Maximum likelihood estimation of export cost thresholds distributions

Guerini Mattia   Musso Patrick   Nesta Lionel

WEHIA 2019

City University, London

Université Côte d'Azur (GREDEG)
OFCE - SciencesPo
Sant'Anna School for Advanced Studies

Research Question

On threshold distributions

Thresholds are unobserved (entry) barriers to be overcome by economic agents

Barriers embody various aspects: sunk and fixed costs, (lack of) capabilities, opportunity costs.

Agent-specific characteristics (the $\theta-$attribute) explain how agents can overcome such hurdles

The combination of thresholds and $\theta-$attributes determines agents' decisions regarding market participation

- Foreign markets: export barriers and productivity
- Technologies: barriers to adoption and capabilities
- Labour markets: reservation wage and unemployment benefit

How economists think about thresholds

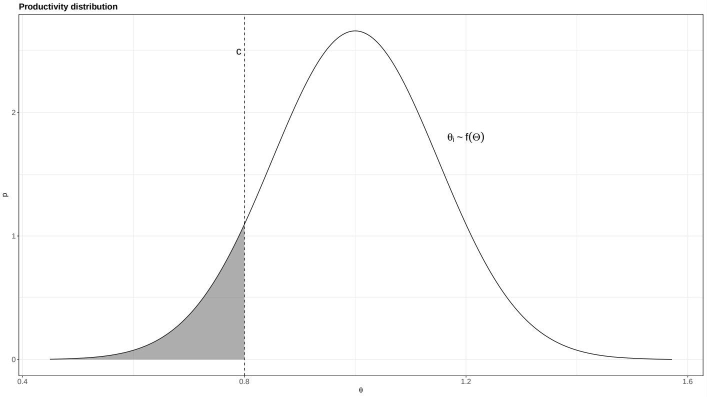

Figure 1: Typical representation of non-participating agents.

Alternative way of thinking about thresholds

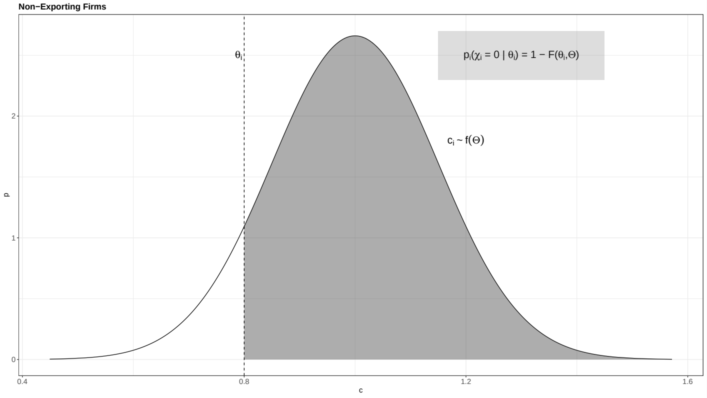

Figure 2: Alternative representation of non-participating agents.

Our Contribution

The contribution of this paper:

- methodological: econometric strategy to estimate threshold cutoffs (general)
- empirical: application to French firms in the context of international trade (export participation)
- economic policy: important to efficiently channel public subsidies to firms that might use them

Econometric strategy

Assumptions

Constraints:

1. Agents are heterogenous in their $\theta-$attribute (θ_i)
2. Thresholds are agent-specific (c_i)

Working assumptions:

A1. Perfect sorting hypothesis

- $\chi_i = 1$ iif $\theta_i > c_i$
- $\chi_i = 0$ iif $\theta_i < c_i$

A2. Agent-specific threshold (c_i) is a random variable which follows the density distribution f with vector of parameter Θ: $c_i \sim f(\Theta)$ to be estimated

We then parametrically estimate the thresholds distribution $\mathbf{c}$.

Maximum likelihood estimation of export costs

Under working assumptions A1 amd A2 the Likelihood function takes the form:

$$L(\Theta) = \prod_{i=1}^{N} [F(\theta_i; \Theta)]^{\chi_i} \times [1 - F(\theta_i; \Theta)]^{1-\chi_i} \tag{1}$$

where Θ is a vector of parameters characterizing the distribution F and can be consistently estimated.

Our preferred choice for F is the Γ distribution, which is characterized by two parameters (α, β) = (shape, scale) and is very flexible.

The case for MonteCarlo simulations

However, a series of questions might emerge:

- **Q1**: under A1 and A2, how reliable is the estimate?
- **Q2**: under deviations of A1, how precise is the estimate?
- **Q3**: under deviations of A2, which distributional assumption reduces the estimation error?

Using a MonteCarlo simulation setting we answer to all these questions!

Results sneak preview

We show that:

- when both A1 and A2 are valid, we precisely estimate $\mathbf{c}$;
- even under some violations of A1, the estimates are good;
- assuming a Γ_{MLE} is in general better (safer) than assuming a $\mathcal{N}_{MLE}$;
- French firms export costs are very skewed;
- the Γ_{MLE} assumption is corroborated by the French evidence.

MonteCarlo Simulations

Baseline: Densities of $(\hat{\alpha}, \hat{\beta})$

In the PM-PS case, our estimation is correct on average. However, there is some variability in the estimates.

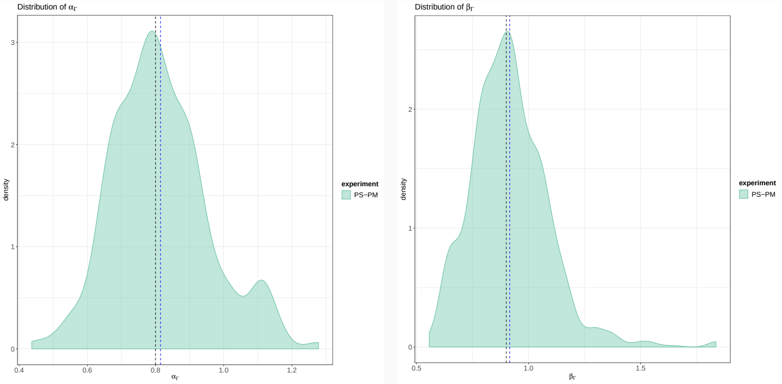

Figure 3: Distribution of the estimates of $(\hat{\alpha}, \hat{\beta})$. The black vertical lines represents the true values (α, β). The blue vertical lines represent the average estimates $(\hat{\alpha}, \hat{\beta})$.

Baseline: first two moments

But what is the performance in terms of the first two moments? In the end, we are concerned in the distributional properties of the export costs, not in $\hat{\alpha}$, $\hat{\beta}$

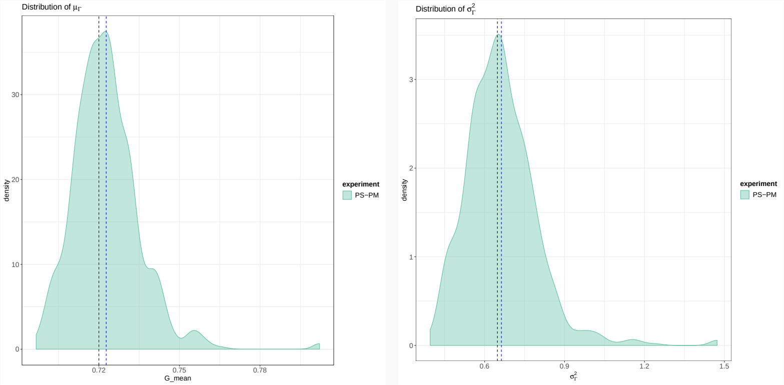

Figure 4: Distribution of the estimates of the first four moments of the distribution.

Baseline: Flexibility of the Gamma and compensation effects

We are sometimes far from estimating the true parameters, but the moments MSEs are lower.

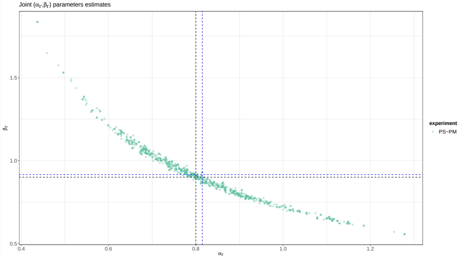

Figure 5: Joint estimates of $(\hat{\alpha}, \hat{\beta})$ in the PM-PS experiment.

There is a *compensation* effect such that when the shape parameters is under-estimated the scale parameter is over-estimated (and vice versa).

This is due to the flexibility of the Γ distribution.

Example of three Γ distributions

This flexibility is extremely good for our practical purposes.

Because even when we make mistakes in terms of parameter estimations, we still obtain a fairly accurate export thresholds density.

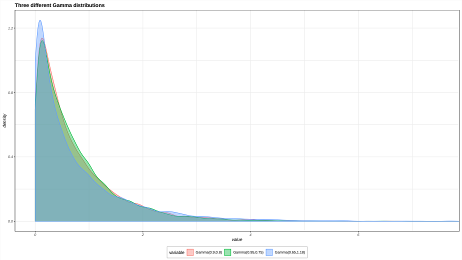

Figure 6: Three different Γ distributions with three different parametrisations.

General MonteCarlo framework

We set the general extensive margin decision problem as:

$$\chi_i = \begin{cases} 1 & \iff \theta_i + \varepsilon_i^\theta > c_i + \varepsilon_i^c \\ 0 & \iff \theta_i + \varepsilon_i^\theta \leq c_i + \varepsilon_i^c \end{cases} \tag{2}$$

where ε^θ and ε^c represent *iid* components that affect respectively the researcher measurement and the agent's decision:

- **IM - Imperfect Measurement**. ε^θ is similar to a measurement error of the $\theta-$attribute (e.g. productivity).
- **IS - Imperfect Sorting**. ε^c is similar to an error by the agent about the actual magnitude of entry barriers.

Test robustness against failures of A1

This general problem give rise to four alternatives:

	PM		IM	
PS	$\varepsilon_i^\theta = 0$	$\varepsilon_i^c = 0$	$\varepsilon_i^\theta \sim \mathcal{N}(0,\sigma)$	$\varepsilon_i^c = 0$
IS	$\varepsilon_i^\theta = 0$	$\varepsilon_i^c \sim \mathcal{N}(0,\sigma)$	$\varepsilon_i^\theta \sim \mathcal{N}(0,\sigma)$	$\varepsilon_i^c \sim \mathcal{N}(0,\sigma)$

Here, for each scenario:

- 10k firms;
- $c \sim \Gamma(\alpha, \beta)$ (i.e. A2 is valid);
- θ with possibly different distributions ($\mathcal{N}, \mathcal{B}, \mathcal{U}, \mathcal{P}_{II}$);
- $\sigma \in \{3\%, 6\%, 9\%, 12\%, 15\%\}$;
- 200 MonteCarlo simulations.

This yields to a total of 16k independent runs

Test A1 violations: shape parameter

But how does the estimation performs when A1 is violated?

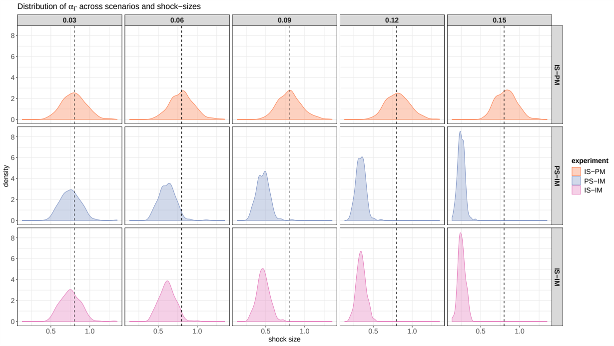

Figure 7: Estimation of the shape parameter $\hat{\alpha}$. Rows represent the three different scenarios. Column represent different sizes of the variance of the disturbance term.

Test A1 violations: scale parameter

But how does the estimation changes when A1 is violated?

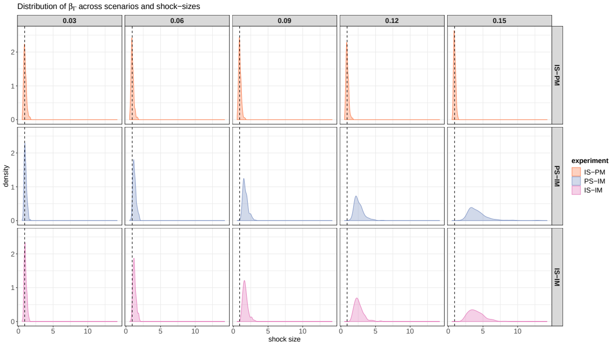

Figure 8: Estimation of the scale parameter $\hat{\beta}$. Rows represent the three different scenarios. Column represent different sizes of the variance of the disturbance term.

Test robustness against failures of A2

In this case we have six alternatives:

	Γ_{MLE}	$\mathcal{N}_{MLE}$
$c \sim \mathcal{N}$	$\nparallel$	$\checkmark$
$c \sim \Gamma_s$	$\checkmark$	$\nparallel$
$c \sim \Gamma_a$	$\checkmark$	$\nparallel$

Here for each scenario:

- 10k agents;
- PS-PM case (i.e. A1 is valid);
- $\theta \sim \mathcal{N}$;
- 200 MonteCarlo simulations.

This yields to a total of 1.2k independent runs.

Test A2 violations: using the $\mathcal{N}_{MLE}$ approach

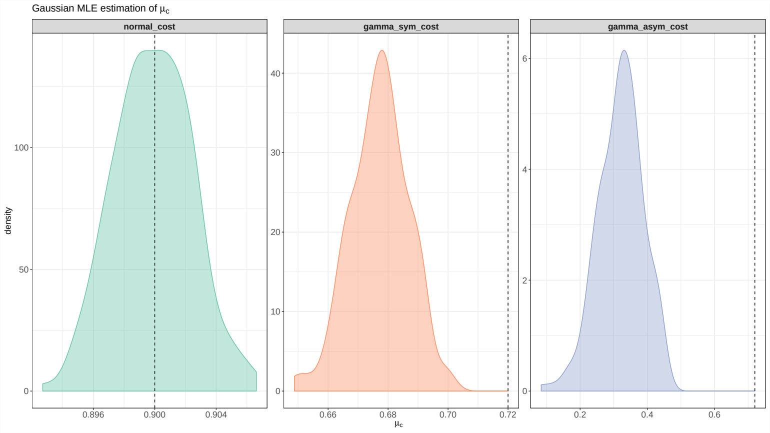

Figure 9: Estimation of the mean of c with Normal-MLE under different cost hypotheses.

Test A2 violations: using the Γ_{MLE} approach

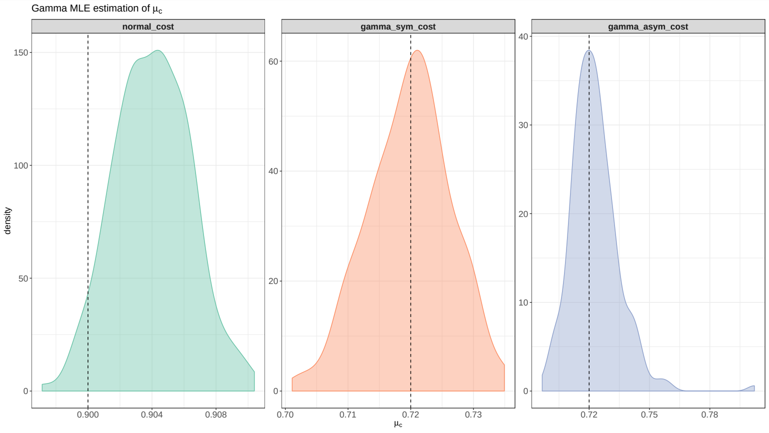

Figure 10: Estimation of the mean of c with Gamma-MLE under different cost hypotheses.

Test A2: comparison of the quality of the estimates

A quantitative comparison:

MLE	costs	MSE_μ	MSE_σ	MSE_{sk}	MSE_k
$\mathcal{N}$	$\mathcal{N}$	0.00000630	0.000000923	0.00	0.00
$\mathcal{N}$	Γ_s	0.00193	0.000576	0.444	0.444
$\mathcal{N}$	Γ_a	0.164	0.190	5.00	56.25
Γ	$\mathcal{N}$	0.0000219	0.00000213	0.104	0.0246
Γ	Γ_s	0.0000428	0.00000930	0.000506	0.00204
Γ	Γ_a	0.000156	0.0183	0.0394	1.84

Table 1: Mean Squared Errors of the estimated first four moments over the different scenarios.

Empirical Application

The export extensive margin problem

In the international trade literature it is typical to assume that most efficient firms self-select themselves into export activity [Melitz, 2003, Melitz and Ottaviano, 2008].

This means that for a firm i, the export extensive margin decision depends upon:

$$\theta_i > c \tag{3}$$

This simple assumption:

- allows one to vastly explain the productivity premium [Bernard and Jensen, 2004];
- clashes with empirical evidence about exporters mismatch [Eaton et al., 2011].

Economic intuition suggests to look for heterogeneity in export costs. But this is difficult to treat theoretically without further assumptions [Mayer et al., 2014].

Data

- Panel database of French manufacturing firms covering the period 1990-2007 (EAE data)
- Firms of at least 20 employees and turnover higher than 5 Millions Euro
- Dataset of about 300k observations

Sector by Sector Estimation

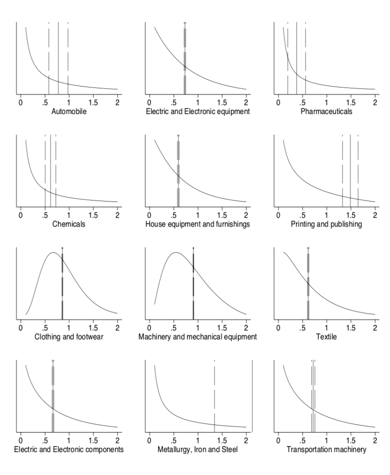

Figure 11: Estimation of the mean of c with Gamma-MLE for French firms in separate sectors.

Explaining participation rate

Table 2: Explaining Participation Rates by Means of Export Premium and Estimated Threshold Average and Variance

	(1)	(2)	(3)	(4)	(5)
TFP Premium	0.239				
	(0.129)*				
Normal μ		-0.018	-0.080		
		(0.006)***	(0.006)***		
Normal σ^2			-0.062		
			(0.005)***		
Gamma μ				-0.205	-0.371
				(0.013)***	(0.009)***
Gamma σ^2					0.019
					(0.001)***
Constant	0.734	0.755	0.893	0.910	1.004
	(0.008)***	(0.005)***	(0.014)***	(0.010)***	(0.006)***
Observations	261	223	223	211	211
Adj. R-squared	0.009	0.128	0.465	0.551	0.900
F Statistics	3.4*	33.5***	97.6***	259.1***	950.9***

Conclusion

Wrap Up

Results

We show that:

- when both A1 and A2 are valid, we precisely estimate **c**;
- even under some violations of A1, the estimates are precise;
- assuming a Γ_{MLE} is in general better (safer) than assuming a $\mathcal{N}_{MLE}$;
- French firms export costs are very skewed;
- the Γ_{MLE} assumption is corroborated by the French evidence.

Future research

We can *easily* extend our approach to:

- entry cost in general in other domain and by employing the time dimension;
- entry cost in different locations or for different products (bivariate Γ distribution);
- comparison also with a Pareto MLE.

Thanks for the attention!

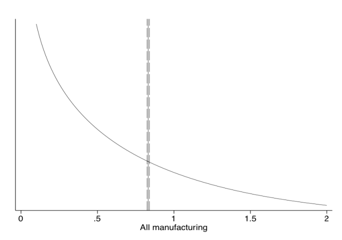

Figure 12: Estimation of the mean of c with Gamma-MLE for French firms. Pooled Manufacturing.

References

References i

Bernard, A. B. and Jensen, J. B. (2004).

Why some firms export.

The Review of Economics and Statistics, 86(2):561–569.

Eaton, J., Kortum, S., and Kramarz, F. (2011).

An anatomy of international trade: Evidence from french firms.

Econometrica, 79:1453–1498.

Mayer, T., Melitz, M. J., and Ottaviano, G. I. P. (2014).

Market size, competition, and the product mix of exporters.

American Economic Review, 104(2):495–536.

Melitz, M. J. (2003).

The impact of trade on intra-industry reallocations and aggregate industry productivity.

Econometrica, 71:1695–1725.

Melitz, M. J. and Ottaviano, G. I. P. (2008).

Market size, trade, and productivity.

The Review of Economic Studies, 75(1):295–316.

Additional Material

Summary statistics

Table 3: Participation Rate and Productivity Export Premium In French Manufacturing - 1990-2007

Industry	♯Obs.	PR	XPP
All manufacturing	353,100	0.733	0.042
Automobile	7,949	0.786	0.023
Chemicals	35,738	0.837	0.025
Clothing and footwear	28,276	0.674	0.143
Electric and Electronic components	14,714	0.776	0.064
Electric and Electronic equipment	19,551	0.756	0.064
House equipment and furnishings	24,054	0.823	0.056
Machinery and mechanical equipment	63,032	0.706	0.043
Metallurgy, Iron and Steel	61,661	0.730	0.000
Mineral industries	15,413	0.587	0.000
Pharmaceuticals	8,270	0.914	0.036
Printing and publishing	29,991	0.617	0.048
Textile	22,113	0.800	0.082
Wood and paper	22,338	0.695	-0.016

PR: Participation rate

XPP: Export Productivity Premium. $XPP = \bar{\theta}_X - \bar{\theta}_D$

All TFP are transformed to wipe out sector-year fixed effects