying beliefs. In order to isolate Z_{ijkt}^q and Z_{ijkt}^p , we need to purge supply side and market specific factors from actual quantities and prices. This is achieved by estimating the following quantity and price equations in logs:¹⁵

$$\ln q_{ijkt} = \mathbf{FE}_{ikt} + \mathbf{FE}_{jkt} + \varepsilon_{ijkt}^q \quad (12)$$

$$\ln p_{ijkt} = \mathbf{FE}_{ikt} + \varepsilon_{ijkt}^p \quad (13)$$

where k is a 6-digit product and t is a year. \mathbf{FE}_{ikt} and \mathbf{FE}_{jkt} represent respectively firm-product-year and destination-product-year fixed effects. Note that we do not have direct price data, so we rely on unit values, defined as S_{ijkt}/q_{ijkt} , where S_{ijkt} denote firms sales, to proxy them. In our

¹⁴Detailed derivations and proofs of all our propositions are relegated to the appendix.

¹⁵We use the Stata routine `reghdfe` developed by Sergio Correia, based on Guimaraes and Portugal (2010).

baseline estimations, we stick to the model and estimate the price equation without the jkt fixed effects, as implied by the CES assumption. In section 5.1 we discuss the implications of relaxing the CES assumption, one of them being that we need to control for market-specific conditions in the price equation.

The estimates of ε_{ijkt}^q and ε_{ijkt}^p are estimates of the Z_{ijkt} terms. Using (7) and (8), we get:

$$\varepsilon_{ijkt}^q = \ln Z_{ijkt}^q = \sigma_k \ln \mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right] \quad (14)$$

$$\varepsilon_{ijkt}^p = \ln Z_{ijkt}^p = \frac{1}{\sigma_k} a_{ijkt} - \ln \mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right]. \quad (15)$$

This identification strategy is possible to implement because we are able to observe the sales of the same product by the same firm in different destination markets, which allows purging market-specific firm dynamics from the evolution of firm productivity through the inclusion of \mathbf{FE}_{ikt} .¹⁶ As we account for all time-varying, market- and firm-product-specific determinants of quantities and prices, our approach could accommodate any underlying dynamic process for the ikt and jkt terms. This includes processes driving the evolution of firm productivity, but also any other time-varying, firm-specific factors that might affect firm dynamics, such as financial constraints, as well as variations in market-specific trade costs.

Consistently estimating the residuals of (12) and (13) however requires some identification assumptions. In particular, ε_{ijkt}^q and ε_{ijkt}^p need to be orthogonal to firm characteristics $\{w_{it}, \varphi_{ikt}\}$, and ε_{ijkt}^q must also be orthogonal to market conditions $\{Y_{jt}, P_{jkt}\}$. This implies that beliefs do not vary systematically with productivity, or, in other words, that initial beliefs must be unbiased also along the firm productivity dimension. This rules out the possibility that firms engage in overall productivity-enhancing investments because they have higher beliefs in a given market. Note however that our identification strategy does not preclude firms to modify a *market-specific* productivity component in response to changes in their information set. In section 5.1, we thus allow productivity to differ across destinations for a given firm-product. The condition on ε_{ijkt}^p also implies that demand signals a_{ijkt} must be orthogonal to firms' overall costs $\{w_{it}, \varphi_{ikt}\}$. Put differently, we make the standard assumption that firms with high productivity do not enjoy higher market specific demand beyond the effect of their productivity on demand through lower prices.

These orthogonality restrictions also reflect our assumption that beliefs are market-specific, i.e. that firms do not adjust their beliefs to information arriving from other markets. As mentioned in the introduction, in theory there could be spillovers taking many different forms: beliefs could depend on the experience accumulated by the firm in selling the same product to other destinations, including the domestic market. They could also vary with the information obtained when selling other products in the same market. Studying such informational spillovers is beyond the scope of this paper. Yet, we are confident that the information we capture is indeed market-specific. The reason is that our identification strategy *de facto* constrains the set of possible determinants of beliefs. For instance, if these are partly determined by past domestic market

¹⁶The reason why we do not model learning about productivity appears more clearly in equations (14) and (15). Identifying demand variations is possible because we are able to control for productivity through the inclusion of ikt fixed effects. On the other hand, we cannot distinguish productivity variations from global demand shocks faced by firms in all the markets, as these would be mixed with unit costs in the \mathbf{FE}_{ikt} .

experience for the same product, or by past experience in other markets for the same product, then the ikt fixed effects will account for them. In other words, ε_{ijkt}^q captures the firms' beliefs net of the effect of experience in other markets at time t .¹⁷

Finally, a note on our interpretation of the residuals (14) and (15). Following the model, we consider that these residuals reflect the demand-side components of prices and quantities. Our identification assumption is that, within a given firm, costs can differ across products but not across products *and* destinations. Note however that we allow variations in costs across markets for a given product. These include in particular trade costs and potential differences in demand for quality and are captured by \mathbf{FE}_{jkt} . In section 5.1 we allow productivity to be market-specific and show that we can still consistently estimate the demand shocks. We do not, however, allow firms to learn about market-specific costs. As discussed in section 6, the evidence we find on the profiles of ε_{ijkt}^q and ε_{ijkt}^p is more consistent with firms learning about demand than about costs, but we cannot exclude that firms are learning about demand shifters such as market-specific trade costs. Such a learning process would be isomorphic to learning about demand. We favor the traditional demand learning formulation, yet what we call demand learning could be encompassing learning about demand-shifters.

Identifying demand shocks. Testing prediction 1 requires getting estimates of the demand signals a_{ijkt} . Because the firm takes its quantity decision before observing the demand realization, $\ln Z_{ijkt}^q$ depends on the firms' beliefs about demand only, while $\ln Z_{ijkt}^p$ is adjusted for the demand shock (an assumption that we discuss in section 5.1). Thus, the residual ε_{ijkt}^q provides a direct estimate of the firms' beliefs. We only need to correct for σ_k . In order to back out the demand shock and get an estimate of σ_k , we regress ε_{ijkt}^p on ε_{ijkt}^q . Using (15) and (14), we get:

$$\left(\frac{1}{\sigma_k} a_{ijkt} - \ln \mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right] \right) = \beta \left(\sigma_k \ln \mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right] \right) + \lambda_{ijk} + v_{ijkt}. \quad (16)$$

We need to include firm-product-destination fixed effects λ_{ijk} to account for the fact that $a_{ijkt} = \bar{a}_{ijk} + \varepsilon_{ijkt}$. Omitting these fixed effects would generate inconsistent estimates of β as both v_{ijkt} and the firm beliefs $\mathbb{E}_{t-1} \left[\exp\left(\frac{a_{ijkt}}{\sigma_k}\right) \right]$ would depend on \bar{a}_{ijk} , which would violate the zero conditional mean assumption.¹⁸ Including λ_{ijk} allows to take out \bar{a}_{ijk} from the error term v_{ijkt} and recover consistent estimates of β . We estimate (17) by 6-digit product to allow σ_k to differ across products and obtain:¹⁹

¹⁷Our methodology does not, on the other hand, take into account the possibility that beliefs depend on the information gathered by the firm while selling other products in the same destination. This would require including ijt fixed effects in equations (12) and (13). We have tried to include these and our estimates were largely unaffected (see section I.5 in the online appendix). This lends support to our assumption that information is indeed mostly product-market specific: if shocks and beliefs were correlated across products within destinations, the firms' response to a demand shock would partly reflect its belief updating behavior on other products, and including ijt fixed effect should dampen the extent of estimated belief updating.

¹⁸We thank an anonymous referee for pointing out this issue.

¹⁹Whenever our estimates of β are statistically insignificant or imply values of σ_k which are lower than 1, we replace \hat{v} by a missing value and do not consider the observation in the estimations. Note that our results are insensitive to such cleaning of the data. σ_k is lower than 1 for only 0.01% of observations, and insignificant β coefficients (at the 5% level) are obtained for 1.6% of observations. See the upper panel of Table A.4 in the online appendix. We also perform a robustness exercise where equation (17) is estimated at the 4-digit instead of 6-digit level to end up with more observations by product and more efficient estimations.

$$\widehat{\beta} = -\frac{1}{\sigma_k} \quad ; \quad \widehat{\lambda}_{ijk} + \widehat{v}_{ijkt} \equiv \widehat{a}_{ijkt} = \frac{1}{\sigma_k} a_{ijkt} \quad ; \quad \widehat{v}_{ijkt} = \frac{1}{\sigma_k} \varepsilon_{ijkt}. \quad (17)$$

Note that the level of uncertainty can be directly inferred from our estimates of demand signals. We define market-specific uncertainty as the standard deviation of a_{ijkt} , computed by product-and-destination, over our data period.

The last variable we need to test our predictions is market-specific firm age, which has been defined in section 2. Age is either constructed as a single discrete variable or as a set of dummies, to allow the learning processes to be non-linear.

Testing prediction #1. We can now derive our testable equation. Equation (11) cannot be tested directly as we do not observe $\widetilde{\theta}_{ijkt-1}$ but only ε_{ijkt}^q . We make use of (11), (5) and (14), to get the following specification (see the appendix):

$$\Delta \varepsilon_{ijkt+1}^q = g_t \left(a_{ijkt} - \varepsilon_{ijkt}^q \right) + g_t \frac{\sigma_\varepsilon^2}{2\sigma_k}. \quad (18)$$

This equation is equivalent to (11), except that it can be tested: our estimates of ε_{ijkt}^q comes from (12), and a_{ijkt} is computed from equation (17) as the product of \widehat{a}_{ijkt} times σ_k . g_t is an inverse function of market-specific age (equation (5)). We estimate:

$$\Delta \varepsilon_{ijkt+1}^q = \sum_{g=2}^G \alpha_g (a_{ijkt} - \varepsilon_{ijkt}^q) \times AGE_{ijkt}^g + \sum_{g=1}^G \beta_g AGE_{ijkt}^g + u_{ijkt} \quad (19)$$

where AGE_{ijkt}^g are dummies taking the value 1 for each age category $g = 2, \dots, 10$ representing the number of years of presence in the export market (e.g. $g = 2$ in the second year of presence). Standard errors are robust to heteroscedasticity and clustered by firm (or, alternatively, bootstrapped). We expect α_g to be positive on average, and β_g to be decreasing with age. Our main prediction is that α_g decreases with age g . Note that equation (18) predicts that $\alpha_g = g_t = \frac{1}{\sigma_\varepsilon^2 / \sigma_{jk0}^2 + t}$ with g_t measuring the speed of learning. Hence, the evolution of the α_g coefficients with firm age allow to assess how firms learn about their demand parameter.

Our test of the passive learning mechanism therefore builds on the evidence that firms adjust their quantities to past demand shocks *and* that such a reaction gets smaller as firms grow older in a market. This decline of the quantity reaction to past demand shocks is a distinctive feature of the learning process. If firms had full information about demand, stochastic *iid* shocks should not generate any quantity reaction beyond the current period as these shocks would not provide any information. In that case, the coefficients α_g should be equal to zero. If instead shocks were persistent, firms would always adjust their next period quantities in the same way: the α_g coefficients would be positive but constant over time.

4 Main results

In this section, we start by providing some descriptive statistics about our final sample, before discussing the results obtained when testing prediction 1. We then study how market uncertainty affect the characteristics of the learning process.

4.1 Sample statistics

Table 1: Sample statistics

	Obs.	Mean	S.D.	Q1	Median	Q3
$\ln q_{ijkt}$	4382989	6.237	2.795	4.277	6.004	8.001
$\ln p_{ijkt}$	4382989	3.115	1.969	1.808	3.058	4.358
$\Delta \varepsilon_{ijkt}^q$	1854141	0.030	1.200	-0.631	0.026	0.687
$\Delta \varepsilon_{ijkt}^p$	1854141	-0.002	0.672	-0.224	-0.000	0.221
$a_{ijkt} - \varepsilon_{ijk,t-1}^q$	1854141	-0.052	3.425	-1.340	-0.041	1.180
a_{ijkt}	1854141	-0.003	0.562	-0.261	0.001	0.256
σ_k	1854141	6.205	4.783	3.566	5.089	6.593
$\text{sd}(a_{ijkt})$	1848126	2.603	1.493	1.895	2.267	2.802
Age_{ijkt}^1	1854141	3.505	1.800	2	3	4
Age_{ijkt}^2	1854141	3.671	1.851	2	3	5
Age_{ijkt}^3	1854141	3.759	1.851	2	3	5

Source: Authors' computations from French Customs data. $\ln q_{ijkt}$ and $\ln p_{ijkt}$ are the logs of quantities and prices sold by a firm i in a market jk a given year t . ε_{ijkt}^q and ε_{ijkt}^p are respectively the belief of the firm about future demand from equation (14) and the residuals of the price equation from equation (15). Age_{ijkt}^1 is the number of years since the last entry of the firm on market jk (reset to zero after one year of exit). Age_{ijkt}^2 : reset after 2 years of exit; Age_{ijkt}^3 : years of exporting since first entry (never reset to zero). a_{ijkt} is our estimate of the demand shock from equation (17). σ_k : elasticity of substitution from equation (17). $\text{sd}(a_{ijkt})$ is the standard deviation of a_{ijkt} , computed by market (product-destination).

Table 1 contains some descriptive statistics about our final sample. Firms are typically young in the markets they serve: the average age is comprised between 3.5 and 3.8 years depending on the definition (note that since we focus on $\Delta \varepsilon_{ijkt}^q$ in the following, firms that exit during the first year are dropped and 2 is the minimum value that our age variable can take). This is evidence of the low survival rates observed during the first years a firm serves a particular market (Figure 1.a). Over the period, the firm-market specific beliefs have been characterized by a positive average growth, while $\Delta \varepsilon_{ijkt}^p$ is slightly negative on average.

Our methodology generates reasonable estimates of σ_k : we get a median value of 5.1 and an average of 6.2 in our final sample. These numbers are comparable to the ones found by the literature, using very different methodologies and data.²⁰ Our estimates of σ_k also follow expected patterns: considering Rauch (1999)'s classification, the median (resp. mean) across products is 5.2 (resp. 6.1) for differentiated goods, 7.3 (resp. 8.6) for referenced priced goods and 8.9 (resp. 10.1) for goods classified as homogenous. These means and medians of σ_k are statistically different across the three groups.²¹

4.2 Baseline results

The results obtained when estimating equation (19) are provided in Table 2. The first column considers separately the effect of demand shocks and age on changes in firms' beliefs. Columns (2)

²⁰See Imbs and Mejean (2015) for a detailed literature review.

²¹See section B of the online appendix for details. Note that these numbers are slightly higher than the means and medians displayed in Table 1 because they are computed across products, while the statistics in Table 1 are based on our final sample, i.e. also reflect the number of French firms selling each product.

Table 2: Prediction 1: demand shocks and beliefs updating

Dep. var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				$\Delta \varepsilon_{ijk,t+1}^q$			
$a_{ijkt} - \varepsilon_{ijkt}^q$	0.064 ^a (0.001)	0.074 ^a (0.001)	0.074 ^a (0.001)			0.047 ^a (0.007)	0.047 ^a (0.006)
× Age _{ijkt}		-0.003 ^a (0.000)	-0.003 ^a (0.000)				
× Age _{ijkt} = 2				0.069 ^a (0.001)	0.069 ^a (0.001)	0.022 ^a (0.007)	0.022 ^a (0.005)
× Age _{ijkt} = 3				0.064 ^a (0.001)	0.064 ^a (0.001)	0.017 ^b (0.007)	0.017 ^a (0.005)
× Age _{ijkt} = 4				0.060 ^a (0.002)	0.060 ^a (0.001)	0.013 ^c (0.007)	0.013 ^b (0.005)
× Age _{ijkt} = 5				0.057 ^a (0.002)	0.057 ^a (0.002)	0.010 (0.007)	0.010 ^b (0.005)
× Age _{ijkt} = 6				0.058 ^a (0.002)	0.058 ^a (0.002)	0.011 (0.007)	0.011 ^b (0.005)
× Age _{ijkt} = 7				0.054 ^a (0.002)	0.054 ^a (0.002)	0.007 (0.006)	0.007 (0.005)
× Age _{ijkt} = 8				0.052 ^a (0.003)	0.052 ^a (0.002)	0.004 (0.006)	0.004 (0.005)
× Age _{ijkt} = 9				0.052 ^a (0.004)	0.052 ^a (0.003)	0.005 (0.006)	0.005 (0.005)
× Age _{ijkt} = 10				0.047 ^a (0.007)	0.047 ^a (0.005)		
Age _{ijkt}	-0.033 ^a (0.001)	-0.033 ^a (0.001)	-0.033 ^a (0.000)				
Observations	1854141	1854141	1854141	1854141	1854141	1854141	1854141

Robust standard errors clustered by firm in parentheses (bootstrapped in columns (3), (5) and (7)). ^c significant at 10%; ^b significant at 5%; ^a significant at 1%. Age dummies included alone in columns (4) to (7) but coefficients not reported. Columns (6) and (7) are the same as column (4) and (5) except that coefficients are estimated relative to the baseline omitted category, age of ten years. a_{ijkt} is our estimate of the demand shock from equation (17); ε_{ijkt}^q is the belief of the firm about future demand from equation (14). Age_{ijkt} is the number of years since the last entry of the firm on market jk (reset to zero after one year of exit).

to (7) study how age affects the reaction of beliefs to demand shocks. Columns (3), (5) and (7) are equivalent to columns (2), (4) and (6) with standard errors being bootstrapped rather than clustered by firm, to account for the fact that the right hand side variables have been estimated.

As predicted, firms update their beliefs positively when they face a higher than expected demand, and the growth in beliefs declines with age on average (column (1)). More importantly, we find support for our key prediction: belief updating following a demand shock is significantly stronger when firms are young (columns (2)-(7)). Including age linearly (column (2) and (3)) or through bins (columns (4) to (7)) leads to the same conclusion. Bootstrapping the standard errors also leaves the results unaffected.

After a decade of presence in the market, the magnitude of belief updating following a given

demand shock is 30 percent smaller than after entry. In columns (6) and (7), we find that, when compared to the benchmark category – age of ten years –, the coefficients of the first four years (first six years with bootstrapped standard errors) are significantly higher. The shape of the learning process is consistent with the theory: age has a strong effect in early years, and matters less for more experienced firms (section C in the online appendix provides a graphical depiction of the result and a discussion of our functional form assumption). Note that most of our estimated coefficients are statistically different from each other up to year seven, which supports the existence of a learning process over this time horizon. After seven years, our results no longer provide clear evidence of learning (note however that the coefficient of the last category is less precisely estimated due to the small number of observations). However, even the most experienced firms in our sample still significantly adjust their quantities following demand shocks. Assuming that part of the demand signals received is persistent would explain this finding: in that case, experienced firms would continue to adjust their quantities to demand shocks even if they have fully discovered their idiosyncratic demand.

4.3 Learning and market uncertainty

Our second prediction is that a higher level of uncertainty in the market (a higher σ_ε^2 in the model) should slowdown the updating process. The underlying intuition is that a demand signal is less informative when uncertainty is higher. It follows that the speed at which firms update their beliefs should decrease with age, but less so when uncertainty is larger (see proof of prediction 1 in the appendix).

We use our theory-based measure of market uncertainty (the standard deviation of a_{ijkt} , computed by product and destination over the entire period). We then add to specification (19) an interaction term between our uncertainty measure and $(a_{ijkt} - \varepsilon_{ijkt}^q)$, and a triple interaction between age, $(a_{ijkt} - \varepsilon_{ijkt}^q)$ and uncertainty (as well as an interaction term between age and uncertainty). Table 3 contains the results. Column (1) shows that, as predicted, the extent of belief updating following a demand shock is smaller in markets characterized by a higher level of uncertainty. On the other hand, the coefficient on the interaction term between age and the demand shocks is virtually unaffected. Quantitatively, the role of uncertainty is non negligible. A standard deviation increase from the mean of the level of uncertainty decreases the response of beliefs to demand shocks from 0.090 to 0.082 in column (1).

Moreover, when uncertainty is large, gaining experience has a lower effect on belief updating, as shown by the coefficient of the triple interaction term in column (2). Another way to represent these results is to separate the sample into high and low uncertainty markets, defined according to the sample median of our uncertainty measure. We run our baseline specification (column (4) of Table 2) separately on each of the two sub-samples. The results are displayed in column (3) and (4) of Table 3. We clearly see that the average extent of belief updating is much larger in markets with low uncertainty levels, and that updating decreases more with age in the least uncertain markets. In the online appendix, section D we use bins of age categories and a more extreme sample split (first and last quartile of uncertainty). In these specifications, we find that the updating coefficient decreases from 0.171 in the second year to 0.128 after ten years in the least uncertain markets, while in the most uncertain markets the relationship is flatter and updating is almost nonexistent as the coefficients decrease from 0.035 to 0.021.

Table 3: Prediction 1: the role of uncertainty

Dep. var. Sample	(1)	(2)	(3)	(4)
		$\Delta \varepsilon_{ijk,t+1}^q$	High	Low
	Uncertainty			
$a_{ijkt} - \varepsilon_{ijkt}^q$	0.102 ^a (0.006)	0.113 ^a (0.005)	0.054 ^a (0.001)	0.163 ^a (0.002)
× Age _{ijkt}	-0.003 ^a (0.000)	-0.006 ^a (0.001)	-0.002 ^a (0.000)	-0.007 ^a (0.000)
× Uncertainty	-0.005 ^a (0.001)	-0.006 ^a (0.001)		
× Age _{ijkt} × Uncertainty		0.001 ^a (0.000)		
Age _{ijkt}	-0.033 ^a (0.000)	-0.027 ^a (0.001)	-0.037 ^a (0.001)	-0.028 ^a (0.001)
Uncertainty	-0.004 ^b (0.001)	0.004 ^c (0.002)		
Age _{ijkt} × Uncertainty		-0.002 ^a (0.000)		
Observations	1848126	1848126	928963	919146

Robust standard errors clustered by firm in parentheses. ^c significant at 10%; ^b significant at 5%; ^a significant at 1%. a_{ijkt} is our estimate of the demand shock from equation (17); ε_{ijkt}^q is the belief of the firm about future demand from equation (14). Age_{ijkt} is the number of years since the last entry of the firm on market jk (reset to zero after one year of exit). Uncertainty is the standard deviation of a_{ijkt} , computed by market jk . High and low uncertainty mean above and below sample median.

5 Robustness

In this section we first assess the implications of several key assumptions of our model for our identification strategy and the interpretation of our results. We then discuss how our results might be affected by endogenous exit, before considering a series of additional sensitivity tests.

5.1 Modelling assumptions

Our model makes three important assumptions. First, firms set their quantities before observing the demand realization, as in Jovanovic (1982). Second, firms face CES demand and monopolistic competition (hence markups are constant). Third, firm productivity is not market-specific. In this subsection we assess the sensitivity of our results to these hypotheses (we discuss the validity of our demand-side modelling of learning in section 6). In particular, we show how they affect (i) the identification of beliefs and demand signals and (ii) our test of prediction 1. Relaxing these assumptions implies in general that the residuals ε_{ijkt}^q can no longer be interpreted as reflecting beliefs only. However – and provided that we control for market-specific firm size in some cases –, these extensions do not alter the qualitative interpretation of our results, in the sense that our baseline estimates of Table 2 can still be viewed as evidence of belief updating. For each extension, we summarize here the main intuitions and refer the reader to the online

appendix E for details.

5.1.1 Fixed quantities

We have assumed so far that quantities are set before firms observe their idiosyncratic demand in each market, while prices adjust to the demand shocks. We relax this assumption in two directions: we start by considering the possibility that prices are set first, with or without a constant price elasticity. Second, we assume that firms can adjust their quantity decision after observing part of the demand shock.

If we completely reverse our assumption and suppose that prices are set *ex-ante* while quantities fully adjust to demand shocks, due to CES demand, prices will only depend on supply side characteristics. They take the form of a constant mark-up over marginal costs and do not vary with the quantity produced, the firm’s beliefs or the demand shock. Quantities on the other hand fully adjust and depend solely on the demand shocks. Regressing ε_{ijkt}^p on ε_{ijkt}^q should therefore generate insignificant $\hat{\beta}$ coefficients, and ε_{ijkt}^q should not vary with age. Both these predictions are clearly at odds with our findings.

Now, assume that prices are set *ex-ante* but the market structure is oligopolistic, which implies variable markups. In this case, prices reflect the firm’s beliefs, as markups depend on its expected market share. Quantities reflect both these beliefs and the demand shocks. We can still estimate demand signals, but our identification strategy should be reversed: ε_{ijkt}^q should be regressed on ε_{ijkt}^p , and the updating process should be observed on $\Delta\varepsilon_{ijkt}^p$. The main prediction of such a model is that a positive demand shock should lead firms to update upwards their beliefs, which would increase their markup and their prices. In the online appendix (section E.2), we follow this alternative methodology and find that prices slightly *decrease* with demand shocks, which is inconsistent with this alternative model of Bertrand competition with a non-constant price elasticity.

Finally, we consider an intermediate case where firms can revise their quantity decision after observing part of the demand shock. In this case, our theoretical predictions still hold, but the identification of the demand shock is affected: ε_{ijkt}^q now also captures part of the demand shock and becomes a noisy measure of the firm’s belief. This may affect our estimates of the demand shocks, although the direction of this bias is unclear. Yet, unless this bias is correlated with age, our main results that young firms update more their beliefs should not be affected. One way to gauge the importance of this possible bias is to focus on sectors or destinations for which quantities are more likely to be rigid – i.e. those for which the demand shocks are more likely to be correctly estimated – and to compare the results with our baseline estimates of Table 2. We expect less quantity adjustment for complex goods (in which many different relationship-specific inputs are used in the production process) and in destinations characterized by longer time-to-ship. In section E.3 of the online appendix, we restrict our sample to sectors or destinations which are above the sample median in terms of time-to-ship or input complexity. The estimated magnitude of belief updating and the coefficient on the interaction terms between demand shocks and age are quantitatively similar to our baseline estimates.²² Altogether, these results suggest

²²The coefficient on the interaction term between demand shocks and age is slightly lower than our baseline in the case of complex goods (col. (5) of Table A.7). In column (6), however, we see that this result is only driven by the effect of the last age category, 10 years of experience, which is itself quite imprecisely estimated.

that our assumption of fixed quantities is not unrealistic and does not lead our identification strategy to artificially generate our results.

5.1.2 Other extensions and control for size

Our next two extensions allow respectively for variable mark-ups and for productivity to be market-specific. We reach similar conclusions in both cases. ε_{ijkt}^q can no longer be interpreted as beliefs about demand only – it is also affected by mark-ups or productivity. $\Delta\varepsilon_{ijkt+1}^q$ therefore reflects changes in beliefs as well as variations in mark-ups or productivity. The key point, however, is that we are still able to interpret the reaction of firms to demand shocks as evidence of belief updating, provided that we control for size. The complete derivations are provided in sections E.4 and E.5 of the online appendix.

Variable mark-ups. The first implication of variable mark-ups for our empirical strategy is that prices could now depend on local market conditions, i.e. the price equation (13) should include a set of jkt fixed-effects. Columns (1) and (2) of Table 4 at the end of this section shows that this modification leaves our results largely unchanged.

Second and more importantly, the quantities residuals ε_{ijkt}^q should now capture the firms' beliefs, but also their expected markups. Hence, changes in expected mark-ups should affect $\Delta\varepsilon_{ijkt+1}^q$. To take into account this possibility, we extend the model to an oligopolistic market structure. Formally, we simply assume that the number of competitors in each sector K , Ω_{kt} , is small enough so that each competitor takes into account the impact of his own decisions on the sectoral price index. As shown in the online appendix, our methodology still produce unbiased estimates of the demand shock. Our main equation however becomes:

$$\Delta\varepsilon_{ijkt+1}^q = g_t \left(a_{ijkt} - \varepsilon_{ijkt}^q \right) + g_t \frac{\sigma_\varepsilon^2}{2\sigma_k} - \sigma_k g_t \ln \left(\frac{\mathbb{E}_{t-1} [\varepsilon(s_{ijkt})]}{\mathbb{E}_{t-1} [\varepsilon(s_{ijkt})] - 1} \right) - \sigma_k \Delta \ln \left(\frac{\mathbb{E}_t [\varepsilon(s_{ijkt+1})]}{\mathbb{E}_t [\varepsilon(s_{ijkt+1})] - 1} \right)$$

where $\mathbb{E}_{t-1} [\varepsilon(s_{ijkt})]$ is the expected elasticity of demand faced by firm i in market jk at the beginning of period t , which itself depends on the expected market share $\mathbb{E}_{t-1} [s_{ijkt}]$. With variable mark-ups, our main equation includes two new terms.

The first term is the level of the expected mark-ups. It comes from the fact that the expected mark-up also affects our measure of beliefs, ε_{ijkt}^q , and in turn $\left(a_{ijkt} - \varepsilon_{ijkt}^q \right)$. We thus need to control for firm size/market share to avoid a standard omitted variable bias.

The second term captures the change in expected mark-ups, and it depends on the updating process through the change in the expected market share. Our measure of belief updating is now underestimated: when firms update positively, they tend to increase their quantities but also their prices, which dampens their overall quantity reaction. It follows that in the case of variable mark-ups, ε_{ijkt}^q becomes an increasing function of firm's beliefs²³ and we only capture the overall reaction of purged quantities to belief updating. Put differently, our results still provide evidence for the updating process, but in a qualitative sense.

Importantly, two firms of different sizes may not have the same mark-up reaction to a given belief update. This is another reason to control for market share: to be able to compare the extent of updating of firms of different age, but with the same market share.

²³Formally, we derive in the online appendix two alternatives sufficient conditions ensuring that the overall quantity response to a positive updating is still positive: either $\sigma_k \geq 2$ or $s_{ijkt} \leq 1/2$.

Product-destination productivity. In the model, we have assumed that productivity was firm-product-specific. Here we relax this assumption and consider the case of product-destination-specific productivity. This again introduces a new source of dynamics in ε_{ijkt}^q . We assume that the unit cost of producing good k for market j at time t is $\frac{w_{it}}{\varphi_{ikt}} \frac{1}{\varphi_{ijkt}}$. This could reflect differences in productivity for the same good across markets, but also differences in product quality. Again, our methodology still produces unbiased estimates of the shock, as shown in the online appendix. But the dynamics of quantities now also reflects the evolution of φ_{ijkt} . We get:

$$\Delta \varepsilon_{ijkt+1}^q = g_t \left(a_{ijkt} - \varepsilon_{ijkt}^q \right) + g_t \frac{\sigma_\varepsilon^2}{2\sigma_k} + \sigma_k g_t \ln(\varphi_{ijkt}) + \sigma_k \Delta \ln(\varphi_{ijkt+1}).$$

As for the case of variable mark-ups, because ε_{ijkt}^q contains a new element, our equation now has two additional terms: one in level because $\ln(\varphi_{ijkt})$ alters our measure of beliefs, and one in difference $\Delta \ln(\varphi_{ijkt+1})$, because $\Delta \varepsilon_{ijkt+1}^q$ also reflects the dynamics of productivity. Again, the first term implies that we need to control for firm size, to avoid a standard omitted variable bias. Second, the dynamics of $\ln(\varphi_{ijkt})$ also affects $\Delta \varepsilon_{ijkt+1}^q$. If this dynamics is uncorrelated with the updating process, the interpretation of our results should be unaffected. If however $\Delta \ln(\varphi_{ijkt+1})$ is positively affected by the updating process – if a positive updating leads firms to invest to improve φ_{ijkt} – then our measure of updating becomes a measure of the overall impact of the updating process on $\Delta \varepsilon_{ijkt+1}^q$: it does not only capture the updating process itself but also how the quantity response is magnified by a change in productivity.²⁴ Again, ε_{ijkt}^q would become an increasing function of firm’s beliefs, and our evidence of the updating process would become qualitative as we would not identify firms’ beliefs *per se*. This productivity response could be size dependent, which again requires to control for firm size. The decline of the overall response of $\Delta \varepsilon_{ijkt+1}^q$ to demand shocks over time, conditional on size, however still provides evidence for an updating process.

Controlling for size. The two extensions of the models discussed above suggest that firm-market size should be included in our regressions, together with its interaction with firm-market age. We do so in Table 4. Columns (1) and (2) are similar to our baseline regressions (Table 2, columns (2) and (4)), except that jkt fixed effects are introduced in the estimation of the price residuals ε_{ijkt}^p , as predicted by models with variable markups. The average level of belief updating is slightly larger than in our baseline estimates, but the effect of age is similar. In columns (3) to (6) we additionally control for firm size, as measured by the value sold by firm i on market jk during year $t - 1$ divided by the total value exported by French firms in market jk during year $t - 1$. Size is introduced either linearly in columns (3) and (4) or through bins computed using market-specific deciles in columns (5) and (6). Our coefficients of interest are extremely stable across specifications.²⁵ In the online appendix E.6, we consider a number of alternative measures of firm size and include interaction terms between size and $a_{ijkt} - \varepsilon_{ijkt}^q$ to account for the fact that age and size are correlated. In all instances the results are similar to our benchmark estimates.

²⁴Note that this possibility does not violate the orthogonality conditions that we need to identify demand shocks. As discussed in section 3.3, we need the beliefs a firm in market jk at time t to be orthogonal to overall firm-product characteristics; yet, beliefs can be correlated with the characteristics of a firm-product in that particular market j .

²⁵The positive coefficient on age in column (5) may appear surprising at first, but this coefficient cannot be directly interpreted as this estimation also includes a full set of interaction terms between age and size.

Table 4: Prediction 1: controlling for size

	(1)	(2)	(3)	(4)	(5)	(6)
Dep. var.			$\Delta \varepsilon_{ijk,t+1}^q$			
Robustness	Controlling for FE _{jk_t} in prices		Controlling for FE _{jk_t} in prices and size			
Size			Linear		Bins	
$a_{ijkt} - \varepsilon_{ijkt}^q$	0.103 ^a (0.002)		0.103 ^a (0.002)		0.102 ^a (0.002)	
× Age _{ijk_t}	-0.003 ^a (0.000)		-0.003 ^a (0.000)		-0.003 ^a (0.000)	
× Age _{ijk_t} = 2		0.096 ^a (0.002)		0.096 ^a (0.002)		0.096 ^a (0.002)
...						
× Age _{ijk_t} = 10		0.074 ^a (0.009)		0.074 ^a (0.009)		0.075 ^a (0.009)
Age _{ijk_t}	-0.034 ^a (0.001)		-0.040 ^a (0.001)		0.019 ^a (0.002)	
Size _{ijk_t}			-1.053 ^a (0.016)	-1.015 ^a (0.017)		
× Age _{ijk_t}			0.109 ^a (0.003)	0.101 ^a (0.003)		
Observations	1870377	1870377	1870377	1870377	1501840	1501840

Robust standard errors clustered by firm in parentheses. ^c significant at 10%; ^b significant at 5%. ^a significant at 1%. a_{ijk_t} is our estimate of the demand shock from equation (17). In this table, jk_t fixed effects are included in the estimation of the price residuals $\varepsilon_{ijk_t}^p$ used to identify demand shocks. $\varepsilon_{ijk_t}^q$ is the belief of the firm about future demand from equation (17). Age_{ijk_t} is the number of years since the last entry of the firm on market jk (reset to zero after one year of exit). Size_{ijk_t} is proxied by the value sold by firm i on market jk during year t divided by the total value exported by French firms in market jk during year t . Columns (5) and (6) include size bins corresponding to the ten deciles of size variable, computed by market-year. Age dummies included alone in columns (2), (4) and (6) but coefficients not reported. See Table A.8 in the online appendix for the full set of coefficients on the interaction terms.

5.2 Survival and selection bias

Our main prediction is tested on the sample of firms which survive in period t . Endogenous sample selection could be a concern in equation (19). The error term u_{ijk_t} might be correlated in particular with demand shocks: the observed sample includes firms with relatively positive demand shocks (as those with negative shocks are more likely to exit), and firms which do not update downward their beliefs too much following a negative signal (otherwise they would exit). In other words, endogenous exit might create a correlation between the error term of (19) and demand shocks.

The predictions of the learning model for survival are discussed in details in section F of the online appendix. We show that exit probability depends (negatively) on demand signals, age, and beliefs, as well as on the dynamics of firm productivity and market conditions. Predicted exit probabilities can therefore be estimated as a function of a_{ijk_t} , $\varepsilon_{ijk_t}^q$, Age_{ijk_t} and fixed effects in the ikt and jk_t dimensions. We use a linear probability model which allows the inclusion of our two high-dimensional fixed effects. Once these survival probabilities have been estimated, we

perform two different types of exercises to check that our results are not affected by endogenous selection.

First, we gauge the importance of this selection bias by estimating (19) on sub-samples defined according to the survival probability. This is an application of the “identification-at-infinity” method (Chamberlain, 1986; Mulligan and Rubinstein, 2008). The general idea is to restrict the estimation sample to firms that are most likely to survive, the selection bias being lower for firms with high survival probability. We allocate firms in 5 bins of survival probability and estimate (19) on sub-samples that include only firms above the 20th, 40th, 60th and 80th percentiles of survival probability. The results are presented in Table A.12 (section H) in the online appendix. Starting from the full sample in column (1), we progressively drop the quintiles of observations with the highest exit probabilities from the sample. Accordingly, column (5) only includes the quintiles of observations with the lowest exit probabilities (i.e. the highest survival probability). If endogenous exits were driving our results, we would expect the patterns of belief updating to substantially differ across samples. On the contrary, we find that the coefficients on $(a_{ijkt} - \varepsilon_{ijkt}^q)$ and its interaction with age are extremely stable across different bins of survival probability.²⁶

Table 5: Demand shocks and beliefs updating: controlling for endogenous exit

Dep. var.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Selection correction		$\Delta \varepsilon_{ijkt+1}^q$ Linear				$\Delta \varepsilon_{ijkt+1}^q$ Semi-parametric		
$a_{ijkt} - \varepsilon_{ijkt}^q$	0.065 ^a (0.001)	0.075 ^a (0.002)	0.075 ^a (0.001)		0.065 ^a (0.001)	0.075 ^a (0.002)	0.075 ^a (0.001)	
× Age _{ijkt}		-0.003 ^a (0.000)	-0.003 ^a (0.000)			-0.003 ^a (0.000)	-0.003 ^a (0.000)	
× Age _{ijkt} = 2				0.069 ^a (0.001)				0.069 ^a (0.001)
...								
× Age _{ijkt} = 9				0.054 ^a (0.007)				0.054 ^a (0.007)
$\widehat{\text{Pr}}(\text{exit}_{ijkt})$	-0.409 ^a (0.005)	-0.409 ^a (0.005)	-0.409 ^a (0.003)	-0.417 ^a (0.005)				
Age _{ijkt}	-0.054 ^a (0.001)	-0.054 ^a (0.001)	-0.054 ^a (0.000)		-0.057 ^a (0.001)	-0.057 ^a (0.001)	-0.057 ^a (0.001)	
Observations	1501766	1501766	1501766	1501766	1501766	1501766	1501766	1501766

Robust standard errors clustered by firm in parentheses (bootstrapped in columns (3) and (7)). ^c significant at 10%; ^b significant at 5%; ^a significant at 1%. Age dummies included alone in columns (4) and (8) but coefficients not reported. a_{ijkt} is our estimate of the demand shock from equation (17); ε_{ijkt}^q is the belief of the firm about future demand from equation (14). Age_{ijkt} is the number of years since the last entry of the firm on market jk (reset to zero after one year of exit). In columns (1)-(4), predicted exit probabilities are obtained from the estimation of Table A.11, column (4) and introduced directly in equation (19). In columns (5) to (8), they are introduced semi-parametrically in the second step, i.e. we included 100 bins corresponding to each percentile of the variable. Online appendix Table A.14 reports the full set of coefficients.

These results suggest that endogenous exit does not bias our results. We can go further and try to account for a potential selection bias by including a correction term in our estimations. Given the structure of our selection equation (which includes two high dimensional sets of fixed

²⁶In the online appendix (Table A.13), we perform a similar analysis but define quintiles based on both exit probability and firm-market size. The results are similar.

effects), we cannot use probit or other maximum likelihood estimators to implement a standard Heckman procedure. Instead, we follow Olsen (1980) and include a correction term constructed from a linear estimation of the selection equation. Crucially, Olsen’s correction term is linear, which implies that the selection equation needs to include variables which do not appear in the second step.²⁷ This is not a problem in our case as ikt and jkt fixed effects can be used as exclusion variables. Results appear in Table 5, columns (1) to (4) and are again close to our baseline estimates. Alternatively, we can relax the linearity assumption of the correction term and use a partially linear approach in the second step. More precisely, as suggested by Cosslett (1991), we replace the linear correction term by a hundred indicator variables constructed from predicted exit probabilities. Results are provided in Table 5, columns (5) to (8). Again our coefficients of interest are largely unaffected.²⁸

5.3 Measurement issues

In section I of the online appendix, we perform some additional robustness checks. In particular: (i) we restrict the sample to extra EU destinations to account for the different treatment of EU trade flows by the customs (section I.1); (ii) we use alternative definitions of firm age (section I.2); (iii) we reconstruct the years, beginning the month of the first entry at the firm-product-destination level, to account for the fact that the first year of export measured over a calendar year is potentially incomplete, as pointed out by Berthou and Vicard (2015) and Bernard et al. (2017), which can affect growth rates in the first period (section I.3); (iv) we replicate the results with equation (17) being estimated at the 4-digit (HS4) instead of 6-digit level, as some 6-digit categories might include few observations, leading to imprecise estimates (section I.4); and (v) we re-estimate ε_{ijkt}^q and ε_{ijkt}^p including ijt fixed effects in equations (14) and (15) to control for the potential informational spillovers from selling other products in the same destination (section I.5). Each set of results is discussed in details in the online appendix. In all cases, they are extremely close to our baseline estimates of Table 2.

Overall, the results presented in this section show that both the magnitude of belief updating and its age dependence are extremely stable across various samples and specifications, which strongly suggests that our findings are not driven by specific sectors, firms or modelling assumptions.

6 Discussion: alternative mechanisms on the demand side?

Several alternative demand side mechanisms have been proposed in the literature to explain firm dynamics. They mainly give rise to demand accumulation, either endogenously or exogenously. A first category of models considers firms engaging in market-specific investment to increase their profitability, or in a costly search for new buyers (see for instance Ericson and Pakes, 1995, Luttmer, 2011, Eaton et al., 2014, Fitzgerald et al., 2016). A second possibility is

²⁷See Vella (1998) for a summary of Olsen (1980) and alternative procedures to correct for endogenous sample selection. More details about the procedure appear in the online appendix, section H.

²⁸Online appendix Table A.15 shows the results of an alternative semi-parametric procedure consisting in using a polynomial expansion of the first-step prediction as a correction term. We also try a standard Heckman procedure, estimating the first step by probit without fixed effects and relying on the nonlinearity of the inverse mills ratio in the second step to identify the selection term.

that firms price low in their first years to build a consumer base (Foster et al., 2016, Gourio and Rudanko, 2014). Finally, demand could simply evolve exogenously over time as in Ruhl and Willis (2017). All these mechanisms would generate the increase in average sales observed over time for surviving firms that we documented in section 2.2 (Figure 2) – and this is precisely the stylized fact that motivated many of these papers. As already underlined, models of demand accumulation, if they do not include some learning about demand, cannot deliver our main prediction, i.e. that firms adjust less and less their quantities to past demand shocks as they grow older in a market. Yet, we cannot exclude *a priori* the possibility that some demand accumulation is at play on top of the updating process. Put differently, our estimates of ε_{ijkt}^q , which we interpret as beliefs, could in theory reflect other types of dynamics of market-specific demand. In this section we first show that our assumption of firms learning about a constant demand parameter is consistent with our data, i.e. that variations in ε_{ijkt}^q and ε_{ijkt}^p can indeed be interpreted as being driven, at least to a first order, by the updating process.²⁹ We then show that the variance of estimated beliefs explains a large part of the observed variance of sales growth within cohort.

6.1 Dynamics of ε_{ijkt}^q and ε_{ijkt}^p

To further check the validity of the model, we study how the quantities and prices residuals ε_{ijkt}^q and ε_{ijkt}^p vary with age within cohorts, as the predictions of the learning model differ from those of demand accumulation theories. In the passive learning model, the dynamics of ε_{ijkt}^q and ε_{ijkt}^p are affected by both within firm-markets dynamics and selection effects. Indeed, conditional on age and fixed effects, the decision to stay or exit the market depends on the firm’s beliefs: there is a threshold value below which firms exit the market.³⁰ Exit decisions thus depend on the beliefs at the beginning of the period and on the demand shocks received. First, for a given demand shock, the smallest firms – firms with the lowest ε_{ijkt}^q – are more likely to exit. Second, for a given level of beliefs, firms that decrease in size – those facing negative demand shocks – exit more. Therefore, survivors are firms that received positive demand shocks on average.

Dynamics of ε_{ijkt}^q . Both effects imply that, conditional on survival, ε_{ijkt}^q should grow on average over time within cohorts. This is due to composition effects: as prior beliefs are unbiased on average, firms have equal probabilities to update upward or downward. Hence, when focusing on within firm-market variations (i.e. controlling for firm-product-destination fixed effects), quantities should become much flatter.³¹ This is indeed what we find in Figure 2 (the complete set of coefficients and standard errors is provided in Tables A.23 and A.24 in the online appendix K). Figure 2.a plots the coefficients obtained when we simply regress ε_{ijkt}^q on firm-market age: ε_{ijkt}^q sharply increases with age. When instead we focus on variation within firms-markets (Figure 2.b), ε_{ijkt}^q becomes almost flat: it only exhibits a slight positive growth in the first years, especially at age 2. This is mostly due to the incompleteness of the first year of export measured over the calendar year; as shown in Figure 2.c, when years are reconstructed to start the month

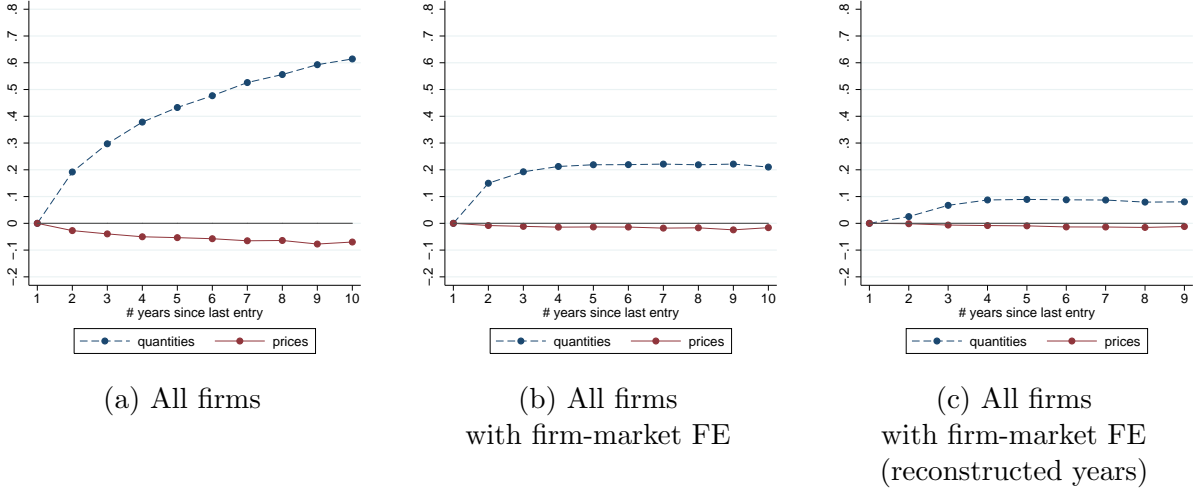
²⁹In online appendix section J, we also implement a test proposed by Pakes and Ericson (1998) to discriminate between models of “active” and “passive” learning. This test mainly shows that firms’ beliefs do not follow a Markov process, a feature that should be natural in models of demand accumulation where the decision to accumulate demand depends on firm size.

³⁰See online appendix section F for details.

³¹The passive learning model actually generates positive *unconditional* growth rates of quantities, i.e. even in the absence of composition effects triggered by selection. As shown in the online appendix G, this is however a weaker prediction, as it is driven by functional form assumptions.

of the first entry,³² the increase observed in the second year almost vanishes. After 3 years, ε_{ijkt}^q is only 9 percent higher than at the time of entry, and remains constant afterwards.

Figure 2: Dynamics of ε_{ijkt}^q and ε_{ijkt}^p



Note: This figure plots the coefficients obtained when regressing the prices and quantities residuals ε_{ijkt}^q and ε_{ijkt}^p on a set of age dummies. Age is defined at the firm-market (firm-product-destination) level. Panel (b) controls for firm-product-destination fixed effects. Panel (c) considers the same specification as panel (b) but on the dataset of reconstructed years (see section I.3 in the online appendix). The complete set of coefficients and standard errors are shown in the online appendix Table A.23 (columns (2) and (6) for panel (a) and (4) and (8) for panel (b)) and Table A.24 (columns (3) and (6) for panel (c)).

These results contrast with the prediction of demand accumulation theories. In these models, we would expect quantities to increase more gradually and more strongly over time. Moreover, such an increase should not only be observed in the pooled regressions of Figure 2.a, but also in the within firms-markets estimations of Figures 2.b and 2.c.³³ We do find some growth at early age even after accounting for composition effects, which is consistent both with demand accumulation theories and with the passive learning model. Yet, this growth is extremely limited in magnitude and in duration, which suggests that the role of demand accumulation processes, if any, seems modest in our data at the firm-market level.

Dynamics of ε_{ijkt}^p . When interpreted through the lens of the learning model, ε_{ijkt}^p represents the difference between demand shocks and the firms' expected demand. Composition effects imply that ε_{ijkt}^p should decrease over time. Because they receive positive demand shocks on average, survivors initially set their price above their optimal pricing rule, to "jump" on realized demand. They next update their beliefs, which progressively become more accurate over time. ε_{ijkt}^p should thus decrease on average and converge toward its steady state value. But again, controlling for firm-market fixed effects, ε_{ijkt}^p should remain constant. These predictions are confirmed in Figure 2: without firm-market fixed effects, ε_{ijkt}^p is decreasing in age, although the effect is quantitatively limited (Figure 2.a). This is what the passive learning model predicts as changes in the firm

³²See section I.3 in the online appendix.

³³Note that the pattern shown in our stylized fact #2 and in Figure A.5 of the online appendix should not be seen as evidence of some sort of demand accumulation: even in the learning model, the subsample of firms-markets surviving the entire period grow over time, as they received positive demand shocks on average.

beliefs are supposed to affect more $\Delta\varepsilon_{ijkt}^q$ than $\Delta\varepsilon_{ijkt}^p$.³⁴ Note that all coefficients statistically differ from zero at conventional levels (Table A.23). More importantly, when composition effects are accounted for, prices become flat (Figures 2.b and 2.c).

While consistent with the learning model, these findings are difficult to reconcile with theories of demand accumulation. In models where such accumulation is driven by firm pricing policy (i.e. pricing low in the first years to attract consumers), prices of young firms should be lower than those of experienced exporters: ε_{ijkt}^p should increase over time in Figures 2.b and 2.c. If demand accumulation is not driven by firm pricing, prices should stay constant over time; they should not decline with age as in Figure 2.a.³⁵ Overall, the results shown in Figure 2 therefore support our interpretation of ε_{ijkt}^q and ε_{ijkt}^p as being mostly driven by the updating process.³⁶

6.2 The variance of firms' growth

We have seen in section 2.2 (Figure 1) that the variance of observed growth rates within market-specific cohorts of firms decline with the age of the cohort conditional on size, a fact that does not arise naturally in models where learning is absent.³⁷ On the other hand, with learning, younger firms update more than older firms and so have larger growth rates in absolute value. It follows that the variance of firms growth decreases with the cohort tenure on a specific market. As formally shown in the appendix, we get the following prediction which is a direct consequence of firm updating:

Prediction # 3 (variance of growth rate): *The within cohort variances of growth rates of Z_{ijkt}^q and Z_{ijkt}^p decrease with cohort age.*

We test this prediction by estimating the following equation:

$$\mathbb{V}(\Delta\varepsilon_{ijkt}^X) = \delta^X \times \text{AGE}_{c_jkt} + \mathbf{FE}_{c_jk} + u_{ijkt} \quad \forall X = \{q, p\} \quad (20)$$

where \mathbf{FE}_{c_jk} represent cohort fixed effects. As in section 2.2, a cohort of new exporters on a product-destination market is defined as all firms entering market jk in year t . We again expect our coefficient of interest δ^X to be negative: because firms update less their beliefs when they gain experience in a market, their quantities and prices become less volatile. Using the estimated coefficients from (20), we can also check whether the variance of the growth in beliefs match the observed variance of sales growth.

Figure 3 shows the results. We plot the variance of the growth of quantities (beliefs) and prices residuals, as well as of the predicted value of sales and compare it with the observed variance of

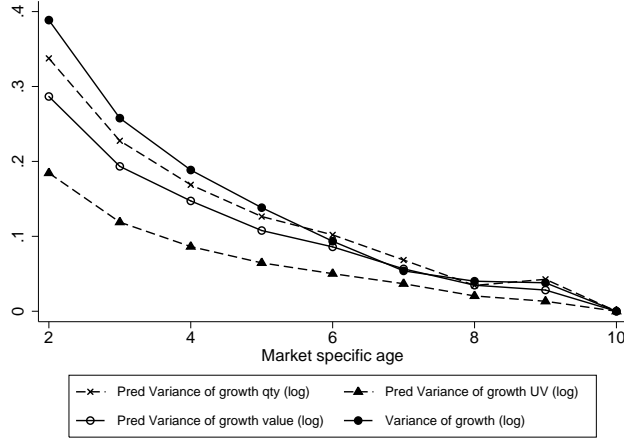
³⁴The magnitude of the difference in growth rates should be a factor σ_k (equations (14) and (15)), which is indeed close to what we find in Table A.23 when comparing the price and quantity equations.

³⁵The price decrease we find in Figure 2.a also suggests that at least part of the updating process we uncover is directly about demand. Indeed, if firms were fully informed about the demand function (and would learn about something else, for instance productivity), they would choose a quantity - prices couple on the demand function and prices should not deviate from the optimal pricing rule.

³⁶This does not imply that demand accumulation processes are not relevant to explain other dimensions of firm dynamics. For instance, firms may accumulate demand due to investment or marketing expenses that affect simultaneously their sales in many markets, or because of product-specific trends in consumer tastes: firms with the "right" product would experience positive growth in demand. Since these elements are purged from our quantities and prices residuals, we cannot infer their importance.

³⁷The literature has however proposed mechanisms allowing to explain the decline in variance of growth rate with size, *conditional* on age (see for instance Luttmer, 2011).

Figure 3: Impact of firm-market specific age conditional on size: predicted patterns



Note: this figure plots the variance of quantities, prices and values residuals within cohort over age on each product-destination market. The coefficients are shown in columns (2) and (6) in table A.25 and column (2) in table A.26 respectively. The coefficients on the growth and variance of growth of sales from figure 1.a.

sales growth. The full set of estimates appear in online appendix L. Within cohort, the variance of the growth rate of both beliefs and prices residuals sharply decreases with age in all columns. Note that this is still true when controlling for the number of observations in the cohort, for average size, or for attrition by concentrating on the firms-markets which survive over the entire period (see Tables A.25 and A.26). The variance of ε_{ijkt}^q follows quite closely the variance of observed sales. Again, given that this decline in variance with age conditional on size cannot be explained by models without learning, this provides further support for the learning model.

7 Conclusion

In this paper we have provided direct evidence that passive learning about demand is an important determinant of firm dynamics. We derived a core prediction from a standard model of market-specific firm dynamics incorporating Bayesian learning about local demand that theories without learning cannot generate: a demand signal leads firms to update their beliefs, especially when they are young. Combining the structure of the model with detailed exporter-level data, we developed a methodology to identify demand shocks and firms beliefs about demand.

The learning process generates the decline in the growth rates and their variance within cohort with firms' age found in the data. Our framework is also consistent with heterogeneous patterns of growth of surviving firms since over-optimistic firms upon entry may experience negative growth. We have focused on a specific dimension of firm dynamics – the post-entry firm behavior at the product-destination level –, yet this dimension explains more than half of the variance in overall firm growth.

Our results open several paths for future research. An implication of the model is that the learning process creates a form of hysteresis: the most experienced firms are less sensitive to demand shocks in terms of sales and exit decisions. This suggests that aggregate uncertainty shocks, thought as an increase in the dispersion of micro-level shocks (Bloom et al., 2014), should

have heterogeneous effects across industries depending on their age structure. Another natural extension of our paper would be to go beyond post-entry dynamics and extend our framework to include explicitly informational spillovers across products, destinations or firms. Such spillovers could affect firms' entry decisions and size upon entry. Quantifying the respective contributions of each of these sources of information to firm dynamics would bear direct policy relevance.

References

- Abbring, Jaap H. and Jeffrey R. Campbell (2005), "A Firm's First Year." Tinbergen Institute Discussion Papers 05-046/3, Tinbergen Institute.
- Albornoz, Facundo, Hector F. Calvo Pardo, Gregory Corcos, and Emanuel Ornelas (2012), "Sequential exporting." *Journal of International Economics*, 88, 17–31.
- Arkolakis, Costas (2016), "A unified theory of firm selection and growth." *Quarterly Journal of Economics*, 131(1), 89–155.
- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott (2009), "The Margins of US Trade." *American Economic Review*, 99, 487–93.
- Bernard, Andrew B., Renzo Massari, Jose-Daniel Reyes, and Daria Taglioni (2017), "Exporter Dynamics and Partial Year Effects." *American Economic Review*, 107, 3211–3228.
- Berthou, Antoine and Vincent Vicard (2015), "Firms' export dynamics: experience vs. size." *World Economy*, 38, 1130–1158.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry (2014), "Really Uncertain Business Cycles." WP 14-18, Center for Economic Studies.
- Buono, Ines and Harald Fadinger (2012), "The micro-dynamics of exporting: evidence from French firms." Working Paper 2012-880, Banca d'Italia.
- Cabral, Luis and Jose Mata (2003), "On the Evolution of the Firm Size Distribution: Facts and Theory." *American Economic Review*, 93, 1075–1090.
- Chamberlain, Gary (1986), "Asymptotic Efficiency in Semiparametric Models with Censoring." *Journal of Econometrics*, 32, 182–218.
- Clementi, Gina Luca and Hugo A Hopenhayn (2006), "A Theory of Financing Constraints and Firm Dynamics." *The Quarterly Journal of Economics*, 121, 229–265.
- Cosslett, Stephen (1991), "Semiparametric estimation of a regression model with sample selectivity." In *Nonparametric and Semiparametric Methods in Econometrics and Statistics* (Barnett. James Powell William A and George Tauchen, eds.), Cambridge University Press.
- Defever, Fabrice, Benedikt Heid, and Mario Larch (2015), "Spatial Exporters." *Journal of International Economics*, 95, 145–156.
- Dunne, Timothy, Mark J Roberts, and Larry Samuelson (1989), "The Growth and Failure of U.S. Manufacturing Plants." *The Quarterly Journal of Economics*, 104, 671–98.

- Eaton, Jonathan, Marcela Eslava, David Jinkins, C. J. Krizan, and James Tybout (2014), “A search and learning model of export dynamics.” Unpublished manuscript.
- Eaton, Jonathan, Marcela Eslava, Maurice Kugler, and James Tybout (2008), “Export dynamics in colombia: Firm-level evidence.” In *The Organization of Firms in a Global Econom* (D. Marin E. Helpman and T. Verdier, eds.), Cambridge, MA: Harvard University Press.
- Eaton, Jonathan, Samuel Kortum, and Francis Kramarz (2011), “An anatomy of international trade: Evidence from french firms.” *Econometrica*, 79, 1453–1498.
- Ericson, Richard and Ariel Pakes (1995), “Markov-Perfect Industry Dynamics: A Framework for Empirical Work.” *Review of Economic Studies*, 62, 53–82.
- Evans, David (1987), “Test of alternative theories of firm growth.” *Journal of Policial Economy*, 95, 657–74.
- Fernandes, Ana and Heiwai Tang (2014), “Learning from Neighbors’ Export Activities: Evidence from Exporters’ Survival.” *Journal of International Economics*, 94, 87–94.
- Fitzgerald, Doireann, Stephanie Haller, and Yaniv Yedid-Levi (2016), “How Exporters Grow.” Staff Report 524, Federal Reserve Bank of Minneapolis.
- Foster, Lucia, John Haltiwanger, and Chad Syverson (2008), “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?” *American Economic Review*, 98, 394–425.
- Foster, Lucia, John Haltiwanger, and Chad Syverson (2016), “The slow growth of new plants: learning about demand?” *Economica*, 83, 91–129.
- Gourio, Francois and Leena Rudanko (2014), “Customer capital.” *Review of Economic Studies*, 81, 1102–1136.
- Guimaraes, Paulo and Pedro Portugal (2010), “A simple feasible alternative procedure to estimate models with high-dimensional fixed effects.” *Stata Journal*, 10, 628–649.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda (2013), “Who Creates Jobs? Small versus Large versus Young.” *The Review of Economics and Statistics*, 95, 347–361.
- Hopenhayn, Hugo A (1992), “Entry, Exit, and Firm Dynamics in Long Run Equilibrium.” *Econometrica*, 60, 1127–50.
- Imbs, Jean and Isabelle Mejean (2015), “Elasticity optimism.” *American Economic Journal: Macroeconomics*, 7, 43–83.
- Impullitti, Giammario, Alfonso A. Irarrazabal, and Luca David Opromolla (2013), “A theory of entry into and exit from export markets.” *Journal of International Economics*, 90, 75–90.
- Jovanovic, Boyan (1982), “Selection and the Evolution of Industry.” *Econometrica*, 50, 649–70.
- Klette, Tor Jakob and Samuel Kortum (2004), “Innovating Firms and Aggregate Innovation.” *Journal of Political Economy*, 112, 986–1018.

- Lawless, Martina (2009), “Firm export dynamics and the geography of trade.” *Journal of International Economics*, 77, 245–254.
- Luttmer, Erzo G. J. (2007), “Selection, Growth, and the Size Distribution of Firms.” *The Quarterly Journal of Economics*, 122, 1103–1144.
- Luttmer, Erzo G. J. (2011), “On the Mechanics of Firm Growth.” *Review of Economic Studies*, 78, 1042–1068.
- Morales, Eduardo, Gloria Sheu, and Andres Zahler (2014), “Extended Gravity.” unpublished manuscript, Princeton University.
- Mulligan, Casey B. and Yona Rubinstein (2008), “Selection, Investment, and Women’s Relative Wages Over Time.” *The Quarterly Journal of Economics*, 123, 1061–1110.
- Munch, Jakob R and Daniel Nguyen (2014), “Decomposing Firm-Level Sales Variation.” *Journal of Economic Behavior and Organization*, 106, 317–334.
- Olsen, Randall J (1980), “A least squares correction for selectivity bias.” *Econometrica*, 48, 1815–1820.
- Pakes, Ariel and Richard Ericson (1998), “Empirical Implications of Alternative Models of Firm Dynamics.” *Journal of Economic Theory*, 79, 1–45.
- Piveteau, Paul (2016), “An empirical model of export entry with customer accumulation.” Unpublished manuscript.
- Rauch, James E. (1999), “Networks versus markets in international trade.” *Journal of International Economics*, 48, 7–35.
- Rossi-Hansberg, Esteban and Mark L. J. Wright (2007), “Establishment Size Dynamics in the Aggregate Economy.” *American Economic Review*, 97, 1639–1666.
- Ruhl, Kim J. and Jonathan L. Willis (2017), “New Exporter Dynamics.” *International Economic Review*, 58, 703–726. Forthcoming.
- Timoshenko, Olga A. (2015), “Product Switching in a Model of Learning.” *Journal of International Economics*, 95, 233–249.
- Vella, Francis (1998), “Estimating Models with Sample Selection Bias: A Survey.” *Journal of Human Resources*, 33, 127–169.

A Appendix - Detailed derivations and proofs

Optimal quantities and prices. Firms choose quantities by maximizing expected profits subject to demand. Using (1), we get:

$$\max_q \int \pi_{ijkt} dG_{t-1}(a_{ijkt}) = \max_q q_{ijkt}^{1-\frac{1}{\sigma_k}} \left(\frac{\mu_k Y_{jt}}{P_{jkt}^{1-\sigma_k}} \right)^{\frac{1}{\sigma_k}} \mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right] - \frac{w_{it}}{\varphi_{ikt}} q_{ijkt} - F_{ijk}$$

The FOC writes:

$$\left(1 - \frac{1}{\sigma_k}\right) q_{ijkt}^{-\frac{1}{\sigma_k}} \left(\frac{\mu_k Y_{jt}}{P_{jkt}^{1-\sigma_k}}\right)^{\frac{1}{\sigma_k}} \mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right] = \frac{w_{it}}{\varphi_{ikt}} \Leftrightarrow q_{ijkt}^* = \left(\frac{\sigma_k}{\sigma_k - 1} \frac{w_{it}}{\varphi_{ikt}}\right)^{-\sigma_k} \left(\frac{\mu_k Y_{jt}}{P_{jkt}^{1-\sigma_k}}\right) \mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right]^{\sigma_k}$$

And from the constraint, we get $p_{ijkt}^* = \left(\frac{\sigma_k}{\sigma_k - 1} \frac{w_{it}}{\varphi_{ikt}}\right) \left(\frac{e^{\frac{a_{ijkt}}{\sigma_k}}}{\mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right]}\right)$

Updating of firm's beliefs about expected demand. First note that firm i has a prior about the demand shock given by $a_{ijkt} \sim \mathcal{N}(\tilde{\theta}_{ijkt-1}, \tilde{\sigma}_{ijkt-1}^2 + \sigma_\varepsilon^2)$ and thus $e^{\frac{a_{ijkt}}{\sigma_k}} \sim LN\left(\frac{\tilde{\theta}_{ijkt-1}}{\sigma_k}, \frac{\tilde{\sigma}_{ijkt-1}^2 + \sigma_\varepsilon^2}{\sigma_k^2}\right)$.

It follows that $\mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right] = \int \left(e^{\frac{a_{ijkt}}{\sigma_k}} \right) dG_{t-1}(a_{ijkt}) = e^{\frac{1}{\sigma_k} \left(\tilde{\theta}_{ijkt-1} + \frac{\tilde{\sigma}_{ijkt-1}^2 + \sigma_\varepsilon^2}{2\sigma_k} \right)}$. Hence:

$$\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right] = \frac{1}{\sigma_k} \left(\Delta \tilde{\theta}_{ijkt} + \frac{\tilde{\sigma}_{ijkt}^2 - \tilde{\sigma}_{ijkt-1}^2}{2\sigma_k} \right)$$

Using the the definition of $\Delta \tilde{\theta}_{ijkt}$, g_t , $\tilde{\sigma}_{ijkt-1}^2$ and $\tilde{\sigma}_{ijkt}^2$ (see (3) and (4)), it is easy to show that $\frac{\tilde{\sigma}_{ijkt-1}^2 - \tilde{\sigma}_{ijkt}^2}{g_t} = \tilde{\sigma}_{ijkt-1}^2$. It follows the expression in the text (11):

$$\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right] = \frac{g_t}{\sigma_k} \left(a_{ijkt} - \tilde{\theta}_{ijkt-1} \right) - \frac{g_t}{\sigma_k} \frac{\tilde{\sigma}_{ijkt-1}^2}{2\sigma_k}$$

But we only observe $\Delta \varepsilon_{ijkt+1}^q = \sigma_k \Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right]$. It follows that $\Delta \varepsilon_{ijkt+1}^q = g_t \left(a_{ijkt} - \tilde{\theta}_{t-1} \right) - g_t \frac{\tilde{\sigma}_{ijkt-1}^2}{2\sigma_k}$. As $\varepsilon_{ijkt}^q = \sigma_k \ln \mathbb{E}_{t-1} \left[e^{\frac{a_{ijkt}}{\sigma_k}} \right]$, we get $\varepsilon_{ijkt}^q = \tilde{\theta}_{ijkt-1} + \frac{\tilde{\sigma}_{ijkt-1}^2 + \sigma_\varepsilon^2}{2\sigma_k}$, or $\tilde{\theta}_{ijkt-1} = \varepsilon_{ijkt}^q - \frac{\tilde{\sigma}_{t-1}^2 + \sigma_\varepsilon^2}{2\sigma_k}$. This leads to the equation we test in the main text:

$$\Delta \varepsilon_{ijkt+1}^q = g_t \left(a_{ijkt} - \varepsilon_{ijkt}^q \right) + g_t \frac{\sigma_\varepsilon^2}{2\sigma_k} \quad (21)$$

Prediction 1. Prediction 1 states that $a_{ijkt} - \varepsilon_{ijkt}^q$, has a larger impact on firms' updating, the younger the firms are. Using (21), we immediately get:

$$\frac{\partial \Delta \varepsilon_{ijkt+1}^q}{\partial \left(a_{ijkt} - \varepsilon_{ijkt}^q \right)} = g_t > 0$$

Updating is larger for younger firms, as g_t decreases with t .

Prediction 2: Impact of market uncertainty. Moreover, the updating process is also affected by the level of market uncertainty σ_ε^2 . Formally:

$$\frac{\partial^2 \left(\Delta \varepsilon_{ijkt+1}^q \right)}{\partial \left(a_{ijkt} - \varepsilon_{ijkt}^q \right) \partial \sigma_\varepsilon^2} = -\frac{g_t^2}{\sigma_{jk0}^2} < 0$$

Updating decreases with uncertainty, as a signal is less informative when market uncertainty is larger. As a consequence, market uncertainty dampens the speed of learning. In other words,

updating decreases less with age, the more uncertain the market. This can be seen noting that:

$$\frac{\partial^2 \left(\Delta \varepsilon_{ijkt+1}^q \right)}{\partial \left(a_{ijkt} - \varepsilon_{ijkt}^q \right) \partial t} = - \frac{1}{\left(\frac{\sigma_\varepsilon^2}{\sigma_{jk0}^2} + t \right)^2}$$

which is larger (less negative) in more uncertain markets (with larger σ_ε^2).

Dynamics of prices and quantities. The model predicts expected growth rates of opposite signs for quantities and prices. This result comes from (14) and (15). Taking the first difference of these equations in expected terms, we directly get the expected growth rates. We find:

$$\mathbb{E} \left[\Delta \ln Z_{ijkt+1}^q \right] = - \frac{1}{\sigma_k} \mathbb{E} \left[\Delta \ln Z_{ijkt+1}^p \right]$$

Given that firms that decrease in size will on average be more likely to exit, the expected growth rate of quantities must be positive for survivors. Hence, the expected growth rate of prices for these firms should be negative and smaller by a factor $-\frac{1}{\sigma_k}$. Quantitatively, this is very close to what we find in table A.23.

Prediction 3. Prediction 3 states that the variance of growth rates within cohort decrease with cohort age. The variance of these growth rates can be expressed as:

$$\mathbb{V} \left[\Delta \ln Z_{ijkt+1}^q \right] = \sigma_k^2 \mathbb{V} \left(\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right] \right) \quad (22)$$

$$\begin{aligned} \mathbb{V} \left[\Delta \ln Z_{ijkt+1}^p \right] &= \left(\frac{1}{\sigma_k} \right)^2 \mathbb{V} (\Delta a_{ijkt+1}) + \mathbb{V} \left(\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right] \right) \\ &\quad - \frac{2}{\sigma_k} \text{Cov} \left(\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right], \Delta a_{ijkt+1} \right) \end{aligned} \quad (23)$$

First, a_{ijkt+1} and a_{ijkt} being drawn from the same distribution, $\mathbb{V}[\Delta a_{ijkt+1}] = 2\sigma_\varepsilon^2$. Second, using (11), we get:

$$\mathbb{V} \left(\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right] \right) = \left(\frac{\sigma_\varepsilon}{\sigma_k \left(\frac{\sigma_\varepsilon^2}{\sigma_{jk0}^2} + t \right)} \right)^2$$

As $\mathbb{E}[\Delta a_{ijkt+1}] = 0$, we get:

$$\text{Cov} \left(\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right], \Delta a_{ijkt+1} \right) = \mathbb{E} \left[\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right] \Delta a_{ijkt+1} \right]$$

Expanding this expression and using the fact that a_{ijkt} and a_{ijkt+1} are independent and that $\mathbb{E}[a_{ijkt}] = \mathbb{E}[a_{ijkt+1}] = \bar{a}_{ijkt-1}$, we get:

$$\text{Cov} \left(\Delta \ln \mathbb{E}_t \left[e^{\frac{a_{ijkt+1}}{\sigma_k}} \right], \Delta a_{ijkt+1} \right) = - \frac{\sigma_\varepsilon^2}{\sigma_k \left(\frac{\sigma_\varepsilon^2}{\sigma_{jk0}^2} + t \right)}$$

Finally, plugging this term into (22) and (23) and after rearranging, we get the following expres-

sions which are both strictly decreasing with t :

$$\mathbb{V} \left[\Delta \ln Z_{ijkt+1}^q \right] = \left(\frac{\sigma_\epsilon}{\left(\frac{\sigma_\epsilon^2}{\sigma_{jk0}^2} + t \right)} \right)^2 \quad (24)$$

$$\mathbb{V} \left[\Delta \ln Z_{ijkt+1}^p \right] = \left(\frac{\sigma_\epsilon}{\sigma_k} \right)^2 \left(\left(\frac{1}{\left(\frac{\sigma_\epsilon^2}{\sigma_{jk0}^2} + t \right)} + 1 \right)^2 + 1 \right) \quad (25)$$