Sudden Stops, Productivity, and the Exchange Rate*

Laura Castillo-Martinez†

December 2025

Abstract

Following a sudden stop, productivity often declines, while real exchange rates adjust through a nominal depreciation, lower domestic prices, or both. Cross-country evidence suggests that productivity declines are larger when nominal depreciation dominates the real exchange rate adjustment. Motivated by this pattern, the paper studies how the nature of exchange rate adjustment shapes productivity dynamics during sudden stops. Using Spanish manufacturing micro-data from two sudden stops under different regimes, it shows that, in a currency union, cleansing through exit is stronger than under a floating regime, with aggregate productivity rising despite weaker firm-level performance. A small open-economy DSGE model with firm dynamics, endogenous markups, and nominal rigidities rationalizes these findings. The model identifies three channels through which a sudden stop affects productivity: pro-competitive, cost, and demand. While only the first operates under a floating regime, all three are active in a currency union. Quantitatively, the model is able to explain about 55 percent of the exit-driven contribution to productivity growth in Spain's 2010–13 episode.

JEL codes: D24, E52, F32, F41, O57.
Keywords: sudden stops, currency union, productivity.

*I am indebted to Ricardo Reis and Silvana Tenreyro for their continual guidance and support throughout this project. I am also grateful to Gianluca Benigno and Gianmarco Ottaviano for their advice during the job market. I would like to thank Mark Aguiar, Miguel Bandeira, Loukas Karabarbounis, Tiloka De Silva, EAGLS, Manu García, Ethan Ilzetzki, Pablo Ottonello, Şebnem Kalemli-Özcan (discussant), Matthias Kehrig, Andrea Lanteri, and all other participants at many seminars and conferences for helpful comments and discussions. Catherine Gonzalez provided excellent research assistance. This research benefited from financial support from Club Financiero Español and the Bank of Spain. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

†Duke University, FRB Minneapolis and CEPR. Email: l.castillo-martinez@duke.edu.

1 Introduction

The role of fixed versus flexible exchange rates during balance-of-payments crises has long been debated in international macroeconomics. The European sovereign debt crisis offers a compelling opportunity to revisit this question. After Greece admitted to having misreported its public debt figures in late 2009, a sudden stop in capital flows hit the periphery of Europe. This episode was unusual not only because it unfolded within a currency union, but also because most affected countries saw measured total factor productivity (TFP) rise, breaking with the historical pattern of productivity contraction during these crises.[1]

I study how exchange rate regimes shape the relationship between sudden stops and productivity dynamics. Cross-country aggregate evidence suggests that productivity growth declines with exchange rate flexibility. To uncover the mechanisms behind this pattern, I move from aggregate data to firm-level manufacturing data from Spain, which experienced two sudden stops under contrasting exchange rate regimes: one within the euro area in 2010–13 and one under a floating currency in 1992–93. The firm-level evidence shows that stronger selection and a larger cleansing effect raise aggregate productivity in the former episode, in contrast to the latter, despite weaker firm-level performance in both cases.

I rationalize these findings with a small open economy model featuring firm dynamics, endogenous markups, nominal rigidities, and a novel link between consumer labor income and firm profits. In a sudden stop in a currency union, the economy gains international competitiveness through lower wages, which improve firm selection via demand contraction. Under a floating regime, the nominal exchange rate bears the adjustment burden, limiting selection and dampening aggregate TFP gains. I use the model to replicate Spain's experience during the 2010–13 episode, quantify the contribution of firm selection to productivity growth, and contrast it with a counterfactual floating regime resembling 1992-93. Overall, the paper shows that firm heterogeneity reshapes how real exchange rate adjustment translates into macroeconomic outcomes during a sudden stop.

Section 2 explores the behavior of macroeconomic variables during sudden stops across a broad set of economies over the 1990–2015 period. Using a standard criterion to identify sudden stops that captures both the episodes emphasized in the earlier literature as well as the most recent European cases, I first confirm the well-established fact that aggregate TFP declines on average. Conditioning on the prevailing exchange rate regime, however, reveals a new pattern: productivity contractions are larger in economies with more flexible exchange rates. This pattern is statistically significant and robust to alternative exchange rate classifications, detrending methods, and controls for crisis characteristics and country-level features. Since all other macroe-

[1] During the crisis period, 2010-2013, TFP rose in Ireland (3.5%), Portugal (1.5%), and Spain (1.8%), and remained unchanged in Italy. For context, since 2000, these economies had seen TFP decline by 2.6%, 2.1%, 2.8%, and 6.4%, respectively (AMECO).

conomic variables exhibit broadly familiar dynamics, the remainder of the paper turns to a closer examination of productivity dynamics.

Section 3 focuses on micro evidence from the Spanish manufacturing sector. Specifically, I exploit firm-level data from the 2010–13 European sovereign debt crisis and contrast it with an earlier sudden stop that hit Spain in 1992–93 during the Exchange Rate Mechanism crisis. The two episodes share a similar crisis onset but differ sharply in their exchange rate policy response. While during the earlier sudden stop the national currency, the *peseta*, depreciated throughout the episode, during the latter Spain was a member of a currency union and could only regain competitiveness by lowering wages. These differences are accompanied by contrasting patterns of aggregate productivity adjustment that mirror the cross-country evidence.

The data granularity allows me to measure TFP at the firm level and aggregate it to an industry-wide measure. A decomposition of aggregate productivity growth shows that, although firm-level productivity declines during both crises, the exit of low-productivity firms, together with reallocation toward surviving incumbents, accounts for the positive aggregate growth observed in the 2010–13 sudden stop. Consistent with this pattern, differences in firm-level productivity changes across episodes are disproportionately concentrated in the lower tail of the productivity distribution. Moreover, a formal regression-based test for cleansing indicates that selection is strengthened in 2010–13 but not in 1992–93. Finally, among other firm characteristics shaping exit dynamics, higher-markup firms are more likely to survive, both in normal times and during 2010–13. This pattern, however, is largely explained by their higher underlying productivity.

I also consider a number of competing explanations for the observed cross-episode differences. While potentially complementary, differences in the severity of the concurrent banking crises, expenditure-switching effects associated with real depreciations, and the geographic scope of the underlying shocks do not fully account for these findings. Taken together, these results suggest that firm dynamics play a key role in shaping the relationship between sudden stops, aggregate productivity, and exchange rate adjustment.

Based on the previous evidence, section 4 develops a small open economy model with firm heterogeneity and exchange rate policy to study the macroeconomic effects of a sudden stop. Preferences are specified to generate endogenous firm selection and variable markups, consistent with the empirical patterns documented in the data. Firms differ in idiosyncratic productivity and face a stochastic option value of exit, giving rise to entry and exit across the entire firm-size distribution. Firm profitability is linked to aggregate conditions through endogenous labor supply, introducing a channel through which income fluctuations affect firms' costs and exit decisions.

To capture the exchange rate regime differences documented in the data, I introduce nominal rigidities in the wage-setting process. The central bank chooses the nominal exchange rate as its main policy tool. I focus on two extreme regimes: a currency union, characterized by a credible commitment to keep the nominal exchange rate constant; and a strict wage inflation targeting regime, where the flexible wage equilibrium is always implemented. A sudden stop is defined as a two-fold shock to the domestic economy. First, it involves an increase in the interest rate that consumers pay when borrowing abroad. By increasing the cost of credit, the domestic economy is forced to deleverage internationally and increase net exports through a real exchange rate depreciation. Second, it simultaneously features a decline in the productivity level of all firms.

Section 5 discusses the effects of a sudden stop shock on aggregate productivity in a simpler version of the model. Abstracting from intertemporal firm decisions—by holding idiosyncratic productivity fixed, shutting down the option value of exit, and simplifying wage dynamics—and in the absence of free entry, the model admits a closed-form solution.The key insight is that aggregate productivity is proportional to a domestic productivity threshold. The threshold represents the minimum productivity level at which a firm can generate positive profits and, thus, select into the domestic market. It therefore suffices to understand how the threshold moves after a sudden stop to learn about its effect on aggregate productivity.

In equilibrium, the domestic threshold is determined by the number of active firms in the market, the wage level and the marginal utility of income. Therefore, there are three endogenous mechanisms through which a shock can affect productivity. First, the threshold increases with the number of active firms, as greater competition lowers profit margins for all firms and, thus, requires a higher level of productivity to remain profitable. This is the pro-competitive channel. Second, higher wages increase the costs of production for all firms, lowering again their profit margin and calling for a higher productivity level. This is the cost channel. Third, higher income increases the demand for overall consumption. This, instead, increases the firm profit margin and relaxes the productivity requirement. This is the demand channel.

The effect of a sudden stop on the domestic productivity threshold will hinge on the relative strength of these conflicting forces. This, in turn, depends on the extent to which the interest rate increase that characterizes a sudden stop is offset by a nominal exchange rate depreciation. If the nominal exchange rate fully adjusts, both the marginal utility of income and wages remain unchanged. As a result, only the pro-competitive channel is active: fewer firms import, leading to an unambiguous decline in productivity. In contrast, if the nominal exchange rate is fixed, the marginal utility of income adjusts accordingly, activating all three channels and making the overall effect quantitatively ambiguous. The simplified model identifies conditions under which the demand channel dominates, allowing a sudden stop to improve productivity in a currency union.

I provide empirical support for the relevance of the demand channel by exploiting sectoral variation in the firm-level data. Intuitively, this channel should be stronger in sectors more reliant on domestic consumers. Consistent with this prediction, I document that during the 2010–13 sudden stop, less tradable sectors experienced stronger productivity growth, with a larger share attributable to firm exit. I also assess the robustness of the main result to relaxing some of the assumptions. Allowing for CES preferences introduces an anti-competitive channel, while incorporating quality-enhancing imported intermediate inputs gives rise to a foreign quality channel. Crucially, both extensions affect firm behavior symmetrically under fixed and flexible exchange rate regimes following a sudden stop, and thus do not alter the differential implications of the core mechanism.

I turn towards studying the quantitative properties of the full model in section 6. To do so, I first parameterize the model using a combination of values from the literature and a moment-matching exercise. The model performs well in replicating key features of the firm size distribution. Second, I solve for the transition dynamics of the model following a shock to the interest rate and a common shifter of firm productivity. Plotting the impulse response function of aggregate TFP confirms that the previous analytical results hold more generally: productivity falls under a floating arrangement and increases in a currency union following a sudden stop. In addition, the model generates the other stylized facts previously documented by the literature: a contraction in output, a reversal in the current account and a real exchange rate depreciation.

Next, I use the quantitative model to simulate the 2010–13 sudden stop in Spain. I begin by filtering the sequences of shocks, the interest rate and the common productivity shifter, required to match the observed evolution of the real exchange rate and output over the 2002–2014 period following Spain's entry into the euro. The model replicates the observed increase in TFP during the sudden stop episode, albeit slightly overpredicting its magnitude. I then perform the productivity growth decomposition exercise described in Section 3 for two samples: one mirroring the coverage of the previously used survey data, and the other comprising the full population of firms in the model. The model accounts for roughly 55% of the exit-driven contribution to productivity growth observed in the data. Comparing the two samples reveals that, while the survey-based sample modestly overstates the magnitude of the selection margin, the overall message is unchanged. As a robustness check, I also examine how much of the observed dynamics the model can explain in the absence of the common productivity shifter.

I then study the mechanisms underlying the 1992–93 sudden stop through a counterfactual exercise. Since a direct simulation is not feasible due to data limitations, I consider a scenario in which Spain operates under a flexible exchange rate regime while experiencing the same real exchange rate adjustment and productivity shock as in the later episode. In contrast to the baseline, productivity declines, driven by a fall in incumbent performance and the absence of a strong cleansing effect. This mirrors key patterns observed in the micro-data during the 1992–93 period. Finally, to assess the broader implications, I compare welfare across regimes using a utility-based measure. The analysis suggests that, during the sudden stop, agents prefer to remain in the currency union, motivating a re-examination of the fixed versus floating debate.

Relation to the literature This paper contributes to several strands of the literature at the intersection of international macroeconomics and firm dynamics.

The starting point of this paper is the literature on sudden stops, first defined by Calvo (1998) as abrupt reversals in foreign capital inflows. This literature documents common empirical regularities across historical episodes (e.g., Calvo, Izquierdo and Mejía (2004), Guidotti et al. (2004), Kehoe and Ruhl (2009)) and develops theoretical frameworks that generate these regularities by modeling sudden stops as adverse shocks to productivity and external borrowing conditions (e.g., Neumeyer and Perri (2005), Meza and Quintin (2007), Mendoza (2010)). This paper contributes to the former by revisiting established stylized facts through the lens of exchange rate policy, and to the latter by incorporating exchange rate policy in a model that departs from the representative-firm paradigm.

A second related line of research embeds firm dynamics into open-economy macroeconomic models to study the transmission of external shocks. Ghironi and Melitz (2005) are among the first to introduce endogenous entry and trade participation in a DSGE framework, while maintaining constant firm heterogeneity. Parallel work adds dynamic firm adjustment through trade states and within-firm margins (Alessandria and Choi, 2007; Alessandria, Kaboski and Midrigan, 2010). This paper instead builds on the Hopenhayn (1992) tradition by allowing for time-varying firm-level productivity.[2] In this sense, the paper is closer to Ates and Saffie (2021), who also study the productivity costs of sudden stops with firm dynamics. Their focus, however, is on the long run effects of entry distortions and how financial selection cushions the fall in endogenous productivity.

Finally, this paper is connected to the literature that studies the contribution of reallocation to TFP growth. In particular, I provide empirical support for Caballero and Hammour (1994)'s cleansing hypothesis and discuss the conditions under which its magnitude is likely to be relevant in the context of a sudden stop.[3] Moreover, this work adds to the recent set of papers that link declining TFP and enhanced misallocation with capital inflows; see Reis (2013), Benigno and Fornaro (2014) and, especially, Gopinath et al. (2017). They study, however, an earlier period.

2 Sudden Stops: Revisiting the Stylized Facts

This section provides cross-country evidence on the behavior of macroeconomic variables during sudden stops. I identify sudden stop episodes using a standard criterion and study the associated dynamics using an event-study approach. Both steps are standard in the literature. I build on them by examining how the resulting dynamics vary across sudden stops occurring under different exchange rate regimes.

[2] Models featuring firm selection and productivity dynamics are well established in closed-economy settings, but have been applied less extensively to open economies, especially in the context of international financial shocks and exchange rate policy.

[3] The cleansing hypothesis is an interpretation of Schumpeter (1939)'s creative destruction argument that emphasizes the role of reallocation among new and incumbent firms at a business cycle frequency.

2.1 Data and Methodology

Following Cavallo and Frankel (2008), I define a sudden stop as an episode in which there is a substantial decline in the capital account surplus together with a recession.[4] In particular, I classify as a sudden stop a period that contains at least one year during which (i) the financial account surplus has fallen at least one standard deviation below its rolling average and (ii) GDP per capita contracts.[5] The start and end of each episode is marked by the first and last year within the period in which the financial account surplus is half a standard deviation below the rolling average.[6] The latter requirement ensures that the capital flow reversals captured by the algorithm strictly qualify as sudden stops; first, by requiring that the financing disruption is accompanied by an appropriate macroeconomic adjustment, and second, by ruling out booming episodes that display similar characteristics, for example a positive trade shock. All data is collected from standard sources and, thus, its description is relegated to Online Appendix A.1

The total number of episodes is 78, representing 5.2% of total available country/year observations in the sample. The full list of episodes per country, plus exchange rate classification, is given by Table A.1. The criterion successfully captures all traditional sudden stop episodes previously discussed by the literature , mostly occurring around the 1994/5 Tequila crisis, the 1997 Asian Financial Crisis, the 1998 Russian default, as well as the most recent balance of payment crisis in the peripheral economies of the European Union.[7]

I build on Ilzetzki, Reinhart and Rogoff (2019) updated *de facto* coding system in order to bin episodes by exchange rate flexibility. In my baseline results, I consider as prevalent the exchange rate regime that is in place during the last year of the sudden stop. There are four different cases: a currency union, a hard peg, a soft peg and a floating arrangement.[8] Out of the 78 episodes identified, 11 occur within a currency union (8 in the euro area and 3 in the West African Economic and Monetary Union), 14 in a hard peg system, 26 in a soft peg regime and 25 in a floating arrangement.

[4] The practice of conditioning on output contraction goes back as far as the canonical Calvo, Izquierdo and Mejía (2004) methodology.

[5] This contrasts with Cavallo and Frankel (2008), who also require an improvement in the current account deficit (or an equivalent decline in foreign reserves). As this is conceptually equivalent to the first condition, I drop it.

[6] Refer to Online Appendix A.2 for further details.

[7] The methodology does not account for changes in TARGET2 balances in the Eurozone and, thus, prevents me from measuring private capital flows accurately. However, this is not problematic for my purposes as the algorithm already identifies the GIIPS episodes.

[8] In terms of the Ilzetzki, Reinhart and Rogoff (2019) fine classification, I deviate as follows: (1) I manually divide code 1 into currency union and no separate legal tender, (2) I group codes 2 to 4 under the hard peg category, (3) I group codes 5 to 11 under the soft peg category, (4) I group codes 12 to 14 under the floating arrangement and (5) I rename group 15 as 5, i.e., other categories.

FIGURE I: PRODUCTIVITY DURING A SUDDEN STOP

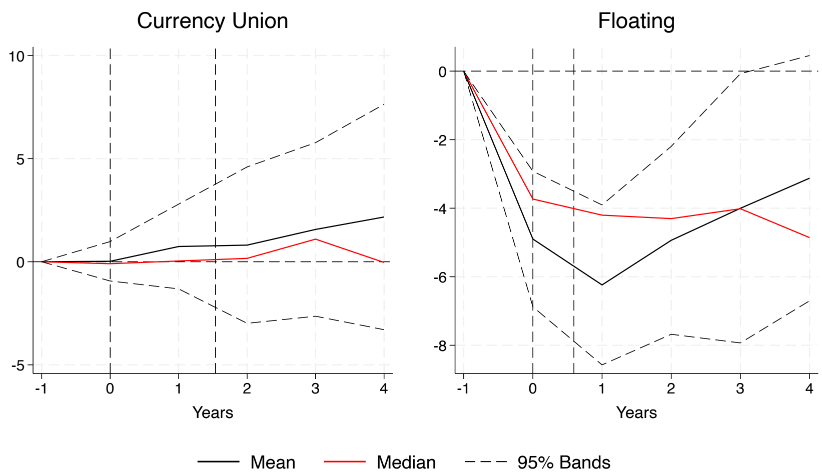

Notes: This figure plots the response of TFP to a sudden stop under a curreny union (left) and under a floating arrangement (right). The black and red solid lines depict the mean and median path of productivity while the black dashed lines represent standard error bands. The two vertical lines show the start and end of an average episode. Productivity is expressed as percentage deviations from an extrapolated linear trend fitted over the five-year pre–sudden-stop window. The data used is collected from IFS, WDI and the Total Economy Database.

2.2 Event Study Analysis

I study the dynamic behavior of macroeconomic variables following the onset of a sudden stop using an event-study framework. For each episode, the start of the sudden stop is normalized to $t = 0$, and outcomes are traced over an event window. The analysis focuses on output, consumption, employment, productivity, wages, the real exchange rate, the current account, exports, and imports. All variables are detrended using a linear pre–sudden-stop trend and expressed as percentage deviations from that trend, except for the current account, which is reported as a share of GDP, and the real exchange rate, which is reported in levels.[9]

I begin by confirming a well-established empirical regularity: productivity falls sharply following a sudden stop. Figure A.1 plots the mean and median path of total factor productivity across all identified episodes, together with standard error bands. On average, productivity declines on impact and remains depressed in the years following the sudden stop, consistent with earlier evidence in the literature.

The analysis then turns to the role of exchange rate regimes. Figure I replicates the event study after binning episodes by their exchange rate arrangement, contrasting sudden stops occurring within a currency union (left panel) with those taking place under a floating regime (right panel).

[9]I choose a five-year window to balance isolating medium-run trends with maximizing the number of episodes for which sufficiently long pre-event data are available. Results obtained using a ten-year pre–sudden-stop window are reported as a robustness check in Online Appendix A.3 and lead to similar conclusions.

TABLE I: MACRO VARIABLES DURING A SUDDEN STOP

	TFP (1)	Output (2)	Consumption (3)	Employment (4)	Wages (5)	CA (6)	RER (7)
$t = 0$	0.023	-2.452	-1.256	-2.869*	-2.973*	3.536*	0.328
	(0.422)	(1.458)	(2.013)	(1.207)	(1.373)	(1.540)	(1.022)
$t = 1$	0.739	-4.583*	-3.948*	-5.868**	-5.375**	4.958*	-1.124
	(0.908)	(1.710)	(1.859)	(1.689)	(1.768)	(2.270)	(0.793)
$t = 2$	0.806	-7.001**	-6.305*	-9.323***	-7.126***	5.195**	-0.701
	(1.675)	(2.401)	(3.105)	(2.022)	(1.972)	(1.609)	(0.810)
$t = 0 \times float$	-4.923***	-4.818	-5.189	2.432	-4.810	0.134	-13.838
	(1.029)	(2.691)	(2.668)	(1.378)	(2.715)	(1.949)	(7.287)
$t = 1 \times float$	-6.978***	-4.338	-6.804*	4.220	-1.892	-0.259	-11.284
	(1.432)	(2.740)	(2.943)	(2.102)	(3.329)	(2.586)	(6.123)
$t = 2 \times float$	-5.700*	-0.399	-4.623	7.248*	0.210	-0.394	-10.915
	(2.093)	(3.555)	(4.136)	(2.699)	(3.765)	(2.155)	(5.649)
N	99	156	148	104	152	159	103
R^2	0.427	0.269	0.355	0.388	0.189	0.221	0.171

Notes: This table reports regression-based event-study estimates of the response of macroeconomic variables to a sudden stop. Regressions include event-time indicators interacted with a floating exchange rate dummy, as well as sudden-stop episode fixed effects. All real variables are expressed as percentage deviations from an extrapolated linear trend fitted over the five-year pre–sudden-stop window. The current account is expressed as a share of GDP, and the real exchange rate (RER), constructed as an index, is reported in levels. Standard errors, clustered at the episode level, are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

The responses differ systematically across regimes. Under a floating arrangement, productivity exhibits a clear and persistent decline. By contrast, within a currency union, productivity remains broadly unchanged and, if anything, displays a slight improvement in the aftermath of the sudden stop, although the associated confidence bands are admittedly wide. [10]

To assess whether the differences across exchange rate regimes are statistically significant, Table I reports regression-based event-study estimates. The table compares the average response of macroeconomic variables across regimes by interacting event-time indicators with a floating-regime dummy. Standard errors are clustered at the episode level. The first column indicates that, relative to a currency union, the post–sudden-stop decline in productivity under floating arrangements is about 5 percentage points larger on impact and about 7 percentage points larger one year after the onset of the episode, with both differences statistically significant. I conduct a battery of robustness checks to evaluate the consistency of this productivity result, including alternative approaches to exchange rate classification, detrending windows, and controlling for crisis- and country-specific characteristics. The results of these exercises are reported in Online Appendix A.3.

Finally, the behavior of the remaining macroeconomic variables is consistent with conventional

[10]For completeness, Figure A.2 reports the corresponding results for hard and soft pegs. Consistent with a negative relationship between flexibility and TFP growth, productivity increases under hard pegs, albeit with wider confidence bands than in currency unions, while under soft pegs there is a noticeable decline, particularly on impact.

accounts of sudden stops. Output, consumption, employment, and wages all decline following the onset of a sudden stop. Capital flow reversals coincide with a depreciation of the real exchange rate, reflected in a decline in the RER index. The current account adjusts rapidly, with deficits narrowing sharply and approaching balance within one year of the sudden stop. These patterns are broadly similar across exchange rate regimes. The dynamic responses are illustrated in Figures A.3 and A.4, and the corresponding estimates are reported in the remaining columns of Table I.

3 Spain: A Tale of Two Sudden Stops

Motivated by the aggregate patterns suggested in Figure I, this section turns to a case-study approach that allows for a closer examination of the mechanisms underlying productivity dynamics during sudden stops. While Section 2 suggests that the relationship between exchange rate flexibility and productivity is not specific to the euro area, a more granular analysis is needed to understand how these patterns arise. This section takes that route, exploiting firm-level data from two sudden stops in Spain's recent economic history: the 1992–93 Exchange Rate Mechanism (ERM) crisis and the 2010–2013 European sovereign debt crisis.

There are clear parallels between these two episodes regarding the onset. Both were preceded by periods of increasing capital inflows, declining international competitiveness and widening current account deficits. Economic growth was fueled by the construction sector, with steep increases in property prices and crawling private debt. Public finances, on the other hand, were in a similar good shape.

Foreign capital inflows abruptly reverted following a confidence crisis affecting the European integration project: the negative outcome of the Danish referendum on the Maastricht Treaty in the first case, and the Greek announcement of substantial upward revisions in the government budget deficit more recently. The flight of international investment led to an urgent correction of misaligned real exchange rates in order to expand net exports. As growth stalled and unemployment rose, austerity measures were put in place in order to curb the rising public deficits generated by automatic stabilizers. In addition, structural reforms aimed at increasing the flexibility of the labor market were passed during both episodes.[11]

The response of exchange rate policy to these events, however, diverged significantly. While the *peseta* was devalued in three occasions during the 1992-93 crisis, Spain already shared a common currency with its largest trading partners since 2002 and underwent a process of internal

[11] There are two stark differences regarding these two sudden stops. First is the magnitude of the shock: Spain's current account surplus as a share of GDP moved from -3.5% to -1.2% between 1991 and 1994 versus -4.3% to 1.0% between 2009 and 2014. However, the duration was longer in the second episode, such that, per year, the reduction was around 1.1% during both episodes. Second, the latter episode coincides with a banking crisis, whereas the former did not. I partially address this concern by examining firms' leverage at the end of the section.

devaluation.[12,13] I take these episodes as examples of sudden stops under floating arrangements and currency unions, respectively, and use firm-level data to explore what is driving the observed aggregate TFP pattern.[14,15]

3.1 Data

I use firm-level data from the Survey on Business Strategies (Encuesta sobre Estrategias Empresariales, ESEE, in Spanish) managed by the SEPI Foundation, a public entity linked to the Spanish Ministry of Finance and Public Administrations. The ESEE surveys all manufacturing firms operating in Spain with more than 200 workers and a sample of firms between 10 and 200 workers, providing a rich panel dataset with over 1,800 firms for the period 1990-2014. It covers around 20 percent of output in Spanish manufacturing and provides information on each firm's balance sheet together with its profit and loss statement.[16]

The primary advantage of the ESEE dataset, especially relative to the ORBIS dataset compiled by Bureau van Dijk Electronic Publishing (BvD), is its long temporal coverage, which extends back to the early 1990s and enables a consistent firm-level analysis of both the 1992–93 Exchange Rate Mechanism crisis and the 2010–13 European sovereign debt crisis.[17] To the best of my knowledge, ESEE is the only Spanish firm-level dataset with reliable financial information spanning both episodes.

A second key advantage is the clean measurement of firm exit. The survey explicitly distinguishes between temporary non-response, permanent exit, and organizational changes such as

[12]In 1992, the *peseta* was first devalued by 5% on September 17th, known as Black Wednesday, when the pound and the lira abandoned the ERM altogether. A further 6% was devalued on November 23rd, with a third devaluation taking place in May 1993.

[13]Figure A.6 applies Engel (1999)'s decomposition to compare real exchange rate dynamics during the 1992–1993 and 2010–2013 sudden stops, using Spain as the domestic economy and Germany as the bilateral foreign partner. While real exchange rate movements in the later episode are entirely due to relative price adjustments (given the euro), over 95% of the variance in real exchange rate changes during the earlier episode is accounted for by nominal exchange rate fluctuations across all horizons.

[14]It can be argued that Spain does not strictly classify as a floating exchange rate regime in 1992-93 as it remains a member of the ERM, a multilateral party grid of exchange rates established in 1979. However, the repeated realignments of its central rate against the *deutsche mark* and the substantial widening of the exchange rate fluctuation bands meant that the overall devaluation of its currency was even larger than that of floating currencies such as the pound. In order words, despite the formal membership of the ERM, the exchange rate effectively behaved as flexible.

[15]Panel A of Figure A.5 shows that regardless of the data source, Spain experienced an improvement in aggregate TFP relative to its pre-crisis trend during 2010-13 sudden stop. While different datasets imply differing paths for TFP in levels—suggesting increases, flat dynamics, or mild declines—this disagreement vanishes once TFP is evaluated relative to trend. Since the focus of this paper is on business-cycle fluctuations rather than long-run growth, trend-adjusted TFP is the appropriate object of analysis.

[16]Table A.4 shows that while the coverage of the ESEE data relative to different releases of the EU KLEMS data for the Spanish manufacturing sector is modest, it is also constant over time.

[17]The other existing firm-level dataset used in the literature, such as in García-Santana et al. (2020), is the Central Balance Sheet Data (Central de Balances Integrada, CBI, in Spanish) owned by the Bank of Spain. This alternative dataset relies on the same underlying source as ORBIS—annual financial statements that firms are obliged to submit to the Commercial Registry—and is therefore subject to similar limitations. See Almunia, López-Rodríguez and Moral-Benito (2018) for details.

mergers, acquisitions, or split-ups. Firms that later resume production or survey participation are re-included and appropriately tracked, substantially reducing spurious exit and re-entry. Entry, however, is observed according to the survey's sampling structure, including all newly created firms with more than 200 employees and a random sample representing approximately 5% of firms with 10 to 200 employees each year.

Finally, ESEE offers several additional features that are particularly valuable for this analysis. It provides firm-level information on exports, which is typically subject to strict confidentiality restrictions in Spain, and is explicitly designed for research purposes, with substantial effort devoted to ensuring internal consistency and accuracy throughout the data collection process.

Details on the cleaning procedure and the deflating of nominal variables are relegated to Online Appendix B.3. I estimate industry output elasticities for capital and labor using Ackerberg, Caves and Frazer (2015)'s algorithm and then compute firm-level productivity as a Solow residual.

Representativeness A well-known shortcoming of the ESEE dataset is its over-representation of large firms. To address this issue, I exploit dynamic sampling weights constructed by the SEPI Foundation based on Census data. Figure A.8 shows that the weighted sample closely matches the actual firm size distribution as documented by Eurostat, substantially improving upon the raw (unweighted) data. In addition, Table A.5 validates the weighted sample along the extensive margin by comparing entry and exit rates for firms with ten or more employees to the Spanish Business Registry (Directorio Central de Empresas, DIRCE), both over time and across three-digit manufacturing sectors.

One remaining challenge of the ESEE dataset is that it provides limited information on micro firms, i.e., firms with nine employees or less. [18] According to administrative data, over 80% of all firms in the Spanish manufacturing sector are micro firms; yet, they account for only 9% of production and 11% of the wage bill.[19] Given the latter, the empirical disregard for micro firms is generally not a concern in macroeconomic analyses that emphasize the intensive margin. However, this paper also examines the role of the extensive margin, making it crucial to understand the potential bias introduced by firm-size truncation.

Before doing so, it's key to determine whether truncation binds differently over time. Figure A.9 shows that the share of micro firms in the Spanish economy has followed a remarkably stable pattern, particularly in the 15 years preceding the 2010–2013 sudden stop. While official data is unavailable before 1995, it is reasonable to expect that the lower end of the firm size distribution remained similar in the early 1990s. This alleviates concerns that the bias introduced by truncation might differ across sudden stops.

To additionally assess potential bias, I implement two strategies —one empirical, the other

[18] Micro firms are present only if they entered the sample with at least ten employees and later shrank below the threshold.

[19] Data corresponds to the Structural Business Statistics for years 2005-2020 as reported by Eurostat.

theoretical. In Online Appendix B.5, I use data from ORBIS to complement the baseline analysis. This dataset improves firm coverage by incorporating all micro firms and a broader share of small and medium sized manufacturing firms, but its time span, starting only in the early 2000s, restricts the analysis to the 2010–2013 sudden stop in Spain. To overcome this limitation, I turn to the 2011–2013 sudden stop in the Czech Republic as a comparative episode; one that unfolded in a similarly advanced economy, within a comparable time frame, but under a floating exchange rate regime and accompanied by a sharp decline in aggregate TFP. Results from this alternative case study reinforce the key asymmetry in firm behavior that I document in the next subsection.

In Section 6, I use the quantitative model to evaluate whether the empirical focus on larger firms, imposed by data limitations, meaningfully distorts the analysis. Specifically, I conduct the analysis twice: once using a sample designed to mimic the empirical dataset (by overrepresenting larger firms, as in the ESEE), and again using the full firm distribution, calibrated to match administrative data on firm size. This comparison allows me to assess whether the key empirical patterns are likely to generalize beyond the selected sample.

Firm Profile To dissect to what extent firms exposed to each of the two episodes are comparable, I include a table of mean differences in observable characteristics. Specifically, Table A.6 examines differences in age, size (both in terms of employment and sales), productivity growth, capital intensity, leverage and trade engagement (propensities to import and export). The results indicate statistically significant differences along two key dimensions: age and openness to trade. Firms in the later episode are older and more frequently engaged in international trade. These findings are reasonable given the 18-year gap between episodes, which naturally accounts for the increased firm age, and the greater integration within the single market following the introduction of the euro, which explains the heightened trade activity.

The bottom section of Table A.6 highlights the industry composition of firms exposed to each of the two sudden stops. While the shifts in industry composition appear minor in magnitude, they are statistically significant for nearly half of the sectors. This underscores the importance of accounting for industry composition in the firm-level analysis that follows.

3.2 Results

Aggregate TFP, defined as the employment-weighted average of firm-level TFP, fell by 12.96% during the 1992–1993 episode but rose by 2.24% in the 2010–2013 period. The granularity of the data allows for a more detailed investigation of the drivers of productivity.

The Lower Tail

I first document changes in the distribution of firm-level productivity before and after each of the crises. A visual inspection of the kernel probability distribution estimate of log TFP before and

FIGURE II: PRODUCTIVITY GROWTH ACROSS THE PRODUCTIVITY DISTRIBUTION

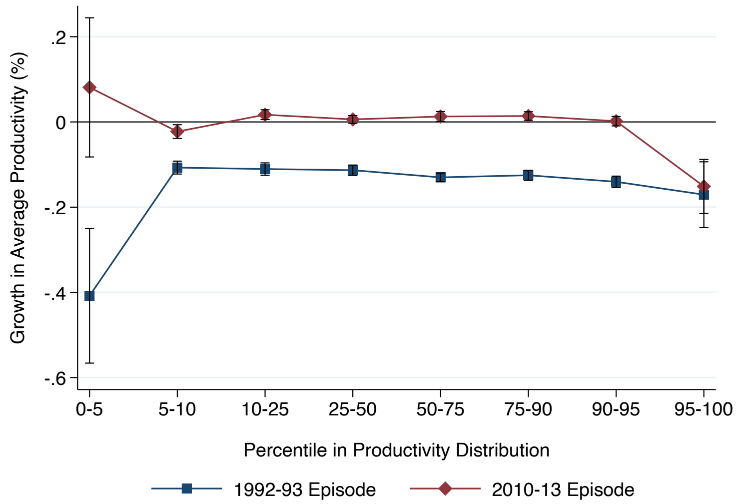

Notes: This graph plots the growth in average TFP by percentile of the productivity distribution. It compares the average TFP of firms in a given percentile before and after each of the two sudden stops. As this is an unbalanced panel, firms are allowed to change percentiles and even exit the sample during the transition. The corresponding base and end years are 1991 and 1993 for the first episode; 2009 and 2013 for the second episode. To account for variability, the vertical lines represent error bands. The data used is collected from the ESEE dataset.

after each of the two sudden stops confirms there is ample heterogeneity in TFP levels among firms in any given year as already highlighted by the literature. More surprisingly, the shape of the distribution is similar and remains unchanged throughout both crisis periods, with no major shifts. In fact, the lower tail concentrates most, if not all, of the action: it lengthens as TFP decreases in the former crisis while shortens as TFP increases in the latter case.

To see this graphically, Figure II presents the percentage change in average productivity for each percentile of the productivity distribution during the two sudden stops. On average, the difference in the change in productivity across episodes, the gap between the red and blue lines, is roughly constant across the entire distribution, with the notable exception of the 5% percentile where TFP decreases by 41% during 1991-1993 while increases by 8% during 2009-2013. Although the error bands are admittedly wide in both cases, the difference relative to other percentiles is large enough to remain relevant - the gap is three times the average.

Estimated moments of the distribution support the predominant role of the lower tail with higher-order moments experiencing the largest swings.[20] During the 1992-93 crisis firms display lower productivity on average and the dispersion of log TFP increases. The increase in dispersion, however, is asymmetric. The distribution of unproductive firms expands while that of productive firms changes little with the coefficient of skewness declining from -0.40 to -1.04. Moreover, increasing kurtosis, 6.12 versus 9.14, is associated with fatter tails as the probability mass moves

[20]Refer to Table A.7 for further details.

TABLE II: DECOMPOSITION OF PRODUCTIVITY GROWTH

	Sudden Stops	
	1992-93	2010-13
Productivity Growth (%)	-12.96	2.24
Contribution to Productivity Growth		
Incumbents' Contribution	**-13.37**	**-4.54**
Within-firm Contribution	-13.89	-7.36
Between-firm Contribution	-1.83	0.32
Cross-term Contribution	2.34	2.50
Net Entry Contribution	**0.41**	**6.79**
Entrants' Contribution	-1.81	-0.40
Exiters' Contribution	2.23	7.19

Notes: Productivity growth refers to accumulated TFP growth for the stated period. Base and final years are 1991 and 1993 for the first episode; 2009 and 2013 for the second episode. Contribution of incumbents and net entrants add up to productivity growth. Contribution of within-firm, between-firm and cross-term components add up to incumbents' contribution. Contribution of entrants and exiters add up to net entry contribution.
The data used is collected from the ESEE dataset.

away from the shoulders of the distribution. Although the behavior of TFP exactly reverses during the 2010-13 crisis - productivity increases while dispersion drops - it is still the tails, and especially, the lower tail, that changes the most. In this case, skewness increases from -2.21 to -0.65 while kurtosis shrinks from 24.98 to 6.40.

Decomposing Productivity Growth

While the above findings support a narrative of shifting productivity cutoffs, it is often the case that firms at the lower end of the productivity scale are small in size and, thus, have negligible effects on the aggregate. A more formal test of growth patterns requires considering weighted measures. Moreover, it should aim at disentangling the role of incumbent, entering and exiting firms in shaping TFP changes. I study this by performing a TFP growth decomposition using the formulation of Dias and Marques (2021):

$$\Delta Z_t = \sum_{i\in C} s^C_{i,t-1}\,\Delta Z_{i,t} + \sum_{i\in C} Z_{i,t-1}\,\Delta s^C_{i,t} + \sum_{i\in C} \Delta s^C_{i,t}\,\Delta Z_{i,t} + s^N_t\left(Z^N_t - Z^C_t\right) - s^X_{t-1}\left(Z^X_{t-1} - Z^C_{t-1}\right),$$

where Z_t denotes employment-weighted aggregate productivity and $Z_{i,t}$ is the productivity of firm i. The term $s_{i,t}$ is its employment share, and $s^C_{i,t} = s_{i,t}/s^C_t$ rescales these shares within the incumbent group. The superscripts C, N, and X indicate incumbents, entrants, and exiters, with s^j_t and Z^j_t denoting the corresponding group-level employment shares and aggregate productivity. Further details are provided in Online Appendix B.4.

Results for the two sudden stops are summarized in Table II.[21] The decline in TFP in the 1992-93 crisis is entirely driven by incumbents. In fact, net entry contributes to positive growth, although the magnitude is small as the positive contribution of exiters narrowly outweighs the negative contribution of entrants. Among incumbents, the between-firm and cross terms jointly contribute positively to productivity through reallocation. Nevertheless, the pronounced decline in within-firm productivity dominates, generating a large negative overall effect.

In contrast, the increase in TFP experienced during 2010-13 is fully driven by net entry, in particular, by unproductive firms exiting the market. The size of the effect is remarkable, especially given that small and medium firms are underrepresented in the sample. Delving deeper into the characteristics of exiting firms shows that during the 2010-13 episode, firms that exit the market were, on average, bigger in terms of market share (7.44% versus 6.10%) and almost 50% more unproductive relative to incumbents (30.18% versus 20.77%) than their 1992-93 counterparts. Moreover, the annualized exit rate more than doubled from 4.63% to 9.82%.[22] To alleviate potential industry composition concerns, Figure A.11 shows this holds broadly at the three-digit industry level too. In sum, there is more and better exit. [23]

Back to Table II, the contribution of incumbents remains negative for the 2010-13 sudden stop. It is still the case that, on average, the productivity of incumbents is procyclical. There is a positive effect of the between and cross terms that cushions the negative effect of the within-firm component. The increase in market share reallocation and a stronger correlation between productivity and market share changes at the individual firm, together with the positive contribution of exiting firms, is consistent with a cleansing effect of the 2010-13 sudden stop which is absent in the 1992-93 episode. The cleansing hypothesis, as discussed by Caballero and Hammour (1994), argues that crises are periods of accelerated productivity-enhancing reallocations, especially as resources are freed by the exit of unproductive firms. I turn to formally testing the firm-level implications of this interpretation in what follows.[24]

[21] Online Appendix B.5 shows that results are robust to alternative decomposition approaches such as Foster, Haltiwanger and Krizan (2001) and Melitz and Polanec (2015) and the use of value-added or sales weights.

[22] The corresponding averages for other years are the following: the annualized exit rate is 7.78%, the employment share of exiting firms is 6.65% and the difference in TFP between exiting firms and incumbents is 9.09%.

[23] Firms entering the market during the 2010-13 episode are, on average, smaller in terms of market share (6.4% versus 11.57%) and slightly less unproductive relative to incumbents (1.65% versus 1.84%). This explains the negative contribution of entrants to aggregate TFP growth during both sudden stops, with a more pronounced effect in the earlier episode. Entry rates are 4% in 1992-93 and 5.55% in 2010-13. In normal times, entrants are more unproductive relative to incumbents (2.21%), and the entry rate is notably higher at 7.7%. Thus, my analysis confirms the main finding in Ates and Saffie (2021): both sudden stops are characterized by fewer but better entrants.

[24] A valid concern is that if firms are forward-looking, they might backload the decision to exit, and, thus, the duration of a crisis might be an important driver of results. I refer the reader to Online Appendix B.5, where I show that exit in the 2010-13 episode is not concentrated on the later years.

The Cleansing Hypothesis: A Formal Test

According to the literature, there is a tight connection between firm exit, input growth and productivity: models of firm dynamics predict that exit is more likely among low productivity firms whereas high productivity firms are expected to grow by more every period. The cleansing hypothesis suggests that recessions accelerate these dynamics. One should therefore observe a stronger correlation between survival, labor growth and productivity levels during crises. To test whether this is the case for the two sudden stop episodes considered, I adjust the empirical specification proposed by Foster, Grim and Haltiwanger (2016) and run the following set of regressions:

$$y_{i,t+1} = \beta \, TFP_{it} + \gamma \, ss^1_{t+1} * TFP_{it} + \theta \, ss^2_{t+1} * TFP_{it} + X'_{i,t} \, \omega + \epsilon_{i,t+1} \, ,$$

where $y_{i,t+1}$ stands for a set of dependent variables. In the exit specification, it is a dummy variable equal to one if a firm is active in period t but reports no activity in period $t+1$ strictly due to market exit as identified by the status indicator. In the regressions for input growth, it is a quantitative variable measuring labor growth between t and $t+1$. The regressor ss^1_{t+1} is a dummy variable for the 1992-93 sudden stop, ss^2_{t+1} is a dummy variable for the 2010-13 sudden stop and TFP_{it} captures the log of firm-level productivity. To abstract from underlying sector-specific trends, the above specification includes industry-year fixed effects at the three-digit level. In addition, $X_{i,t}$ controls for firm characteristics. For the baseline specification, I follow Foster, Grim and Haltiwanger (2016) in accounting for firm size and age effects.[25] However, the role of other firm characteristics is also explored in the section that follows.

For the exit specification, the relationship between survival probability and productivity is expected to be positive and, thus, $\beta < 0$. Under the cleansing hypothesis, this correlation should strengthen during a sudden stop episode and one would anticipate $\gamma < 0$ and $\theta < 0$. For the input growth specification, the exact opposite applies.

Results of these regressions are summarized in Table III. The first column shows the relationship between productivity and the probability of exit. Consistent with earlier findings, firms that exit the market tend to feature lower productivity levels. Focusing on the interaction terms, there is evidence of a cleansing effect only during the second episode. In terms of quantitative significance, the predicted difference in probability of exit between a firm one standard deviation below and a firm one standard deviation above average is 4.7 percentage points in normal times but almost 11.3 percentage points during the latter sudden stop.

The second and third columns support further the predictions of the cleansing hypothesis for the 2010-13 episode. First, note that there is a positive impact of productivity on labor growth as predicted by the literature. Of greater interest, this correlation is even higher during the second

[25] For firm size effects, I use a categorical variable: firm size class =1 if firm employment < 20; =2 if $20 \leq$ firm employment ≤ 50; =3 if $50 \leq$ firm employment ≤ 200; =4 if firm employment > 200. To control for the life cycle of the firm, I use a dummy variable $Young_{it}$ that equals one if the firms is five years old or younger.

TABLE III: REALLOCATION AND PRODUCTIVITY

	Exit (1)	Labor Growth (Incumbent & Exiters) (2)	Labor Growth (Incumbents Only) (3)
TFP_{it}	-0.038***	0.037***	0.033***
	(0.005)	(0.005)	(0.004)
$ss^1_{t+1} * TFP_{it}$	0.007	-0.002	-0.003
	(0.013)	(0.011)	(0.012)
$ss^2_{t+1} * TFP_{it}$	-0.052***	0.022*	0.016*
	(0.011)	(0.010)	(0.009)
Observations	36,252	32,262	29,679
Industry-Year FE	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes

Notes: Regression for exit is a linear probability model where exit=1 if a firm that operates in period t exits the market in period $t+1$. Labor growth is measured from period t to period $t+1$. TFP_{it} is the log firm-level TFP at time t, ss^1_{t+1} is a dummy equal to one for years 1992-93 and ss^2_{t+1} is a dummy equal to one for years 2010-13. Firm controls in period t account for age and size effects. Standard errors (in parentheses) are calculated using bootstrapping methods; $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

sudden stop. In fact, the predicted difference in labor growth between a firm one standard deviation above and a firm one standard deviation below average increases from 4.7 (4.1) percentage points in normal times to 7.5 (6.1) percentage points in 2010-13 according to coefficients reported in the second (third) column.

3.3 Alternative Explanations

Though so far the focus has been on the marked divergence in the exchange rate policies implemented during the two sudden stops, there are a number of additional dimensions along which the Spanish economy differed in 1992 versus 2010 that could also explain the contrast in firm dynamics documented in the previous section.[26] While it is unfeasible to fully rule out all alternative explanations, this section explores to what extent they might be driving the results. More specifically, I investigate the role of the banking crisis, expenditure switching effects of a real depreciation and the geographic scope of the crisis. Table IV augments the above empirical model for exit by adding relevant firm-level controls and interactions to test whether the coefficients of interest, especially θ, remain significant and stable when considering alternative explanations. To ease comparison, the first column of Table IV reiterates results for the baseline specification.

The most important difference across the two sudden stops, besides the exchange rate, is that

[26] As mentioned above, Online Appendix B.5 presents complementary evidence based on ORBIS, contrasting Spain and the Czech Republic at the same point in time. This exercise serves as a further check for readers who prefer a comparison across countries rather than across periods.

TABLE IV: FIRM EXIT AND PRODUCTIVITY WITH ADDITIONAL CONTROLS

	(1)	(2)	(3)	(4)	(5)
TFP_{it}	-0.038***	-0.033***	-0.037***	-0.038***	-0.032***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
$ss^1_{t+1} * TFP_{it}$	0.007	-0.002	0.007	0.007	-0.003
	(0.013)	(0.014)	(0.013)	(0.012)	(0.014)
$ss^2_{t+1} * TFP_{it}$	-0.052***	-0.040***	-0.051***	-0.050***	-0.038***
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
$leverage_{it}$		0.052***			0.052***
		(0.013)			(0.013)
$ss^1_{t+1} * leverage_{it}$		-0.008			-0.006
		(0.037)			(0.037)
$ss^2_{t+1} * leverage_{it}$		0.054*			0.055*
		(0.032)			(0.032)
$importer_{it}$			-0.012***		-0.012**
			(0.003)		(0.004)
$ss^1_{t+1} * importer_{it}$			0.002		0.012
			(0.013)		(0.014)
$ss^2_{t+1} * importer_{it}$			-0.010		0.005
			(0.009)		(0.010)
$exporter_{it}$				-0.001	0.002
				(0.003)	(0.004)
$ss^1_{t+1} * exporter_{it}$				-0.020**	-0.025*
				(0.009)	(0.012)
$ss^2_{t+1} * exporter_{it}$				-0.020**	-0.016*
				(0.007)	(0.009)
Observations	36,252	34,279	36,252	36,252	34,279
Industry-Year FE	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes

Notes: All regressions are linear probability models where exit=1 if a firm that operates in period t exits the market in period $t+1$. TFP_{it} is the log firm-level TFP at time t, ss^1_{t+1} is a dummy equal to one for years 1992-93 and ss^2_{t+1} is a dummy equal to one for years 2010-13. $leverage_{it}$ is captured by the bank debt-to-assets ratio. $importer_{it}$ is a dummy equal to one if the firm reports imports above 50,000 € in value. $exporter_{it}$ is a dummy equal to one if the firm reports exports above 50,000€ in value. Firm size classes in period t are used to control for firm size effects. Standard errors (in parentheses) are calculated using bootstrapping methods; $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Spain simultaneously experienced a banking crisis only during the latter. While intuition works in the opposite direction: highly leveraged firms, with a higher propensity to exit during a credit crunch, feature higher productivity levels on average; the second column of Table IV considers the role of leverage explicitly. In particular, the empirical model is augmented to account for the bank debt to assets ratio and the corresponding interactions. As expected, leverage is positively correlated with exit and this correlation is strengthened during the banking crisis, in line with findings in Bornstein and Castillo-Martinez (2023) for other European countries. However, the productivity results remain mostly unchanged. While the 2010-13 disruption of credit contributed to the overall increase in firm exit, it cannot fully explain the extend of the cleansing effect that followed.

Other well-known effects of a real exchange rate depreciation include (i) an expenditure switching effect on imported intermediate inputs and (ii) balance sheet effects resulting from liability currency mismatches. While in the absence of a model it is ex-ante unclear whether these effects should be different across episodes, most economists tend to expect a greater impact whenever the currency depreciates. This would involve more exit in the first episode, which does not hold in the data. Although the ESEE dataset does not provide information on debt denomination, the third column of Table IV provides some evidence on the role of imported intermediate inputs by featuring the import status of the firm. As theory predicts importers have a lower propensity to exit. Interestingly, this correlation remains unchanged during both sudden stops. As in the previous column, the productivity coefficients remain significant and stable.

Next, I explore the role of export status—not only for completeness but also as a way to account for the geographic scope of each sudden stop. A global crisis implies a larger contraction of foreign demand, which could increase firm exit rates relative to a more localized or regional crisis.[27] As extensively documented in the literature, exporters are, on average, significantly more productive (138.15%) and, consequently, less likely to exit (19.05%) than non-exporters.[28] However, once I control for firm productivity, the primary channel through which exporter status influences exit likelihood is foreign demand. The extent to which a crisis is global determines how insulated exporters are from local downturns relative to non-exporters, thereby shaping their lower propensity to exit.

The fourth column of Table IV confirms that being an exporter reduces the probability of exit during a sudden stop, all else equal. This suggests that foreign demand remains relatively stronger than domestic demand during such episodes and exporters are better able to substitute between the two, in line with results by Almunia et al. (2021). More importantly, the effect's magnitude is similar across both sudden stops, indicating that the underlying shocks had comparable impact on foreign demand.

Online Appendix B.2 discusses additional contrasts between the two episodes that may concern the reader, such as long-run shifts in the manufacturing sector, the aftermath of the 2009 construction bust, the presence or absence of European-wide policy support, and the evolution—both in level and composition—of Spain's net foreign asset position.[29]

Finally, the literature has suggested market power as another firm-level characteristic that may

[27] The 2010–13 sudden stop clearly occurred during a global crisis. However, the first episode was also part of a broader economic upheaval: the Exchange Rate Mechanism crisis triggered a recession across the continent. For context, annual GDP growth for the Euro area during 1992–93 was only 0.57%, compared to 1.09% during the 2010–13 period.

[28] During both episodes, firm-level data show an increase in the propensity to export (9.51% in the first episode; 17.65% in the second) and a decline in the average productivity of exporters (23.31% in the first episode; 4.45% in the second). This pattern aligns with the model's prediction of a lower productivity threshold for exporting during sudden stops.

[29] A popular complementary channel through which reallocation contributes to productivity growth is increased allocative efficiency. In fact, Gopinath et al. (2017) and García-Santana et al. (2020) have shown that increasing capital misallocation is responsible for the slowdown of productivity growth prior to the 2010-13 crisis. Figure A.10 explores whether a reversal in factor misallocation underlies the recent improvement in aggregate TFP. The evidence suggests that this is not the case. That said, accurately measuring dispersion in marginal revenue products requires comprehensive coverage of small firms, which the ESEE dataset underrepresents, even when using sampling weights.

shape exit behavior. Table A.11 shows that firms charging higher markups, estimated following De Loecker and Warzynski (2012), are significantly less likely to exit. This negative correlation mirrors the patterns documented for productivity and remains robust after controlling for other firm attributes. Notably, the relationship is stronger during the second sudden stop. However, the effect does not survive once productivity is included as a control. This suggests that while there is meaningful heterogeneity in markups across firms, it is closely tied to productivity differences; at least in terms of how these characteristics influence the propensity to exit.

As a brief summary, the above findings call for a theory of sudden stops that features heterogeneously productive firms, selection into production and variable firm-specific markups. All of these elements, together with the exchange rate dimension, are featured in the theoretical framework that I develop next.

4 A Small Open Economy with Firm Dynamics

Consider an infinite-horizon small open economy. Time is discrete and indexed by t. The economy is populated by a continuum of households, $i \in (0,1)$, that consume goods, provide specialized individual labor types and engage in financial transactions with foreign investors. A large number of firms, domestic and foreign, indexed by j, produce differentiated consumption goods using labor. Finally, a monetary authority sets the nominal exchange rate as the policy instrument.

4.1 Consumers

Households derive utility from leisure and the consumption of differentiated varieties. Labor decisions are outsourced to type-specific labor unions as discussed below. Under the assumption of complete contingent claims markets for consumption, consumption and saving decisions are identical across households. Thus, I drop the household specific index i and, thus, consider a representative consumer for the rest of this section.[30] Its lifetime utility is given by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t U\left(q_{jt}, L_t\right) \right], \tag{1}$$

where $\mathbb{E}_t$ is the expectation operator conditional on the information set available at time t, β is the discount factor, q_{jt} is the consumption level of variety j, and L_t is labor supplied by the household.

[30] Note that this is a slight abuse of notation. There is technically not a representative household in this model due to labor specialization. However, all households face the same budget, choose the same consumption and, thus, have the same marginal utility of income.

The period utility function is based on Melitz and Ottaviano (2008):

$$U\left(q_{jt}, L_t\right) = \alpha \int_0^{N_t} q_{jt}\, dj - \frac{1}{2}\gamma \int_0^{N_t} q_{jt}^2 dj - \frac{1}{2}\eta \left(\int_0^{N_t} q_{jt} dj\right)^2 - L_t\,,$$

where N_t is the measure of available consumption varieties, γ determines the level of product differentiation and α and η jointly capture the substitutability between the consumption of differentiated goods and leisure. All three demand parameters are strictly positive. [31,32]

Households can only engage in financial transactions with foreign investors by trading in risk-free bonds, B_t, denominated in units of foreign currency, which pay a gross nominal interest rate $R_t = R \exp\left(\varepsilon_t^R\right)$, where ϵ_t^R is an interest rate shock around a mean value $R = \beta^{-1}$. [33]

The budget constraint is given by:

$$\int_0^{N_t} p_{jt} q_{jt}\, dj + e_t B_t = W_t L_t + \Pi_t + e_t R_{t-1} B_{t-1}\,, \tag{2}$$

where p_{jt} is the price of variety j, W_t is the nominal wage, and Π_t is profit received from firms, all expressed in units of domestic currency. [34] In addition, e_t denotes the nominal exchange rate, defined as units of domestic currency per unit of foreign currency.

Each period, the representative household maximizes (1) by choice of q_{jt} and B_t, subject to (2).

4.2 Labor Market

As is standard in sticky wage New Keynesian models, a central authority competitively aggregates units of differentiated labor input, L_t^i, into aggregate employment services, L_t, using a CES technology, $L_t = \left(\int L_t^{i\,\frac{\varepsilon_w - 1}{\varepsilon_w}} di\right)^{\frac{\varepsilon_w}{\varepsilon_w - 1}}$, where ε_w is the elasticity of substitution across labor types. Firms can purchase aggregate employment services at a rate W_t, where $W_t = \left(\int W_t^{i\,1-\varepsilon_w} di\right)^{\frac{1}{1-\varepsilon_w}}$. Nominal wages, W_t^i, are set by unions, each of which represents households specialized in a given

[31] Melitz and Ottaviano (2008) preferences are appealing for three reasons. First, they capture love of variety. As γ increases, consumers place higher weight on the distribution of consumption across varieties. Second, the quadratic form gives rise to a linear demand function which ensures the existence of a choke price and an extensive margin of production even in the absence of fixed costs of production. Third, they generate endogenous variable markups, which capture the effect of market competition on firm sales (the so-called pro-competitive effect) as opposed to standard CES preferences.

[32] Melitz and Ottaviano (2008) also consider a second consumption good, which is homogeneous, non-tradable, and serves only to close the model. I drop non-tradable consumption and instead feature leisure in the utility function. This follows the convention in small open economy models to account for labor supply decisions, and enables the introduction of nominal rigidities in the wage-setting process later on. Note that while linearity in leisure simplifies the algebra, it is not essential to the analysis.

[33] Households are not allowed to trade in domestic bonds in the baseline model for the sake of simplicity. However, extending the model to include domestic bonds would be trivial as these would be in zero net supply.

[34] All prices and costs in the model are expressed in units of domestic currency. Unlike in the CES case, total consumption under Melitz-Ottaviano preferences is non-homogeneous, making it unfeasible to derive a price index that can be used to deflate nominal variables as it is standard.

type of labor. Wage setting is staggered à la Calvo (1983): a union is able to reset its wage with a probability $1-\theta$ every period. [35]

All unions that reset their wage in a given period choose the same wage since they face an identical problem. As a result, I drop the index i and simply distinguish wages based on the time of their last adjustment. In particular, unions choose W_t^* to maximize $\mathbb{E}_t\left[\sum_{k=0}^{\infty}(\beta\theta)^k U\left(q_{j\,t+k|t}, L_{t+k|t}\right)\right]$, where $q_{j\,t+k|t}$ and $L_{t+k|t}$ denote consumption and labor at time $t+k$ for a household that last updated the wage at time t. This is subject to the sequence of household budget constraints that are effective while W_t^* remains in place, as well as the corresponding labor demand schedules, as given by: $L_{t+k|t} = \left(\frac{W_t^*}{W_{t+k}}\right)^{-\varepsilon_w} L_{t+k}$.

Following this wage setting structure, the evolution of the aggregate nominal wage is given by

$$W_t = \left[\theta W_{t-1}^{1-\varepsilon_w} + (1-\theta)(W_t^*)^{1-\varepsilon_w}\right]^{\frac{1}{1-\varepsilon_w}} . \tag{3}$$

4.3 Domestic Firms

There is an endogenous measure, M_t, of domestic firms producing the differentiated varieties. Labor is the only factor of production and the unit nominal cost is a concave function in the factor price, i.e., $\frac{W_t^{\sigma}}{Z_{jt}}$ where $0 < \sigma \leq 1$ is the labor income share.[36] Firm productivity $Z_{jt} = Z_t^A z_j^P z_{jt}^T$ is the product of a common productivity shifter, Z_t^A, an idiosyncratic permanent component, z_j^P, and an idiosyncratic transitory effect, z_{jt}^T. Permanent idiosyncratic productivity is drawn from a distribution $F(z^P)$ while transitory idiosyncratic effects, z_t^T, follow a Markov process with transition function $H(z_{t+1}^T|z_t^T)$. The common productivity shifter evolves according to $Z_t^A = exp(\epsilon_t^A)$, where ϵ_t^A is a productivity shock with mean zero.

The timing of events for an incumbent firm at period t is as follows. At the start of the period, aggregate uncertainty is resolved. The firm observes its stochastic nominal value of exiting, x_{jt}, which is positive and drawn from a common time-invariant distribution $\xi(x_t)$ and decides whether to keep operating. If it exits, it receives the exit value. If it decides to stay, its idiosyncratic productivity is revealed and the firm then sets prices and chooses labor. Production takes place and the period concludes.

Every period there is an endogenous mass of potential entrants, M_t^s. Entry is a two-step procedure and takes place once aggregate uncertainty is realized. First, a potential entrant comes up with an idea, which involves paying a fixed cost $c_s > 0$. The quality of the idea, measured by $s_j \sim S(s)$, is informative of the entrant's productivity. Conditional on entry, the distribution of the

[35] The model could alternatively feature Rotemberg wage adjustment costs, sticky information à la Mankiw and Reis (2002) or downward wage rigidities, as long as the latter are incorporated in the form of a constraint on the wage setting problem of the worker (or of the union that represents them).

[36] To rationalize this functional form, suppose there is a second factor of production, which is inelastically supplied by households and the production function is Cobb-Douglas. If the price of this second input, κ, is assumed to be constant, the marginal cost is given by $\frac{1}{z}\left(\frac{W_t}{\sigma}\right)^{\sigma}\left(\frac{\kappa}{1-\sigma}\right)^{1-\sigma}$.

idiosyncratic shocks in the first period of operation is $H^s(z_t|s_t)$, strictly decreasing in s. Potential entrants that decide to implement their idea must pay an additional cost $c_e > 0$ to enter. Both fixed costs are nominal and expressed in units of domestic currency. The measure of actual entrants is denoted by $\mathcal{E}_t$.

Firms with the same level of productivity make the same decisions. Thus, in what follows, I index firms (and varieties) by their idiosyncratic productivity, $z_t = z^P z_t^T$, rather than by their index j.

Incumbents' static problem Domestic firms can sell their varieties in both the domestic and the export market. Markets are segmented and selling abroad requires incurring a per-unit trade cost $\tau > 1$.

While domestic demand for a differentiated good, $q_t^H(z_t)$, is given by the household's optimization problem above, I assume the foreign demand for a domestic variety is: $q_t^X(z_t) = A - Bp_t^X(z_t)$, where A and B are exogenous given a small-open economy setting. Domestic firms set prices for goods sold domestically, $p_t^H(z_t)$, and abroad, $p_t^X(z_t)$, to maximize flow profits given by

$$\pi_t(z_t) = p_t^H(z_t)\, q_t^H(z_t) + e_t\, p_t^X(z_t)\, q_t^X(z_t) - \frac{W_t^\sigma}{Z_t z_t}\left(q_t^H(z_t) + \tau q_t^X(z_t)\right) .$$

Incumbents' recursive problem At all $t \geq 0$, the distribution of operating firms is denoted by $\Gamma_t(z_t)$, such that $\int d\Gamma_t(z_t) = M_t$. Let $\mu_t \in \mu$ denote the set of aggregate state variables and $J(\mu_{t+1}|\mu_t)$ its transition operator. Note that $\mu_t = \{W_{t-1}, e_t, \Gamma_t, R_t, Z_t\}$.

Given the timing of events, I refer to $z^{T'}$ as the current realization of the transitory productivity component. Then, $z = z^P z^{T'}$ denotes the idiosyncratic productivity of a firm operating in the current period.

The start-of-period value of an incumbent firm is given by $V(\mu, z^P, z^T)$ which solves

$$V(\mu, z^P, z^T) = \int \max\left\{E_z\left[V^c(\mu, z^P, z^{T'})\middle|z^T\right], x\right\} d\xi(x)\,,$$

where $E_z\left[V^c(\mu, z^P, z^{T'})\middle|z^T\right]$ is the expected value of a continuing firm

$$E_z\left[V^c(\mu, z^P, z^{T'})\middle|z^T\right] = \int V^c(\mu, z^P, z^{T'})\; dH\left(z^{T'}|z^T\right),$$

and

$$V^c(\mu, z^P, z^{T'}) = \begin{cases} \pi(\mu, z^P, z^{T'}) + \beta \int V(\mu', z^P, z^{T'})\; dJ\left(\mu'|\mu\right) & \text{if } \pi(\mu, z^P, z^{T'}) > 0 \\ 0 & \text{otherwise} \end{cases}$$

In other words, the firm will choose to exit if its expected flow of profits is zero or if the exit value is high enough.

Entry decision For an aggregate state μ, the value of a new firm with permanent idiosyncratic component z^P that enters after receiving signal s is

$$V^e(\mu, z^P, s) = \int V^c(\mu, z^P, z^{T\prime})\; dH^s\;(z^{T\prime}|s)\,.$$

It follows that only potential entrants with sufficiently high values of $s(z^P) \geq s^*(z^P)$ find it profitable to enter, where $s^*(z^P)$ is determined by $V^e(\mu, z^P, s^*(z^P)) = c_e$. Finally, the free-entry condition ensures that potential gains from entry are exhausted. Potential entrants keep generating ideas right until $V^p(\mu) = c_s$, where the value function of a potential entrant is given by

$$V^p(\mu) = \int\int \max\left\{V^e(\mu, z^P, s) - c_e,\, 0\right\}\, dS(s)\, dF(z^P)\,.$$

4.4 Foreign Importers

Foreign firms might choose to serve the domestic market. While the distribution of operating foreign firms, $\Gamma^*(z)$, is exogenous given the small open economy assumption, the mass of importers, M_t^M, is endogenous. Similarly to the domestic pricing problem above, flow profits for a foreign firm, indexed by z_t, are given by

$$\pi_t^M(z_t) = p_t^M(z_t)\, q_t^M(z_t) - \frac{e_t \tau}{z_t} q_t^M(z_t)\,,$$

where $p_t^M(z_t)$ is the price set by the foreign importer and the demand for imports, $q_t^M(z_t)$, follows from the household's optimization problem above. A foreign firm will select into importing provided that $\pi_t^M(z_t) > 0$.

4.5 Exchange Rate Policy

To pin down the nominal variables of the model, I need to determine exchange rate policy. Suppose the central bank implements monetary policy by setting the nominal exchange rate according to the following rule:

$$(\Pi_t^w)^{\phi_W}\,(e_t)^{1-\phi_W} = 1\,, \tag{4}$$

where $\Pi_t^w = \frac{W_t}{W_{t-1}}$ is nominal wage inflation and $0 \leq \phi_w \leq 1$ is the weight that the central banks puts on wage stabilization. I explore two extreme versions of this rule. Under a currency union, the central bank perfectly commits to a currency peg in which $e_t = 1$ at every period t, i.e. , $\phi_W = 0$. Under a flexible exchange rate arrangement, the central bank offsets all the distortions that originate from nominal rigidities by implementing the flexible wage equilibrium, i.e. , $\phi_W = 1$.

4.6 Stationary Competitive Equilibrium

I now define the recursive stationary competitive equilibrium of this economy. The law of motion for the measure of operating (domestic) firms $\Gamma'(z^P, z^{T\prime})$ is defined as follows. For all Borel sets $\mathcal{Z}^P \times \mathcal{Z}^T \in \Re^+ \times \Re^+$,

$$\Gamma'(\mathcal{Z}^P \times \mathcal{Z}^T) = \int \int_{B(\mu,x)} \int_{\mathcal{Z}^T} dH(z^{T'}|z^T)\, d\xi(x)\, d\Gamma(z^P, z^T) + \mathcal{E}'(\mathcal{Z}^P \times \mathcal{Z}^T)\,,$$

where $B(\mu, x) = \left\{x \text{ s.t. } E_z\left[V^c(\mu, z^P, z^{T\prime})\Big|z^T\right] \geq x\right\}$ and $\mathcal{E}'(\mathcal{Z}^P \times \mathcal{Z}^T)$ is the measure of actual entrants, such that, for all Borel sets $\mathcal{Z} \in \Re^+$,

$$\mathcal{E}'(\mathcal{Z}^P \times \mathcal{Z}^T) = M^s \int_{s^*}^{\infty} \int_{B_p(\mu, z^P, z^{T\prime})} F(z^P)\, dH^s(z^{T\prime}|s)\, dS(s)\,,$$

where $B_p(\mu, z^P, z^{T\prime}) = \{z^{T\prime} \in \mathcal{Z}^T \text{ s.t. } V^c(\mu, z^P, z^{T\prime}) > 0\}$.

Clearing in the goods market requires:

$$e\,(1-R)\,B = e \int p^X(\mu,z)\, q^X(\mu,z)\ d\Gamma'(z) - \int p^M(\mu,z)\, q^M(\mu,z)\ d\Gamma^*(z) - (M^s c_s + M^e c_e - X)\ , \tag{5}$$

where $M^e = \int d\mathcal{E}'$ is the number of actual entrants and X is the total value of exiting given by $X = \int\int \mathbf{1}\{E_z\left[V^c(\mu, z^P, z^{T\prime})\Big|z^T\right] < x\}\, x\, d\xi(x)\, d\Gamma'(z)$. The left-hand side is the net foreign asset position of the economy while the right-hand side is the trade balance after accounting for net entry costs.

Finally, note that the exchange rate rule above, equation (4), implies that $e = 1$ in the stationary equilibrium. The definition of the recursive stationary competitive equilibrium is as follows.

Definition 1. *A* ***Recursive Stationary Competitive Equilibrium*** *consists of value functions* $V(\mu, z)$, $V^c(\mu, z)$, $V^e(\mu, s)$, *and* $V^p(\mu)$; *firm policy functions* $\pi^k(\mu, z)$ *and* $p^k(\mu, z)$ *for* $k = \{H, X, M\}$; *consumer's demand schedule* $q(z)$; *a vector of aggregate prices* $\{W^*, W\}$; *foreign bonds* B; *the number of varieties* N; *the number of potential entrants* M^s, *and the distribution of operating (domestic) firms* $\Gamma'(z)$, *such that:*

1. $V(\mu, z)$, $V^c(\mu, z)$, $\pi^H(\mu, z)$, $\pi^X(\mu, z)$, $p^H(\mu, z)$ *and* $p^X(\mu, z)$ *solve the incumbent's problem*
2. $V^e(\mu, s)$, *and* $V^p(\mu)$ *solve the entrant's problem*
3. $\pi^M(\mu, z)$ *and* $p^M(\mu, z)$ *solve the importer's problem*
4. $q(z)$ *and* B *solve the household's problem*
5. W^* *solves the union's problem*
6. W *follows equation* (3)

7. $N = \int \mathbf{1}\{\pi^H(\mu, z) > 0\}\ d\Gamma'(z) + \int \mathbf{1}\{\pi^M(\mu, z) > 0\}\ d\Gamma^*(z)$

8. *Free entry condition holds:* $V^p(\mu) = c_s$

9. *Goods market clears - i.e., equation 5 holds*

10. *The distribution of operating (domestic) firms is stationary*

4.7 Characterizing the Equilibrium

Here I show the properties of the representative household, labor unions and firms' optimality conditions and provide some intuition.

Given quadratic preferences, domestic consumption is governed by the existence of a choke price, p_t^{max}, i.e., the maximum price at which the household is willing to consume any differentiated variety. If $p_{jt} > p_t^{max}$, optimal consumption of variety j falls to zero. If instead $p_{jt} \leq p_t^{max}$, optimal consumption of variety j is given by the corresponding first order condition:

$$\alpha - \gamma q_{jt} - \eta \int_0^{N_t} q_{jt} dj = \lambda_t p_{jt}\,,$$

where λ_t is the time t Lagrange multiplier on the budget constraint, equation (2). The demand for variety j can be written compactly as

$$q_{jt} = \max\left\{\frac{\lambda_t}{\gamma}\left[p_t^{max} - p_{jt}\right], 0\right\}\,, \tag{6}$$

where $p_t^{max} \equiv \frac{\frac{\alpha\gamma}{\lambda_t} + \eta \int_0^{N_t} p_{jt} dj}{\gamma + \eta N_t}$. As in Melitz and Ottaviano (2008), a larger number of competing varieties, induce a decrease in the choke price. In addition, higher income, lower λ_t, increases the choke price.

The optimal decision for the purchase of the foreign asset, B_t, delivers a standard Euler equation:

$$\lambda_t = \beta R_t E_t\left[\frac{e_{t+1}}{e_t}\lambda_{t+1}\right]\,. \tag{7}$$

A higher interest rate and expectations of nominal exchange rate depreciation both increase the cost of borrowing internationally and, thus, encourage consumer savings.

The optimality condition associated with the union's problem can be written as

$$\sum_{k=0}^{\infty}(\beta\theta)^k \mathbb{E}_t\left[\lambda_{t+k} L_{t+k}\left(\frac{W_t^*}{W_{t+k}}\right)^{-\varepsilon_w}\left(W_t^* - \frac{\varepsilon_w}{\varepsilon_w - 1}\frac{1}{\lambda_{t+k}}\right)\right] = 0\,. \tag{8}$$

In the limiting case of full wage flexibility, $\theta = 0$, $W_t = W_t^* = \frac{\varepsilon_w}{\varepsilon_w - 1}\frac{1}{\lambda_t}$. Higher wages increase household income all else equal. Given diminishing marginal utility, the Lagrange multiplier, i.e.

, the increase in utility of an extra dollar of income, falls. With staggered wages, the goal is to close the discounted sum of expected wedges between actual and desired (flexible) wages.

Finally, optimal price setting by domestic firms and importers delivers expressions for prices and output such that:

$$
\begin{aligned}
p_t^H(z_t) &= \min\left\{\frac{1}{2}\left[p_t^{max} + \frac{W_t^\sigma}{Z_t z_t}\right], p_t^{max}\right\},\ q_t^H(z_t) = \max\left\{\frac{\lambda_t}{2\gamma}\left[p_t^{max} - \frac{W_t^\sigma}{Z_t z_t}\right], 0\right\}, \\
p_t^X(z_t) &= \min\left\{\frac{1}{2}\left[\frac{A}{B} + \frac{\tau}{e_t}\frac{W_t^\sigma}{Z_t z_t}\right], \frac{A}{B}\right\},\ q_t^X(z_t) = \max\left\{\frac{B}{2}\left[\frac{A}{B} - \frac{\tau}{e_t}\frac{W_t^\sigma}{Z_t z_t}\right], 0\right\}, \\
p_t^M(z_t) &= \min\left\{\frac{1}{2}\left[p_t^{max} + \frac{e_t\tau}{z_t}\right], p_t^{max}\right\},\ q_t^M(z_t) = \max\left\{\frac{\lambda_t}{2\gamma}\left[p_t^{max} - \frac{e_t\tau}{z_t}\right], 0\right\}.
\end{aligned}
$$

5 Sudden Stops and Productivity in a Simple Example

Before proceeding to the full characterization of the model's solution, it is useful to build intuition of the mechanism by providing some analytical results. In order to do this, I focus on a simpler version of the model with (i) constant idiosyncratic firm productivity, (ii) no free entry of firms and (iii) zero exit value at all times. There is an exogenous number of potentially active firms at home, M, and abroad, M^*, with productivity z that follows a Pareto distribution $1-\Gamma(z) = \left(\frac{1}{z}\right)^k$ with shape parameter k and minimum level equal to one. I also simplify the wage dynamics of the model by assuming that, after a shock, only a fraction $1-\theta$ of labor types adjust their wages immediately, while the rest remain fixed until the next period.

In the following, I define aggregate productivity and discuss the channels through which shocks can potentially affect it. I then show how the effects of a sudden stop on productivity depend on the exchange rate regime. All proofs are provided in the Appendix. I provide empirical support for one of these channels using firm-level data.

5.1 Aggregate Productivity

Given the new set of assumptions above, the domestic firm's optimization problem is static. Firms' decision to operate can be written in terms of a productivity threshold: a firm will choose to serve market i provided that $z \geq z_t^i$ where $i \sim \{H, X, M\}$. Thus, the variable of interest, domestic aggregate productivity, is given by:

$$
Z_t^H = M\int_{z_t^H}^{\infty} \Omega(z)\, z\, Z_t\, \frac{\gamma(z)}{1-\Gamma(z_t^H)}\ dz\,,
$$

where $\Omega(z)$ is the weight used in the aggregation and $\gamma(z)$ is the probability distribution function of the Pareto distribution. I normalize $\Omega(z)$ so that weights sum to one over the set of firms that operate domestically (i.e., $z \geq z_t^H$).

The empirical analysis in Section 3 uses labor weights, $\Omega(z) = \ell(z)/L_t^H$, whereas alternatives include equal weights $\Omega(z) = 1/\big(M(1-\Gamma(z_t^H))\big)$, output weights $\Omega(z) = q(z)/Q_t^H$, and revenue weights $\Omega(z) = r(z)/R_t^H$. [37] The following Lemma establishes that $z_t^H Z_t$ is the key statistic for measuring aggregate productivity independent of the weights used in the aggregation. [38]

Lemma 1. *Domestic aggregate productivity, Z_t^H, is an increasing function of the domestic productivity threshold, z_t^H and the common productivity shifter, Z_t.*

In other words, changes in productivity in this model are partly governed by firms' extensive margin decisions. This is in contrast to alternatives in the literature that either model productivity exclusively as an exogenous shock to the economy, allow for variable capacity utilization or consider R&D decisions.

5.2 Pro-competitive, Cost and Demand Channels

In the absence of shocks to the common productivity shifter, $Z_t = 1$, the productivity threshold is determined by the number of firms in the market, the cost of production and the level of consumer demand; all three are potentially subject to change during a sudden stop episode.

Proposition 1. *In the stationary equilibrium, i.e. , $Z_t = 1$:*

$$dz_t^H = \underbrace{F_N\, dN_t}_{\textit{Pro-competitive}} + \underbrace{F_W\, dW_t}_{\textit{Cost}} + \underbrace{F_\lambda\, d\lambda_t}_{\textit{Demand}},$$

where $F_N > 0$, $F_W > 0$ and $F_\lambda > 0$.

The intuition follows next. In the first place, a larger number of active firms in the market, $dN_t > 0$, implies greater competition. Given the preferences considered, enhanced competition lowers individual firm demand. This forces less productive firms out of the market as profit margins shrink across the productivity distribution. This *pro-competitive effect* was first introduced by Melitz and Ottaviano (2008), with a focus on competition in the goods market.

Second, a higher wage, $dW_t > 0$, reduces firms' profit margins by raising production costs. As a result, only more productive firms remain profitable and select into production, leading to an increase in aggregate productivity. This is what I denote the *cost effect*, which is the underlying mechanism in the canonical Melitz (2003) model, which focuses on competition in the labor market.

[37] $Q_t^H = M\int_{z_t^H}^{\infty} q_t^H(z)\,\gamma(z)\,dz$ is total domestic output, $L_t^H = M\int_{z_t^H}^{\infty} l_t^H(z)\,\gamma(z)\,dz$ is total employment in domestic production, and $R_t^H = M\int_{z_t^H}^{\infty} p_t^H(z) q_t^H(z)\,\gamma(z)\,dz$ is total domestic revenue.

[38] The reader may wonder how this measure relates to other definitions of aggregate productivity. First, the (model-consistent) Solow residual commonly reported by statistical agencies maps directly to the employment-weighted average. Second, under non-homothetic preferences the welfare-relevant measure of productivity is not uniquely defined; in Melitz and Ottaviano (2008), for example, welfare statements are tied to the unweighted average of firm productivities.

Finally, higher aggregate demand from consumers, $d\lambda_t < 0$, raises individual firm demand across the productivity distribution and relaxes the productivity threshold for profitability. As a result, less productive firms are more likely to enter or remain in the market. This final channel, a novelty of this model, is referred to as the *demand effect*.

5.3 Effects of a Sudden Stop

I define a sudden stop as a one-time unexpected increase in the interest rate R_t, while keeping the common productivity shifter constant, i.e., $Z_t = 1$ for all t. The shock is a purely transitory MIT deleveraging shock that forces an expansion of net exports and an improvement in international competitiveness. The following proposition considers its effect on productivity under the two alternative exchange rate regimes.

Proposition 2. *Given a sudden stop,*

1. *In a floating arrangement, only the pro-competitive channel operates and productivity falls:*

$$dN_t < 0,\ dW_t = 0 \ \text{ and } \ d\lambda_t = 0 \ \text{ so that } \ dz_t^H < 0\,.$$

2. *In a currency union, all three channels operate and the effect on productivity is ambiguous:*

$$dN_t < 0,\ dW_t < 0 \ \text{ and } \ d\lambda_t > 0 \ \text{ so that } \ dz_t^H \gtrless 0\,.$$

A higher interest rate must be matched by an expected appreciation of the domestic currency, stronger future demand growth, or both. Holding future variables constant, this implies a depreciation of the currency today (higher e_t) and/or a contraction in current demand (higher λ_t). Suppose, for now, that the nominal exchange rate depreciates just enough to fully offset the increase in the interest rate. Under this assumption, both the cost and demand effects are muted, as the marginal utility of income, and thus the wage level, through the union's optimality condition, remains unchanged. However, there is a decline in the number of active firms in the domestic economy, driven by a fall in the number of importers. The depreciation makes foreign firms less competitive, shifting expenditure toward domestic varieties. This negative pro-competitive force leads to an unambiguous decline in aggregate productivity.

Now suppose instead that the nominal exchange rate remains unchanged, so the adjustment occurs entirely through an increase in λ_t. This constitutes a negative demand effect. At the same time, wages fall following the optimal wage-setting condition, reducing production costs and triggering a negative cost effect. The pro-competitive channel remains negative, following the logic above. All three channels are active but exert opposing forces on the productivity threshold: the demand effect raises it, while the cost and pro-competitive effects push it downward. As a result, the overall impact of a sudden stop on productivity in a currency union is ambiguous and depends on parameter values. Nevertheless, it is possible to characterize conditions under which the demand effect dominates, resulting in an increase in aggregate productivity.

Corollary 1. *For a small enough ϵ_t^R, a sufficient condition for $dz_t^H > 0$ is that*

$$(1-\theta)\sigma(k+1) < \left(\frac{W_t}{W_t^*}\right)^{1-\epsilon_w} \approx 1 \,.$$

There are three key parameters for this condition to hold: the share of labor income, σ, the degree of wage rigidities, θ, and the shape parameter of the productivity distribution, k. The share of labor income governs the mapping between the wage level and the unit cost. As σ increases, labor represents a greater share of the optimal input bundle and falling wages cheapen production costs by more. This reinforces the cost effect of a sudden stop. In the Melitz (2003) model, the cost channel is at its strongest featuring a production function which is linear in labor, $\sigma = 1$.

The degree of wage rigidities determines the extent to which changes in the marginal utility of income translate into wage adjustments. As wage stickiness increases, a decline in income results in a smaller reduction in aggregate wages. Consequently, the absolute ratio of the demand effect to the cost effect rises, amplifying the relative strength of the demand channel in response to a sudden stop.

The shape parameter measures the concentration of firms at the lower end of the productivity distribution. This represents the inverse of dispersion in firm-level productivity. As firms only differ in their productivity levels, if k increases, they become more homogeneous and, thus, more reliant on their relative cost advantage to survive. This implies that changes in the economy's international competitiveness will lead to larger swings in the number of importers, thus, increasing the size of the pro-competitive effect.

5.4 Demand Channel in the Data

Given the central role of the demand channel in shaping the productivity response to a sudden stop, I carry out a simple validation exercise. I return to the firm-level data from Section 3, focus on the 2010-13 sudden stop exclusively and use sectoral variation in tradability to assess whether a stronger demand channel is associated with lower aggregate productivity. The intuition is the following: industries more reliant on domestic consumers are more exposed to the demand contraction resulting from a sudden stop under a currency union. As a result, one should observe a larger increase in the productivity threshold together with more cleansing in less tradable industries.

I adopt the tradability index from Mian and Sufi (2014), which classifies industries based on geographic concentration under the idea that industries dependent on national demand (more tradable) tend to be more spatially concentrated, while those serving local markets (less tradable) are more uniformly distributed. It is key to distinguish between traded goods, whose classification

TABLE V: 2010-13 PRODUCTIVITY GROWTH DECOMPOSITION

	All Sectors	High Tradability Sectors	Low Tradability Sectors
Productivity Growth (%)	2.24	0.18	2.41
Contribution to Productivity Growth			
Incumbents' Contribution	**-4.54**	**-3.73**	**-6.07**
Within-firm Contribution	-7.36	-4.10	-9.80
Between-firm Contribution	0.32	-1.76	1.59
Cross-term Contribution	2.50	2.13	2.15
Net Entry Contribution	**6.79**	**3.91**	**8.47**
Entrants' Contribution	-0.40	0.76	-0.59
Exiters' Contribution	7.19	3.14	9.06

Notes: Productivity growth refers to accumulated TFP growth for the 2010-13 sudden stop. Columns (2) corresponds to the baseline decomposition as explained in the main text. Columns (3) restricts the sample to sectors with a tradability index over the median while Column(4) restricts the sample to sectors with a tradability index below the median. The data used is collected from the ESEE dataset.

is endogenous in this model and can shift in response to a sudden stop, and tradability, which I take as a structural characteristic that remains stable over time and across countries.

Using this classification, I conduct two exercises to test whether, indeed, less tradable sectors feature a larger increase in selection at exit in 2010-13. First, I repeat the TFP decomposition for industries with high and low tradability. Results are summarized in Table V. The first column repeats the baseline decomposition for 2010–13 to ease comparison. The key patterns that emerge are as follows: firm-level productivity declines across all three samples, and while the magnitude of the incumbents' contribution varies, it is consistently negative. This implies that observed differences in overall productivity growth are primarily driven by net entry dynamics—specifically, the contribution of exiting firms. As predicted by the theory, exit contributes most positively to productivity growth in low-tradability sectors and least in high-tradability sectors.

Second, I reestimate the empirical specification for cleansing effects of a sudden stop separately for each industry and visualize the results in Figure III. In particular, I plot the difference in the exit propensity gap between a high- and low-productivity firm during the 2010–13 sudden stop relative to normal times against my measure of sectoral tradability. The fitted line (in red) shows a negative relationship, with a slope of -0.8, confirming that sectors producing less tradable goods experience stronger selection at exit.

The analysis above holds the common productivity shifter constant, making the framework well-suited to capturing the forces at play in a currency union. However, it cannot fully account for productivity dynamics under a floating regime, where the intensive margin is empirically more relevant. I now turn to a fully-fledged quantitative model that reintroduces shocks to Z_t and relaxes the earlier simplifications, in order to quantify how much of the observed difference in the

FIGURE III: CHANGES IN EXIT SELECTION BY INDUSTRY

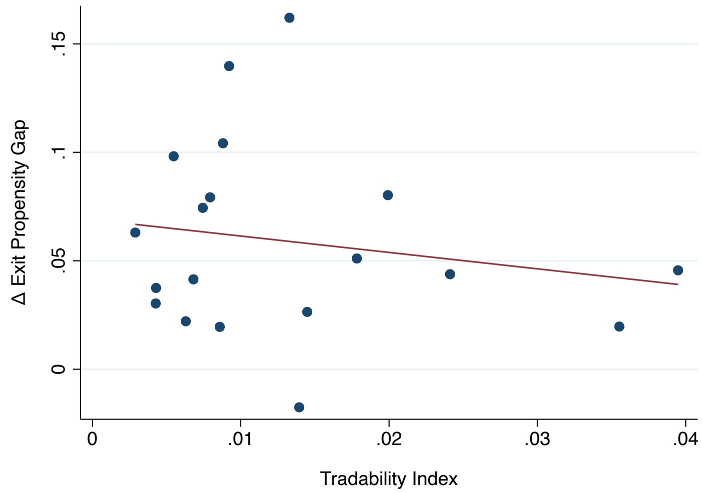

Notes: This figure plots the relationship between a sector's difference in the exit propensity gap between a high- and low-productivity firm during the 2010–13 sudden stop relative to normal times and the tradability index. Each point represents a sector, with the red line depicting the fitted relationship. The data used is collected from the ESEE dataset.

aggregate productivity response across exchange rate regimes is accounted for by the mechanisms proposed in this paper.

6 Quantitative Analysis

This section begins with a discussion of the model parameterization. I then solve for the transition dynamics following an aggregate shock. Finally, I discuss how I use the model to simulate the 2010-13 sudden stop in Spain and evaluate its quantitative performance.

6.1 Parameterization

I start by discussing how I parameterize firm technology and choose the parameters of the model. I divide the parameters into two groups. For the first group, I use commonly used values in the literature. For the second group, I implement a moment-matching exercise to assign values. To do so, I use the steady state of the model with constant $Z_t^A = 1$ and $R_t = \beta^{-1}$. Details on the algorithm that solves the stationary competitive equilibrium are relegated to Online Appendix C.2.

TABLE VI: PARAMETER VALUES

Panel (a): Externally Calibrated Parameters					
Parameter	Description	Value	Source		
β	Discount factor	0.99	Annual real return on bonds is 4%		
σ	Labor share	0.64	National Accounts Spain		
ε_w	Elasticity of substitution (labor)	4.3	Galí and Monacelli (2016)		
θ	Index of wage rigidity	0.8	Galí and Monacelli (2016)		
τ	Iceberg trade cost	1.3	Ghironi and Melitz (2005)		
ρ_z	Transitory productivity - persistence	0.59	Gopinath et al. (2017)		
σ_z	Transitory productivity - st. dev. of innovations	0.13	Gopinath et al. (2017)		
Panel (b): Internally Calibrated Parameters					
Parameter	Description	Value	Moment	Data	Model
z_L	Permanent productivity parameter	0.611	Firm size - average	15.17	18.37
z_H	Permanent productivity parameter	0.793	Firm size - st. dev.	45.98	56.77
π	Permanent productivity parameter	0.99993	Employment share: <10	16.03	20.33
γ	Preference parameter	0.00040	Employment share: 10-249	53.46	57.59
α	Preference parameter	1.216	Average size of entrants	5.21	6.62
e_0	Prob. of zero exit value	0.914	Average size of exiters	4.87	4.46
k_e	Exit shape parameter	2.744	Exit rate	0.081	0.078
c_e	Fixed cost of entry	8.03×10^{-11}	Exit rate: <10	0.095	0.079
c_s	Fixed cost of ideas	4.48×10^{-6}	Export share	0.16	0.16
A	Foreign demand parameter	4388	Export intensity	0.27	0.35
B	Foreign demand parameter	3269	W/W^*	1.30	1.30

Functional Forms I assume a discrete process for transitory idiosyncratic firm productivity shocks that approximates the autoregressive process

$$\log z_t^T = -\frac{\sigma_z^2}{2(1+\rho_z)} + \rho_z \log z_{t-1}^T + \sigma_z u_{it}^z \quad \text{with} \quad u_{it}^z \sim N(0,1)\,,$$

where ρ_z captures persistence and σ_z denotes the standard deviation of innovations u_{it}^z. The constant term normalizes the unconditional mean of transitory idiosyncratic productivity to one for any choice of ρ_z and σ_z. The permanent component of productivity, z_i^P is drawn from the following distribution:

$$z_i^P = \begin{cases} z^L\,, & \text{with probability } \pi \\ z^H\,, & \text{with probability } 1-\pi \end{cases}$$

The exit value is assumed to be zero with probability e_0. With probability $(1-e_0)$, the exit value is drawn from a Pareto distribution characterized by shape parameter k_e. The entry transition function is assumed to be identical to that of incumbents: $H^q(z|q) = H(z'|z)$. The distribution of ideas follows $Q(q) = B\,exp(-q)$ over the lower part of the transitory productivity distribution where B is a scale parameter to ensure $Q(q)$ adds up to one. Γ^* is assumed to be equal to the stationary distribution of the domestic operating firms.

Externally Calibrated Parameters Panel (a) of Table VI provides a summary of the model parameters that are externally calibrated, their baseline values and the source or the empirical target. This first set of parameters are standard and, thus, values are chosen in line with the literature and, when possible, consistent with Spanish statistics taking the 2002-08 period as a reference. The time period of the model is a quarter. Accordingly, the discount factor β is chosen to be 0.99. The output elasticity parameter σ is set to 0.64, roughly the average labor share and within the range that is common in the literature. For the elasticity of substitution for labor types, ϵ_w, and the index of wage rigidities, θ, values are taken from Galí and Monacelli (2016) which are based on empirical studies on European countries conducted by the OECD. In terms of trade costs, τ is equal to 1.3 following Ghironi and Melitz (2005) and many others. Finally, for the parameters that govern the evolution of firm productivity, I borrow from Gopinath et al. (2017) that exploit ORBIS data for Spanish manufacturing firms. In particular, ρ_z and σ_z, are estimated by fitting a firm fixed-effects model on firm idiosyncratic productivity to be 0.59 and 0.13 respectively.

Internally Calibrated Parameters I calibrate the rest of the parameters to match informative moments of the data. I target a total of eleven moments.

The first six moments include the mean and standard deviation of firm size; the employment shares of firms with fewer than 10 employees and those with 10–249 employees; and the average size of entering and exiting firms, all based on the Structural Business Statistics from Eurostat. The next two moments are the average exit rate and the exit rate among micro firms, drawn from the Spanish Business Registry (Directorio Central de Empresas, DIRCE). I also target the share of exporting firms and the exports-to-sales ratio, calculated using the 5th vintage of the CompNet dataset and restricted to manufacturing firms from 1998 to 2008.[39] The final moment is the relative wage level, measured as the ratio of the unit labor cost in Spain vis-à-vis Germany, according to the ECB.

The model is highly nonlinear, and all parameters affect all moments. However, some parameters are more determinant for certain statistics. The relative wage level is pinned down by the free entry condition and, thus, largely driven by the fixed cost of ideas. Selection in entry, as reflected by the average size of entrants, is mostly governed by the fixed cost of entry. Exit in the model is due to the choke price or the arrival of the exit shock. As the choke price is largely determined by the preference parameter α, there is a close relationship between this parameter and the exit rate of micro firms. Holding this constant, the overall exit rate is then determined by the probability of a non-zero exit value. The average of exiters is controlled by the maximum exit value. The export share and the export to sales ratio are affected by the foreign demand parameters A and B. Finally, the mean and standard deviation of firm size, as well as the employment shares, are predominantly affected by the preference parameter γ and the parameters governing permanent productivity.

Panel (b) of Table VI reports the estimated parameter values as well as the target moments in the data and the model. In addition, note that the preference parameter η is normalized to one

[39] The CompNet dataset is an initiative by the Competitiveness Research Network, established by the ECB, to compile detailed firm-level data across several European countries.

FIGURE IV: FIRM SIZE DISTRIBUTION - DATA AND MODEL

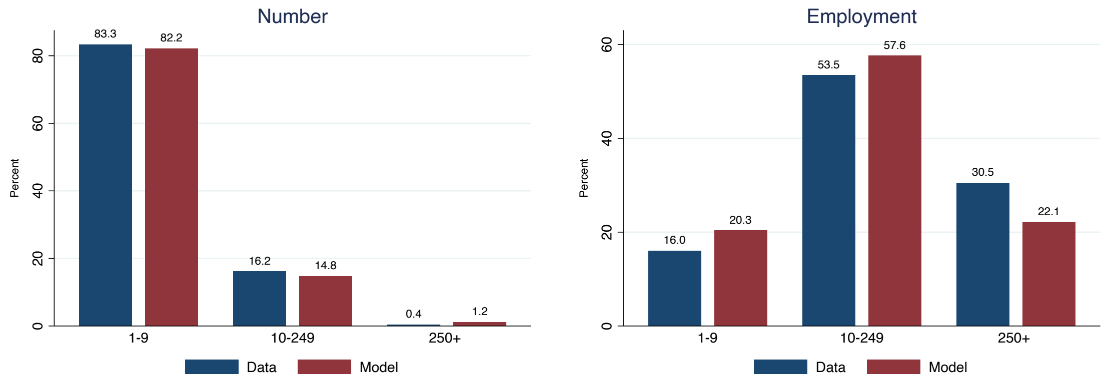

Notes. Each panel compares data and model shares across firm size classes. The left panel shows the distribution of firms, where each bar represents the share of firms in that size category. The right panel reports employment shares by firm size. In both panels, the values represent fractions of the total—each bar group (Data and Model) adds up to one across size classes. Data moments are based on the Structural Business Statistics from Eurostat.

without loss of generality and the steady state value of debt is set to ensure that trade is balanced in steady state, $\bar{B} = 5.6 \times 10^{-6}$.

Model Fit As a measure of model fit, Figure IV compares the population and employment shares by firm size in the model and in the data. The calibration captures key features of the firm size distribution. While the employment shares are directly targeted, the population shares are not—yet the model closely replicates them. Most firms are micro, with fewer than 10 employees, but they account for only a modest share of total employment. In contrast, medium and large firms, though fewer in number, employ the bulk of the workforce.

6.2 Adding Aggregate Dynamics

Before I turn towards studying the quantitative relevance of the mechanism, I explore the dynamic properties of the quantitative model. The goal is to show that the analytical results in 5.3 hold more broadly. I model a sudden stop as the result of a positive shock to the interest rate and a decline in the common productivity shifter in the spirit of Mendoza (2010).[40] The latter serves as a reduced-form approach to capturing the decline in the within-firm productivity component that is common to both the 1992–93 and 2010–13 sudden stop episodes, as shown in Table II. [41]

[40]Note that Mendoza (2010) also features a shock to the price of imported intermediate inputs, which are absent from my model.

[41]While endogenizing the common productivity shifter is beyond the scope of this paper, incorporating imported inputs and variable capacity utilization alongside credit constraints would establish a financial link between the common productivity shifter and the risk premium shock. This financial link would operate similarly across regimes, as firms substitute foreign intermediate inputs with domestic alternatives following the real exchange rate depreciation, leading to a decline in measured productivity at the firm level.

FIGURE V: MACROECONOMIC EFFECTS OF A SUDDEN STOP

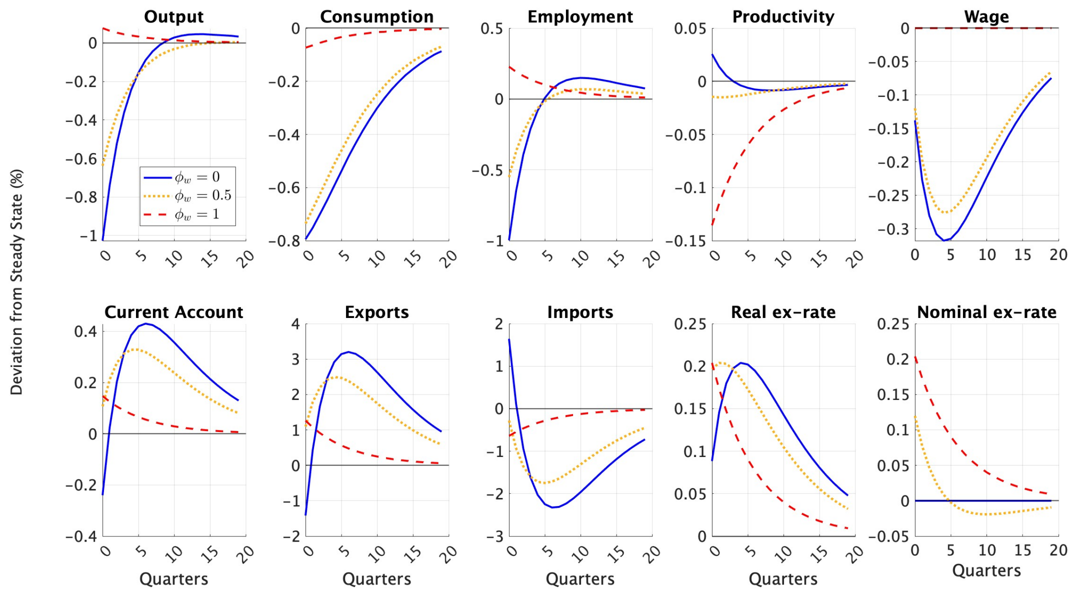

Notes. These figures plot the impulse response functions of key macroeconomic variables to a simultaneous increase in the interest rate and a decline in the common productivity shifter under different exchange rate regimes, as predicted by the model described in Section 4. Under the currency union ($\phi_w = 0$), the shocks correspond to a one percent cumulative increase in the country-specific risk premium and a one percent cumulative decrease in the common productivity shifter. For the floating regime ($\phi_w = 1$) and the intermediate case ($\phi_w = 0.5$), the interest rate shock is recalibrated such that the peak real exchange rate depreciation matches that of the currency union case. All variables except the current account are expressed in log deviations from steady state. The current account is expressed as a share of output and shown in percentage point deviations from steady state. Exports and imports are denominated in domestic currency; output and consumption are expressed in real terms.

When the economy is out of steady state, the value functions, the distribution of operating firms, and all aggregate variables change over time. Hence, I solve for the sequences $\{V_t, V_t^c, V_t^e, V_t^P, \Gamma_t\}_{t=0}^T$ and $\{W_t, W_t^*, \lambda_t, p_t^{\max}, R_t, B_t, N_t,\}_{t=0}^T$ consistent with (i) household, labor union and firm (incumbent, entrant and importer) optimization, (ii) aggregate wage and interest rate dynamics, (iii) law of motion for the distribution of operating firms and (iv) goods market clearing, where $t = 0$ is the period in which shocks hit and agents learn about them, and T is the period in which the economy is back in steady state. Online Appendix C.3 describes the numerical algorithm.

Figure V summarizes the model-implied responses of key macroeconomic variables to a simultaneous increase in the interest rate and a decline in the common productivity shifter. The shocks are initially specified under a currency union ($\phi_w = 0$) as a one percent cumulative increase in the interest rate and a one percent cumulative decrease in the productivity shifter. For the floating regime ($\phi_w = 1$) and an intermediate case ($\phi_w = 0.5$), the interest rate shock is recalibrated so that the size of the real exchange rate adjustment is comparable to that of the currency union. [42] All

[42]See Section 6.3.2 for further details.

variables, except the current account, are expressed in log deviations from steady state.

As expected, a sudden stop is characterized by a depreciation of the real exchange rate and a current account surplus. The model predicts a slight delay in the adjustment within a currency union. [43] The path of TFP clearly diverges across the two polar cases. On the one hand, under the baseline calibration, the positive effect of a lower aggregate demand offsets the negative effect of lower production costs and fewer competing firms on the domestic productivity cutoff and, thus, TFP improves in the currency union. On the other hand, productivity falls unambiguously in the floating regime.

Output and consumption are measured in real terms. The model predicts a fall in both variables under a currency union. In a floating regime, there is a decline in consumption, although less pronounced, and a slight increase in output. Following output fluctuations, employment only falls in the currency union. [44] The current account surplus, expressed in domestic currency, reflects a combination of rising exports and falling imports.

The intermediate regime is included to illustrate how macroeconomic responses vary with the degree of nominal exchange rate adjustment. As more of the real exchange rate depreciation is absorbed through the nominal exchange rate, the contraction in aggregate demand is weaker, and productivity begins to fall. At the same time, the decline in output and employment in this case illustrates that once some wage adjustment begins, as would likely occur under more realistic interpretations of floating regimes, domestic activity contracts.

6.3 Simulating the 2010-13 Sudden Stop

With the calibrated model in hand, I simulate the dynamics of the Spanish economy during the 2010–2013 sudden stop episode. Specifically, I filter the interest rate shock, ϵ_t^R, and the productivity shifter, ϵ_t^A, such that the model-generated paths for the real exchange rate and output replicate the actual Spanish data over the 2002–2014 period. Agents are assumed to have perfect foresight but are continuously surprised by the realization of these two aggregate shocks, which they interpret as following AR(1) processes.

Figure VI plots the time series for the real exchange rate and output in the data (solid blue) alongside the model-implied series (dashed red). The real exchange rate is based on relative unit labor costs and constructed using National Accounts data from Eurostat. Output corresponds to real GDP at constant prices, as reported in the World Bank's World Development Indicators. The

[43] This is entirely driven by nominal rigidities as the model disregards additional policy instruments available under a currency union, such as public capital inflows, that might directly cushion the adjustment in the data.

[44] The slight increase in output and employment under a floating regime arises because exchange rate policy in the model fully offsets the nominal rigidity and replicates the flexible-wage allocation. This makes the equilibrium effectively frictionless, and in such environments it is difficult to generate both a current account reversal and a contraction in domestic activity. Any small departure from strict zero wage-inflation targeting—as in the intermediate case considered below—restores the model's ability to generate the empirically observed decline in output and employment during sudden stops.

FIGURE VI: TARGETS FOR MODEL SIMULATION - SPAIN, 2002–2014

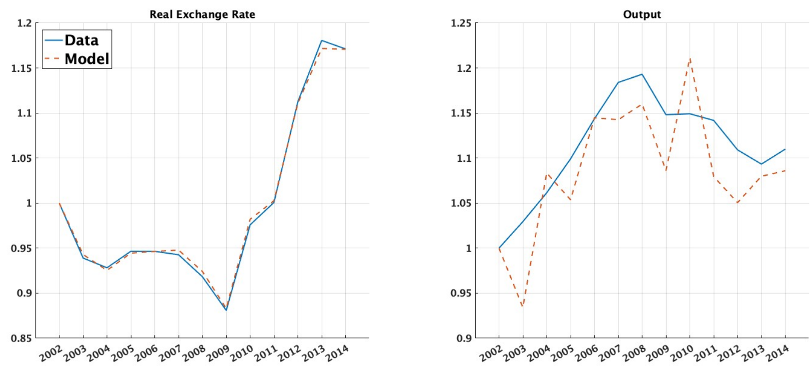

Notes. This figure plots the data (solid blue) and model-implied series (dashed red) for the real exchange rate (left panel) and output (right panel) in Spain over the period 2002–2014. The real exchange rate is based on relative unit labor costs, constructed using National Accounts data from Eurostat. Output corresponds to real GDP at constant prices, taken from the World Bank's World Development Indicators. Most series are expressed as indices normalized to 1 in the initial year.

model matches the real exchange rate dynamics almost perfectly, including the sharp appreciation observed between 2010 and 2013. While it is less precise in replicating output, it generates a contraction during the crisis years that is broadly in line with the magnitude observed in the data. The fitted time series for the underlying shocks are shown in Figure A.12 in the Online Appendix.

Having simulated the Spanish economy, I return to the aggregate productivity decomposition exercise to evaluate the quantitative performance of the model. For convenience, the first column of Table VII reproduces the 2010–2013 productivity growth decomposition based on the ESEE data, while the second column summarizes the results generated by the model under the baseline calibration.

The baseline results are computed using a firm sample constructed to mimic key characteristics of the ESEE dataset. Specifically, I restrict attention to all firms with more than 200 employees and a random 5% sample of those with 10 to 200 employees, drawn once at the beginning of the simulation to mimic the sampling strategy used in assembling the ESEE dataset. Similarly, all new entrants with more than 200 employees are included each year, along with a 5% sample of entrants with 10 to 200 employees.

The model performs well in replicating the empirical decomposition of productivity growth. It attributes the gains primarily to firm exit, consistent with the data: in the empirical decomposition, exiters contribute 7.19 percentage points to an overall growth of 2.24, while in the model they contribute 7.00 out of 3.97. Relative to total productivity growth, the importance of exit in the model is about 55% of that observed in the data.The model broadly replicates the negative total contribution of incumbents, although it somewhat overstates aggregate productivity growth. This overestimation arises in part because the model assigns a modestly positive contribution to entrants, whereas the firm-level data show a small negative effect. At the same time, while

TABLE VII: PRODUCTIVITY GROWTH DECOMPOSITION 2010-13 - DATA AND MODEL VARIANTS

	Data	Baseline	Full Sample	No Shifter	Flexible
Productivity Growth (%)	2.24	3.97	3.50	8.41	-2.63
Contribution to Productivity Growth:					
Incumbents' Contribution	**-4.54**	**-3.54**	**-0.96**	**2.17**	**-4.33**
Within-firm Contribution	-7.36	-34.08	-8.12	-19.44	-29.11
Between-firm Contribution	0.32	55.03	8.53	38.03	41.00
Cross-term Contribution	2.50	-24.49	-1.37	-16.41	-16.21
Net Entry Contribution	**6.79**	**7.51**	**4.45**	**6.23**	**1.70**
Entrants' Contribution	-0.40	0.52	-1.43	-0.20	1.70
Exiters' Contribution	7.19	7.00	5.89	6.44	0.00

Notes: This table reports the decomposition of productivity growth in Spain over the period 2010–2013. Column 1 repeats the empirical results shown in Table II, based on the ESEE data. Column 2 shows the model's baseline simulation using a sub-sample generated to match the characteristics of the ESEE dataset. Column 3 extends the baseline simulation to the full firm sample, representing the entire economy. Column 4 presents a variant of the simulation in which the common productivity shifter is shut down. Column 5 reports a counterfactual scenario in which Spain operates under a flexible exchange rate regime during the 2010–2013 simulation. All figures are expressed in percentage points.

the model matches the total incumbent contribution, the individual within-firm and reallocation components are much larger than in the data, even if they offset in the aggregate.[45]

6.3.1 Robustness

To assess the robustness of the baseline findings, I consider two complementary exercises. First, I evaluate the implications of the sample restrictions in the empirical analysis by extending the model-based decomposition to the full firm sample. This helps assess the extent to which the patterns documented in Section 3 may be affected by selection bias due to data limitations. Second, I explore how much of the observed productivity gains can be solely accounted for by financial tightening alone by running the simulation without the shock to the common productivity shifter.

The results are summarized in Columns 3 and 4 of Table VII respectively. When the decomposition is performed using the full firm sample, the model still predicts positive productivity growth (3.50%), and the overall message remains consistent with the baseline: firm exit remains the primary driver of improvements in productivity, overturning the negative contribution of incumbents. The magnitudes, however, differ slightly. The exit-driven contribution is about 5% lower in the full sample than in the baseline. Interestingly, including the full set of entrants also introduces a negative contribution from entry, consistent with the firm-level evidence, though not large enough to overturn the positive effect of net entry. Taken together, these findings suggest

[45]Non-convex adjustment in the model amplifies conditional responses among surviving firms, generating large and offsetting within-firm productivity changes and reallocation effects among incumbents despite a realistic aggregate contribution.

that while the empirical analysis focuses on a selected subset of firms, the main patterns it reveals are most likely not an artifact of sample selection and would extend to the full economy.

Turning to Column 4, I re-filter the interest rate shock so that the model continues to match the same real exchange rate depreciation as in the baseline simulation, but now with the common productivity shifter shut down. This leads to a sharper increase in predicted productivity growth (8.41%), primarily because the contribution of incumbents turns positive. In contrast, the contribution of net entry, the model's core mechanism, remains robust and does not depend on shocks to common productivity. The role of the shifter is primarily to discipline the behavior of incumbent productivity and to prevent overstating the aggregate productivity growth.

6.3.2 A Flexible Exchange Rate Counterfactual

Extending the simulation backward to explicitly capture the 1992–1993 sudden stop is not feasible due to data limitations prior to 1995. Instead, I construct a counterfactual scenario in which Spain retains a flexible exchange rate regime, i.e., does not join the euro, but is exposed to the same real exchange rate fluctuations and common productivity shocks as in the baseline simulation.[46]

In this flexible regime, productivity declines rather than rises, driven entirely by worsening incumbent performance. The contribution of net entry is limited, and in particular, exiting firms do not play any role in offsetting the aggregate decline. Although this simulation does not directly replicate the conditions of the 1992–1993 crisis, the patterns it generates are broadly consistent with the firm-level evidence for that period, as reported in Table II.

From productivity to welfare Productivity gains are not always welfare-enhancing. To assess the broader implications of the counterfactual, I compute a consumption-equivalent variation measure using the full firm sample, which quantifies the percentage change in lifetime consumption that would make households indifferent between the two regimes. The findings suggest that agents would be willing to give up approximately 1% of lifetime consumption to remain in the currency union during the sudden stop. This result may appear surprising, given the conventional wisdom that floating exchange rates perform better. However, in this model, the flexible equilibrium is also inefficient due to the presence of variable markups, which distort firm behavior and undermine the selection process. This highlights the importance of incorporating firm selection and market structure into the analysis of optimal exchange rate policy in future work.

[46] An alternative approach would be to simply feed the original sequence of structural shocks. The counterfactual I construct here instead holds constant the real exchange rate fluctuations while varying only the exchange rate regime. This disentangles the role of the regime from the size of the real exchange rate adjustment, ensuring that the mechanism operates through the regime itself rather than through differences in the real devaluation. Retaining the primitive shocks would mechanically produce larger real exchange rate movements under a float, making differences across regimes appear even more pronounced.

7 Conclusion

This paper revisits a classical question in international macroeconomics: how does exchange rate policy affect macroeconomic performance after a shock? While the literature provides many attempts at answering this issue, it has mostly overlooked the effect on firm dynamics. I study the question anew in the context of a sudden stop, emphasizing the divergence in TFP patterns that emerges across exchange rate regimes in the aggregate data and relating them to observed differences in firm exit at the micro level.

Taking the firm-level analysis of two sudden stops in Spain as a starting point, the paper argues that documented differences in the reallocation of resources from unproductive exiting firms to productive survivors might be related to the degree of currency appreciation vis-à-vis wage devaluation. A small open economy DSGE model featuring firm selection, variable markups and elastic labor supply formalizes the mechanism. Productivity is determined by the number of firms (pro-competitive channel), the marginal utility of income (demand channel) and the unit cost of production (cost mechanism). The relative magnitude of these forces depends on the exchange rate policy with a currency union generating quantitatively more cleansing because of a larger demand effect. Systematic analysis of the behavior of macroeconomic variables during sudden stops under different exchange rate regimes confirms that the model's implications hold for a wide set of economies.

This paper provides a primarily positive account of how exchange rate policy affects short-term productivity growth, while offering a utility-based welfare comparison as a first step toward understanding its normative implications. A full welfare analysis, however, remains pending. Looking ahead, an important question is how productivity gains translate into welfare improvements. Evaluating the trade-off between enhanced resource reallocation and the persistence of nominal rigidities will be key to assessing the desirability of different policy responses. In particular, determining the optimal weight that policy should place on each margin remains an open avenue for future research.

Appendix

Proof of Lemma 1

For a Pareto productivity distribution with shape parameter k, the unweighted average productivity among active domestic producers is

$$\tilde{Z}_t^H = \frac{1}{1-\Gamma(z_t^H)} \int_{z_t^H}^{\infty} z\, Z_t\, \gamma(z)\, dz = \frac{k}{k-1} z_t^H Z_t.$$

Using the zero-profit condition for the marginal firm, $p_t^{\max} = W_t^{\sigma}/(Z_t z_t^H)$, domestic policies can be written as

$$q_t^H(z) = \frac{\lambda_t}{2\gamma} \frac{W_t^\sigma}{Z_t} \left(\frac{1}{z_t^H} - \frac{1}{z} \right), \qquad l_t^H(z) = \frac{W_t^{\sigma-1}}{zZ_t} q_t^H(z), \qquad r_t^H(z) = \frac{\lambda_t}{\gamma} \left(\frac{W_t^\sigma}{2Z_t} \right)^2 \left[\left(\frac{1}{z_t^H} \right)^2 - \left(\frac{1}{z} \right)^2 \right].$$

Evaluating the corresponding aggregates under the Pareto distribution yields

$$\hat{Z}_t^H = \frac{k+1}{k} \tilde{Z}_t^H, \qquad \check{Z}_t^H = \frac{(k-1)(k+2)}{k^2} \tilde{Z}_t^H, \qquad \bar{Z}_t^H = \frac{k+2}{k+1} \tilde{Z}_t^H.$$

Proof of Proposition 1

Using the zero-profit condition,

$$p_t^{\max} = \frac{W_t^\sigma}{Z_t z_t^H},$$

the optimal domestic and import prices are

$$p_t^H(z) = \frac{1}{2} \frac{W_t^\sigma}{Z_t} \left(\frac{1}{z_t^H} + \frac{1}{z} \right), \qquad p_t^M(z) = \frac{1}{2} e_t \tau \left(\frac{1}{z_t^M} + \frac{1}{z} \right).$$

The choke-price condition implies

$$z_t^M = \frac{e_t \tau}{W_t^\sigma} Z_t z_t^H.$$

Aggregating prices yields

$$P_t = \frac{2k+1}{2k+2} \frac{W_t^\sigma N_t}{Z_t z_t^H}.$$

Substituting into the choke-price definition gives

$$Z_t z_t^H = \frac{\lambda_t W_t^\sigma}{\alpha \gamma} \left[\gamma + \frac{\eta}{2k+2} N_t \right].$$

Setting $Z_t = 1$ and totally differentiating delivers the result.

Proof of Proposition 2

From the flexible-wage condition,

$$W_t^* = \frac{\varepsilon_w}{\varepsilon_w - 1} \frac{1}{\lambda_t}.$$

Rewriting the wage Phillips curve yields

$$1 = \left[\theta \left(\frac{W_{t-1}}{W_t} \right)^{1-\varepsilon_w} + (1-\theta) \left(\frac{W_t^*}{W_t} \right)^{1-\varepsilon_w} \right]^{\frac{1}{1-\varepsilon_w}}.$$

The Euler equation implies

$$\lambda_t e_t = \beta R_t \mathbb{E}_t[\lambda_{t+1} e_{t+1}],$$

so a sudden stop raises $\lambda_t e_t$ on impact. Under a floating regime, the exchange rate absorbs the adjustment, implying $d\lambda_t = dW_t = 0$; under a currency union, $d\lambda_t > 0$ and $dW_t < 0$.

Using

$$N_t = \left(\frac{1}{z_t^H}\right)^k \left[M + M^* \left(\frac{W_t^\sigma}{\tau e_t}\right)^k\right]$$

together with the expression for z_t^H yields an implicit equation in N_t. Under a float, $dN_t < 0$; under a currency union, the sign is ambiguous.

Proof of Corollary 1

Combining Propositions 1 and 2 yields

$$z_t^H - \frac{\eta}{2k+2} \frac{\lambda_t W_t^\sigma}{\alpha\gamma} \left(\frac{1}{z_t^H}\right)^k \left[M + M^* \left(\frac{W_t^\sigma}{\tau e_t}\right)^k\right] = \frac{\lambda_t W_t^\sigma}{\alpha}.$$

A sufficient condition for $dz_t^H > 0$ is $d(\lambda_t W_t^{\sigma(k+1)}) > 0$, which holds provided

$$(1-\theta)\sigma(k+1) < \left(\frac{W_t}{W_t^*}\right)^{1-\varepsilon_w}.$$

A first-order approximation implies the condition holds for small shocks.

References

Ackerberg, Daniel A, Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica*, 83(6): 2411–2451.

Alessandria, George, and Horag Choi. 2007. "Do Sunk Costs of Exporting Matter for Net Export Dynamics?" *The Quarterly Journal of Economics*, 122(1): 289–336.

Alessandria, George, Joseph P Kaboski, and Virgiliu Midrigan. 2010. "Inventories, Lumpy Trade, and Large Devaluations." *American Economic Review*, 100(5): 2304–2339.

Almunia, Miguel, David López-Rodríguez, and Enrique Moral-Benito. 2018. "Evaluating the Macro-Representativeness of a Firm-Level Database: an Application for the Spanish Economy." Banco de España Occasional Paper 1802.

Almunia, Miguel, Pol Antràs, David Lopez-Rodriguez, and Eduardo Morales. 2021. "Venting Out: Exports during a Domestic Slump." *American Economic Review*, 111(11): 3611–3662.

Ates, Sina T, and Felipe E Saffie. 2021. "Fewer but Better:Sudden Stops, Firm Entry, and Financial Selection." *American Economic Journal: Macroeconomics*, 13(3): 304–356.

Baily, Martin Neil, Charles Hulten, David Campbell, Timothy Bresnahan, and Richard E Caves. 1992. "Productivity Dynamics in Manufacturing Plants." *Brookings Papers on Economic Activity: Microeconomics*, 1992: 187–267.

Benigno, Gianluca, and Luca Fornaro. 2014. "The Financial Resource Curse." *The Scandinavian Journal of Economics*, 116(1): 58–86.

Bornstein, Gideon, and Laura Castillo-Martinez. 2023. "Firm Exit and Financial Frictions."

Caballero, Ricardo J., and Mohamad L. Hammour. 1994. "The Cleansing Effect of Recessions." *American Economic Review*, 84(5): 1350–1368.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics*, 12(3): 383–398.

Calvo, Guillermo A. 1998. "Capital Flows and Capital-Market Crises: the Simple Economics of Sudden Stops." *Journal of Applied Economics*, 1(1): 35–54.

Calvo, Guillermo A., Alejandro Izquierdo, and Luis-Fernando Mejía. 2004. "On the Empirics of Sudden Stops: the Relevance of Balance-Sheet Effects." NBER Working Paper 10520.

Cavallo, Eduardo A, and Jeffrey A Frankel. 2008. "Does Openness to Trade Make Countries More Vulnerable to Sudden Stops, or Less? Using Gravity to Establish Causality." *Journal of International Money and Finance*, 27(8): 1430–1452.

De Loecker, Jan, and Frederic Warzynski. 2012. "Markups and Firm-Level Export Status." *American Economic Review*, 102(6): 2437–71.

Dias, Daniel A., and Carlos Robalo Marques. 2021. "Every Cloud Has a Silver Lining: Micro-Level Evidence on the Cleansing Effects of the Portuguese Financial Crisis." *Oxford Bulletin of Economics and Statistics*, 83(2): 352–376.

Engel, Charles. 1999. "Accounting for US Real Exchange Rate Changes." *Journal of Political Economy*, 107(3): 507–538.

Foster, Lucia, Cheryl Grim, and John Haltiwanger. 2016. "Reallocation in the Great Recession: Cleansing or Not?" *Journal of Labor Economics*, 34(S1): S293–S331.

Foster, Lucia, John C Haltiwanger, and Cornell John Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*. 303–372. University of Chicago Press.

Galí, Jordi, and Tommaso Monacelli. 2016. "Understanding the Gains from Wage Flexibility: the Exchange Rate Connection." *American Economic Review*, 106(12): 3829–3868.

García-Santana, Manuel, Enrique Moral-Benito, Josep Pijoan-Mas, and Roberto Ramos. 2020. "Growing like Spain: 1995-2007." *International Economic Review*, 61(1): 383–416.

Ghironi, Fabio, and Marc J. Melitz. 2005. "International Trade and Macroeconomic Dynamics with Heterogeneous Firms." *The Quarterly Journal of Economics*, 120(3): 865–915.

Gopinath, Gita, Şebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez. 2017. "Capital Allocation and Productivity in South Europe." *The Quarterly Journal of Economics*, 132(4): 1915–1967.

Griliches, Zvi, and Haim Regev. 1995. "Firm Productivity in Israeli Industry 1979–1988." *Journal of Econometrics*, 65(1): 175–203.

Guidotti, Pablo E., Federico Sturzenegger, Agustin Villar, José de Gregorio, and Ilan Goldfajn. 2004. "On the Consequences of Sudden Stops." *Economia*, 4(2): 171–214.

Hopenhayn, Hugo A. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica*, 1127–1150.

Ilzetzki, Ethan, Carmen M. Reinhart, and Kenneth S. Rogoff. 2019. "Exchange Arrangements Entering the 21st Century: Which Anchor Will Hold?" *The Quarterly Journal of Economics*, 134(2): 599–646.

Kalemli-Özcan, Şebnem, Bent E Sørensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yeşiltaş. 2024. "How to Construct Nationally Representative Firm-level Data from the Orbis Global Database: New Facts on SMEs and Aggregate Implications for Industry Concentration." *American Economic Journal: Macroeconomics*, 16(2): 353–374.

Kehoe, Timothy J., and Kim J. Ruhl. 2009. "Sudden Stops, Sectoral Reallocations, and the Real Exchange Rate." *Journal of Development Economics*, 89(2): 235–249.

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: a Proposal to Replace the New Keynesian Phillips Curve." *The Quarterly Journal of Economics*, 117(4): 1295–1328.

Melitz, Marc J. 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica*, 71(6): 1695–1725.

Melitz, Marc J., and Gianmarco I.P. Ottaviano. 2008. "Market Size, Trade, and Productivity." *The Review of Economic Studies*, 75(1): 295–316.

Melitz, Marc J., and Sašo Polanec. 2015. "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit." *The Rand Journal of Economics*, 46(2): 362–375.

Mendoza, Enrique G. 2010. "Sudden Stops, Financial Crises, and Leverage." *American Economic Review*, 100(5): 1941–66.

Meza, Felipe, and Erwan Quintin. 2007. "Factor Utilization and the Real Impact of Financial Crises." *The BE Journal of Macroeconomics*, 7(1): 1–41.

Neumeyer, Pablo A., and Fabrizio Perri. 2005. "Business Cycles in Emerging Economies: the Role of Interest Rates." *Journal of Monetary Economics*, 52(2): 345–380.

Olley, G Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica*, 64(6): 1263–1297.

Prados De la Escosura, Leandro. 2017. *Spanish Economic Growth, 1850–2015.* Springer Nature.

Reis, Ricardo. 2013. "The Portuguese Slump and Crash and the Euro Crisis." *Brookings Papers on Economic Activity*, 2013(1): 143–210.

Schumpeter, Joseph A. 1939. *Business Cycles.* McGraw-Hill New York.