Uncertainty and Export Status: Theory and Evidence

Eliav Danziger
Simon Fraser University, Burnaby, BC V5A 1S6, Canada

Abstract

I develop and analyze a model of uncertainty and international trade in which firms can allocate output across markets in response to market-specific shock realizations. The model's main prediction is that uncertainty in home shocks increases the propensity to export, while uncertainty in foreign shocks as well the correlation between home and foreign shocks decrease the propensity to export. I test the model's prediction using French firm-level data and find strong empirical support for the theory's predictions across a wide range of assumptions and specifications. Moreover, the effect of uncertainty on firms' export probability is nontrivial. For instance, I find that for a typical manufacturing firm moving up one standard deviation in the home uncertainty distribution increases the probability of exporting by a factor of 7.5 percentage points, while a similar shift in the foreign uncertainty distribution or in the distribution of the correlation between home and foreign shocks decreases the probability of exporting by 3.5 and 3.6 percentage points, respectively.

Keywords: international trade; export status; uncertainty; reallocation

JEL Classifications: D21; F12; F61

1 Introduction

All firms face uncertainty, whether about demand for their goods, production costs or countless other economic circumstances. Beyond these uncertainties, exporters potentially face even greater uncertainty due to unpredictable transportation costs or changing trade policies. How, then, does uncertainty affect a firm's decision to export? Does uncertainty increase or decrease the likelihood that a firm will choose to export? These are the questions which I address in this paper.

Despite the ubiquity of uncertainty, the role of uncertainty has been relegated to relative insignificance in the vast international trade literature. The standard justification for ignoring uncertainty is that firms are risk neutral. The argument is that, while it is true that uncertainty may lead to profit volatility, it should not affect a firm's decision making. However, this conclusion belies a voluminous literature specifically dedicated to studying the impact of uncertainty and risk on firms' decision making, not to mention the substantial investment that many firms make in reducing their exposure to risk. Indeed, the conclusion that risk-neutral firms remain unswayed by uncertainty does not hold when firms make irrevocable decisions that incur sunk costs prior to the resolution of uncertainty. For instance, if firms must decide on their output level prior to learning what demand for their product will be, the uncertainty will surely affect the firm's output decision, whether or not the firm is risk-neutral.

That firms must choose their output level prior to the resolution of demand uncertainty is sensible. After all, there is some space of time between the time a firm produces its good and the time it brings it to market. Recent empirical findings also point in the same direction. In particular, several studies have demonstrated that among exporters there is a negative correlation between domestic and home sales (my data replicates these findings as well (see Appendix ???)). This result is unlikely to be due to a negative correlation of demand across markets because the majority of exporters export to countries that are similar to their home market. For example, there tends to be a strong positive correlation between the aggregate growth rate in a firm's home country and the aggregate growth rate in a firm's export markets. This makes negatively correlated demand shocks highly unlikely. Indeed, in the Section 5.2 I show that there is a strong positive correlation between demand shocks at home and abroad, at least for French exporters. Instead, a natural interpretation of the negative correlation between domestic and foreign sales is that when conditions are favorable in one market relative to the other a firm increases its sales in that market by diverting sales from other markets rather than by increasing production. These findings suggest that the decision on how much output to produce precedes the decision of where to sell the output, and these decisions, at least to some extent, are independent of one another.

In the model, I introduce uncertainty as a firm-market-specific shock. In line with the aforementioned empirical findings, I assume that, in addition to choosing their export status, firms also decide how much to produce before demand uncertainty is resolved. Importantly, however, firms need not commit to the final destination of their output until after the resolution of demand uncertainty. Not surprisingly, firms will sell a greater proportion of their output in the destination with more favorable demand shock realizations. Under such circumstances, the decision to become an exporter provides the firm with the option value of of reallocating output across markets. In particular, a nonexporting firm must sell all its output at home even if it can only fetch a low price due to an unfavorable demand shock. However, an exporting firm in the same situation can profitably divert output abroad where more favorable conditions may enable the firm's good to command a higher price. The reverse is true as well. A nonexporting firm benefits from high demand at home, but it cannot sell more than it originally intended. However, an exporter can divert output, otherwise slated for the foreign market, to the home market to take advantage of the high demand at home.

The theoretical mechanism just described, which I refer to as the reallocation motive for exporting, constitutes a novel motive for exporting that has yet to be explored in the literature. I show theoretically that the probability with which a firm exports increases with home uncertainty and decreases with either foreign uncertainty or the correlation between home and foreign shocks. Indeed, the greater is uncertainty at home, the greater is the ability of an export destination to mitigate the negative impact of uncertainty. Similarly, greater foreign uncertainty increases the ability of home to mitigate the negative impact of foreign uncertainty. However, the benefit to reducing uncertainty through reallocation cannot possibly exceed the loss from introducing this uncertainty in the first place. Finally, the benefit to reallocating output across markets is greatest when the shocks are least positively (or most negatively) correlated.

RESULTS UNDER CONSTRUCTION

The rest of the paper is organized as follows: In Section 2 I develop a simple model of exporting and uncertainty. In Section 3 I derive the model's predictions for the relationship between the distribution of demand shocks faced by firm's and the probability that a firm exports. In Section 5 I test the model's predictions using French firm-level data and confirm the predictive power of the model. In Section **??** I assess the quantitative importance of the mechanism in the model with some counterfactual exercises. In Section 6 I describe some robustness checks. Finally, Section 7 concludes.

2 Model

In this section, I propose two versions of an export-decision model. The firm-level data I will use to test the alternative versions of the model contain information on a firm's total exports, but not on destination-specific exports. Therefore, to maintain consistency with the data, I model the world as consisting of a home country and one foreign country that can be thought of as the rest of the world. The two countries will be indexed by $c \in \{h, f\}$. In what follows, I develop the model from the perspective of a home-country firm. Total demand for firm i's good in period t (q_{cit} is demand in country c)

$$Q_{it} = q_{hit} + \mathbf{1}_{X_{it}=1} q_{fit} = A_{it}\left[\gamma_{hit} p_{hit}^{-\sigma} + \mathbf{1}_{X_{it}=1} B_{if} \gamma_{fit} p_{fit}^{-\sigma}\right], \tag{1}$$

where B_{if} is the size of the foreign market relative to the home market; A_{it} is a time-varying firm-specific demand factor; $X_{it} \in \{0, 1\}$ indicates whether (1) or not (0) firm i is an exporter; γ_{cit} is a time varying firm-specific demand shock in country c with a unit mean; p_{cit} is the country c consumer price of the good produced by firm i; and $\sigma > 1$ is the price elasticity of demand.[1]

Access to the foreign market requires payment of a fixed cost, f_x. In addition to the fixed cost, an exporter incurs an iceberg cost so that it must ship $\tau > 1$ units of its good for each unit sold abroad. Firms differ in their productivity, φ_i, and the marginal cost of production for firm i is $1/\varphi_i$. Therefore, the total cost to a firm of producing a total quantity of $Q_{it} = q_{Hit} + \tau q_{Fit}$ is $Q_{it}/\varphi_i + \mathbf{1}_{X_{it}=1} f_x$.

In every period, the firm draws a pair of demand shocks, $\left(\gamma_{hit}, \gamma_{fit}\right)$, for the home and foreign countries from the joint cumulative distribution function, $G(\cdot, \cdot)$. However, unlike the demand factor A_{it} whose value the firm knows all along, the firm only learns the realization of $\left(\gamma_{hit}, \gamma_{fit}\right)$ after choosing both its export status and the quantity it wishes to produces. Thus, while A_{it} as well as $\left(\gamma_{hit}, \gamma_{fit}\right)$ are time-varying, the former represents predictable changes in demand, while the latter represents unpredictable changes in de-

[1] The assumption that the shocks, γ_{hit} and γ_{fit}, have a unit mean is made for simplicity and is without loss of generality because any mean different than one in country c can be absorbed by A_{it} and B_{if}.

mand, i.e., uncertainty. This uncertainty introduces risk from the point of view of the firm since the firm must make irreversible decisions prior to the resolution of this uncertainty. At this point, I consider two alternative version of the model:

Version I - Ex-Ante Allocation In this version the firm must allocate its output between the home and foreign countries prior to learning the realization of shocks, that is, q_{hit} and q_{fit} are allocated ex ante.

Version II - Ex-Post Allocation In this version the firm need only commit to total production, Q_{it}, prior to learning the realization of shocks. After learning the realization of shocks the firm may choose to allocate its output as it sees fit subject to the constraint $Q_{it} = q_{hit} + \tau q_{fit}$, that is, q_{hit} and q_{fit} are allocated ex post.

For the most part I will refrain from using subscripts to differentiate the results from each version of the model. However, when clarity necessitates it, I will introduce subscripts $v \in \{I, II\}$ to distinguish results from the two versions of the model. Also, since in what follows I consider a particular firm, i, in a given period, t, I drop the subscripts i and t for the remainder of the theoretical portion of the paper. Indeed, the purpose of their inclusion thus far was to ease the transition to the empirical portion of the paper, where these subscripts will be reintroduced as they serve an important role in the empirical analysis.

2.1 Ex-Ante Allocation

I begin by analysing the ex-ante allocation version of the model. After having chosen export status and quantities for each market it serves, the firm is entirely at the mercy of the shock realizations, γ_h and γ_f. Therefore, given export status, X, the firm's realized revenue for any choice of q_h and q_f and any pair of realized shocks is determined entirely by the demand function (Equation (1)), which yields

$$R\left(q_h, q_f, X; \gamma_h, \gamma_f\right) = A^{1/\sigma}\left[\gamma_h^{1/\sigma} q_h^{1-1/\sigma} + \mathbf{1}_{X=1} {B_f}^{1/\sigma} {\gamma_f}^{1/\sigma} q_f^{1-1/\sigma}\right], \tag{3I}$$

Taking this result into account, the firm chooses q_h and q_f to maximize expected variable profit given export status

$$\pi\left(X\right), = \max_{q_h, q_f} E\left[R\left(q_h, q_f, X; \gamma_h, \gamma_f\right) - \frac{\left(q_h + \mathbf{1}_{X=1}\tau q_f\right)}{\varphi}\right].$$

Standard profit maximization yields

$$q_h = A\rho^\sigma \varphi^\sigma \left[E\left(\gamma_h^{1/\sigma}\right)\right]^\sigma \qquad \text{and} \qquad q_f = AB_f \tau^{-\sigma} \rho^\sigma \varphi^\sigma \left[E\left(\gamma_f^{1/\sigma}\right)\right]^\sigma, \tag{4I}$$

where $\rho = (\sigma - 1)/\sigma$. The firm's expected variable profit function is therefore

$$\pi(X) = \frac{A\rho^{\sigma-1}\varphi^{\sigma-1}}{\sigma}\left[\left[E\left(\gamma_h^{1/\sigma}\right)\right]^\sigma + \mathbf{1}_{X=1}\mathscr{O}\left[E\left(\gamma_f^{1/\sigma}\right)\right]^\sigma\right], \tag{5I}$$

where $\mathscr{O} = \tau^{1-\sigma}B_f$ is a measure of the firm's openness taking into account both the iceberg cost (lower iceberg cost implies greater openness) and the size of the foreign market relative to the home market (larger market implies greater openness). A key feature of this version of the model is that there is no interaction between the two markets, and therefore the firm's total profit is equal to the sum of the profits earned in each market. Therefore, the firm chooses to export if and only if the expected profit attributable to the Foreign country exceeds the fixed cost of exporting:

$$\frac{A\rho^{\sigma-1}\varphi^{\sigma-1}}{\sigma}\mathscr{O}\left[E\left(\gamma_f^{1/\sigma}\right)\right]^\sigma > f_x. \tag{6I}$$

This condition defines a cutoff productivity φ^* above which firms export and below which they serve only the Home market. So far the analysis is almost identical to the standard analysis in the existing literature with the exception of the introduction of uncertainty. The ex-post allocation version of the model represents a more substantial departure from the existing literature.

2.2 Ex-Post Allocation

In the ex-post allocation version of the model the firm need only commit to the total quantity it produces, Q, prior to the learning the realization of shocks. However, the ultimate destination of that output can be chosen after shocks are realized. The timing of this version of the model can be divided into three stages. In the first stage (export-decision stage) the firm chooses its export status. In the second stage (production stage) the firm chooses the total quantity it produces, Q. Finally, in the third stage (allocation stage) the firm learns the realization of the demand shocks, γ_h and γ_f, and, if the firm is an exporter, decides how to divide its output, Q, chosen in the second stage, between the home and foreign countries subject to the constraint that $Q \geq q_h + \tau q_f$. The firm's problem can be solved backwards in three steps corresponding to the three stages of the model outlined above.

Allocation Stage The firm finds the maximal profit for any potential choices of Q and export status, as well as any possible realizations of the demand shocks, γ_h and γ_f. Since at this point all costs are sunk (apart from the iceberg cost), the firm maximizes its revenue by choosing q_h and q_f as follows

$$\begin{aligned} R\left(Q,X;\gamma_h,\gamma_f\right) &= \max_{q_h,q_f}\left[q_h p_h + \mathbf{1}_{X=1}\left\{q_f p_f\right\}\right] \\ &\quad \text{subject to } Q \geq q_h + \mathbf{1}_{X=1}\left\{\tau q_f\right\} \text{ and Demand given by Equation (1).} \end{aligned}$$

This is a trivial decision for nonexporters, and their solution is $q_h = Q$. Exporters, however, achieve the optimal division of output by equalizing marginal revenue across the two markets, which yields:

$$q_h\left(Q,1;\gamma_h,\gamma_F\right) = \frac{\gamma_h}{\Gamma}Q \qquad \text{and} \qquad q_f\left(Q,1;\gamma_h,\gamma_f\right) = \frac{\mathscr{O}\gamma_f}{\Gamma}Q, \tag{2II}$$

where $\Gamma = \gamma_h + \mathscr{O}\gamma_f$ and as in Model I $\mathscr{O} = \tau^{1-\sigma}B_f$ is a measure of openness of the firm taking into account both the foreign relative to home market size as well as the additional cost associated with supplying the foreign relative to home country. The solution for the optimal division of output across markets is to sell output in each market in proportion to its realized market size, that is, when the shock is favorable in the home country it should sell proportionally more of its output domestically, and when it is favorable in the foreign country it should sell proportionally more abroad. The solution to the firm's problem implies realized revenues in each destination given by

$$R\left(Q,X;\gamma_h,\gamma_f\right) = A^{1/\sigma}\left(\gamma_h + \mathbf{1}_{X=1}\mathscr{O}\gamma_f\right)^{1/\sigma}Q^{1-1/\sigma}. \tag{3II}$$

Production Stage Taking into account the results of the allocation stage, the firm chooses the total quantity to produce, Q, that maximizes its expected profit given export status. Thus, the firm solves the maximization problem

$$\pi\left(X\right) = \max_Q E\left[R\left(Q,X;\gamma_h,\gamma_f\right) - Q/\varphi\right].$$

The solution is

$$Q(X) = \rho^{\sigma}\varphi^{\sigma}A\left[E\left(\gamma_h + \mathbf{1}_{X=1}\mathscr{O}\gamma_f\right)^{1/\sigma}\right]^{\sigma}, \tag{4II}$$

The firm's variable profit is then given by

$$\pi(X) = \frac{A\rho^{\sigma-1}\varphi^{\sigma-1}}{\sigma}\left[E\left(\gamma_h + \mathbf{1}_{X=1}\mathscr{O}\gamma_f\right)^{1/\sigma}\right]^{\sigma}. \tag{5II}$$

Export-Decision Stage Using the computations in the allocation and production stages, the firm chooses whether or not to incur the fixed cost of exporting to the foreign country. The firm will choose to export if and only if the expected export profit, i.e., the expected profit from exporting less the fixed cost of exporting, exceeds the profit from serving only the domestic market. That is, the firm will export if and only if

$$\frac{A\rho^{\sigma-1}\varphi^{\sigma-1}}{\sigma}\left(\left[E\left(\gamma_h + \mathscr{O}\gamma_f\right)^{1/\sigma}\right]^{\sigma} - \left[E\left(\gamma_h\right)^{1/\sigma}\right]^{\sigma}\right) > f_x. \tag{6II}$$

As was the case in the ex-ante allocation model, this condition defines a cutoff, φ^*, such that a firm exports if and only if its productivity exceeds this cutoff.

3 The Diversification Motive for Exporting

A comparison of Inequalities (6I) and (6II) shows one of the key differences between the two versions of the model. In Version I (ex-ante allocation) the export profit is simply the profit attributable to the foreign country, i.e., foreign revenue less the cost of producing the output shipped abroad. However, in Version II (ex-post allocation) attributing profit solely to the foreign country is no longer possible. This is because in addition to the profit associated with sales in the foreign country, an exporting firm also profits from the flexibility of adjusting the location of its sales to account for the shock realizations. Because, as will become evident, this latter component of the profit stems from the fact that exporting diversifies a firm's sources of demand, I define the diversification profit as the difference between the expected profit of a firm in the ex-post allocation version of the model and an identical firm in the ex-ante allocation version of the model.

For a nonexporter there is no difference between the two versions of the models since it must sell all its output domestically, and therefore for such a firm the diversification profit is zero. For an exporter, the diversification profit is

$$\pi^{diversification} = \frac{A\rho^{\sigma-1}\varphi^{\sigma-1}}{\sigma}\left\{\left[E\left(\gamma_h + \mathscr{O}\gamma_f\right)^{1/\sigma}\right]^{\sigma} - \left[E\left(\gamma_h\right)^{1/\sigma}\right]^{\sigma} - \left[E\left(\mathscr{O}\gamma_f\right)^{1/\sigma}\right]^{\sigma}\right\}. \tag{7}$$

The diversification profit must be nonnegative, since a firm cannot be worse off when it can allocate its output ex post. After all, nothing prevents a firm from foregoing the option of ex post allocation and commiting to an ex ante allocation if it were desirable to do so. Indeed, Equation 2II shows that only when shocks are proportional, i.e., there exists an $\alpha > 0$ such that $\gamma_f = \alpha\gamma_h$ for all possible $\left(\gamma_h, \gamma_f\right)$ pairs, do the particular shock realizations have no bearing on a firm's choice of q_h and q_f. However, so long as home and foreign shocks are not proportional to one another, the firms ex post allocation depends on the particular shock realizations. It therefore follows that as long as shocks are not proportional, the diversification profit is strictly positive.

Since nonexporters earn the same expected profits in the two versions of the models, the diversification profit is also equal to the excess export profit in Version II as compared to Version I. Since the diversification profit is positive (assuming nonproportionality of shocks), it follows that $\varphi_I^* > \varphi_{II}^*$ (see Inequalities (6I) and (6II)). This means that there are exporters that would export even without the ability to allocate output ex post ($\varphi > \varphi_I^*$) and there are exporters that would export only when given the ability to allocate output ex post ($\varphi \in [\varphi_I^*, \varphi_{II}^*]$). This, in turn, implies that if firms do, in fact, have the abilty to allocate output ex post, then taking advantage of this ability constitutes a motive for exporting. I refer to this as the diversification motive for exporting.

The prediction that $\varphi_I^* > \varphi_{II}^*$, i.e., if ex post allocation is possible then firms are more likely to export than if allocation is only possible ex ante, is a major theoretical findings of this paper. However, since the finding only predicts a quantitative difference between the two versions of the model, i.e., more firms export in one version than the other, it is not useful for testing which of the two versions of the model the data supports. Fortunately, the two versions of the model offer other predictions that are better suited to empirically testing which of the two versions of the model provides a better fit for the data. I now turn to these predictions.

4 The Impact of Uncertainty on Exporting

In this section, I address the following question: what characteristics of the distribution of the home- and foreign-country shocks affect a firm's incentive to export? Since the answer to this question depends on whether or not ex-post allocation of output is possible, the answer will provide a useful testable hypothesis, which I will exploit in the empirical portion of the paper. A simple numerical example will serve to illustrate.

4.1 A Simple Numerical Example

Suppose the two countries are symmetric and that there is no iceberg trade cost, i.e., $A = B_f = \tau = 1$, and that $\sigma = 2$ and $\varphi = 4$. These numbers are chosen for simplicity and result in an expected variable profit in Version I of the model given by (see Equation 5I)

$$\pi_I(X) = \left[E(\gamma_h)^{1/2}\right]^2 + \mathbf{1}_{X=1}\left[E\left(\gamma_f\right)^{1/2}\right]^2, \tag{8I}$$

and in Version II of the model given by (see Equation 5II)

$$\pi_{II}(X) = \left[E\left(\gamma_h + \mathbf{1}_{X=1}\gamma_f\right)^{1/2}\right]^2. \tag{8II}$$

The function $g\left(\gamma_h, \gamma_f\right)$ denotes the probability that a firm receives the shock-realization pair $\left(\gamma_h, \gamma_f\right)$. All five cases are constructed so that the expected shock realization is one in both countries, that is, $E(\gamma_h) = E\left(\gamma_f\right) = 1$. This ensures that the results in the different cases depend on higher moments of the shock distribution rather than on the first moment. The relevant results for the cases discussed below are summarized in Table 1 and can be easily calculated from Equations (8I) and (8II) and the information in the table (see table notes).

Case 1 is a baseline case with no uncertainty where both γ_h and γ_f alway equal one. In this scenario a firm always sells half of its output in each destination and there is no advantage to having the ability to allocate output ex post. Indeed, as Table 1 shows there is no diversification profit, and there is no difference in profits in the two versions of the model.

Case 2 introduces uncertainty in the home country (γ_h equals either zero or two with equal probability), while maintaining certainty in the foreign country (γ_f always equals one). In Version I, where the firm has to commit to q_h and q_f before learning the realization of shocks, it would always have a profit of one in the foreign country, just as in

Table 1: The Impact of Uncertainty on Exporting

	Probabilities $g(\gamma_h, \gamma_f)$	Profit I & II: Nonexporter	Profit I: Exporter	Profit II: Exporter	Export Profit I	Export Profit I	Diversification Profit
Case 1: No Uncertainty	$g(1,1) = 1$	1	2	2	1	1	0
Case 2: Home Uncertainty	$g(0,1) = \frac{1}{2}$ $g(2,1) = \frac{1}{2}$	$\frac{1}{4}$	$1 + \frac{1}{4}$	$1 + \frac{\sqrt{3}}{2}$	1	$\frac{3}{4} + \frac{\sqrt{3}}{2} > 1$	$\frac{\sqrt{3}}{2} - \frac{1}{4}$
Case 3: Foreign Uncertainty	$g(1,0) = \frac{1}{2}$ $g(1,2) = \frac{1}{2}$	1	$1 + \frac{1}{4}$	$1 + \frac{\sqrt{3}}{2}$	$\frac{1}{4}$	$\frac{\sqrt{3}}{2} < 1$	$\frac{\sqrt{3}}{2} - \frac{1}{4}$
Case 4: Positive Correlation	$g(0,0) = \frac{1}{2}$ $g(2,2) = \frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{4}$	0
Case 5: Negative Correlation	$g(0,2) = \frac{1}{2}$ $g(2,0) = \frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{2}$	$\frac{7}{4}$	$\frac{1}{4}$	$\frac{3}{2}$	$\frac{5}{4}$

The profit for nonexporters is identical in the two models and are reported in the second column. Profits for exporters differ in Versions I and II of the model and are reported in columns 3 and 4, respectively. The export profit is the profit of an exporter less that of a nonexporter, i.e. column 3 (4) minus column 2 in Version I (II). This is reported in column 5 (6) for Version I (II). Finally, the reallocation profit is the difference in the export profit between the two versions of the model (column 6 minus column 5), and it is reported in the last column.

the no uncertainty case. However, in Version II the firm can do better by adjusting the allocation of its output in response to shock realizations. In particular, when $\gamma_h = 0$, the firm will sell all of its output abroad, and when $\gamma_h = 2$, it will only sell one third of its output abroad. As Table 1 shows this ex-post allocation leads to an additional profit of $\sqrt{3}/2 - 1/4$, bringing the total export profit in Version II of the model to $\sqrt{3}/2 + 3/4$ which is greater than the export profit in Case 1. The introduction of uncertainty in the home country makes exporting more likely even while reducing the overall profit of the firm.

When ex-post allocation is possible, not having access to a foreign market hurts a nonexporter both when the home-country shock is favorable and when it is unfavorable. On the one hand, a nonexporting firm is forced to sell all of its output at home even when $\gamma_h = 0$. On the other hand, when $\gamma_h = 2$ a nonexporting firm cannot sell more than it initially intended at home to take advantage of the high demand. However, with access to the additional market the firm benefits in both scenarios. When $\gamma_h = 0$, the firm can sell the output abroad instead of being forced to sell it for nothing at home. Moreover, when $\gamma_h = 2$, the firm can divert more of its output away from the foreign market and towards the more favorable home market.

Case 3 reverses the roles of yhr home and foreign countries. In this case, exporting remains more likely in Version II than in Version I of the model because the firm recoups some of the loss from uncertainty via the diversification profit, which, as a consequence of the symmetry between the two countries, is the same as in Case 2. However, even when ex-post allocation is possible, a firm is less likely to export compared to the full-certainty case. This is because the benefit from mitigating foreign risk, i.e. the diversification profit, cannot possibly make up for the loss from introducing the risk in the first place, i.e. the drop in the Version I export profit from one to one quarter.

Cases 4 and 5 illustrate the impact of the correlation between home and foreign shocks on a firm's incentive to export. In both cases both γ_h and γ_f are either zero or two with equal probability. In Case 4 the shocks are perfectly positively correlated, so that always $\gamma_h = \gamma_f$ and the diversification profit is zero. Indeed, the firm cannot do better than always sells half of its output in the home country and half of its output abroad regardless of the shock realization, even when ex-post allocation is possible. Therefore, just as in Case 1, there is no diversification profit, and profits are the same in the two models.

Finally, in Case 5 the home and foreign shocks are perfectly negatively correlated so that the firm knows that the shock will be zero in one market and two in the other.

It just does not know which until after the shocks are realized. With no possibility of ex-post allocation (Version I) an exporter will sell half of its output in each destination leading to results identical to the case of positive correlation. In Version II, however, an exporter always sells all of its output in the market that receives the favorable shock thereby eliminating all demand uncertainty. Thus, the diversification profit makes up for the uncertainty-induced losses at home and abroad suffered by an equivalent firm in Version I of the model. As a consequence, while the correlation between the home and foreign shocks have no impact on the incentives of firms in Version I of the model, in Version II of the model firms are more likely to export when shocks are less (positively) correlated as this increases the benefits of diversification.

4.2 The General Case

The numerical example highlights the way in which the probability that a firm exports depends both on uncertainty and on the ability to mitigate this uncertainty through ex-post allocation of output. In so doing, the numerical example demonstrates the key theoretical finding that distinguishes the two versions of the model. The following proposition generalizes the results in the numerical example and shows that the intuition borne out in the numerical example is not contingent on its specifics (proof in Appendix **??**):

Proposition 1. *The probability that a firm exports:*

1. *increases with the variance of home-country shocks if only if the firm has the ability to allocate output ex post.*
2. *decreases with the variance of foreign-country shocks regardless of the firms ability to allocate output ex post.*
3. *decreases with the covariance between the home- and foreign-country shocks if only if the firm has the ability to allocate output ex post.*

While uncertainty abroad decreases the likelihood of exporting in both models, the models differ in their predictions regarding the impact of both home uncertainty and the correlation between home and foreign shocks on a firm's likelihood of exporting. In the ex-ante allocation version of the model neither has any impact on exporting, but in the ex-post allocation version of the model the former increases the likelihood of exporting while the latter decreases it. In other words, uncertainty at home creates an incentive for the firm to export as a means of diversifying its sources of demand, and the smaller the correlation between the home and foreign shocks the greater is the diversification achieved. Thus, the proposition makes clear that the diversification motive for exporting is aptly named. Importantly, because its predictions diverge for the two versions of the models, Proposition 1 provides a testable hypothesis that I will exploit in the empirical portion of the paper to show that the data support the ex-post allocation version of the model.

5 Empirics

The remainder of the paper will be devoted to testing the two alternative versions of the model proposed in the theoretical portion of the paper. I begin by introducing the French data I use in Section 5.1. In Section 5.2, I show how to use this data to infer model-consistent measures of the second moments of the shock distribution, i.e., the moments implied by the data if the model were the true data-generating process. In Section 5.3, I provide some descriptive statistics of the model-consistent shock distributions faced by firms. In Section 5.4, I use these shock distributions to show that the data rejects the ex-ante allocation version of the model in favor of the ex-post allocation version of the model. Finally, in Section 6, I provide a series of robustness checks on the results, none of which change the conclusion that the data support Version II of the model.

5.1 Data

I carry out the empirical analysis using French firm-level data from the ORBIS database. This database contains worldwide firm-level data, but the availability of particular variables varies widely across countries. The French subset of this data is particularly useful because, unlike the data for other countries, it contains firm-level export data. In particular, the database contains annual data on French firms from 2006 ($t = 1$ in what follows) to 2014 ($t = T$ in what follows), including information on each firm's NAICS industry, revenue and proportion of that revenue that comes from exporting. Although the dataset does not contain a wide range of variables, the structure of the models is rich enough that the aforementioned variables suffice for my purposes.

I drop firms that are missing data for any of the nine years from 2006 to 2014. Because the model is not intended to capture temporary or experimental entry and exit into export

Table 2: Data on Exports

	Manufacturing	Trade and Transport	Financial....	Total
Number of Firms	17,248	44,381	38,717	100,346
Percentage that Export	36.2%	20.8%	9.0%	18.9%
Export Share among Exporters	26.1%	21.6%	25.8%	23.8%

The export share among exporters is calculated as the mean export share across exporters without weighting by firm size. This is the relevant statistic in the scenario that all firms within a sector draw their openness from the same distribution.

markets, I consider only firms that either export in all nine years or never export.[2] This also ensures that each firm's status as exporter or nonexporter is unambiguous. This leaves me with a balanced panel of 171,644 French firms, which represents almost a fifth of all French employing firms. Of these firms 13.3% export in every year, while the remainder never export. Due to data limitations, which I discuss in Section 5.3, I use only firms in manufacturing; trade and transport; and information and financial activities, corresponding to the 1-digit NAICS supersectors 3, 4 and 5, respectively. Table 2 provides information of the number of firms and exporters in these sectors.

5.2 From Theory to Data

My strategy for testing the two alternative theories laid out in Section 2 is to exploit variations in firms' propensity to export in conjunction with variations in the distribution of home and foreign shocks faced by firms. This presents an immediate challenge because, while firms' export status is readily available in the data, the distribution of shocks faced by firms is only indirectly observed via firm-level revenue volatility. In what follows, I show how to derive model-consistent measures of uncertainty from the revenue data, that is, I show how to use the theoretical model to extract information on the shock realizations (and, by extension, the shock distributions) from revenue data. Importantly, as I will show, because the two versions of the model differ in their prediction of firm behavior, the same revenue data implies different model-consistent measures of uncertainty.

Substituting Equations (4I) and (4II) into Equations (3I) and (3II), respectively, yields, for $v \in \{I, II\}$, realized revenue in country c for firm i in period t:

[2] Stuff about experimental entry. Stuff about switchers. Sunk costs.

$$r_{cit}^{m}=K_{ci}^{m}A_{it}\psi_{c}^{m}\left(\gamma_{Hit},\gamma_{Fit}\right) \tag{9}$$

where variance

$$\begin{aligned} K_{ci}^{I} &= \mathcal{O}_{ci}\left[\rho\varphi_{i}E\left(\gamma_{c}\right)^{1/\sigma}\right]^{\sigma-1} \qquad \text{and} \qquad \psi_{c}^{I}\left(\gamma_{Hit},\gamma_{Fit}\right)=\gamma_{cit}^{1/\sigma} \\ K_{ci}^{II} &= \mathcal{O}_{ci}\left[\rho\varphi_{i}E\left(\Gamma\right)^{1/\sigma}\right]^{\sigma-1} \qquad \text{and} \qquad \psi_{c}^{II}\left(\gamma_{Hit},\gamma_{Fit}\right)=\Gamma_{it}^{1/\sigma-1}\gamma_{cit}, \end{aligned} \tag{10}$$

where $\Gamma_{it}=\gamma_{hit}+\mathcal{O}_{fi}\gamma_{fit}$, $\mathcal{O}_{hi}=1$ and $\mathcal{O}_{fi}=\tau^{1-\sigma}B_{f}$. The term K_{ci}^{m} is a firm-destination specific term that is time invariant and therefore does not contribute to variations in revenue over time. Rather, variations in revenue over time are due to the two remaining terms, A_{it} and ψ_{i}^{m}. The former leads to predictable variations in revenue, while the latter is a function of shocks leading to unpredictable variations in revenue .

Since the goal is to use the variations in the revenue data to infer the shocks realizations faced by the firm, it is necessary to purge the revenue data of variations in revenue unrelated to the shocks, i.e., variations caused byA_{it}. In order to do so, I assume that A_{it} can be approximated by a firm-level trend:

Assumption 1. $A_{it}=\left(1+a_{i}\right)^{t-1}$*, where* a_i *represents the trend of firm i.*

This assumption does not preclude the possibility of a sector- or economy-wide trend that is predictable from a firm's point of view since such a trend would be captured by the firm-level trend. In Section 6, I consider an alternative in which there is a sector-level component which is not constrained to follow a trend, but that is nevertheless predictable to the firm. As I show, this has almost no bearing on the results.

Given Assumption 1, Equation (9) and the fact that total realized revenue is $R_{it}^{m}=r_{hit}^{m}+r_{fit}^{m}$, the log difference (with respect to time) of R_{it}^{m} is

$$\log R_{it+1}^{m}-\log R_{it}^{m}=\alpha_{i}^{m}\left(t-t_{1}\right)+\varepsilon_{it}^{m}, \tag{??a}$$

where $\exp\left(\alpha_{i}\right)=\left(1+a_{i}\right)$.[3] The OLS estimate of a_i represents changes in demand (and , as a result, aso revenue) that the firm foresees prior to making its output and export-status

[3] The error term is equal to $\varepsilon_{it}^{m}=\log\sum_{i\in\{H,F\}}\psi_{i}^{m}\left(\gamma_{Hjt+1},\gamma_{Fjt+1}\right)-\log\sum_{i\in\{H,F\}}\psi_{i}^{m}\left(\gamma_{Hjt},\gamma_{Fjt}\right)$. Since the shocks are independently and identically distributed (with respect to different periods), it follows that so too is $\psi_{i}^{m}\left(\cdot,\cdot\right)$. As a consequence, $E\left(\varepsilon_{it}\right)=0$.

decisions. After controlling for the firm-level trend, the remaining revenue volatility is attributable to home and foriegn shocks.

The remaining variables that need to be backed out of the data are the firm-destination-specific productivity factors (K_{ci}^{m}) which determine the level of a firm i's revenue, and shock realizations (γ_{cit}^{m}), which determine the volatility around this level. For a given pair $\left(K_{hi}, K_{fi}\right)$, Equation 9 together with data on firm revenue in each market yields shock realizations $\left(\gamma_{hit}, \gamma_{fit}\right)$ for each period, t. Since the best estimate for the sample mean of shocks is the population mean, the pair $\left(K_{hi}, K_{fi}\right)$ is chosen so that the mean of shock realizations for each firm is one, i.e., $\sum_{t=1}^{T} \gamma_{hit}/T = \sum_{t=1}^{T} \gamma_{fit}/T = 1$ for all i. While, the assumption of a unit mean is without loss of generality (see footnote 1), it has the advantage that the standard deviations and coefficient of variations of shocks are equal.

The procedure outlined above yields the shock realizations for every firm in every year from 2006 to 2014 in each market served by the firm as well as K_{hi}^{m} for every firm and K_{fi}^{m} for every exporter. However, to test the theory I need data on the shock distributions faced by firms rather than the particular shock realizations. To infer a shock distribution from the shock realizations, I assume that all firms within a 6-digit NAICS industry face the same shock distribution. I therefore compute the second moments of the shock distribution for each 6-digit industry by pooling all the realization of all the firms in that 6-digit industry and computing statistics from this pooled sample.[4] In Section 6, I consider a specification in which every firm has its own unique shock distribution. However, since the hypothetical shock that nonexporters would have received had they exported are unobserved, this method only allows me to infer a home shock distribution for all firms.

Importantly the results for the inferred shock realizations and by extension the inferred shock distributions (as well as destination-specific factors, K_{ij}^{m}) differ across the two versions of the model. This is because when firms can only allocate output ex ante, revenue in the home (foreign) country is determined solely by home (foreign) shock realizations. In contrast, when firms can allocate output ex post, home and foreign revenues are both determined by both shock realizations simultaneously. Formally, this can be seen from the definitions of $\psi_i^m(\cdot,\cdot)$. While $\psi_i^I(\cdot,\cdot)$ depends only on the shock realization in market i, $\psi_i^{II}(\cdot,\cdot)$ depends on the shock realizations in both markets. As an illustrative example,

[4]For concreteness, consider a typical manufacturing industry, NAICS sector 312130 that has a total of 237 firms of which 87 export. Since I have nine years of data, this given me $237 \times 9 = 2133$ observed draws from the home shock distribution and $87 \times 9 = 783$ observed draws from the foreign shock distribution from which I can directly estimate the variance of home and foreign shocks. In addition, from the exporting firms I have 783 pairs of home and foreign shocks from which I can estimate the covariance between home and foreign shocks for this 6-digit industry.

suppose that we observe revenue at home which is above the firm's revenue trend. With ex-post allocation ruled out, the only possible explanation is a favorable home shock. However, when ex-post allocation is allowed, the high revenue at home could alternatively be the result of an unfavorable foreign shock which led the firm to allocate more output towards the home market. Thus, in Version I of the model, revenue in a particular market depends only on the shock realization in that market, while in Version II of the model revenue in each market depend both on the shock in that market and the favorability of that shock relative to the shock in the other market.

5.3 Model-Consistent Second Moments

Before testing the theory, I present some descriptive statistics of the second moments of the shock distributions implied by the model and the data. As mentioned above, I focus on NAICS sectors beginning in 3, 4 and 5, which contains 370 6-digit NAICS industries, but not all of them have active exporters. Of course, I have no choice but to drop industries with no exporters since I have no information on foreign shocks in those industries. After dropping these industries, I am left with 351 6-digit manufacturing industries with a total of 100,346 (see Table 2 for more detail).

Table 3 reports descriptive statistics of the second moments of the shocks distributions derived in Section 5.2. These results, together with the results in Table 2, show that not only is there substantial variation in the proportion of exporting firms across industries, but also in the second moments of the distribution of home and foreign shocks. It is this variation that I will exploit to test which of the versions of the model accords better with the data.

The results in Table 3 provide some insight into the difference between the shock distributions implied by the two versions of the model. Specifically, the results make clear that the same revenue data implies higher variance of shocks in Version I as compared to Version II of the model. To see why this is so, consider the shock realization implied by a given revenue increase in a particular market for an exporter. In Version I, this can only be explained by an increase in the shock realization by the same factor as the increase in revenue, i.e., if revenue doubled then so too did the shock realization. However, in Version II, in addition to the direct effect of the shock realization on revenue, the revenue gets an additional boost as a result of ex-post allocation of output towards this market, i.e., a doubling of revenue implies less than a doubling of the shock realization. Analogously, a decrease in revenue in Version I must be due to an equally large decrease in the shock

Table 3: Descriptive Statistics for 6-Digit NAICS Manufacturing Industries

		NAICS	Min.	Max.	Median	Mean	St. Dev.
Version I	H. Variance	31-33	0.05	4.05	0.27	0.33	0.33
		42-49	0.03	0.92	0.23	0.29	0.19
		51-56	0.08	1.69	0.35	0.39	0.25
	F. Variance	31-33	0.10	2.67	1.06	1.09	0.50
		42-49	0.05	6.04	1.38	1.40	0.71
		51-56	0.17	4.13	1.47	1.59	0.77
	Covariance	31-33	-0.84	0.47	-0.05	-0.05	0.12
		42-49	-0.31	0.30	-0.04	-0.05	0.08
		51-56	-0.66	0.53	-0.07	-0.08	0.16
Version II	H. Variance	31-33	0.05	0.73	0.19	0.21	0.10
		42-49	0.03	0.57	0.20	0.22	0.11
		51-56	0.08	1.69	0.31	0.35	0.23
	F. Variance	31-33	0.07	1.31	0.38	0.40	0.19
		42-49	0.02	1.45	0.49	0.51	0.21
		51-56	0.07	1.88	0.56	0.61	0.31
	Covariance	31-33	-0.17	0.64	0.10	0.12	0.09
		42-49	-0.03	0.40	0.12	0.12	0.08
		51-56	-0.04	0.90	0.15	0.18	0.16

realization, while in Version II part of the decrease in revenue is attributable to ex-post allocation away from the market. In other words, compared to the ex-ante allocation version of the model, the ex-post allocation version of the model requires a smaller (larger) shock to rationalize a given increase (decrease) in revenue. As a result, to rationalize the same revenue data requires a greater variance in shock realizations in the ex-ante allocation version of the model compared to the ex-post allocation version of the model. A key lesson from Table 3 is that the shocks inferred by revenue data depend heavily on the specifics of the model, and in particular ignoring firms' ability to allocate output ex postleads to a substantial overstatement of the uncertainty firms face.

Besides being useful for the testing the theory, the shock distributions are interesting in their own right. For example, some authors have conjectured that firms face greater uncertainty abroad than domestically because firms that export a large fraction of their output tend to also have more volatile revenue streams than either similar nonexporting firms or firms that export only a small fraction of their output.[5] My results confirm this view more directly. Indeed, according to the computations for ex-post allocation version of the model, the variance of foreign shocks are on average 0.19, 0.29 and 0.26 higher than the variance of home shocks in the NAICS supersectors beginning with 3, 4 and 5, respectively. Moreover, the variance of foreign shocks is greater than the variance of fome shocks in all but 3 out of the 351 6-digit industries. The results also show that home and foreign shocks are strongly positively correlated. The correlation coefficient between home and foreign shocks overall is 0.37, and it is positive in 95.7% (336 out of 351) of 6-digit industries. The results for the shock distributions implied by ex-ante allocation version of the model are similar, but I do not present them here because as I will presently show the data rejects Version I of the model in favor of Version II of the model.

5.4 Evidence on the Reallocation Motive for Exporting

Regardless of the specifics of the model, firms export if and only if $\pi(1) - \pi(0) > f_x$, that is, if the difference between the expected profit when exporting and when not exporting exceeds the fixed cost of exporting. Taking logs of both sides of this inequality yields:

$$X_i = \begin{cases} 1 & y_i > 0 \\ 0 & \text{otherwise} \end{cases}, \tag{12}$$

[5] See, for example, Vanoorenberghe (2012).

where $y_i = \log[\pi_i(1) - \pi_i(0)] - \log f_{ix}$. The index on the fixed cost indicates that fixed exporting costs may potentially differ across firms. In Appendix **??**, I show that the theoretical model motivates the following linear approximation of the latent variable, y_i:

$$y_i = \beta_0 + \beta_1 \log K_{Hi} + \sum_j \beta_j D_{ij} + \sum D_{ik} \alpha_k Z + \varepsilon_i, \tag{13}$$

where D_{ij} (D_{ik}) is a dummy variable which equals one if firm i is in NAICS sector j (k) and

$$\alpha_k Z = \alpha_{1k} \mathrm{var}_i(\gamma_H) + \alpha_{2k} \mathrm{var}_i(\gamma_F) + \alpha_{3k} \mathrm{cov}_i(\gamma_H, \gamma_F), \tag{14}$$

where the seconds moments for firm i are determined by its 6-digit NAICS industry as explained in Section 5.2. This specification, with interaction dummies for sector k, implies that the effect of the second moments of the shock distribution faced by a firm on the firm's export probability depends on the firms sector, and, in particular, the effect is common to all firms within sector k.[6] The interpretation of the β_j's is that they account for any sector-specific productivity factors and exporting fixed costs at the 2-digit (3-digit) industry level, while the ε_i's account for any firm-specific exporting fixed cost.[7] These productivity factors and fixed costs are unobserved in my data and assumed to be time invariant (see Appendix ??? for more details).

I test Version I (the null hypothesis) against Version II (the alternative hypothesis) of the model by running a logit regression defined by Equations (12), (13) and (14). As per Proposition 1 the null and alternative hypotheses are

$$\begin{array}{llll} H_0: & \alpha_2^I < 0 & \text{and} & \alpha_1^I = \alpha_3^I = 0 \\ H_a: & \alpha_1^{II} > 0 & \text{and} & \alpha_2^{II}, \alpha_3^{II} < 0. \end{array}$$

In this case it is insufficient to run just one regression to test the alternative hypothesis

[6] The k indexes 1-digit NAICS supersectors. For each of these supersectors, I estimate three coefficients (one for each second moment of the shock distribution). Since all firms within a 6-digit industry are assumed to have the same shock distribution, in order to reliably estimate α_k, industry k must contain a sufficient number of 6-digit NAICS industries. In particular, since I require at least forty industries per supersector, I cannot use supersectors other than 3,4,5. This limitation also precludes estimating these coefficients at the 2-digit level except in the manufacturing sector, which I do in Section 6.

[7] The j indexes 2-digit (or in an alternative specification 3-digit) NAICS industries. Dummies at more disaggregated levels are inappropriate because at the 4-digit NAICS level there are 171 industries as compared to 351 at the 6-digit NAICS level with 71 of the 4-digit industries containing only a single 6-digit industry. This is problematic because in those cases the dummy variables will invariably pick up variations in the second moments of the shock distributions, which are common to all firms within a 6-digit industry.

against the null hypothesis, since, as explained in Section 5.2, the inferred shock distributions and therefore the data inputs into the regressions are model dependent. Rejecting the ex-ante allocation version of the model based on data implied by the ex-post allocation version would be inconsistent as would as accepting the ex-post allocation version based on data implied by the ex-ante allocation version of the model. Therefore, in Table 4 I present two sets of results for the logit regression, one for each of the two versions of the model.

The results for the ex-ante version of the model (positive first row, negative second and third rows) are at odds with the null hypothesis (zero first and third rows, negative second row) for all three 1-digit supersectors regardless of whether the fixed effects are at the 2- or 3-digit level. Indeed, all of the estimates are significant at the 99% level. The data, therefore, strongly rejects the null hypothesis. For the ex-post version of the model, however, the results are exactly as the model predicts (positive first row, negative second and third rows) for all three supersectors when fixed effects are the 2-digit level. When fixed effects are at the 3-digit level the results are as the model predicts for two out of the three supersectors, and in the third supersector, two out of three coefficients have the correct sign. Again, all these results are significant at the 99% level. Moreover, the regression results also deliver a significant and positive effect of productivity on export status, which is also consistent with the model's predictions. On the basis of these results, therefore, I conclude that the data support the hypothesis that firms have the ability to allocate output across markets in response to shock realizations and that the opportunity to do so constitutes a motive for exporting.

Finally, the results in Table 4show that the impact of uncertainty and in particular the ability to allocate output across destinations in response to shock realizations has a substantial impact on a firm's export probability. For instance, in the manufacturing sector moving up the distribution of home-shock variances by one standard deviation increases the odds of exporting by 91%, which, for a firm with an initial 10% probabilty of exporting translated into a 7.5 percentage point increase in the probability of exporting.[8] Similarly, moving up by one standard deviation in the distribution of foreign-shock variances or covariances between home and foreign shocks leads to a decrease in the odds of exporting by 60% and 63%, respectively.[9] This translates into 3.5 and 3.6 percentage point decreases, respectively, in the probability of exporting for a firm with an initial 10%

[8]This corresponds to an increase of 0.1 in the variance of home shocks faced by the firm, (see Table 3).

[9]This corresponds to an increase of 0.19 in the variance of foreign shocks and an increase in 0.09 in the covariance between home and foreign shocks faced by the firm, (see Table 3).

Table 4: Logit Estimates for Ex-Ante & Ex-Post Allocation Versions of the Model

		Ex-Ante Allocation		Ex-Post Allocation	
		2-Digit FE	**3-Digit FE**	**2-Digit FE**	**3-Digit FE**
Manufacturing	$\mathrm{var}(\gamma_H)$	3.98^{**}	2.68^{**}	6.52^{**}	5.25^{**}
		(0.189)	(0.25)	(0.336)	(0.436)
	$\mathrm{var}(\gamma_F)$	-0.93^{**}	-1.10^{**}	-2.47^{**}	-2.51^{**}
		(0.050)	(0.063)	(0.158)	(0.203)
	$\mathrm{cov}(\gamma_H, \gamma_F)$	-3.79^{**}	-4.55^{**}	-5.43^{**}	-7.11^{**}
		(0.330)	(0.377)	(0.467)	(0.589)
Trade and Transport	$\mathrm{var}(\gamma_H)$	4.21^{**}	3.95^{**}	7.81^{**}	7.01^{**}
		(0.114)	(0.153)	(0.239)	(0.288)
	$\mathrm{var}(\gamma_F)$	-0.02	-1.01^{**}	-0.52^{**}	-4.29^{**}
		(0.035)	(0.072)	(0.122)	(0.217)
	$\mathrm{cov}(\gamma_H, \gamma_F)$	-1.43^{**}	-1.36^{**}	-1.64^{**}	2.11^{**}
		(0.298)	(0.499)	(0.381)	(0.505)
Finance	$\mathrm{var}(\gamma_H)$	1.55^{**}	2.69^{**}	1.79^{**}	3.53^{**}
		(0.091)	(0.123)	(0.129)	(0.184)
	$\mathrm{var}(\gamma_F)$	-0.52^{**}	-0.46^{**}	-1.00^{**}	-1.01^{**}
		(0.044)	(0.048)	(0.138)	(0.150)
	$\mathrm{cov}(\gamma_H, \gamma_F)$	-2.28^{**}	-1.49^{**}	-1.37^{**}	-2.19^{**}
		(0.223)	(0.249)	(0.283)	(0.338)
	$\log K_i$	0.55^{**}	0.57^{**}	0.65^{**}	0.68^{**}
		(0.007)	(0.007)	(0.007)	(0.007)

Statistical significance at the 99 percent level is denoted by ** and at the 95 percent level is denoted by *. Standard errors are in brackets.

probability of exporting.

6 Robustness

In this section, I provide some robustness checks on the results reported in Section 5. Here I report the major findings across various alternative specifications and assumptions, but I relegate the tables with the particular point estimates to Appendix **??**. The overarching theme in this section is that the findings are remarkably consistent across all specifications, with only minor deviations from the results reported in the previous sections. Thus, the results that follow provide additional empirical support for the Model II and the reallocation motive for exporting.

Alternative Measure of Firm Productivity and Controlling for Firm-Level Openness

In the main specifications in Section 5, I control for firm productivity using a model-consistent measure of productivity, i.e., the firm's productivity if the model were the true data-generating process. To test the robustness of the results with respect to this measure of productivity, I instead use the log of the mean domestic revenue for each firm as that firm's measure of its productivity. The results are little changed by this alternative measure of productivity. Indeed, the measure of productivity implied by the model and this alternative measure of productivity are highly correlated (correlation coefficient equal to 0.97).

One of the inherent limitations of the data is the lack of information on the hypothetical openness of nonexporters, that is, it is impossible to know how open a nonexporter would be in the event that it had chosen to export. As a result, in the main specification, I assume that firms draw their openness only after deciding whether or not to export. Nevertheless, it is worthwhile considering whether or not the results are robust to controlling for firm-level openness. I consider two ways of addressing the lack of data on openness for nonexporters. The first way is to compute the mean openness of all exporters within a 6-digit NAICS level and assign all firms in the 6-digit industry, including nonexporters, that level of openness. An alternative approach is to assign to exporters the model-consistent value of openness ($\mathscr{O}_i$) and use those values to impute openness for nonexporters by assuming a linear relationship between the log of productivity and openness.[10] Both approaches are problematic due to selection issues, i.e., nonexporters may

[10] I estimate this linear relationship at the 6-digit level by regressing log productivity on openness of

choose not to export because they know (and the researcher does not) that they would have a low level of openness had they chosen to export. However, this is unlikely to be a significant issue because openness is uncorrelated or very weakly correlated with the second moments of the shock distributions.[11] Even so, neither method of controlling for firm-level openness changes the results, and this is true regardless of whether the measure of openness used is the model-consistent measure ($\mathscr{O}_i$) or the more commonly used foreign share of firm revenue (r_{iF}/R_i).

Alternative Assumption to Firm-Level Trend Next, I consider an alternative to Assumption 1, that the predictable element of volatility follows a time trend. Here, in addition to a firm-level time trend, I allow for a sector-level component at either the 3-digit or 6-digit NAICS level which is predictable to the firm even though it does not necessarily follow a time trend. Formally, this assumption is:

Assumption 2. *$A_{it}w_{it}^{1-\sigma} = A_{st}\,(1+a_i)^{t-t_0}$, where $A_{st} = \sum_{i\in NAICS} r_{it} / \sum_{i\in NAICS} r_{it_0}$, a_i represents the trend of firm i and t_0 is the first period in the data.*

Thus, the sectoral factor, A_{st} is normalized to unity in t_0 (2006), which means that this term captures variations in firm revenues attributable to sectoral changes relative to the initial year in the data. In this formulation, A_{st} captures all changes that affect every firm equally in sector s, such as sector-wide changes in demand or cost of inputs, while the firm-specific element captures each firm's trend relative to the rest of their sector.[12] The procedure for inferring shocks is identical to the one outlined in Section 5.2, except that Equation **??**a is replaced with Equation **??**b:

$$(\log R_{it+1} - \log A_{st+1}) - (\log R_{it} - \log A_{st}) = \alpha_i\,(t - t_0) + \varepsilon_{it}, \qquad \text{(??b)}$$

where, as before, $\exp(\alpha_i) = (1+a_i)$ and $\exp(\varepsilon_{it}) = \Gamma_{it+1}^{1/\sigma} - \Gamma_{it}^{1/\sigma}$.

This alternative assumption regarding the structure of the predictable element of volatility has almost no effect on the specific values of the shock realization that I infer for each

exporters.

[11] Among exporters, the correlation coefficients between openness and the variance of home shocks, the variance of foreign shocks and the covariance between home and foreign shocks are 0.019, -0.002 and 0.009, respectively.

[12] The term A_{st} can be thought of as corresponding to the revenue factor in a general-equilibrium version of a Melitz model, $E\,(Pw)^{1-\sigma}$, where E is total consumer expenditure in the sector and P is the price index and w is the input price. Thus, this specification is consistent with a full general equilibrium version of the model, where changes in A_{st} correspond to changes in sectoral input prices (w), consumer demand (E) or sectoral the price level (P).

firm in each year, and hence on the inferred shock distributions faced by firms. Indeed, the correlation coefficient between the shock realizations using Assumption 1 versus Assumption 2 is 0.99 and 0.96 when the sectoral component, A_{st}, is at the 3-digit and 6-digit levels, respectively. As a result, none of my conclusions are altered by this alternative specification, and, in fact, the point estimates are almost identical to the point estimates in the main specification.

Alternative Values for the Demand Elasticity

The values for the shock realizations and subsequent results in Section 5were obtained using $\sigma = 2.7$, which is the median value of the elasticity of substitution reported in Broda and Weinstein (2006). Since the problem of estimating the price elasticity of demand is a complex one, I do not take a stand on the correct estimates. Instead, I check the robustness of my results with respect to alternative values of the demand elasticity. Importantly, because σ enters the computations only in Equation **??**, the choice of σ does not affect the theory in any qualitative way. Rather the choice of σ affects only the particular values of the shock realizations inferred from the data. The results remain largely unchanged when shocks are inferred using demand elasticities ranging from slightly greater than one to five. For larger values of the demand elasticity some of the results become statistically insignificant because, as explained below, as σ increases the variation in the data on second moments of the shock distribution decreases.

To see how the value of the demand elasticity affects the values of the shock realizations inferred from the revenue data, note that as $\sigma \to 1$, the revenue in each destination ceases to depend on the quantity sold in that destination, and therefore there is no incentive to reallocate output after learning the shock realization even when the reallocation option exists. In fact, the shocks inferred from Model I and Model II converge as σapproaches one (see Equation (10). Conversely, as $\sigma \to \infty$, a firm will reallocate as much as possible by selling all of its output in the destination with the more favorable shock. The upshot is that, as σ increases, the variance of shocks needed to explain the variance in the revenue data decreases. As a result, for high values of σ there is very little variation in the second moments of the shock distributions, which leads some of the results to become statistically insignificant.

Effect of Second Moments Interacted at Different NAICS Levels

In the specifications considered in Section 5, the coefficients on the second moments of the shock distribution interacted with dummies at the 1-digit level, i.e., the effect of the second moments of the shock distributions on the probability of exporting depended on the firm's 1-digit NAICS supersector. Here I consider the possibility that these dummies interact at the 2-digit level instead, i.e, the effect of the second moments depends on the firm's 2-digit NAICS sector. However, since I estimate three parameters for each 2-digit sector and all firms within a 6-digit industry face the same shock distribution, I consider only 2-digit sectors which contain at least forty 6-digit sectors. In practice, this limits me to the three 2-digit manufacturing sectors (sectors 31, 32 and 33).[13] The results are consistent with Model II and statistically significant at the 99% level for two out of the three sectors (31 and 33), while in the other sector the results are statistically insignificant.

As an alternative, I consider a scenario in which there are no interacting dummies on the coefficients on the second moments of the shock distribution, i.e., the effect of the second moments of the shock distribution on a firm's export probability is independent of the firm's sector. The advantage of this approach is that it allows me to incorporate more sectors into the analysis, since I am not forced to discard sectors with too few 6-digit industries. Therefore, in addition to considering results including only the firms in the three 1-digit supersectors that I have considered thus far (sectors beginning in 3, 4 and 5), I also consider results including all firms in the data set (a total of 157,934 firms spread over 480 6-digit industries).[14] The results are exactly as the model predicts and statistically significant regardless of which sectors are included in the analysis.

Firm-Specific Shock Distributions

The final assumption that I relax is that all firms in a given 6-digit sector face the same shock distribution. Here I assume that each firm faces a firm-specific shock distribution. The second moments of each firm's shock distribution can then be estimated from the same shock realizations that were derived in Section 5.2.[15] The main problem with

[13] It impossible to consider scenarios in which the effect of the second moment on the export probability depends on the firm's 3-digit sector because the 3-digit sector that nests the most 6-digit sectors (in the data) has only 21 6-digit sectors, with the majority having fewer than 10.

[14] This includes all firms which meet the criteria discussed in Section 5.1, i.e., firms that either export in all years in the data or never export.

[15] The difference is that in this case I do not pool all the realizations of firms in the 6-digit industry to estimate the second moments of the shocks distributions.

this approach is that it is impossible to estimate the variance of foreign shocks or the covariance between home and foreign shocks for nonexporters, and, indeed, one of the motivations for the assumption that the shock distribution was common to all firms at the 6-digit level was to circumvent this problem. Thus, I estimate Equation 13 with only the variance of home shocks while omitting the variance of foreign shocks and the covariance between home and foreign shocks. These omissions introduce omitted variable bias into the estimate of the coefficient on the variance of home shocks. Because both the omitted variables are positively correlated with the variance of home shocks and have a negative effect on the probability of exporting, the estimate of the coefficient on the variance of home shocks will be biased downward.[16] However, since this militates against the model's prediction, finding a positive coefficient despite the downward bias constitutes evidence in favor of the model.

One advantage of this approach is that the effect of the variance of home shocks can now differ across more disaggregated levels, rather than only at the 1- or 2-digit levels as discussed above. The reason is that since each firm now has its own shock distribution there is sufficient variation even within 6-digit industries to estimate an industry-specific effect of the home variance, provided that there are enough firms in that industry. For instance, when the effect of home variance is allowed to differ at the 3-digit level, the coefficient of interest is positive in twenty five of the thirty two sectors in which the results are statistically significant at the 99 percent level. At the 2-digit level, eight out of the nine sectors with statistically significant coefficients are consistent with the model's prediction.[17] Considering that the estimates are downward biased, these results strongly support the model.

7 Conclusion

This paper explores the impact of uncertainty on a firm's decision to export. I introduce uncertainty as a market-specific demand shock whose realizations firms learn only after choosing their export status and level of output. Importantly, however, firms need not

[16]The variance of home shocks has a correlation coefficient of 0.33 and 0.55 with variance of and foreign shocks and the covariance between home and foreign shocks, respectively.

[17]Not surprisingly, at more disaggregated levels fewer of the sectors have statistically significant coefficients. At the 1-digit level all sectors are statistically significant at the 99% level, at the 2-digit level 64% of sectors are statistically significant and for 3-digit sectors this drops to 54%. However, given the downward bias of the estimates, it should be expected that fewer estimates will be statistically significant than if no variables were omitted.

commit to the final destination of their output until after shocks are realized. Thus, firms may reallocate output ex post by diverting a greater proportion of their output towards the market with favorable shock realizations and away from the market with unfavorable shock realizations. I show that in such a setting firms may be incentivized to export as a means of mitigating their overall risk. I refer to this novel mechanism as the reallocation motive for exporting.

The model yields three testable implications: (1) the propensity to export increases with home uncertainty; (2) the propensity to export decreases with foreign uncertainty; and (3) the propensity to export decreases with the correlation between home and foreign shocks. I use French firm-level data to test the model's predictions. The data strongly support the model's predictions. Moreover, the results are economically significant. Indeed, my estimates imply that for a manufacturing firm moving from the 25-th to the 75-th percentile in the home uncertainty distribution increases the probability of exporting by a factor of 0.62, while a similar move in the foreign uncertainty distribution leads to a decrease in the probability of exporting by a factor of 0.24.

Finally, I assess the importance of uncertainty and the reallocation motive for exporting in explaining firms' export status. I find that while eliminating uncertainty would increase the probability that a typical manufacturing firm will export by a factor of 0.13, while eliminating the ability to reallocate output after learning the realization of shocks would decrease this probability by a factor of 0.17. These results point to the importance of both uncertainty and the ability to reduce uncertainty through reallocation of output in shaping a firm's export-status decision.

As this is the first paper to explore the reallocation motive for exporting, my goal has been to explore some of its theoretical implications and to confirm its existence empirically. The findings in this paper raise many interesting avenues for future research. How does uncertainty affect not only whether or not a firm exports but also the number of export destinations? What are the general-equilibrium and welfare implications of the reallocation motive for exporting? These are all fruitful directions for future research.