Trade Dynamism and Import Market Maturity

David Hummels*, Purdue University and NBER

Kan Yue**, Xavier University

February 12, 2024

(Preliminary and Incomplete: Please do not cite or distribute)

Abstract:

Recent work documents declining business dynamism in the United States and European countries, with concerning implications for markups, innovation and productivity. Using import data for 146 countries over four decades we document a set of new stylized facts describing market dynamism. The extent of variety "churn" is an order of magnitude larger than net growth in varieties, but falls significantly over time. Markets exhibit falling entry rates and increased price and market share advantages for incumbents relative to new entrants. Further, incumbents appear to exhibit greater resilience than newer entrants, showing lower volatility in prices, market shares, and market duration, as well as a smaller response to cost shocks such as rising tariffs. These patterns hold for 90 percent of countries and product types, and so are difficult to ascribe to unique industrial organization or country characteristics. We compare various models of firm entry to see which are able to explain these facts.

Acknowledgements: We thank seminar participants at Purdue University, and especially Mario Crucini, Andrew Greenland, John Lopresti, Kim Ruhl, Anson Soderbery, and Chong Xiang for helpful comments on earlier drafts.

Author info: David Hummels hummelsd@purdue.edu Kan Yue ** yuek@xavier.edu

1. Introduction

Recent work documents declining business dynamism in the last few decades within markets in the United States and a few European countries.[1] Among the key factors highlighted in this literature[2], the entry rate of new firms and the share of young firms in economic activity have declined. The labor productivity gap between frontier and laggard firms has widened. The reallocation of jobs and churn in the labor market have slowed. The dispersion of firm growth rates has decreased. Market concentration has increased and in some instances markups have increased.

The consequence of declining dynamism depends primarily on the reason dynamism has declined. If declining dynamism results from rising barriers to entry for new firms or perhaps overt exercise of market power by dominant firms, then concentration translates to a reduction in new varieties that consumers would value, and rising markups. (Covarrubias et al 2020) It may also have implications for slowing of aggregate productivity growth (Decker et al 2017, Alon et al 2018) or changing labor shares of income. If on the other hand, dominant firms are simply getting better faster, then rising concentration is the result of growing productivity advantages in these firms, with consumers enjoying falling prices. Similarly, if there are costs associated with establishing new business relationships, then a reduction in churn (between input suppliers to firms, between firms and their consumers) may generate welfare gains from stability. (Baqaee et al 2023)

One of the challenges in sorting out the causes and consequences of declining dynamism is that the literature to date has primarily focused on firm-level micro data within a series of single country case studies, either explaining aggregate macro facts or in some cases exploiting cross-industry or cross-regional variation.[3] This makes it difficult to know whether changes in business dynamism are a widespread phenomenon, or to exploit useful cross-country variation in market behavior or policy, and difficult to untangle multiple causes that might arise within a single time period within a country.

In this paper we offer a new approach to measuring business dynamism comprehensively worldwide by focusing on import markets. While a focus on importing behavior prevents us from seeing the activity of domestic firms in each market, it provides insights into competitive market dynamics that

[1] While the majority of this work has focused on the US, related results have been demonstrated for Belgium (Bijnens and Konings 2020) and Turkey (Akcigit et al 2020).
[2] Akcigit and Ates (2021) provide an excellent overview of the literature and its key findings.
[3] Bijnens and Konings 2020, Akcigit et al 2021 and Covarrubias et al 2020 draw some limited cross country comparisons.

we show bear a strong resemblance to those seen when studying domestic firms only. Import data also provides certain advantages. One, the consistency of measurement and long time series coverage of the trade data series allows us to make comparisons across different markets and product types to provide insights into causes of dynamism. Two, the level of disaggregation available in product level data enables a better comparison of market shares and prices for like products, as opposed to using relatively broad industries that incorporate firms producing dissimilar products as well as multiproduct firms. Three, international comparisons enable us to examine the behavior of firms and countries as they sell their product into markets of varying age, as well as to study the impact of cost shocks such as freight costs and tariffs that provide a rich set of time varying and time invariant advantages to subsets of firms.

We use both US import data (for comparison to firm level findings on business dynamism in the US) as well as import data of 146 countries from 1991-2020. We document a sharp decline in the dynamism of import markets (defined as an importer-HS6 product pair) over time. We show:

- The "churn" of varieties entering/exiting markets is an order of magnitude larger than the net growth of variety into import markets.
- Entry rates for new exporters falls sharply over time, and as exit rates remain roughly constant, this leads to a corresponding reduction in the amount of market churn.
- The market share for new entrants (first year competitors, as well as those with less than 3 or 5 years tenure) falls relative to longer-tenured incumbents over time.
- The product prices for long-tenured incumbents falls relative to new entrants over time. However, this is caused not by incumbents getting better, but rather a steady worsening of the price competitiveness of new entrants.

These patterns hold with and without controlling for market characteristics such as size, income, and levels of protection. They are remarkably robust across importers and products, suggesting the inadequacy of explanations that point to specific policy regimes in particular countries or the unique industrial organization of particular products. However, one pronounced difference across countries is the date at which an importing country first registers positive imports of a particular HS6 product from any source country. Using this subset of the data we can we characterize the "maturity" of an import cycle as the length of time since first import, and show that the reduction in import dynamism is more closely associated with "maturity" rather than the calendar date.

We next consider measures of dynamism that are focused on the persistence of productivity advantages. We develop four stylized facts.

- Like the firm-based literature on business dynamism, market share volatility declines over calendar years and with market age.
- Related, the durability of trade relationships (i.e. the likelihood that an exporters exits an import market in a given year) is greater for incumbents than new entrants, and the durability gap grows steadily larger over calendar years and market age.
- Price volatility also declines with market age and calendar year, with volatility declining the most for the most experienced firms.
- Related, the extent to which established firms absorb cost shocks (in the form of tariff or freight increases) rather than passing them onto consumers is greater than new firms and grows steadily over calendar years and exporter expierence.

We next explore a set of canonical models from the literature to better understand how elements of these models generate particular outcomes for market dynamism. Broadly speaking, these are all dynamic models of firm entry in which firms choose whether to enter import markets as a function of their productivity, and fixed and variable costs of serving the market. They differ primarily in the nature of the productivity process, the persistence of productivity and nature of innovations to the productivity time series. Using these models, we simulate the evolution of an import market to show which productivity dynamics generate import-market dynamics that match the broad patterns in the data.

While a modified version of Melitz (2003) with permanent productivity draws and a growing market can generate many of the stylized facts described in the business dynamism literature, explaining facts about price and market share volatility requires a model where productivity mixes persistent and time varying factors, and where established firms face progressively lower variance in their time varying factors.

This paper adds to several literatures. We follow the literature on business dynamism, including papers by Akcigit 2020, Akcigit and Ates 2021, Akcigit et al 2021, Alon et al 2018, Baqaee 2023, Bijnens and Konings 2018, Covarrubias 2020, and Decker et al 2017. Many of our measures of business dynamism are inspired by this work, but we use imports data rather than domestic production data. A weakness of our approach is that we do not observe behavior by domestic firms or how it is

affected by the import maturity cycle. A strength is that we are able to provide broader coverage relative to country-time specific studies and more disaggregated product coverage. We can observe the dynamics for the same HS6 product across multiple importing countries and across the same calendar date corresponding to different points in the import maturity cycle for different importing countries. This enables us to show that patterns are robust to countries and products and are therefore not obviously linked to unique policy or other characteristics in these markets.

A second literature relates to variety growth in markets. Since Feenstra (1994) and Romer (1994), authors in the international trade field have placed great emphasis on the net growth in variety entering an import market as central to the welfare gains from trade. This works ignores which exporters are providing the variety or any costs associated with switching sources. However, very recent work by Baqaee et al 2023 demonstrates that churning – changing source relationships – can have large welfare consequences because of the costs of disrupting existing business relationships. We show that the churn in variety for a typical importer-product-time is an order of magnitude larger than the net growth in variety but that falls entry rates and increasing incumbency advantages cause the rate of churn to fall both over time and with market maturity.

The paper proceeds as follows. In Section 2 we describe our data and measures. Section 3 provides stylized facts on import market dynamism. Section 4 provides a model of productivity dynamics with serial dependence, simulates versions of this model to draw contrasts between its behavior and stylized facts developed in Section 3. Section 5 concludes.

2. Data and Measures

2.1 Data

Our primary data is bilateral import data drawn from two sources. The first is annual US Imports of Merchandise data for 1991-2020, aggregated to the level of exporter x HTS-10 digit x year. Because many of our measures focus on entry and exit behavior we follow the Pierce-Schott () algorithm, creating synthetic codes to replace HTS codes that are reclassified over time. This results in just over 10,000 unique product codes.

Our second is all bilateral imports worldwide at the HS-6 level extracted from the UN COMTRADE database for the period 1991-2020, using the earliest version of the Harmonized System for consistent product coding throughout. The HS-6 data (over 5000 product lines) provide sufficiently

granular detail to see significant entry/exit behavior, unlike the older SITC data which is available for longer time frames but much less detailed. The HS data does present a challenges in that some countries are slow to switch from reporting their data from the SITC to the HS nomenclature, which creates incomplete coverage in the early years and some challenges with "dating" markets. The details on the number of countries each year reporting imports in the Hs6 nomenclature is provided in Appendix Table A1, along with the maximum number of possible importer i – HS6 h combinations reported, the number of actually reported ih combinations in that year, and the number of newly reporting countries and products in each year. More on this below.

For some exercises we employ additional data. GDP per capita and population are collected from the World Bank. For the COMTRADE data, tariff (bilateral and product-level) and gravity variables are obtained from TRAINS and CEPII, respectively. For the US data, tariff and shipping costs data are reported as part of Imports of Merchandise.

2.2 Measures

Some of the measures used to capture market dynamism in the literature, such as market shares and prices, are straightforward. We next discuss some measures that are slightly more complicated.

2.2.1 Import Market Maturity

In various exercises we explore different notions of market age or maturity. The simplest of these is calendar year. The second, when using exporter-specific data, is the length of time that an exporter has been in a particular market. We will use the first two measures when looking at US imports data. The third – market age – is relevant for the COMTRADE World imports data and requires some explanation.

Different importers *i* begin importing specific HS6 products *h* at different points in time. An important challenge is distinguishing between an i-h import flow occurring for the first time, versus being reported for the first time. This is most clear when looking at the data in the 1990s where many countries are not yet reporting trade data using the HS6 nomenclature (nearly all are reporting SITC imports in this period).

To show this, we use the COMTRADE data to calculate the date of first import, the year in which a country first reports an importing a particular HS-6 product from any exporter.[4] Focusing on the first year that an importer reports any data, we can ask, how many of the available HS6 codes did it utilize in that first year, or in subsequent years? Appendix Figure A1 is a histogram of the share of i-h flows that appear in an importer's first or subsequent years of reporting HS6 data, and a similar histogram by calendar year. While many i-h flows appear for the first time after the first year of COMTRADE data, most of this is because the importer in question wasn't reporting any data until a later year. When they do begin reporting, roughly 75% of possible i-h flows are reported in year 1 of reporting for that importer. The remaining 25% are distributed across the subsequent two decades. (This phenomenon is much less important for the US data once we have used the Pierce-Schott algorithm to create consistent HTS10 codes across time.)

When using the US data and the full COMTRADE data sample we will simply use the calendar year, noting that year 2020 is distinct from 1991 because this market has experienced an additional 40 years of competition. For some of our COMTRADE analysis we will focus on an AFTER sample. It is the 25% of observations that an importer reports beginning in a year after they first begin reporting HS data to COMTRADE, and where no other country reports selling that HS product to that importer. For these observations we have information about the "age" of this market that are distinct from calendar year and a more precise look at the behavior of i-h markets in their first years. That is to say, market age varies significantly over both importers (for a given Hs6) and products (for a given importer) and we can use this variation to compare "new" and "old" markets in the same calendar year. Note that one limitation of this approach is that "first import" is not the same as "first purchase from any source", as a country could buy that HS6 good from its own producers. Still, we will show that dating markets in this manner has significant explanatory value for understanding the evolution of the import markets themselves, and changing dynamism within them.

2.2.2 Entry, Exit, and Churning.

For each importer *i* - HS6 product *h* - year *t* "market", we measure entry, exit and gross and net churning of varieties into the market, where a "variety" is sales from a unique exporting country within a

[4] Note that in a typical year, about 75-80% of potential importer-HS6 pairs have positive trade, e.g. in 2003, there were 133 importers reporting trade in 5001 HS6 categories, for a total of 665,133 potential "import markets" or importer-HS6 pairs. Of these, 527,640 or 79% reported positive trade with at least one exporter.

HS 6-digit product.[5] We explore unweighted and weighted measures. For a given market, the unweighted entry rate and exit rate are as follows:

$$Entry\ rate_{iht} = \frac{Nentrant_{iht}}{(N_{iht} + N_{ih,t-1})/2} \quad Exit\ rate_{iht} = \frac{Nexit_{iht}}{(N_{iht} + N_{ih,t-1})/2}$$

$Nentrant_{it}$ is the number of **exporters** that were not on the market of importer *i* in year *t*-1 but enter in year *t*. $Nexit_{it}$ is the number of **exporters** that were on the market of importer *i* in year *t*-1 but exit in year t. N_{it} and $N_{i,t-1}$ are the total number of exporters that trade with importer *i* in year *t* and *t*-1, respectively. We call the sum of entry and exit rates "Gross churning" and the difference in entry and exit rates "net churning".

The unweighted measures treat all exporters as equally important. An alternative approach is to assign each exporter some weight. We use the exporter's world export share of the product in that period, i.e. if China has a 50% world market share of some product and Zambia has a 1% share, then China is weighted 50X compared to Zambia for entry, exit, and churning measures.

More formally, define the world export share (*wes*) of each exporter *j* for product *h* in year *t* as $wes_{jht} = \frac{EX_{jht}}{\sum_{j \in J_{Ht}} EX_{jht}}$, where EX_{jht} is country *j*'s export to the world. J_{Ht} is the set of countries that export the product in year *t*. Then we can characterize weighted entry and exit rates as

$$weighted\ entry\ rate = \frac{\sum_{j \in J_{i,entry,t}} wes_{jht}}{(\sum_{j \in J_{iht}} wes_{jht} + \sum_{j \in J_{ih,t-1}} wes_{jh,t-1})/2}$$

$$weighted\ exit\ rate = \frac{\sum_{j \in J_{i,exit,t}} wes_{jt}}{(\sum_{j \in J_{iht}} wes_{jt} + \sum_{j \in J_{ih,t-1}} wes_{j,t-1})/2}$$

where $J_{i,entry,t}$ is the set of exporters that enter market *i* in year *t*. $J_{i,exit,t}$ is the set of exporters that were on import market *i* in year *t*-1 but exit in year *t*. J_t and J_{t-1} are the sets of exporters that serve importer *i* in year *t* and *t*-1, respectively.

[5] There is no firm-level exports data available that enables one to study large numbers of importers, products and time periods. In the model section we explore the properties of aggregating firm-level dynamics to country level exporting behavior.

Weighted gross and net churning are simply the sum and difference of weighted entry and exit rates. Note that in the trade literature many papers refer to the extensive margin of trade growth. Widely used measures of the extensive margin used for cross country comparisons (e.g. Hummels-Klenow) and trade growth (Feenstra) are closely related to our weighted net churning measure.

In Table 1 we report summary statistics on (weighted) entry, exit and gross and net churning measures, calculated for each of the four samples described above. A few things are clear from the data. First, there is substantial heterogeneity in all measures across the i-h-t observations. Second, net churning, very closely related to the commonly employed "extensive margin of trade" that calculates the number of varieties present in an import market, is positive at the mean in all samples, indicating net variety growth throughout the i-h-t sample. Third, gross churning is an order of magnitude or more larger than net churning in all samples. This is particularly relevant if supplier dislocation is itself a source of utility or productivity losses as in Baqaee et al 2023. Fourth, in the fourth panel (AFTER FIRST IMPORTER REPORT) , where we have an untruncated age variable, gross churning is much higher than in samples where the age variable is potentially truncated or where we are missing the early years of the flow.

3. Stylized Facts on Import Dynamism

In this section we characterize import market dynamism over calendar years and across import market maturity. Because we have over a half-million i-h markets that are distinctive in a variety of ways, we use simple descriptive regressions that relate each market dynamism measure to calendar year or maturity as follows

$$D_{iht} = \alpha_{ih} + \sum_{t=1993}^{2018} \alpha_t YRDUM_t + \epsilon_{iht} \qquad (1)$$

where D_{iht} include measures of dynamism as described below. We include importer-product fixed effects α_{ih} , as well as year dummies from 1993 to 2018. Many of our dynamism measures are undefined in 1991, so each α_t measures the difference in a measure of market dynamism for year t relative to 1992 for an i-h, averaged over all i-h. (Results from subsamples where we look only at cohorts present for all or a substantial portion of the calendar years are similar.)

We also relate changing dynamism to the maturity of the market in two ways. The first repeats equation 1, but replacing calendar year dummies with market age dummies based on market age as defined above. The second augments the vector of fixed effects for calendar year as in (1) by including market age and age squared as continuous variables. (We cannot include full vectors of calendar year and market age dummies because, once we include an i-h fixed effect, these vectors are collinear.)

$$D_{iht} = \alpha_{ih} + \sum_{T=2}^{28} b_T AGEDUM_{ihT} + \epsilon_{iht} \qquad (2)$$

$$D_{iht} = \alpha_{ih} + b_1 AGE_{iht} + b_2 (AGE_{iht})^2 + \sum_{t=1993}^{2018} \alpha_t YRDUM_t + \epsilon_{iht} \qquad (3)$$

The appendix also reports regressions where we experiment with simple time varying control variables for each market including importer population, per capita income, and MFN tariff rates.

3.1. Entry, Exit, and Churning

Figure 1 plots the coefficients of year dummies (α_t) in Eq. (1) using the US and the two COMTRADE world samples described above. In all samples, entry rates and gross churning relative to 1992 decline steadily over calendar years, and the magnitudes are substantial: in the full COMTRADE sample, entry and gross churning is 50% below 1992 levels by 2000, reaching 70% below 1992 levels in 2020. They are much larger in the AFTER FIRST REPORT sample where products appear after an importer begins reporting data, and where we have more accurate market dating. In contrast, exit rates are flat, or nearly so, in all samples, which explains why entry and gross churning track each other closely. Similarly, the change in the US entry and gross churn numbers numbers are smaller, but show a clear decline of 15% over calendar years, while exit rates also decline somewhat.

Figure 2 plots the coefficients of market age dummies from eqn (2) for the COMTRADE "After" sample where we have clear age data. The message is similar to that for calendar years: entry and churn track each other and decline relative to the early years in the market. Here the drop in exit rates is also more clear.

It is natural to ask whether the declining entry and churning rates we observe are due to market age or something trending that is specific to calendar year. To evaluate this, we estimate equation (3) for our AFTER samples and plot the year dummies in Figure 3. Here we see that, across all samples, entry and churning are declining with market age, but once we include age in the regression, market dynamism over calendar years looks markedly different, with entry rising, and exit falling over time. Loosely speaking this says, once we account for market age, markets experience greater entry dynamism in later calendar years. This is a very different pattern than is emphasized in the single country studies of business dynamism that focus only on calendar year changes.

How robust are these patterns? One approach is to include simple controls such as market size and tariff rates that might affect the profitability of import markets and therefore entry/exit dynamics. Appendix Figure A2 plots calendar year coefficients, and appendix Figure A3 plots market age coefficients, after including simple controls. The basic message of declining entry and churn and constant exit rates hold up after including controls.

A second approach is to ask whether the patterns we see here occur broadly or only in select subsets of markets. To evaluate this while producing statistics that are reportable in a compact way, we estimate a version of equation (2) in which market age dummies are replaced with a simple indicator variable that takes a value of 1 for years 2 through six that the market is active (in year 1, exit and churn rates are undefined, and entry rates always take a value of 2). We estimate this separately for every product (pooling over all countries) and for every country (pooling over all products). Table 2 shows the number of cases where the indicator is significantly negative, or positive, or insignificant (at 5% levels).

Looking across products, entry rates in the first five years are significantly higher than in remaining years in 4553 out of 5023 products, and significantly lower in only 101 products. Looking at countries, entry rates in the first five years are significantly higher than in remaining years in 132 of the 146 countries, and significantly lower for only 4 countries. Patterns are similar for gross churning, but considerably more mixed for exit rates. Including controls shifts makes more of these coefficients insignificant, but the basic message holds with or without controls.

Note that all our estimates include i-h fixed effects so we are already absorbing large differences across markets in important (time invariant) characteristics like market size, openness to trade, and competition policies. It may be that there could be much to learn from exploring variation across i-h markets, both in time varying and time invariant characteristics. We could, for example, look to explain

differential speeds of decline in entry rates by collecting detailed i-h-t varying data on barriers to entry. But the patterns here are sufficiently robust that we will eschew that approach in the remainder of the paper to focus on the first order fact that dynamism declines sharply with market age in almost all products and importing countries.

A simple explanation for declining market dynamism as measured in this way might be that markets are increasingly saturated. That is, once a large majority of possible exporters are present in the market entry rates would have to mechanically slow down as there are few source countries left to draw on. This is more problematic when measuring entry and exit at the level of countries as opposed to individual firms. To address this concern, we calculate a weighted measure of market saturation, with the weight being the world market share of each exporter for the given product ($wes_{jt} = \frac{EX_{jt}}{\sum_{j \in J_{Ht}} EX_{jt}}$).

$$saturation_{iht} = \sum_{j \in J_{iht}} wes_{jht}$$

Where J_{it} is the set of exporters that export product *h* to importer *i* in year *t*. Simply put, saturation = 50% for an i-h-t if exporters representing 50% of world sales of h in year t are actively selling on this market.

Using this measure, we can provide several exercises. First, we regress saturation on market age dummies in the manner of equation 2 and report these appendix Figure A4. It is clear saturation rises with market maturity but the changes relative to the first year are not large – from first year to peak saturation only rises by 15 percentage points. Second, we experiment with repeating the regressions reported in Figure 3 using only those i-h markets whose saturation never exceeds 50%. Results in the appendix Figure A5 and A6 show that the basic patterns of declining entry rates and declining gross churning in markets hold regardless of the level of maximum saturation of the market, and in fact the declines are larger in this sample. That is to say, to the extent that high levels of market saturation matters for the results, it weakens rather than strengthens them.

Still we might worry that a combination of imperfectly measured market age and the limitations associated with calculating entry and exit statistics from national trade data rather than firm level data might limit the conclusions one is willing to draw from looking at the entry, exit, and churning statistics alone. We next turn to the behavior of new and old exporters, comparing market shares and prices, as

well as measures of volatility and persistence in each, to see whether import markets can help us understand additional elements of business dynamism found in the firm-level literature.

3.2. Market share volatility

Entry and exit rates are one way to measure competitive pressures on firms, but firms may experience large changes in their competitive position while remaining in active in the market. To capture this intensive margin change in dynamism, the firm-level literature looks at the volatility of firms' market shares. We can provide a closely related measure by examining the over-time volatility of exporters' market shares in a particular i-h. This tells us whether reallocations across firms are occurring over time, and to see whether aging markets become more or less stable as the market ages.

For each exporter *j* on import market *i*, we calculate its share in *i*'s total import of each product *h*, $share_{ijht}$. Then, we calculate the absolute value of the year-on-year change in $share_{ijht}$ to measure volatility in shares.

$$share\_gr_{ijht} = \left|share_{ijht} - share_{ijh,t-1}\right|$$

We follow equation (2) above and regress market share volatility on i-h fixed effects and either calendar year (for the US market and whole WORLD) sample, or market age dummies (for the AFTER World sample). For compactness of presentation, we will show only the regressions without including simple controls; results including controls are very similar qualitatively and quantitatively.

Figure 4 focuses on calendar year or market age and shows pronounced declines in market volatility in all samples: market shares are becoming more stable as the market ages. The coefficients here are smaller than with entry/exit behavior. Focusing on the COMTRADE ALL SAMPLE, in the oldest markets the average year on year change in market shares is 2 percentages points smaller than in the initial years. To put this in perspective, the mean (median) market share in our sample is .087 (.0084), and the mean (median) year-on-year absolute value of the change in the market share is .048 (.005). The 2 percentage point drop in average market share change from the newest to the oldest markets is equivalent to reducing market share changes at the mean by 42% (=.02/.048). If we focus on the AFTER sample where we have the best information on market age, the fall in market share volatility is much more pronounced – 12 percentage points. In this sample market shares are larger (mean 0.316) and

more volatile (mean year on year change 0.20). The reduction in volatility over time is then equivalent to reducing market share changes at the mean by 60% (=.12/.20).

3.3 Comparing Young and Old Exporters

The literature on business dynamism that focuses on single-country studies shows a particular pattern to the loss of dynamism. It is not merely that entry is slowing and market share volatility is declining. It is that older firms are increasingly dominant, enjoying both growing market share advantages and productivity advantages relative to newer firms.

We can use imports data to measure relative import market shares within a given i-h import market. New entrants' import market share is measured as the combined market share of exporters who have been selling in this import market for less than 3 (or less than 5) years. Established incumbents market share is the combined market share of exporters who have been selling for six years or longer.

Single country studies of dynamism with firm level data often incorporate measures of firms' productivity in order to compare the productivity of frontier versus laggard firms or to compare productivity of entrants and incumbents, and generally shows a growing gap between older and newer firms. This approach has certain limitations inherent in calculating TFP, which are particularly problematic in cases where firms are producing a (potentially changing) mix of products and industrial classifications group together potentially dissimilar firms. We do not have firm level data, or productivity data that is i-j-h-t specific.

However, we can measure prices as trade unit values and do so for a great many more source countries and products than is possible in the dynamism literature. Unit values should presumably be correlated with firm productivity in a way that is similarly instructive.[6] We calculate the (trade share weighted average) of prices for new firms (less than three and less than five years old), and for incumbents in the market six or more years and express these in relative terms. We also provide a sample with first year entrants compared to all firms two or more years in the market. In our base

[6] There is a literature in trade that looks at whether unit values reflect only firms' productivity or also variation in product quality. We pass by this issue, noting only that we are focused on changes in prices for incumbents relative to new entrants. Differences in quality that are i-h market specific (reflecting, e.g. income levels of consumers) are absorbed into fixed effects, and level differences in quality across exporters that separate incumbents from entrants are absorbed into base year differences.

specification we focus on relative prices. In several specifications we also examine absolute price levels for the incumbent and entrant groups, as well as regressions that capture the time path of prices for each individual exporter.

We follow equation (2) above and regress these measures (market shares and prices) on i-h fixed effects and either calendar year (for the US market and whole WORLD) sample, or market age dummies (for the AFTER World sample). Note that we only include year dummies starting at year y+1, where y is the minimum number of years to have a complete cohort of firm (and similarly for age dummies in that sample.) For market shares n=1, for relative market share and relative prices n=5 (i.e. to get a complete set of importers who entered in age 1 through 5, the first market age dummy included is age 6.) In the case of relative market share and relative price measures, by expressing these as a ratio we also control for i-h-t varying information such as macro demand factors or openness to trade that affect all exporters. For compactness of presentation, we will show only the regressions for market age without including simple controls. The results using calendar year and including controls are very similar qualitatively and quantitatively.

Figure 5 shows the market shares of new entrants and those of established incumbent firms, regressed on market age dummies. In all samples we see a sharply declining market share for younger firms as the market ages and a rising market share for older firms. In the full sample, young firms' market share has fallen by as much as 70 percentage points relative to the initial year, while firms at least 6 years in the market see their market share rise by 20 percentage points. This basic pattern holds whether we look at calendar year or age of market, and holds when controlling for market size, income and tariffs, differing only in how quickly the rate of decline in the market share of young firms flattens out.

Figure 6 shows regressions of the relative price of new/older firms on calendar year and market age. Whether we use first year new entrants, or firms less than three or less than five years old, we see the same pattern. The price of new firms is initially similar to that of older incumbents (log relative price close to zero), but as the market ages, younger partners are selling at prices 40-60 percent higher than established firms.

An important question, for both welfare and for understanding what is generating these patterns, is whether relative prices of new/older firms are rising because the numerator is increasing or the denominator is decreasing. To see this, we separately regress numerator and denominator on year

or age dummies and share results in Figure 7. Here we clearly see that older firms are not getting better. Rather, the set of new firms are progressively worse. This suggests the possibility that the composition of these baskets is confounding our ability to understand price movements over time. To address this, we regress prices for each exporter (i.e. exploiting the full i-j-h-t variation) on dummies for calendar year or market age. These results (shown in Appendix Table A3 show that for a given exporter, prices are rising over time, but not as quickly as the average prices shown in Figure 7.

How then do we understand this pattern? One possibility is simply that entry into markets is ordered: the most productive exporters enter first and as time goes on and the market size expands, there is room for progressively less productive entrants. The challenge is that we cannot typically see productivity or prices for firms who have not entered a market. But using the trade data we can address this in two ways. First, we can regress enter order of exporters into a market i-h using the productivity distribution for that same market some years into the future. Second, recalling that importers are typically accessing only a small fraction of the available exporters of product h in a given year t, we can use the price distribution in other import markets for product h in the same year t.

Both sets of regressions (shown in Appendix Table A4) show a strong positive correlation between a country's position in the price distribution and the order of its entry into a market. Exporters with high prices enter later. While we are wary of giving this correlation any causal interpretation, it is consistent with a world in which productivity (and therefore prices) exhibit persistence, and where that productivity selects the order of entry for firms into particular import markets. We next address persistence directly.

3.4 Persistence:

An advantage of using imports data relative to single country firm-level data is that we can use market participation, market shares, and product prices to assess a longer and richer source of variation in the persistence of advantage in the face of common or idiosyncratic shocks.

We began our examination of persistence by measuring the duration of trade, specifically focused on how long exiting firms stay in the market prior to exit, and whether this is changing over time. There are a number of papers that characterize the duration of trade, beginning with the seminal work of Besedes and Prusa (). But this measure takes on new importance in light of the work by Baqaee et al () which shows first order welfare gains from supplier stability.

Our results comparing young and old firms shows rising market share and price advantages for older firms, conditional on continuing to participate in the market. However, this misses an important selection effect if one of these groups experiencing a growing likelihood of exit. In each year we calculate two numbers: the average time in market i-h for all exporters who exited that market in year t, and the average time in market i-h for all exporters in a market i-h who continue selling. We then express this as a ratio, and regress this on calendar year or market age dummies. Note that if we focus on only the numerator or denominator, we would build in a mechanical correlation between the age of the market and the average year in the market. That is, in year 3 of our data, the maximum duration of continuing sellers is 3 years, whereas in year 30 the maximum duration is 30. However, by expressing duration of exiters relative to duration of stayers, we sweep out this mechanical correlation.

Figure 8 shows the relative duration of exiters to stayers in each calendar year, and by market age. Relative to the baseline year, the relative duration of exiters falls by 50 percent over calendar years and by 25 percent over market age. Put another way, with each passing year exiting firms are younger and younger relative to the rest of the market and stayers are older and older.

We next examine market share volatility at the level of the exporter. Recalling our results above, we showed that the average volatility within market i-h falls considerably over calendar year and market age. We can address the source of that decline by running the same regression, but incorporate the age of the exporter (i.e. the length of time they have been active in market i-h). This enables us to incorporate fixed effects that control for the average volatility for that exporter-product, and for any trends in volatility for that product and time period. Given the size of our data, and the computational burden to calculate j-h, h-t fixed effects alongside age dummies, we group our age dummies into six: ages 6-10, 11-15, 16-20, 21-25, 26-30. Table 3 shows these results: market share volatility drops by two percentage points relative to the baseline for the exporters longest in the market. This is the same order of magnitude as the overall decline in volatility corresponding to calendar years. Put another way, markets are aging because the firms who increasingly dominate them are aging. And as they age, their market shares are increasingly stable.

Why exactly are these market shares increasingly stable? An obvious possibility is that prices are increasingly stable. To address this, we look at year to year variability in prices for a given ijht. We begin by regressing (ln) prices on their one year lag, then capture the residuals from that regression. We then regress the residuals on calendar year or market age. Results are reported in Table 4. The top panel of the table reports the results of the lagged price regression: the coefficient is around 0.3,

depending on the specification, indicating a modest degree of persistence. The top panel reports the result of regressing residuals from this price regression on year or age dummies, as well as on age dummies for the exporter's age. Here we see a clear pattern: price residuals are getting much smaller as markets age, and this effect is concentrated in the oldest exporters. The world sample incorporates a rich set of exporter-product-time fixed effects, and so controls for any shocks that are hitting this exporter-product in the rest of the world, price variability drops 25% relative to baseline for the oldest exporters.

Now, one might worry that this variability is simply noise in the data. Unit values are considerably noisier than scanner data or the carefully vetted price series constructed by, e.g. the Bureau of Labor Statistics when calculating the CPI. Perhaps it is the case that older exporters are simply better at accurately reporting both value and quantity, yielding increasingly accurate and therefore less variable unit values. Two things argue against this. One, these regressions control for price variability in the rest of the world with j-h-t fixed effects. And two, the price variability matches the market share variability shown in the previous results.

Still, it would be useful to see whether and how these prices respond to measurable cost shocks. Happily the trade data allow us to see shocks that hit both young and old exporters and to see how their delivered prices respond to these shocks. We join our trade data to detailed tariff data that provide importer-exporter-product-time varying shocks to delivered prices, constructed as the unit value multiplied by (1 + ad-valorem tariff rate). We regress delivered prices on the ad-valorem tariff rate, and an interaction between the tariff rate and the age of the exporter.

Results are in Table 5. When using a level specification, the coefficient on the tariff rate shows that between 60 and 90 percent of the tariff rate is passed onto consumers for exporters in their first year in market i-h. (When using a first differences specification, around 27 percent of the one-year change in tariffs is passed onto one-year change in prices.) This passthrough elasticity drops dramatically for the oldest exporters, as low as 20 percent passthrough for the oldest firms in the US market. For the world as a whole, the passthrough elasticity is cut in half, from around 85 percent to 42 percent. Put another way, early in their tenure exporters experience price changes induced by tariff shocks and pass most of these changes onto consumers. Late in their tenure, exporters absorb most of the shock, leaving delivered prices much less variable.

It may be tempting to interpret these changes in pass through as the result of increasingly dominant world firms. However, several things argue against this interpretation. First, average import market shares in this sample are less than 10 percent. Second, even though entry is slowing, it is still occurring and a rising number of aging exporters means their average market share is actually falling. Third, we control for market share in the regressions. Fourth, these effects are dramatically larger than what one would expect from rising market share alone. So something other than just exercise of market power in pricing is at work here.

Of course, this is just one cost shock that exporters experience, and it is not immediately clear why older firms are better able to absorb these tariff cost shocks. But the difference between older and younger firms provides a new insight into a unique source of advantage as markets age: resilience in the face of shocks.

3.5 Summarizing the Evidence

We now summarize the facts on market dynamism developed in this section. We show that there is a substantial reduction the rate of entry and variety churning over calendar years and as markets age. Exit rates, in contrast, change little over time or age. These patterns are robust to inclusion of country-product fixed effects, time varying measures of market size and openness, and occur in 90% of products and 90% of countries. The fall in market dynamism is also reflect in market shares, market share volatility, and prices. Older markets exhibit less market share volatility on average, smaller market shares for newer entrants, and a pronounced and growing price disadvantage for newer entrants compared to incumbents.

We then develop new stylized facts on persistence that import data are uniquely capable of delivering. We show that the duration of time in market for exiting firms relative to stayers has fallen substantially over time and market ageMarket share volatility declines observed at the market level are concentrated primarily among firms that have been in the market longest. Price variability also declines with market age, with declines concentrated primarily among the oldest firms. And finally, the oldest firms are significantly better able to absorb cost shocks without passing them onto consumers in the form of delivered prices.

4. Which theories match the facts? (to be revised…)

In this section we explore models in which the age of a market affects the dynamics of market entry. Our goal is to match the facts described in section 3, and then to identify testable implications of the model in addition to these facts. The models we have chosen incorporate elements suggested in the business dynamism literature as potential explanations for falling dynamism, including rising barriers to entry, rising substitutability between firms, and changes over time in the distribution of productivity shocks. A difference is that we explore these effects in the context of equilibrium entry and exit of firms, which we show changes the implications of these changes significantly.

Our basic setup is similar to that of a Melitz model, in which firms differ in their productivity, and face fixed and marginal costs of entering foreign markets. We depart from Melitz in a number of ways. First, we experiment with time varying productivity processes so that the time path of firm productivity, prices, and entry, depend on initial conditions. Two, because our focus is on import market dynamics rather than the general equilibrium of firm supply, we examine these productivity dynamics in the context of a small open economy, sales into which do not affect wages in supply markets. Three, since we want to match country-product level facts from the preceding section, we will aggregate firm level export data into national aggregates for our analysis.

Suppressing product subscripts, consumers in import market i have CES utility over varieties of a product with elasticity of substitution σ_t. When selling into this import market, firm g faces iceberg trade costs of τ_{gt}, a fixed cost of entry f_t, and productivity φ_{gt} . Firms will enter this market if revenues R_{gt} and profits are sufficient to cover fixed entry costs

$$\frac{R_{gt}\tau_{gt}}{\sigma_t} > f_t \qquad (6)$$

We assume both productivity and trade costs are a composite of a constant and a time varying component. Trade costs are ad-valorem and common to all firms g in a country c, and depend on a time invariant component dictated by factors like geography, and time varying components associated with factors like tariffs.

$$\tau_{ct} = (1 + TAR_{ct})(1 + DIST_c^{\delta}) \qquad (7)$$

Current productivity is firm specific and depends on lagged productivity and an innovation term.

$$\varphi_{gt} = \left(\varphi_{g,t-1}\right)^{k1}\left(\varepsilon_{gt}\right)^{k2} \quad (8)$$

We are primarily interested in the nature of the productivity dynamics needed to generate the patterns we see in Section 3. This model nests several approaches in the literature. If k1=0 and ε_{gt} is a random draw, we have the canonical Melitz model. If k1=1 and k2=0, we have a "permanent productivity" Melitz model where productivity remains unchanged for all time. If $0 < k1 < 1$, and ε_{gt} is a random draw, we have an AR(1) productivity process of the sort emphasized in the business dynamism literature. Within these settings we can explore additional explanations for rising concentration found in the literature, including increased fixed costs of entering markets, increasing the substitutability between firms, and varying the relative strength of the persistent and time-varying components of productivity (i.e. increasing the value of k1 relative to k2, or changing the variance of ε_{gt}). Finally, we examine models in which the time varying innovation term depends on market exposure.

Many of the papers in the business dynamism literature that explore simple models of the sort above are able to derive analytical solutions. However, they tend to abstract from the problem of entry and exit, and instead focus on intensive margin changes in the relevant aspects of dynamism. However, entry/exit is an essential part of what we wish to explain, and as we show below, titrating this margin significantly impacts all others. Because we want to explore time dependence in the market, the solution for which firms enter and exit the market, the prices they charge, the evolution of their market shares, and the overall price index, all will depend on past values and initial draws. This means that there is not a unique solution to the model, but rather one that depends on the draws obtained by firms in the initial and subsequent periods. This becomes particularly relevant as the number of firms gets small and so reliance on the large number properties of statistical distributions becomes problematic.

To make progress, we use simulation. When assessing each model, we assume there are 1000 potential entrant firms, divided between 100 exporting countries, in proportion to the size of that exporter. We generate an initial productivity distribution as described below, solve the initial equilibrium of potential prices and revenue for each firm, and using equation (6) determine which firms enter and the first period equilibrium. We then redraw values of φ_{gt} according to different models as described below. We iterate on this for 30 periods (as in our data) to complete one round of the simulation, where we also allow market size to grow annually. We recover from this process the key summary statistics that correspond to the exercises in section 3. We then restart the simulation with a

different initial set of productivity draws and repeat for 30 periods, then repeat 50 times for each model.

To explore similarities between the model simulation and the actual empirics we run the same regressions as reported above and compare, e.g. how entry rates change with age of the market. We can think of the distinct model runs as essentially creating a separate i-h market 50 times, corresponding to that markets we analyze in the data. Each will be distinct in levels of key variables, just as in the empirics, but the regressions will recover the average evolution of markets as they age.

Using this setup we can explore productivity dynamics from canonical models, and explore departures from each. In the results below we experiment with three main models with variants in each.

1. No-memory Melitz: k1=0, k2=1 and productivity each period is iid, φ_{gt} pareto, mean 1.
    a. Base model: all parameters constant, τ_{gt}=1. Exogenous death rate 5%.
    b. Increase fixed costs of entry over time.
    c. Increase the elasticity of substitution over time.
    d. Increase the variance of the pareto distribution draws over time (holding mean fixed)
2. Permanent Melitz: k1=1, k2=0. Initial productivity is iid, φ_{gt} pareto, mean 1
    a. Variants as above
3. AR(1). Initial and subsequent productivity draws are iid, φ_{gt} pareto, mean 1
    a. Base model k1 = 0.5, k2=(1-k1), τ_{gt}=1.
    b. Increase k1 (weight on persistent productivity) from 0.5 to 0.8.

We collect the data generated from the simulations and run the regressions as in Section 3. In Table 4 we summarize the findings of the key models. For each of 6 stylized facts (number of entrants, entry rates, gross churning, market share volatility, market share of old/young firms, and prices for old/young firms) we capture the change in that variable over time: increasing (+), decreasing (-), or no change (0). In addition, we show changes over time in the price index generated by the model and the threshold productivity needed to enter the market.

The first row summarizes our stylized facts: While the number of firms is rising over time, entry rates, gross churning, and the volatility of market shares are all falling over time, while the relative price of old/young firms is falling, and the market shares of old/young firms are rising.

The memory-less Melitz model can explain none of these facts, which is not at all surprising since that model was parameterized to generate a stable productivity equilibrium with constant entry/exit rates. Notably this model generates extremely high levels of churning: because productivity for each firm is independently draw in each period, an unusually high draw (and successful entry) in one period is typically followed by lower draws in subsequent periods and exit. This also means that the model generates no "old" (age greater than 6 years firms) to use for studying the evolution of old v. young firms prices and market shares.

We can tweak some of these results by changing model parameters over time. Increasing fixed cost barriers to entry reduces the number of competing firms, which is not at all surprising given equation (6). We can also see this in the price index (rising because the number of competitors is falling) and the threshold level of productivity needed to enter. Interestingly, increasing f actually raises entry *rates* (the number of entrants at time t, relative to the average number of firms at t and t-1). This is a bit subtle: both the number of new entrants and the total number of firms are falling over time, it is just the case that the total number is falling slightly faster. Further, it increases the churning rate: with higher barriers to entry it is more likely that a firm with high productivity in period one will fall below threshold productivity in period 2 and exit. This also creates more market share volatility. With fewer active firms and higher average sales per firm, swinging from high to zero sales and back again across a sequence of draws generates more volatility over time. Once accounting for entry/exit dynamics, increasing barriers to entry gets every stylized fact wrong.

Authors focused on business dynamism have argued that an increase in substitutability across firms should yield greater market concentration – for a given productivity distribution, higher sigma means that highly productive firms get a larger market share. And inspection of equation (6) suggests sigma should have an effect similar to rising barriers to entry. However, the effects are subtly different, with increased substitutability having no effect on the number of firms, reducing entry rates and churning, and having little effect on volatility. Why are these effects different? The key is that the revenue term in equation (6), as well as the price index also depends on the elasticity of substitution, with the net effect being that the threshold productivity needed to enter the market is unchanged.

Finally, we experiment with increasing the variance of the pareto draw (holding mean constant) in each period. Increasing the variance yields more firms with very high productivity and very low prices, lowers the price index and the threshold productivity to enter the market. This decreases the number of firms, and increases both churning and market volatility, as the increasingly large year to year movements in firm productivity make them more likely to cross the entry threshold in a given year and for revenues swings from positive to zero and back to get larger.

Clearly, the memory-less Melitz model is not a suitable candidate for explaining the behavior of young versus old firms. But it has been a useful starting place for understanding how incorporating equilibrium entry/exit thresholds into models of firm dynamics can our ideas of what changes in barriers to entry and substitutability might do to measurable quantities for firms who do enter. We turn next to a model that does have memory to see how our insights change.

The Permanent Melitz model does much better in generating facts. The key aspect of this model is that firms are ordered from most to least productive in their initial draw, with the most productive entering the market. Increases in market size enable more firms to exceed the profitability cutoff (equation 6) and enter. However, because firms at the margin are always less productive than the firms who previously entered, this causes a slowing entry rate over time. Further, this pattern generates many of the key facts comparing longer lived incumbents relative to new entrants. As time goes on new entrants are worse and worse, which means their relative price rises and market share falls. What this model does not generate are any of the key facts about imperfect persistence in prices, market shares, or trade duration. This suggests the need for a hybrid model.

The most common model invoked in papers on business dynamism has an AR(1) process in productivity in the manner of equation (8), with $0<k1<1$, and $k2=1$. We choose $k1=k2=0.5$ to begin, then experiment with raising k1 and lowering $k2=1-k1$, as well as raising F and sigma as in the Melitz experiment. Our base AR(1) model matches none of the stylized facts. Because the productivity distribution is stationary, there are no changes in entry behavior, volatility or prices. We are able to observe the behavior of old versus young firms: old firms tend to receive good initial draws that persist in the face of adverse subsequent shocks and so have lower prices than young firms. But their relative price is not changing over time.

This insight begins to change as we allow the weight on innovation to fall, which is equivalent to creating greater persistence in the model. Threshold productivity falls, and with it the number of firms

rises. The reason is subtle. Taking the product of a series of iid pareto shocks results in a distribution whose standard deviation shrinks substantially relative to both the mean productivity and the productivity threshold, and this effect grows stronger the more shocks one hits the firm productivity distribution with. The price index falls (because it overweights extremely productive firms, which steadily disappear from the market) and revenue is more evenly distributed across firms. That enables more firms, with lower productivity, to enter the market. Prices for both young and old firms are rising – old firms because the repeated pareto draws are bringing their productivity down, young firms because the increasingly lax entry standard allows in worse firms. In this particular parameterization, the productivity threshold is falling fast enough that young firms prices are rising faster than old, and their market shares falling.

Returning to constant innovation weights, and manipulating entry barriers or substitutability in the AR(1) model enables us to extend the Melitz model insights to think about old v. young firms. The effect of rising barriers to entry has effects similar to Melitz rising f: increasing the productivity threshold, shrinking the number of firms, and raising the entry rate. Market shares are rising for old firms relative to young firms, despite the fact that there is not change in their relative prices. This is not particularly robust: it has more to do with the paucity of old firms when the persistence weight is 0.5.

Finally, we turn to the learning model. This model duplicates all of the key features of the stylized facts, with the exception of an increasing number of firms competing in the market. Entry rates, churning and market volatility fall over time. The prices of old firms fall relative to young firms, and their market shares rise.

It is also useful in context to note that we have been discussing model predictions on firm level entry and exit dynamics while our data from section 3 is based on country level data. We next aggregate the data to the level of countries (recall: firms are assigned to countries in proportion to those countries share in world GDP) and reproduce the statistics. Here a country "enters" if any one firm in that country enters and exits at year t if all firms present at year t-1 are gone from the market at time t. Measures of volatility, market shares, and prices are (revenue-weighted) aggregates of all firms from that country active in the market. We see in the last row of Table 4 that all the stylized facts except for rising overall entry are present in the learning model.

We pause here to focus on two key mechanisms and discuss their empirical magnitudes. First, learning increases the threshold productivity needed to enter the market. We can describe this effect

by examining how high productivity must be to enter in year 1 compared to a market that has largely matured by year 20. Recalling that firms are drawing initial values from a pareto distribution with mean 1, in the initial year the minimal productivity to enter was 1.13, which is in the 86th percentile of productivity. By year 20 the minimal productivity to enter was 1.44, which is in the 99.9% percentile of productivity.

Second, the price index. In this model the price index is a sufficient statistic for consumer welfare, as it captures the number of entrants, average prices, and their distribution (as a quantity weighted average with elastic demand the lowest priced varieties are the largest part of the consumption basket). The price index at the firm level falls from (mean across simulations) .28 to .18, and at the country level from .285 to .175.

6. Conclusion

Using imports data for 146 countries and 5000 products from 1991-2018 we develop new approaches to systematically measure market dynamism that bear a strong resemblance to findings from single country case studies. Declining dynamism, whether measured by entry rates, market share volatility, market shares for young firms, prices for young firms relative to old firms, all show sharp reductions over time and with the age of the market. These patterns are more strongly related to market age than to calendar year, and hold in the vast majority of countries and products. This suggests measures of declining dynamism are a robust pattern of market aging that holds regardless of the regulatory or market organization details of individual products or countries.

We develop additional stylized facts on persistence and volatility enabled by the richness of data coverage provided by imports data. We show that the duration of time in market for exiting firms relative to stayers has fallen substantially over time and market age. Market share volatility declines observed at the market level are concentrated primarily among firms that have been in the market longest. Price variability also declines with market age, with declines concentrated primarily among the oldest firms. And finally, the oldest firms are significantly better able to absorb cost shocks without passing them onto consumers in the form of delivered prices. Further, the patterns are broadly consistent with a view that the drop in dynamism is a potentially positive outcome, particularly if

increased market stability lowers costs of supplier dislocation as recently emphasized by Baqaee et al 2023.

We show that these results are broadly consistent with models in which firms that establish early strong positions due to initial productivity advantages are able to consolidate those advantages and enjoy further productivity growth relative to entrants. There are further testable implications from these models, in particular the notion that time invariant advantages such as those conferred by geographic proximity can be magnified over time while smaller time varying disruptions such as those induced by tariff changes have much weaker effects in older markets. Both are borne out in the data.

In our view, this way of viewing market age through the lens of import markets has the potential to significantly augment the literature on dynamism that focuses on single country case studies and requires highly detailed firm level data. It also suggests that work focused on trade determinants, especially trade costs, should pay close attention to market age when measuring responsiveness to shocks. Finally, the model implication that established firms lock in market shares has potential implications for the success of late entrants into markets, and the prospect for export-led development in lagging countries.

References

Akcigit, Ufuk, Akgunduz Yusuf Emre, Cilasun Seyit Mumin, Ozcan-Tok Elif, Yilmaz Fatih, 2020 “Facts on Business Dynamism in Turkey” European Economic Review 128.

Akcigit, Ufuk and Ates, Sina 2021 “Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory” *American Economic Journal: Macroeconomics* 2021, 13(1): 257–298

Akcigit, Ufuk, Alp Harun, and Peters Michael, 2021, “Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries” American Economic Review 111(1): 231-275.

Alon, Titan, Berger David, Dent Robert, Pugsley Benjamin 2018 Older and Slower: The Startup Deficits Lasting Effects on Aggregate Productivity Growth” Journal of Monetary Economics 93: 68-85.

Baqaee, David, Burstein Ariel, Duprez Cedri, Farhi Emmanuel, “Supplier Churn and Growth: A Micro-to-Macro Analysis” NBER Working Paper 31231.

Bijnens, Gert and Konings, Jozef 2020, Declining Business Dynamism in Belgium, “Small Business Economics 54: 1201-1239.

Covarrubias, Matias, Gutierrez German, Philippon Thomas, “From Good to Bad Concentration? US Industries over the Past 30 Years”.

Decker, Ryan, Haltiwanger John, Jarmin, Ron, and Miranda, Javier 2020 “Changing Business Dynamism and Productivity: Shocks versus Responsiveness” American Economic Review 2020, 110(12): 3952–3990

Feenstra, Robert 1994, “New Product Varieties and the Measurement of International Prices” American Economic Review 84(1) 157-177.

Romer, Paul 1994, “New Goods, Old Theory and the Welfare Costs of Trade Restriction” Journal of Development Economics 43(1) 5-38.

Table 1: Summary statistics on trade dynamism measures (by sample)

	(1) mean	(2) p25	(3) p50	(4) p75	(5) min	(6) max	(7) N
WORLD ALL							
Gross churning	0.592	0.014	0.112	1.134	0	2	13877676
Net churning	0.0537	-0.045	0.0001	0.089	-2	2	13877676
Entry rate	0.323	0.001	0.017	0.196	0	2	13877676
Exit rate	0.269	0.0001	0.011	0.132	0	2	13877676
WORLD AFTER							
Gross churning	1.287	0.203	2	2	0	2	1053806
Net churning	0.075	-1.333	0	1.897	-2	2	1053806
Entry rate	0.681	0	0.068	1.937	0	2	1053806
Exit rate	0.606	0	0.023	1.557	0	2	1053806
US							
Gross churning	0.326	0.03	0.099	0.318	0	2	335785
Net churning	0.004	-0.043	0.0001	0.051	-2	2	335785
Entry rate	0.165	0.005	0.03	0.118	0	2	335785
Exit rate	0.161	0.004	0.027	0.112	0	2	335785

Notes: This table reports the summary statistics on the weighted measures of entry, exit and gross and net churning, calculated for each of the four samples described in Section 2.2.1: ALL includes all iht observations. AFTER includes observations where a product h appears for the first time in a year after the importer i reports data.

Table 2: Robustness on trade dynamism trend over tenure by importer and product

	Positively significant	Positive but insignificant	Negatively significant	Negative but insignificant
By product				
Entry rate	4553	354	10	101
Exit rate	1666	2130	286	936
Gross churning	4495	341	37	144
By country				
Entry rate	132	2	7	4
Exit rate	112	8	20	5
Gross churning	139	1	3	2

Note: This table lists the number of cases where the coefficient of early age dummy (from running Eq. 2) is significantly negative, or positive, or insignificant (at 5% levels).

Table: Market Share Volatility by Exporters' Age

Table 4: Price residuals and Exporter Age

dep var: residuals from reg of price on lagged price

	US imports			World imports		
Exporter Market share	-0.7865***		-0.7844***	-0.1754***		-0.1762***
	(0.0043)		(0.0043)	(0.0006)		(0.0006)
Exporter Age 6-10		-0.0518***	-0.0225***		-0.0719***	-0.0728***
		(0.0026)	(0.0026)		(0.0003)	(0.0003)
Exporter Age 11-15		-0.0511***	-0.0064**		-0.1095***	-0.1103***
		(0.0029)	(0.0029)		(0.0004)	(0.0004)
Exporter Age 16-20		-0.0657***	-0.0084***		-0.1403***	-0.1407***
		(0.0031)	(0.0031)		(0.0004)	(0.0004)
Exporter Age 21-25		-0.0695***	0.0001		-0.1922***	-0.1924***
		(0.0034)	(0.0034)		(0.0005)	(0.0005)
Exporter Age 26-30		-0.1086***	-0.0221***		-0.2447***	-0.2451***
		(0.0039)	(0.0039)		(0.0006)	(0.0006)
N	4864002	4864002	4864002	128034692	128034692	128034692
R2	0.1232	0.1168	0.1232	0.1523	0.1533	0.1539
Product-time FE	YES	YES	YES			
Exporter-product-time FE				YES	YES	YES

Notes: Dependent variable is residuals from regression of $\ln(p_{ijht}) = \widehat{0.3}\ln(p_{ijht-1}) + \varepsilon_{ijht}$ for US and World (all sample). US sample has exporter-product-time variation; World sample has importer-exporter-product-time variation.

Table 5: Tariff Passthrough by Exporter Age

	US imports				World Imports			
	Delivered price (levels)		Delivered price (first differences)		Delivered price (levels)			
Tariff	0.6079***	0.6021***	0.2739***	0.2767***	0.8626***	0.9539***	0.8582***	0.9476***
	(0.0205)	(0.0205)	(0.0329)	(0.0329)	(0.0027)	(0.0029)	(0.0027)	(0.0029)
Tariff*exporter age	-0.0130***	-0.0127***	-0.0184***	-0.0188***	-0.0145***	-0.0116***	-0.0143***	-0.0112***
	(0.0009)	(0.0009)	(0.0021)	(0.0021)	(0.0001)	(0.0002)	(0.0001)	(0.0002)
Exporter age	0.0251***	0.0251***	0.0253***	0.0253***	0.0075***		0.0075***	
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)		(0.0001)	
Exporter Market Share		-0.1844***		-0.2007***			0.0859***	0.1330***
		(0.0056)		(0.0085)			(0.0008)	(0.0008)
constant	3.1076***	3.1203***	-0.3788***	-0.3650***	2.4932***	2.5894***	2.4851***	2.5780***
	(0.0015)	(0.0015)	(0.0015)	(0.0016)	(0.0003)	(0.0002)	(0.0003)	(0.0002)
N	5563350	5563350	4998019	4998019	136899975	133995139	136899975	133995139
R2	0.8472	0.8472	0.0385	0.0386	0.7342	0.7588	0.7342	0.7589
FE								
Exporter-product	YES	YES	YES	YES				
Importer-exporter-Product					YES	YES	YES	YES
Exporter-product-time							YES	YES

Table 6: Changes in trade dynamism measures over tenure (by simulation model)

	Entry Rate	Number of firms	Gross churn	Market Shares		Prices	
				Volatility	Old/Young	Old/Young	In
Stylized Facts	**-**	**+**	**-**	**-**	**+**	**-**	
Melitz							
Base	0	0	0	0			
Rising F	+	-	+	+			
Rising sigma	-	0	-	0			
Rising variance	0	-	+	+			
AR(1)							
innovation weight = 0.5	0	0	0	0	0	0	
Innovation weight falling	-	+	-	0	+	-	
Rising F	+	-		+	+	0	
Rising sigma	0	-		0	0	0	
Learning							
Firms	-	-	-	-	+	-	
Aggregated to countries	-	-	-	-	+	-	

Figure 1: Trade dynamism trend over calendar year (weighted)

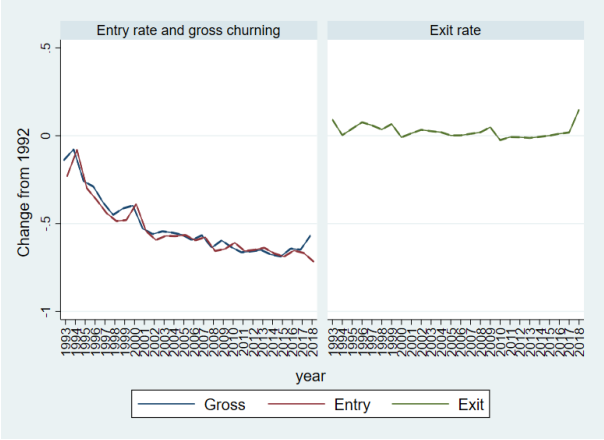

World ALL

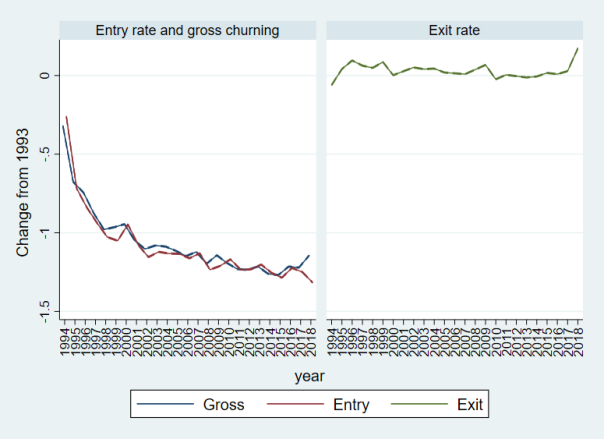

World AFTER first import

US Imports

Notes: ALL includes all iht observations. AFTER includes observations where a product h appears for the first time in a year after the importer i reports data.

Figure 2: Trade dynamism trend over market age (weighted)

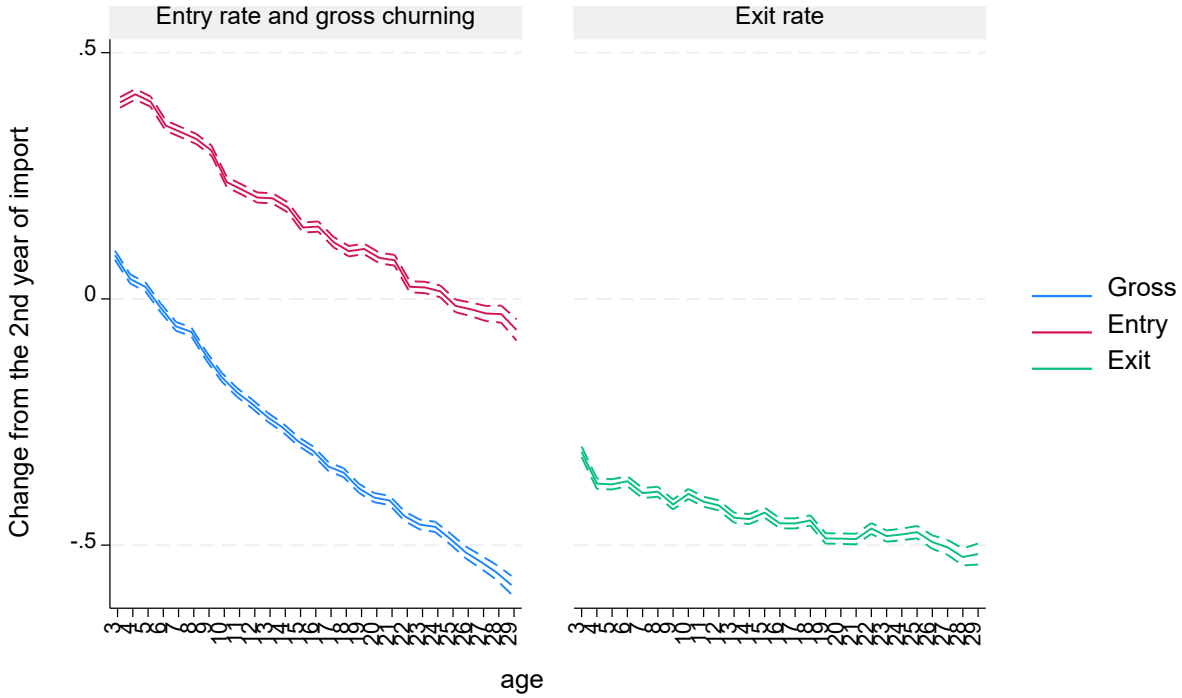

Notes: . AFTER includes observations where a product h appears for the first time in a year after the importer i reports data.

Figure 3: Trade dynamism trend over calendar year (including both age and year in the regression)

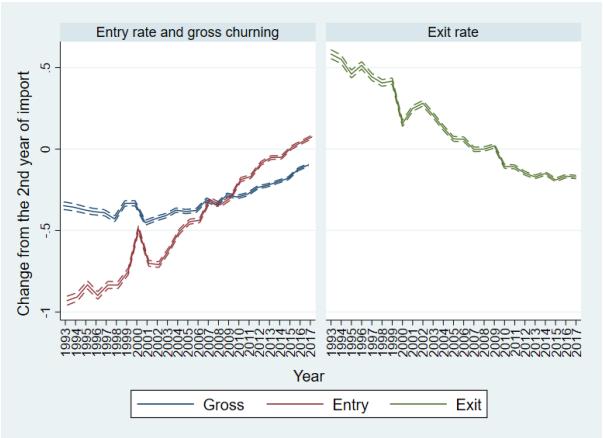

Notes: AFTER includes observations where a product h appears for the first time in a year after the importer i reports data.

Figure 4: Trend of market share volatility over years and market age

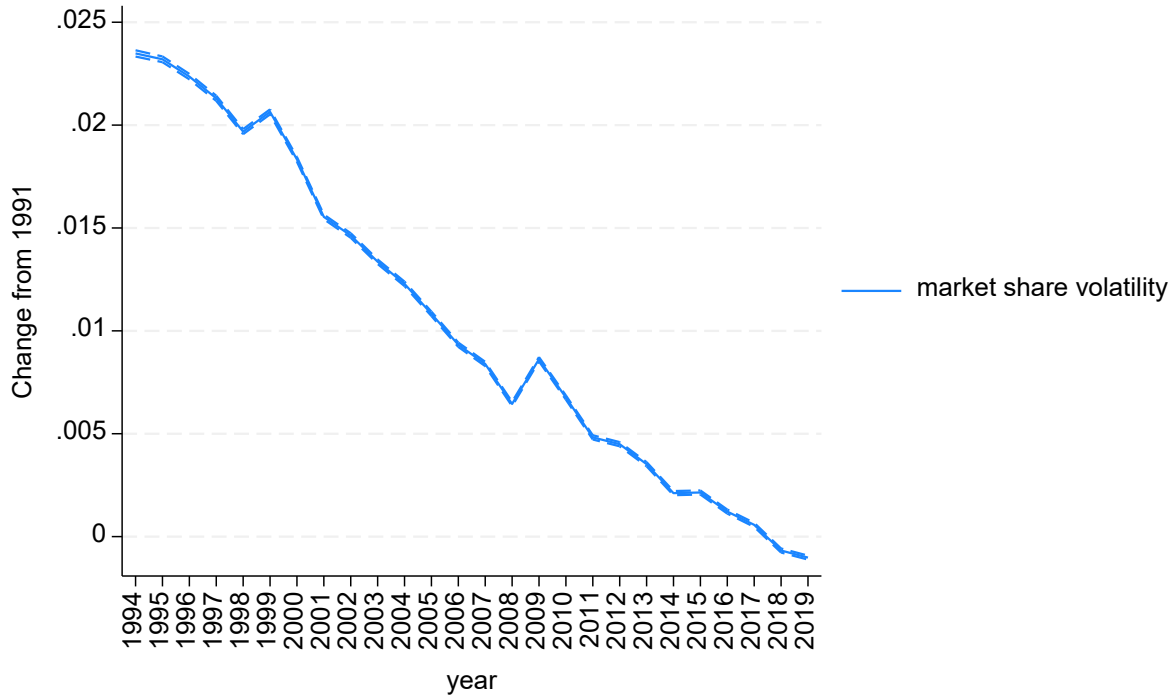

Notes: ALL includes all iht observations. AFTER includes observations where a product h appears for the first time in a year after the importer i reports data.

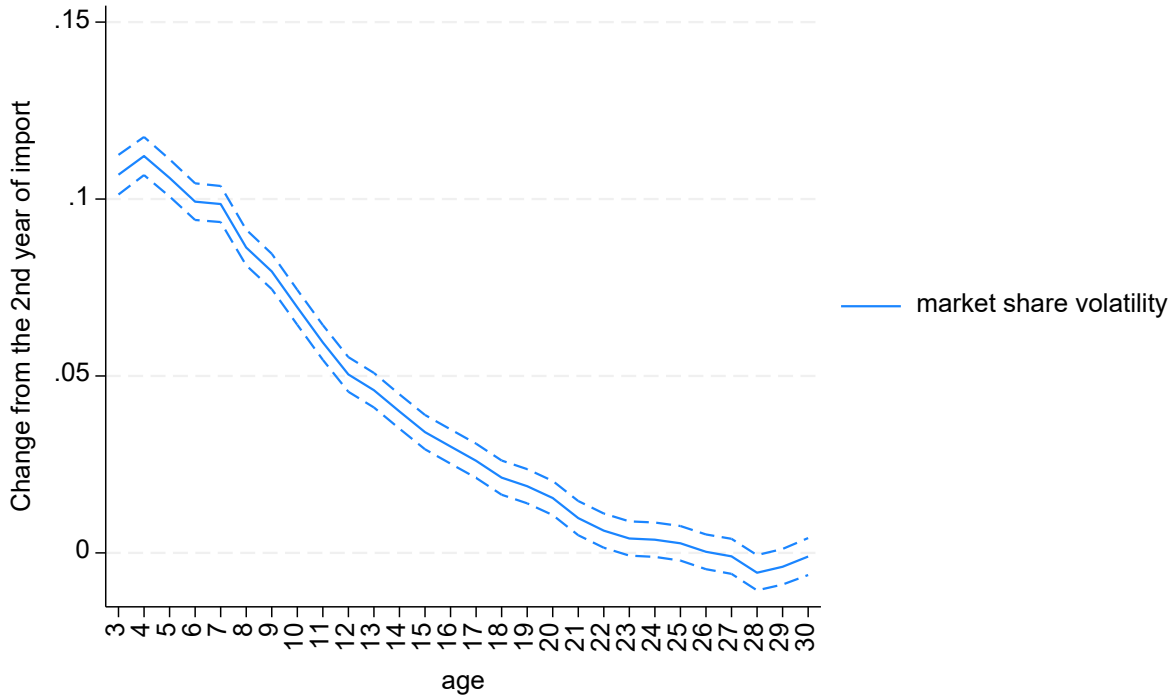

Figure 5: Market share of young and old partners over years and market age

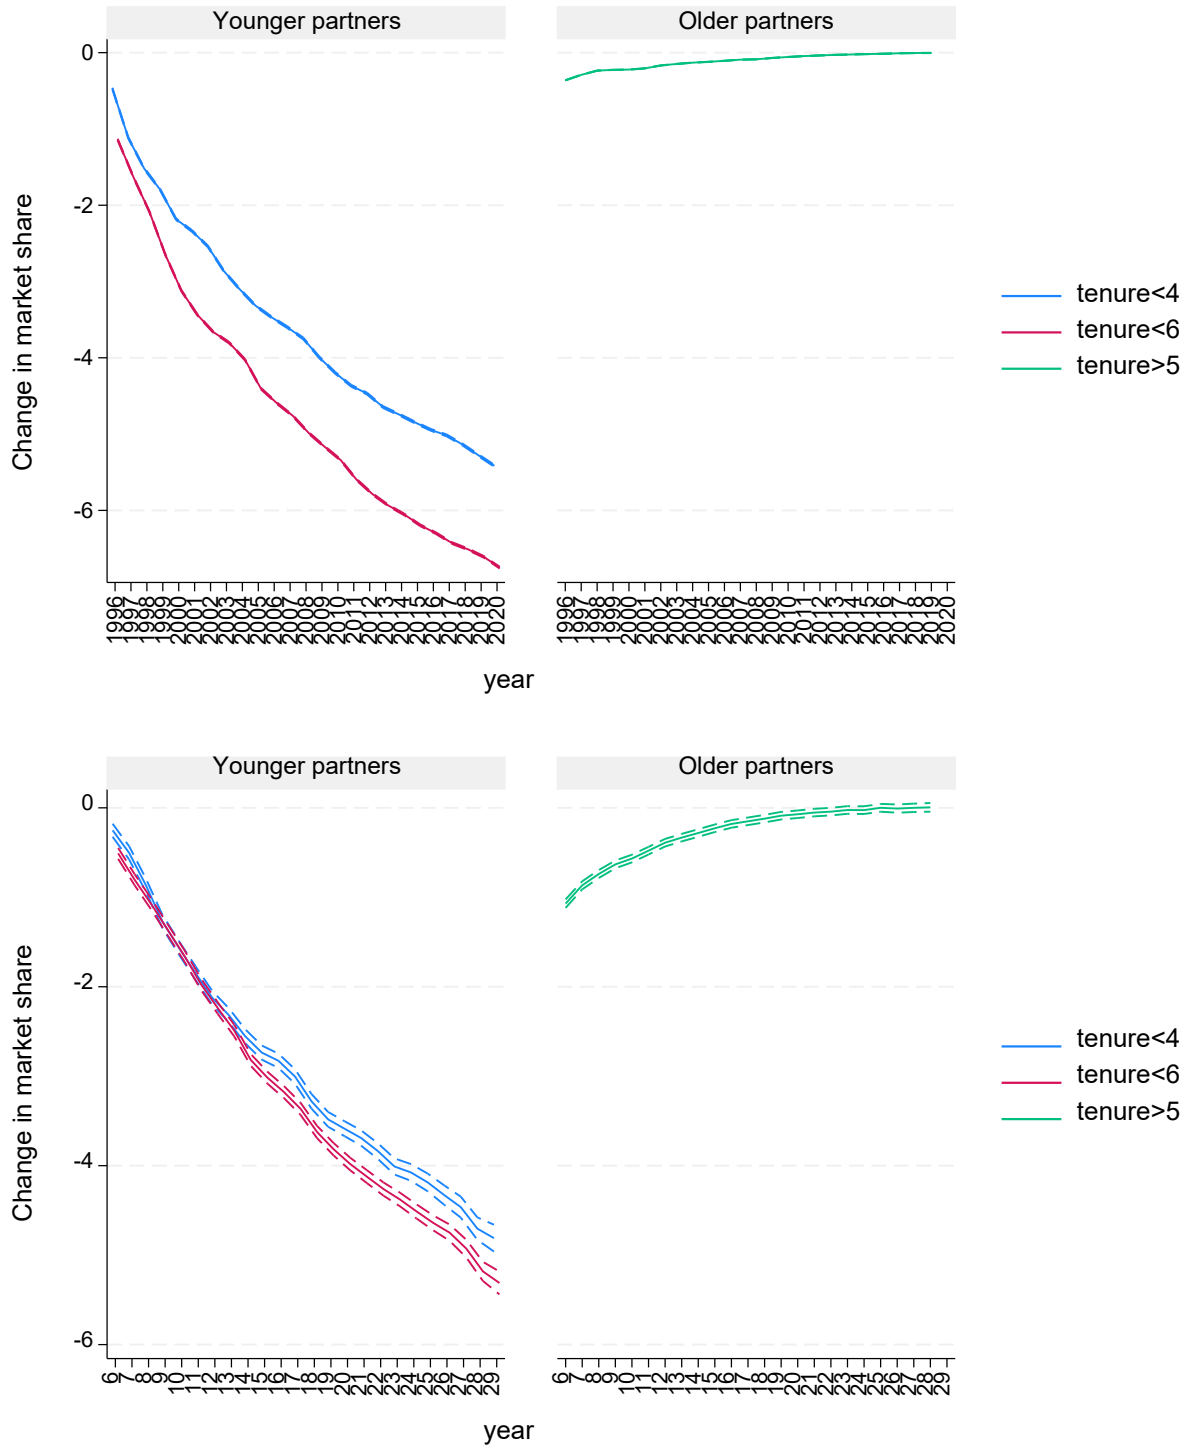

Figure 6: Relative prices: young/old firms

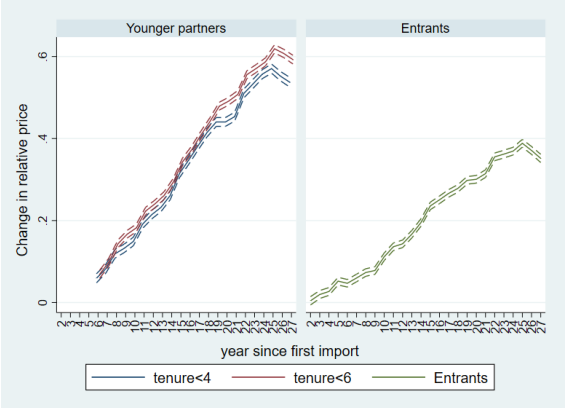

ALL

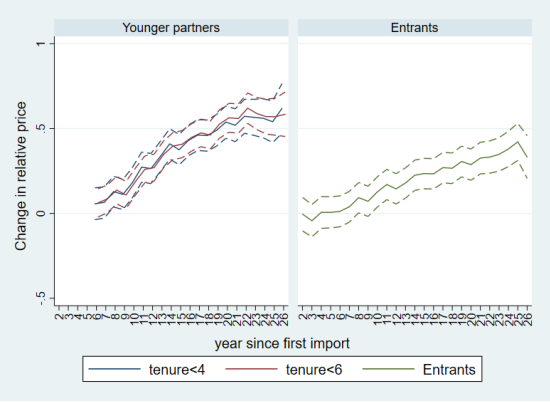

AFTER first import

Notes: Three relative prices are included. First, the average price of varieties (i.e. exporters) that have been on the ih market for less than 4 years relative to the average price of varieties that are more than 5 years old on the same market. Second, the average price of varieties that have been on the ih market for less than 6 years relative to the average of price of varieties that are more than 5 years old on the same market. Lastly, the average price of entrants in each versus exporters that were on the market in the year before.

Figure 7: Absolute Prices

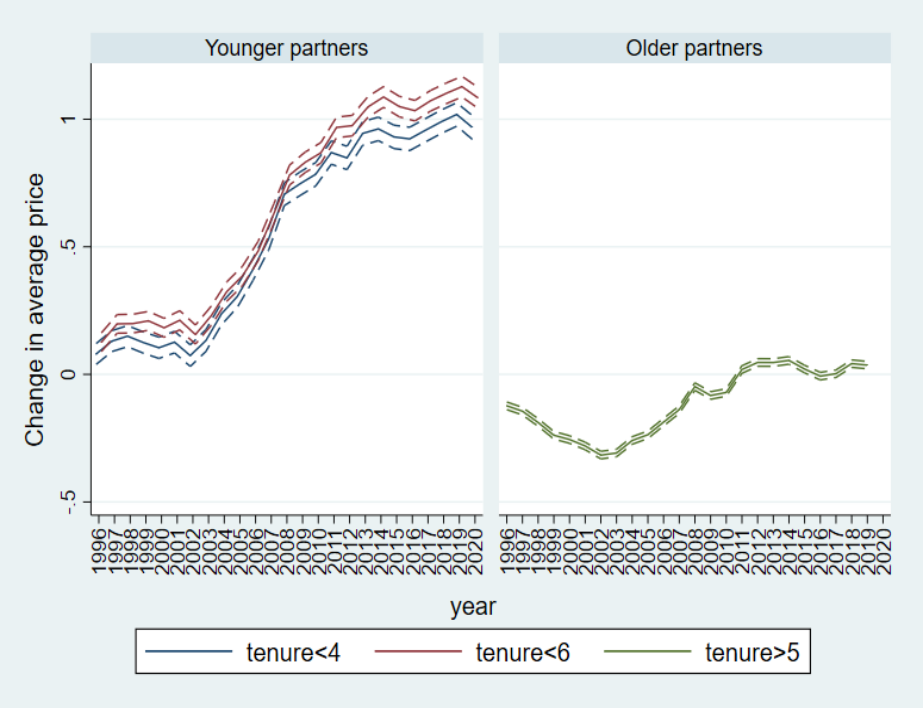

FIGURE 8 Duration

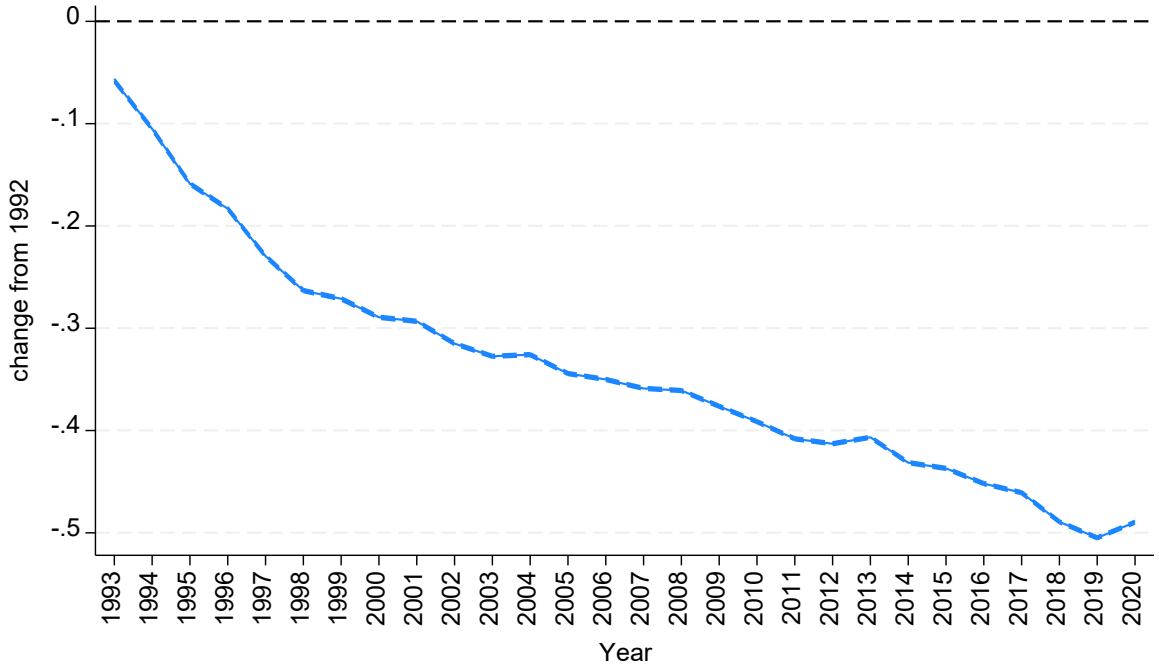

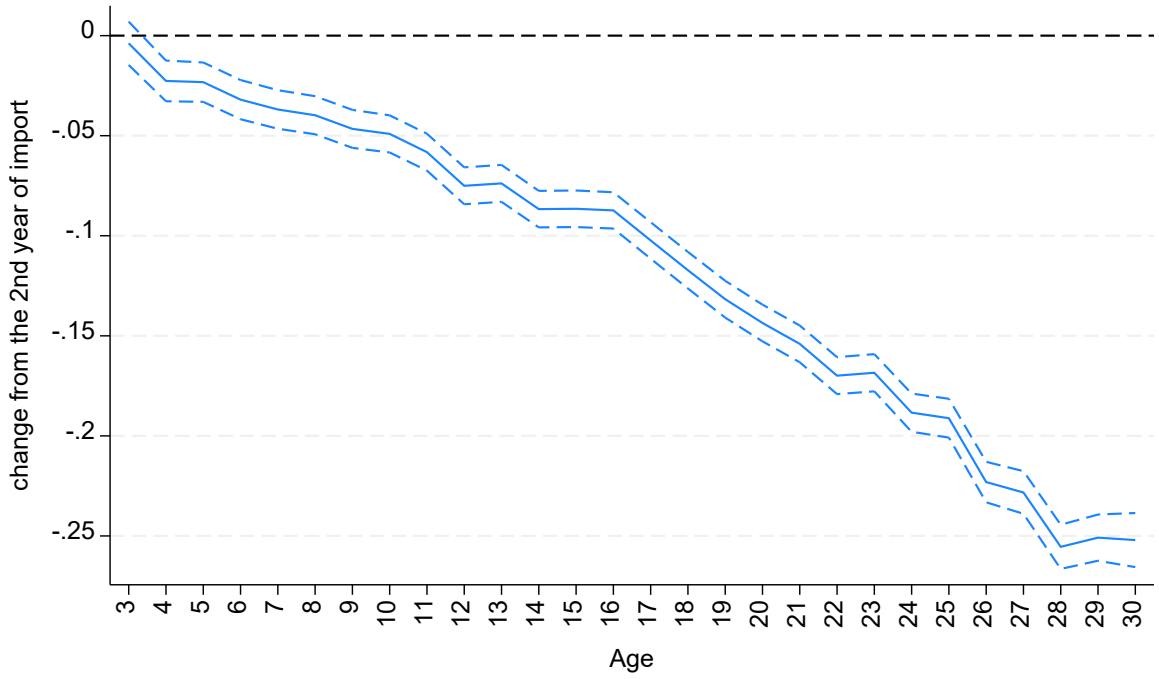

Notes: Figure plots time in market of exiting exporters relative to time in market for stayers, over calendar years and over the age of the import market

Appendix 1: Market Dating

Different countries begin importing specific HS6 products at different points in time. To show this, we calculate the date of first import, the year in which a country first reports an importing a particular HS-6 product from any exporter. Note that in a typical year, about 75-80% of potential importer-HS6 pairs have positive trade, e.g. in 2003, there were 133 importers reporting trade in 5001 HS6 categories, for a total of 665,133 potential "import markets" or importer-HS6 pairs. Of these, 527,640 or 79% reported positive trade with at least one exporter.

An important challenge is distinguishing the difference between an i-h import flow occurring for the first time, versus being reported for the first time. This is most clear when looking at the data in the 1990s where many countries are not yet reporting trade data using the HS6 nomenclature (nearly all are reporting SITC imports in this period).

Focusing on the first year that an importer reports any data, we can ask, how many of the available HS6 codes did it utilize in that first year, or in subsequent years? Appendix Figure A1 is a histogram of the share of i-h flows that appear in an importer's first or subsequent years of reporting HS6 data, and a similar histogram by calendar year. Roughly 75% of possible i-h flows are reported in year 1 of reporting, the remaining are distributed across the subsequent two decades.

COMTRADE includes both importer and exporter reports of trade flows. This means that we can, in many cases, identify a first year of import for an i-h even when that importer has not yet begun reporting Hs6 trade data. E.g. Russia begins reporting HS6 imports in 1996, but by looking at data from exporting countries reports, we can see that they were purchasing products such as Iron, Yarn spun and heterocyclic compounds in years 1991-1995.

For our analysis we are interested in the behavior of market dynamism variable both across calendar years and across the maturity or "age" of the market. The "age" of each market i-h is just the number of years between an observation i-h in year t and the year that imports first occurred for an i-h market. To determine the year of first import for each i-h, we use the following procedure.

1. If an importer i begins reporting in 1991 and an i-h appears in 1991, we treat 1991 as year 1.

2. If an importer i begins reporting in year t>1991, and its first report of an i-h flow occurs in that same year t, we look for instances where any exporter reports selling h to i in a year prior to t.

   a. If there is no such earlier report, then we treat year t as year 1.

   b. If there is such a report in a prior year, we use that prior year as year 1.

3. If an importer i begins reporting in year t, and its first report of an i-h flow is in a year y occurring after t, then we treat y as year 1.

This approach has some limitations. The most obvious is that some of the i-h flows in case 1 were presumably initiated prior to 1991. Another is that in case 2a, we might miss some instances of flows that initiated in early years because neither the importer nor their export partners were reporting data in earlier years. In these cases, the age variables we examine will be measured with error, and likely truncated for flows beginning in 1991. A third is that in case 2b, we know that the market began before the importer's first report of data, but we don't have complete information on trade (varieties, market shares, prices, and measures derived therefrom) until the importer begins reporting data.

We are interested in how market dynamism behavior changes over calendar year and with market maturity. If our aging procedures are excessively noisy, attenuation bias will presumably weaken any estimated relationship between dynamism and age. We might also expect to see a different relationship between dynamism and age in data samples where we are more likely to have a truncated age measure or where we are missing data from the earlier years in the sample.

To address potential biases we will split our samples for all our analysis into four sets.

1. ALL. All i-h-t observations, using age variables as above. Total: 13.9 million obs.

2. SAME AS FIRST IMPORT REPORT. Cases 1 and 2a above: All i-h-t observations where a product h appears for the first time in the first year that importer i reports data, and there is no exporter report of an early flow. Total: 5.6 million obs

3. BEFORE FIRST IMPORT REPORT. Case 2b. All i-h-t observations where an exporter reports selling h to importer i in a year before the importer begins reporting data. Note that here we only correct the age of the market, we don't replace the trade data itself, and so market dynamism data will be missing for early market ages. Total: 7.2 million obs

4. AFTER FIRST IMPORT REPORT. Case 3: All i-h-t observations where a product h appears for the first time in a year after the importer i begins reporting data. This is a smaller sample, but it is also the cleanest in our ability to age. Total: 1 million obs.

Market maturity varies significantly over both importers (for a given Hs6) and products (for a given importer). We can use this variation to compare "new" and "old" markets in the same calendar year. Note that one limitation of this approach is that "first import" is not the same as "first purchase from any source", as a country could buy that HS6 good from its own producers. Still, we will show that dating markets in this manner has significant explanatory value for understanding the evolution of the import markets themselves, and changing dynamism within them.

Appendix 2: Additional Tables and Figures

Table A1: Number of importers and products (Potential and actual)

Year	Importer	Product	Importer-products (actual)	Importer-product (potential)
1991	38	5036	150418	191368
1992	51	5038	205900	256938
1993	59	5037	242020	297183
1994	79	5037	321419	397923
1995	93	5037	364178	468441
1996	102	5037	394805	513774
1997	106	5037	414377	533922
1998	111	5036	432734	558996
1999	115	5035	449364	579025
2000	131	5031	511596	659061
2001	133	5030	527382	668990
2002	132	5019	524324	662508
2003	133	5001	527640	665133
2004	134	4997	530392	669598
2005	132	4998	543538	659736
2006	134	4992	550587	668928
2007	139	4992	557127	693888
2008	137	4943	541525	677191
2009	136	4908	532404	667488
2010	140	4878	546197	682920
2011	138	4877	540695	673026
2012	138	4877	540811	673026
2013	138	4859	546462	670542
2014	139	4869	542335	676791
2015	136	4860	533230	660960
2016	135	4869	530402	657315
2017	131	4789	522147	627359
2018	117	4773	468488	558441

Notes: This table lists the number of importers, products, importer-products that are recorded in the data set. The last column lists the number of importer-products that could potentially show up in the data set if every country imports all products in any given year.

Figure A1. Share of first import (ih) over tenure and calendar year

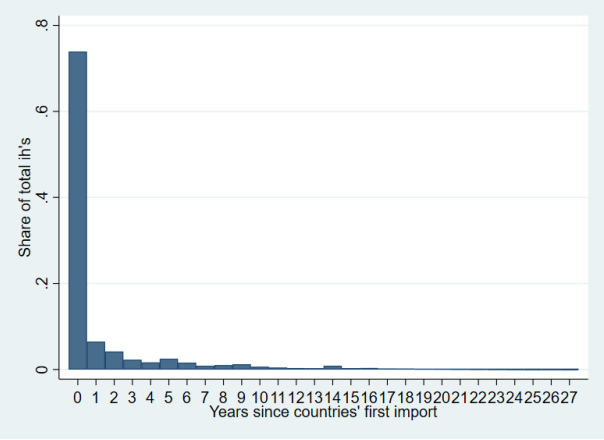

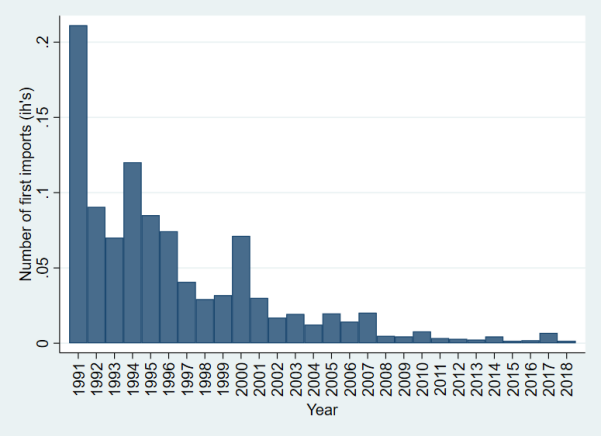

Figure A2: Trade Dynamism trend by calendar year (with controls)

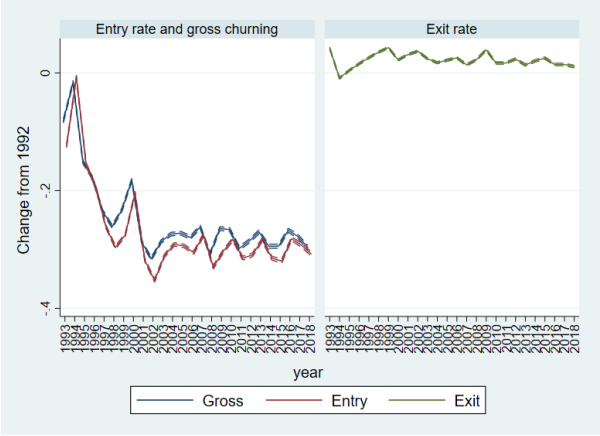

ALL

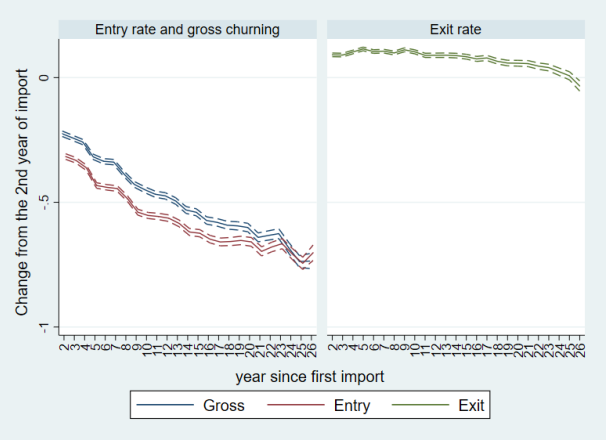

AFTER first import

Figure A3: Trade dynamism trend over age (with controls)

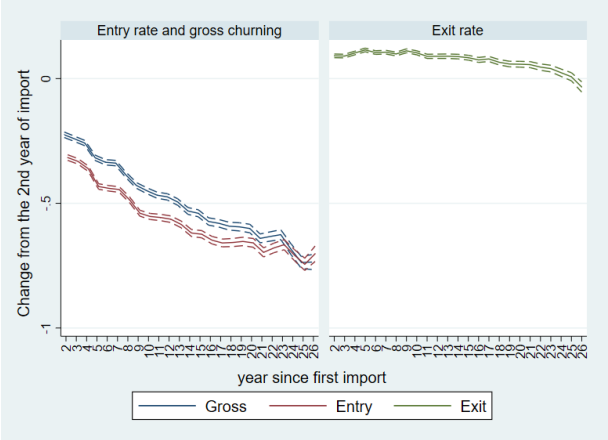

Figure A4 Trend of market saturation over age (weighted, no control)

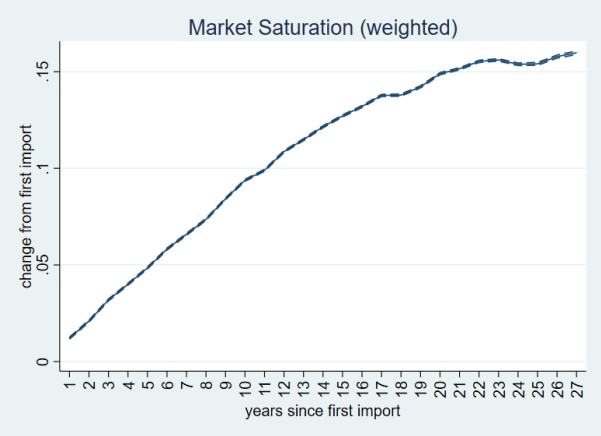

Figure A5: Trade dynamism trend over calendar year (no control, max saturation<0.5).

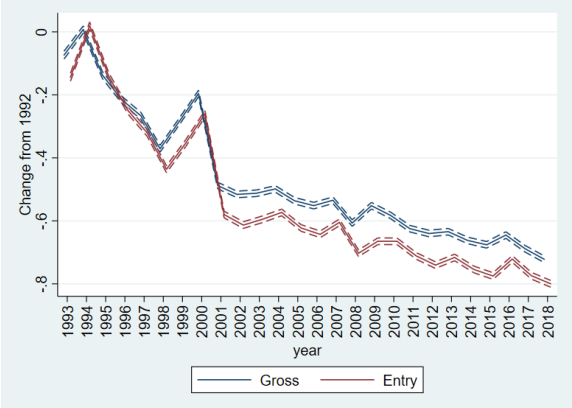

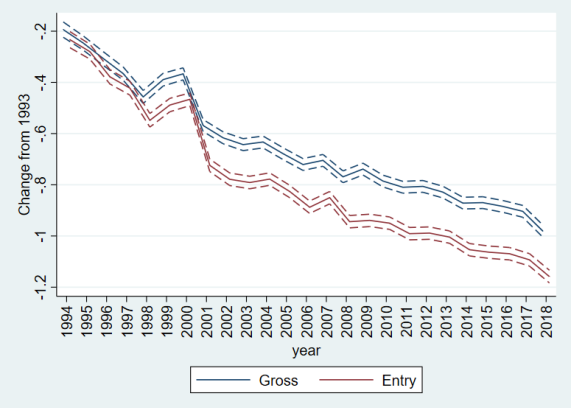

Figure A7: Trade dynamism trend over calendar year

(including both tenure and year the regression, with controls)

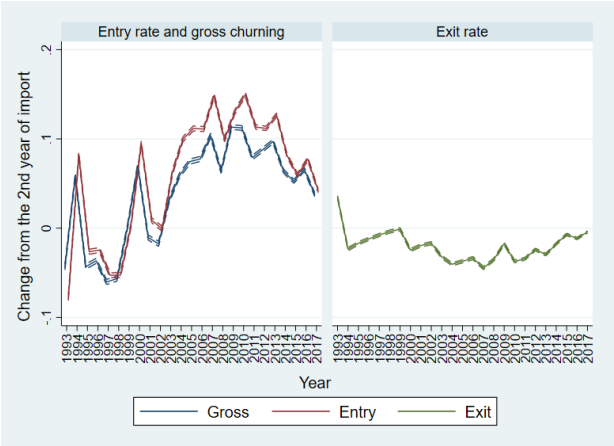

ALL

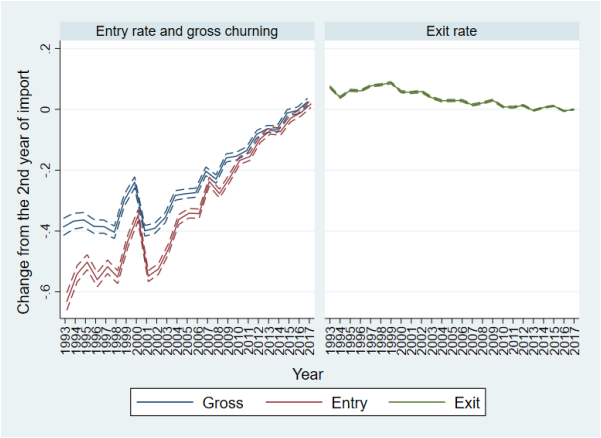

AFTER import

Notes: ALL includes all iht observations. AFTER includes observations where a product h appears for the first time in a year after the importer i reports data.