NBER WORKING PAPER SERIES

ON THE UBIQUITY OF DECLINING BUSINESS DYNAMISM

David Hummels
Kan Yue

Working Paper 32637
http://www.nber.org/papers/w32637

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2024, Revised October 2025

We thank seminar participants at Purdue, Colorado, Tennessee, North Carolina State, the European Trade Study Group, and the Midwest International meetings, and particularly Mario Crucini, Ryan Decker, Andrew Greenland, James Lake, Andrew Liang, John Lopresti, Keith Maskus, Jim Markusen, Kim Ruhl, Georg Schaur, Anson Soderbery, and Chong Xiang for helpful comments on earlier drafts. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

On the Ubiquity of Declining Business Dynamism
David Hummels and Kan Yue
NBER Working Paper No. 32637
June 2024, Revised October 2025
JEL No. D22, F14, L23, O40

ABSTRACT

Recent work documents declining business dynamism in establishment data, and points to difficulty in reallocating resources across firms as an important cause. Using import data, we document declining market dynamism that bears a striking resemblance to establishment-level findings and holds for the vast majority of countries and products. Entry rates and the reallocation of market shares fall over time. Young exporters experience rising prices, falling market shares, and reduced skewness in growth rates relative to older exporters. A heterogeneous firm model with additive productivity shocks and growing expenditures, but no reallocation of resources across exporters, can generate all these facts.

David Hummels
Purdue University
and NBER
hummelsd@purdue.edu

Kan Yue
Xavier University
Department of Economics
yuek@xavier.edu

1. Introduction

Recent work documents declining business dynamism over the last few decades in the United States and a few European countries.[1] Among the key factors highlighted in this literature, 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 across firms and churn in the labor market have slowed. And firms' employment growth and exit probability have become less responsive to profitability shocks.

The consequences of, and the remedies for, declining dynamism depend primarily on the reason dynamism has declined. A series of important recent papers has focused on shocks versus reallocation. Put simply, a lack of dynamism could result either from a reduction in the variance of productivity shocks hitting firms, or from rising adjustment costs that prevent reallocation of market share and inputs toward more productive firms.

Decker et al (2020) provide evidence that the variance of productivity shocks in the US has risen, as measured by the increased variance of TFP level and shocks across firms and a decreased persistence in productivity in the time series. That paper and Cooper et al (2024) point to an inability to reallocate labor across firms as the culprit behind declining dynamism. This is bolstered by evidence from Akcigit et al (2021) showing that an inability to delegate managerial functions (i.e. effectively reallocate managerial labor) prevents higher productivity firms from expanding. Related, Hopenhayn et al (2022) point to slowing population growth and the associated labor supply constraints as an important factor slowing firm entry rates and concentrating growth among older firms.

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- or establishment-level micro data within a single country, either explaining aggregate macro facts or in some cases exploiting cross-industry or cross-regional variation.[2] 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.

[1] Akcigit and Ates (2021) provide an excellent overview of the literature and its key findings.

[2] 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). These papers and Covarrubias et al (2020) draw some limited cross-country comparisons.

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 we show bear a strong resemblance to those seen when studying domestic firms only.

We use US import data at the HTS10 digit level from 1991-2020 as well as world-wide bilateral import data at the HS6 digit level from 1991-2020. We show for the US and the world as a whole:

- Entry rates for new exporters and the reallocation of market shares across exporters decline significantly over time.
- The behavior of young and old exporters differs markedly, with those differences becoming larger over time. Market shares for young exporters fall relative to longer-tenured incumbents, while product prices for young exporters rise relative to longer-tenured incumbents over time.
- Fast growing young exporters are less common over time – the skewness in exporter growth rates among young exporters is declining quickly.
- The variance and skewness of price shocks is rising over time.

All these facts closely parallel establishment-level findings on business dynamism in the US, and are consistent with a sharp decline in the dynamism of import markets over time. We also introduce a new fact not previously investigated in the establishment level data: when looking at changes in prices for a given importer-exporter-product, the magnitude of price shocks is inversely related to (lagged) market share and decreasing in exporter age. This fact will be key to our model explaining the ubiquity of the phenomenon.

These patterns 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. It is also unlikely that explanations for business dynamism that focus on slowing labor reallocation across firms (Decker et al 2020, Cooper et al 2024), or on demographics and the aggregate supply of labor (Hopenhayn et al 2024), will tell us much about ubiquitous patterns of import market dynamism we highlight, because we are comparing outcomes across different exporting countries.[3]

[3] International migrants represented 2.9% of the world population in 1990 and 3.6% in 2020, meaning that reallocation of labor across countries is a small portion of employment but is increasing. (UN World Migration Report 2024, Table 1.) Further, sales of a specific product to a specific destination tend not to involve large shares

Rather than resource reallocations, one needs to look either at the consumption side of the market or at the evolution of productivity itself. We focus on productivity, and show how a simple extension of Melitz (2003) allows the age of a market to affect the dynamics of productivity and market entry and can generate all the business dynamism facts described above. There are two key features of this model.

One, productivity shocks are normally and identically distributed across firms, but shocks are additive rather than multiplicative, similar to the Philippon (2022) model for aggregate TFP. Unlike multiplicative shocks, additive shocks are consistent with our evidence on the distribution of price innovations. Additive shocks imply that the impact of a given shock is smaller the more productive is a firm and the larger its market share. Moreover, even though the shocks are identically distributed in each period, the accumulation of normally distributed shocks over time leads to a rising variance in productivity, consistent with the facts demonstrated in Decker et al (2020). Combined with a fixed cost of entry, firms that experience a series of positive shocks are more likely than firms receiving a series of negative shocks to stay in the market. Even though the model has no learning per se, the average productivity of incumbents rises over time.

The second key feature is that markets are growing over time. Combined with initial differences in firm productivity, this leads to an entry order effect. Early entrants have high productivity while growing market size enables later entrants to succeed despite lower productivity.

The relative strength of these two effects (rising productivity variance, market growth) determines whether market dynamism, as captured by standard measures in the literature, is rising or falling. Faster growth leads to a stronger entry order effect and more rapid declines in entry rates and market share reallocations because faster growth allows progressively worse firms to successively enter. We confirm that entry order effects are an important feature of world trade data, and that a number of our business dynamism measures in imports data drop off most quickly in the fastest growing markets.

This paper adds to several literatures. The underlying questions in this paper and many of our measures are inspired by the literature on business dynamism using US firm- and establishment-level data, with important contributions by Autor et al (2017), Akcigit et al (2020), Akcigit and Ates (2021),

of aggregate employment for a given exporting country, and age demographics differ markedly from other countries.

Akcigit et al (2021), Alon et al (2018), Bijnens and Konings (2018), Cooper et al (2024), Covarrubias (2020), Davis and Haltiwanger (2014), Decker et al (2014), Decker et al (2016), Decker et al (2017), Decker et al (2020) and Hopenhayn et al (2024).

Relative to this literature we show that patterns of declining business dynamism revealed in establishment data are also found in imports data for the large majority of importers and products. We provide a model of this process that abstracts from many of the theoretical channels highlighted in that literature. Shocks to productivity are not changing over time, though their impact may be as a function of the changing distribution of incumbents' productivity in the market. There is no role for changes in market power and no change in the substitutability of products. Most importantly, there is no role through which factor markets, either ease of reallocation or aggregate tightness, contribute to declining dynamism in this model.

To be clear, we are ***not*** saying that labor reallocation or aggregate labor supply is unimportant as a source of declining US business dynamism at the establishment level. Rather, we are saying that dynamism declines are worldwide in scope, and occur in contexts where labor reallocation or aggregate labor supply constraints are unlikely to be a major factor. It may be that the mechanism we highlight could also play some role in slowing dynamism across establishments in the US, but we have no direct evidence for that.

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. Recent work by Baqaee et al (2023) demonstrates that churning (changing source relationships) can have large welfare and productivity 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 falling entry rates and increasing incumbency advantages cause the rate of churn to fall over time.

A third large literature beginning with Melitz (2003) uses heterogeneous firm models of monopolistic competition to study how changes in trade openness lead to entry and exit and the reallocation of resources between firms within a given producing country. In the canonical model, productivity has no memory. Other papers such as Aw et al (2011) and Sanghamitra et al (2007) model richer firm productivity dynamics to describe, from a single exporting country's perspective, which firms

enter or exit exporting (to any market). We use imperfectly persistent productivity dynamics to focus on entry, exit, and reallocation of market shares across multiple exporting countries selling into a given consumption (import) market. Other papers such as Arkolakis et al (2019) and Amiti and Heise (2025) focus on how import competition affect market concentration among domestic firms. We do not examine domestic firms but rather focus on how productivity dynamics affect prices and market shares among the set of exporters to a given market.

The paper proceeds as follows. In Section 2 we describe our data. Section 3 provides stylized facts on entry, exit and market reallocation. Section 4 compares the behavior of young and old exporters. Section 5 examines the distribution and persistence of price shocks. Section 6 provides an extended Melitz (2003) model with growing market size and additive productivity shocks and simulates its behavior for comparison with stylized facts developed in Sections 3, 4, and 5. Section 7 concludes.

2. Data and Measures

2.1 Data

Our primary data are bilateral imports drawn from two sources. The first source is all bilateral imports worldwide at the HS6 level extracted from the UN COMTRADE database for the period 1991-2020, using the 1992 version of the Harmonized System for consistent product coding throughout. The HS6 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. We will refer to exporters as *j*, importers as *i*, and HS products as *h*. A "market" is an importer-HS product (ih). We have as many as 550,000 *ih* markets in the world data.

Our second import source is annual US Imports of Merchandise data for 1991-2020, aggregated to the level of exporter- HTS10 digit product-year. Because many of our measures focus on entry and exit behavior, we follow the Pierce and Schott (2012) algorithm, creating synthetic codes to replace HTS codes that are reclassified over time. This results in just over 10,000 unique product codes that are consistently present in the data over time. We provide the US data as a more disaggregated check on the world data, and to enable a more direct comparison to dynamism measures based on US establishment data.

2.2. Comparing Establishment- and Exporter-Level Data

A weakness of our approach is that we do not observe output, employment or sales by domestic producers and our data aggregates at the exporting country level over firms/establishments. This means that we cannot directly address patterns of entry/exit, the distribution of productivity, or the reallocation of production or employment that occur across firms producing *within* a given country. Instead, we focus on closely related measures that capture similar patterns of dynamism *across* multiple exporting countries within a given import market. A strength is that we are able to observe prices and market shares for foreign sellers with much broader coverage relative to country-time specific studies and with more disaggregated product coverage.

It is useful here to comment on the differences in entry and exit rates as well as other measures we employ based on calculations using imports data as opposed to the more common establishment-level data employed in US-centric measures of market dynamism. The establishment data from the Longitudinal Business Database (or its public use form, the Business Dynamics Statistics, or BDS) defines entry on the first year of employment in an establishment and exit as the closure of the establishment. For papers that focus on firm-level aggregations of these establishments, entry of firms corresponds to the year of first entry by an affiliated establishment and exit corresponds to the closure of the last affiliated establishment. The LBD and BDS do not capture entry by these establishments into any market outside the US, employment or sales by any foreign competitor producing outside the US but selling into the US, or mergers/acquisitions that change the ownership of an establishment. Establishment age is left censored in 1977.[4]

In contrast, the imports data are based on location of production, i.e. they are aggregated over firms/establishments to the level of exporting country. Entry is based on first export sales into a market *ih* by any firm producing in exporter *j*, and exit is based on all firms from exporter *j* ceasing sales. If there is variation over time in the number of active firms from an exporter *j* in market *ih*, it does not show up in entry/exit, but does potentially show up in variables such as market share and (weighted average) prices for that exporter.

The imports data do not capture sales by domestic firms, but do capture entry (for a given exporting country) into new destination markets and new product lines. That is, an exporter already present in market *i* can newly enter *i* with a new product *h*. Because the data are based on location of production not ownership, a multinational firm that switches production location for sales into a given

[4] Left censored establishments account for 65% (employment share) of the total establishments in 1991, 51% in 2005 and 44% in 2020. Authors calculations based on the BDS public use data.

market can count as entry/exit, if it is the first/last firm from an exporting location. The data are left-censored to a varying degree depending on when a market first appears in the data. For robustness, we also estimate on a subsample of data where the first year is not left-censored.[5]

While it would be desirable to conduct our exercises using firm or establishment data, there do not exist matched firm-level data sets that enable a comprehensive view of export entry into multiple import markets from all exporting firms and countries world-wide. That said, there are four reasons to think that our exporting country x product level data still capture interesting and meaningful variations in firm-level activity that is informative about firm-level dynamism in markets.

First, while exports between two large countries such as China and the US likely incorporate many distinct firms for each HS code[6], smaller country pairs are far more likely to involve a single exporting firm for a given trade flow. For example, Hummels et al (2014) using data on Danish imports and exports report that in two-thirds of cases, a given partner-HS6 trade flow involves only a single Danish firm. If the pattern for Denmark, which is in the upper quartile of trading countries worldwide by trade volume, holds for countries with smaller trade volumes, this would mean that a "typical" entry/exit as we measure it at the exporting country x product level would also correspond to entry/exit at the firm level.

Second, entry and exit rates are weighted by the share of each exporter in world trade for that HS code. If a country's world export share is proportional to the number of firms from that country exporting into world markets, the weighted measures are effectively a count of the number of potential firms from a particular exporter selling into a given import market.[7]

Third, many of our measures of market dynamism relate to market shares or prices at the exporter level. When the imports data capture a single firm participating in a given importer-product-exporter trade flow, variation in an exporter's market share and prices *is* firm variation, and reflects intensive margin variation in market participation. When the imports data aggregate over multiple firms participating in a given importer-product-exporter trade flow, price and market share measures mix intensive margin changes alongside (invisible to us) entry by new firms. That is, a market share could

[5] A market is not left-censored if a country *i* begins reporting imports in year *t*, a product *h* appears in *i*'s imports for the first time in a year $Y > t$, and no exporting country reports selling that product to *i* before year *Y*.
[6] See Blonigen and Soderbery (2010) for a specific example of mapping varieties at the firm level to varieties at the country level in the case of automobile imports.
[7] In a very simple model where all firms are symmetric, variation in world export shares would exactly reflect variation in numbers of firms in each exporter.

rise either because existing firms' market shares go down or because new firms from that same exporter have entered and contributed to aggregated country sales. But if the underlying story of rising or falling dynamism holds regardless of the measure employed, we can have more confidence that our insights are not overly plagued by the inability to see individual firms.

Four, for our US sample we can compare some of our measures of dynamism in imports data to the corresponding measures taken from public use sample of establishment data captured in the Business Dynamics Database. In the appendix, we show that dynamism measured using imports data strongly correlates with dynamism measured using the establishment data.

3. Entry, Exit, and Market Share Reallocation

In this section we characterize import market dynamism over calendar years and across import market age for measures that examine entry, exit, and market reallocation. Because we have over a half-million *ih* markets that are distinctive in a variety of ways, we use simple descriptive regressions that relate each market dynamism measure to calendar year as follows

$$D_{iht} = \alpha_{ih} + \sum_{t=1993}^{2020} \alpha_t YRDUM_t + \epsilon_{iht} \qquad (1)$$

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 2020. 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 a specific market *ih*, averaged over all markets. (Results from subsamples where we look only at cohorts present for all or a substantial portion of the calendar years are similar.)

3.1 Entry, exit, churning

For each importer *i* - HS product *h* "market", we measure entry, exit, net entry, and gross churning of exporters *j* into the market in each year *t*. One approach would be to treat all exporters symmetrically and simply count the number of countries who enter or exit as a share of countries in each import market. Our preferred approach assigns each exporter a weight equal to that 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. (The main patterns we highlight look qualitatively similar using both weighted and unweighted 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

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

$$exit\ rate = \frac{\sum_{j \in J_{i,exit,t}} wes_{jt}}{(\sum_{j \in J_{iht}} wes_{jt} + \sum_{j \in J_{ih,t-1}} wes_{j,t-1})/2} \qquad (3)$$

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.[8] The denominators are also weighted: e.g. a denominator of 0.4 means that averaging between periods t and t-1, this market *i-h* purchased goods from exporters *j* representing 40% of the world's total exports of product *h*.

Weighted net entry and gross churning are simply the difference and the sum of weighted entry and exit rates, respectively. Our net entry measure is very closely related to the commonly employed "extensive margin of trade" measures used for time series comparisons by Feenstra (1994) and cross-country comparisons by Hummels-Klenow (2005).

In Table 1, we report summary statistics on (weighted) entry, exit, net entry and gross churning measures, calculated for US and World samples. The mean value of 0.369 for entry rate in the World sample means that in an average *iht*, the world market share of newly entering exporters is equal to 36.9% of total market shares of exporters in that market in years *t* and *t-1*. The mean value of .301 for exit rates in that World sample means that in an average *iht*, the world market share of exiting exporters is equal to 30.1% of the total market shares of exporters in that market in years t and t-1. Adding these together gives the Gross Churning rate of 0.67 (the combined market shares of new entrants and exits is equal to 67% of total market shares for active exporters), and taking the difference gives the net entry rate of 0.068 (market shares of entrants exceeds that of exits by an amount equal to 6.8% of total market shares for active exporters).

[8] For the WORLD samples, weights and entry/exit are all based on HS6 data. For the US sample, we take the weights from HS6 data worldwide, and map them into HTS10 digit entry/exit data, i.e. applying the same weights to all HTS10 flow that have a common first six digits.

A few things are notable. First, there is substantial heterogeneity in all measures across importer-product-time observations. Second, net entry is close to zero for the median *iht* observation, meaning that roughly half of *iht* observations experience negative net entry throughout the period. Third, net entry is overshadowed by gross churning, which is an order of magnitude larger than net entry in the World sample and 80 times larger than net entry in the US sample. This is particularly relevant if supplier dislocation is itself a source of utility or productivity losses as in Baqaee et al (2023).

Figure 1 plots the coefficients of year dummies (α_t) in Eq. (1) using the World sample. Entry rates and gross churning relative to 1992 decline steadily over calendar years, and the magnitudes are substantial: entry and gross churning are 50% below 1992 levels by 2001, reaching 70% below 1992 levels in 2020. Exit rates are either flat or decline much less than entry rates, which explains why entry and gross churning track each other closely.

Results using the US-only sample (for this and other dynamism measures) are captured in the Appendix. The change in the entry rate and gross churning are smaller in the US sample, but show a clear decline of 10% over calendar years. Exit rates are similar. Further, the appendix Table A1 comparison of imports and establishment data for the US show that industries that have seen the greatest decline in import entry rates have also seen the greatest decline in establishment entry rates (and similarly for exit).

3.1.1 Are declining entry rates a mechanical artefact? No.

A simple explanation for declining entry rates as measured in this way might be that markets are increasingly saturated. That is, once a large majority of possible exporters are present in a specific import market, entry rates would have to mechanically slow down as there are few source countries left to draw on. This is potentially more problematic when measuring entry and exit at the level of exporting countries as opposed to individual firms.

To show that this possibility is *not* driving our results we provide several exercises. First, note that the number of exporters in a given product *h* is not fixed over time. We calculate the number of exporters who sell to at least one importer for each product and year, and show (25^{th}, 50^{th}, 75^{th}) percentiles of this value over time in Figure 2, first panel. At all percentiles, the number of exporters of a product *h* doubles (or more) over the sample.[9]

[9] At the 25^{th} percentile, the number of exporters rises from 30 to 67; at 50^{th} from 41 to 96; at 75^{th} from 52 to 124.

We can then ask whether the number of exporters that a particular importer buys a product from is rising or falling as a share of the total number of exporters of that product worldwide. We calculate this ratio for each *iht* and plot percentiles of this value in each year in Figure 2, second panel. At all percentiles, the fraction of potential sellers that an importer buys from is falling throughout the 1990s, then remains flat through the rest of the sample. This shows that the number of available exporters is growing faster than the number of exporters a particular importer buys from. This is the opposite of what one would expect if declining entry rates were driven by importers running out of exporters to purchase from.

These statistics are based on unweighted counts since weighted counts by definition sum to one in each period, and so do not capture the growth in suppliers over time. In a final exercise we look at whether entry behavior looks different for markets that purchase from a large fraction of available suppliers in a given year, where the fractions are calculated on a trade-weighted basis.

Define $saturation_{iht} = \sum_{j \in J_{iht}} wes_{jht}$, where J_{iht} is the set of exporters that export product *h* to importer *i* in year *t*.[10] We repeat the regressions reported in Figure 1 using only those *ih* import markets whose saturation never exceeds a particular threshold (70%, 50%, 30%, 10%) and report the results in Figure 3. The basic pattern of declining entry rates holds in all subsamples, and markets with the lowest saturation rates exhibit the greatest decline in entry rates. This is the opposite of what one would expect if entry rates were being slowed mechanically by running out of exporters to buy from.

3.1.2 Additional robustness

How robust are these patterns? One 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 Eq (1) in which year dummies are replaced with a simple indicator variable that takes a value of 1 for years 2-6 for which we have data for that *ih* , and value 0 for all subsequent years. (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 importers) and for every importer (pooling over all products).[11] Table 2 shows the number of cases where the indicator is

[10] Saturation = 50% for a market *ih* if exporters representing 50% of world sales of *h* are selling on this market

[11] In these regressions we restrict our country samples to *ih* markets where there are at least ten years of data, and where there are four or more exporters into the market. The latter restriction is most relevant for calculating the other dynamism measures reported in Table 2.

positive or negative, with additional breakouts using a p-value threshold of 0.05. Note that a positive estimate means that in the first five years of this market, the relevant dynamism measure has higher values than the remaining years, i.e. that the measure is declining over time.

Looking across products, entry rates in years 2-6 are higher than in remaining years in every single product, with 4992 out of 5038 products having p-values below .05. Looking across countries, we are able to run importer-specific regressions for 100 importers, with the remainder lacking a sufficient number of products with a sufficient number of partners or lacking a consistent time series to produce the pooled estimate. These 100 countries represent 93% of world trade and GDP. Of these, 96 countries (representing 91% of world trade) have entry rates that are higher in years 2-6 than subsequently. For both product- and country-pooled regressions, exit rates are mixed between positive and negative signs in early years, but because changes in exit are small, the gross churning patterns reflect the entry rates.

We also experimented with including time varying controls for market *ih* (income and population), and using only those markets where the first year of data are not left-censored. Results in these cases are similar to, or when eliminating left-censoring even stronger than, results in Figure 1.

3.2 Market share reallocation

Entry and exit rates are one way to measure dynamism in markets, but firms may experience large changes in their competitive position while remaining active in the market. To capture this intensive margin change in dynamism, the firm- and establishment-level dynamism literature looks at reallocations of employment across establishments or firms.[12] We can provide a closely related measure by examining reallocation of market shares, that is, the over-time change in exporters' market shares in a particular import market. This tells us whether reallocations across exporting countries are occurring over time, and to see whether markets become more or less stable as they age.

[12] See Davis et al (2014) and Decker et al (2014).

For each exporter *j* on import market *ih*, we calculate its (value) share in *i*'s total import of product *h*, sh_{ijht}. Then, we calculate the absolute value of the year-on-year change in sh_{ijht} to measure volatility in shares.[13]

$$Dsh_{ijht} = \left|sh_{ijht} - sh_{ijh,t-1}\right| \qquad (4)$$

We employ a variant of Eq (1) in which we regress this *ijht* varying measure of market share volatility on calendar year and incorporate ijh fixed effects (World sample) or jh fixed effects (US sample).

Figure 4 provides changes in market share volatility across calendar year and shows pronounced declines in reallocation. The coefficients here are smaller than with entry/exit behavior. By 2019 the average year on year change in market shares is .025 (2.5 percentage points) smaller than in the initial year. 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 2.5 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 52% (=.025/.048). Focusing only on the sample where the first year of data are not left-censored, the drop in volatility is similar to Figure 4 but the reduction in volatility over time is equivalent to reducing market share changes at the mean by 60%.

Results in the Appendix are similar for US, with market share volatility declining over time. Further, those industries that see the largest declines in market share volatility in the imports data also see the largest declines in job reallocation in the establishment data.

This pattern appears for most products and countries. In Table 2, we report sign tests for whether market share volatility is rising or falling across many products and markets. In 88% of products, and in 78 out of 80 countries representing 91% of world trade, the first five years for which we have data for an import market have higher market share volatility than in subsequent years.

[13] We include cases where a share turns from 0 to positive or positive to zero in consecutive years, but not cases where shares are zero in consecutive years. We have also experimented with using the absolute value of log change in shares (which requires positive shares in consecutive years), and found similar changes over time.

4. Prices, Market Shares and Growth Rates for Young and Old Exporters

The establishment-data literature on business dynamism shows a particular pattern to the loss of dynamism. It is not merely that entry is slowing and employment reallocation is declining. It is that older firms are increasingly dominant, enjoying both growing shares of employment and growing productivity advantages relative to newer firms.[14]

We define "young" exporters as those who have been selling in this import market for 5 or fewer years (results are similar looking at exporters on the market for 3 or fewer years). "Old" exporters are those who have been selling on this market for six years or longer. We do not see employment or productivity in our data, but import data provides related measures: market shares and prices.

Market share is simply the combined market share of all exporters in the "old" and "young" groupings, respectively. For prices, we capture the trade-weighted average unit value for each group, dividing total value of import from all exporters in the group by total quantity of import. We follow Eq (1) and regress these measures (market shares and prices) on importer-product fixed effects and calendar year.[15] Note that we only include year dummies starting at year 1997, since we need minimum of 5 years to have a complete cohort of old exporters, and year 1996 becomes the base year.

Figure 5 shows the market shares of young and old exporters in the World sample, regressed on calendar year dummies. We see a rising market share for older exporters and a sharply declining market share for younger exporters over time. Young exporters' market share falls continuously and by a final total of 70 percentage points relative to the initial year.

Figure 6 shows regressions of prices for old and young exporters on calendar year for the World sample. Relative to the base year, the price of young exporters rises steadily over time, so that in the last year younger partners are selling at prices 40 percent higher than in the initial year. In contrast, average prices for older exporters are 10-35 percent below the initial year but fluctuate with no clear trend.

[14] See Decker et al (2014, 2016) on the share of young firms as a share of firms, job creation and employment. See Decker et al (2020) on productivity.

[15] We focus on absolute price levels for the "young" and "old" groups but have experimented with specifications where we measure young/old relative prices. By expressing these as a ratio we also control for importer-product-time varying information such as macro demand factors or openness to trade that affect all exporters. Ratios and absolute levels of prices show similar patterns.

Prices reflect a combination of productivity, input costs, and markups; while we cannot separate these individual influences, the differential changes in prices between old and young paired with changes in market shares are instructive.[16] For example, it is unlikely that the rising prices for young exporters reflect growing markups at a time when their market shares are falling rapidly. A simpler explanation would seem to be that younger exporters are at a rising cost or productivity disadvantage relative to older exporters.

Patterns shown in Figures 5 and 6 are very similar for the US sample reported in the Appendix. The public use establishment data do not contain information on the distribution of prices or productivity, but they do have the employment shares of young and old establishments. In those industries where the import market shares of young exporters are declining the most, the employment share of young establishments is also declining the most. And similarly for import market shares and employment shares for establishment data.

Results on prices and market shares for young and old exporters are similar for the sub-sample of World markets that are not left-censored. They also appear for most products and countries. In Table 2, we report sign tests for whether market shares and relative prices for young exporters are rising or falling across many products and markets. In 99% of products and in 79 out of 82 countries representing 90% of world trade, the first five years for which we have complete data for an import market (i.e. years 6-10) young exporters have higher market shares than in subsequent years. In 84% of products and in all 78 countries where we have estimates, the first five years for which we have complete data for an import market young exporters have lower relative prices than in subsequent years.[17]

Decker et al (2016) emphasize that not only are market shares for young firms falling but that fewer young firms are exhibiting rapid growth necessary to become large firms over time. They show

[16] There is a literature in trade that looks at whether unit values reflect only firms' productivity or also variation in product quality. Undoubtedly quality change will affect unit values for some observations. However, we are examining time period averages over all *ih* markets, and are focused on changes in prices for older versus newer exporters. Differences in quality that are *ih* market specific (reflecting, e.g. income levels of consumers) are absorbed into fixed effects, and level differences in quality across exporters that separate old and young exporters are absorbed into base year differences. For quality change to drive the results in this picture it would need to be the case that quality change is happening systematically (to drive averages), is continuously trending, and is doing so at different rates for old and new exporters.

[17] Note that there are fewer significant estimates in the regressions that pool over countries for a given product. This is due to unit values being a somewhat noisier measure, particularly when we split samples on age so that the set of exporters in the young/old categories are fewer.

this by calculating the growth rate for each firm, grouping them into young and old firms, and then comparing 90th to 50th percentiles of growth for young and old firms. They show that the 90-50 ratio of growth rates for young firms is dropping over time, absolutely, relative to the 50-10 ratio for young firms and relative to the 90-50 ratio of old firms.

We provide a similar calculation based on import data. We calculate year-on-year growth rates for each *ijht* import flow, group these into old and young exporters for each *ih* as above. We then examine whether 90-50 growth rates for young exporters are falling over time, both absolutely and relative to 50-10 growth rates for young exporters and 90-50 growth rates for old exporters. Figure 7 shows that all three measures are falling over time. As with establishment data, over time fewer and fewer young exporters are exhibiting the rapid growth necessary to become large exporters.

5. Shocks and Residuals

Recent work in the establishment dynamism literature emphasizes the idea that markets can become less dynamic if either the volatility of shocks (to productivity or profitability) declines, or firms' response to those shocks declines. Decker et al (2020) provide a variety of productivity measures at decadal frequencies and show that within industry total factor productivity (for manufacturing firms) and labor productivity (for other firms) is more dispersed in the 2000s than in the 1980s or 1990s. Similarly, innovations to TFP are also more dispersed in the last decade, while the persistence of productivity measured as the AR(1) parameter is declining. This leads them to conclude that declining dynamism must be due to declining response to shocks. Cooper et al (2024) provide additional evidence linking this declining responsiveness to labor market adjustment costs.

What do our trade data tell us about shocks? We do not see productivity but we do see detailed price time series that are specific to importer x exporter x products. We can use this to measure persistence in the price series as well as moments of the distribution of shocks to prices.

We begin by estimating an AR(1) model for prices and studying the behavior of the residuals.

$$\ln p_{ijht} = \alpha_{ij} + \beta \ln p_{ijht-1} + \varepsilon_{ijht} \qquad (5)$$

Using our World sample, our estimate of the AR(1) persistence parameter from Eq (5) is 0.53. For each *iht*, we calculate the mean, standard deviation, and skewness of the price residuals across exporters *j*, and then regress the moment statistics for each *iht* on calendar year in the manner of Eq (1). Results are

reported in Figure 8. The mean size of the residuals varies year to year but does not show any clear trend. In contrast, the standard deviation of the residuals is rising steadily (the coefficient of variation shows a similar pattern), as is the skewness. The rising standard deviation of the residuals is consistent with the Decker et al (2020) finding of increasing dispersion in TFP shocks. Rising skewness indicates that there is an increasing mass of price innovations in the right tail, i.e. price increases.

To further explore these patterns, we regress the price residuals from Eq (5) on the (lagged) market share of exporter *j* in market *ih*, as well as a group of dummies that indicates the age of an exporter *j* in market *ih*. This corresponds to the total number of years the exporter has been active in this market, grouped into five year intervals corresponding to ages: 6-10, 11-15 and so on through 25-30 (age 1-5, the "young" exporters in previous estimations, is the omitted group). We also include *jht* fixed effects for the World sample and *ht* fixed effects for US sample.

Results are reported in Table 3. Here we see two clear patterns in both the US and world data. Within a market, price residuals are negatively correlated with lagged market shares, and price residuals are falling over time for older exporters relative to younger exporters. Note that the time dimension of the fixed effects eliminates any year-specific movement in the price residuals, such as those shown in Figure 8. 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. In this case, price residuals for the oldest exporters drop 15% relative to the "new exporter" baseline. This pattern holds when conditioning on the time varying market share of the exporter in each market.

To put this more plainly, these results show a marked difference between the price dynamics of older and younger exporters and between exporters with larger and smaller market shares. This is inconsistent with a standard model in which shocks to prices are multiplicative and identically distributed across suppliers. In the next section we provide a simple alternative model of productivity dynamics that generates this key fact and has important implications for understanding business dynamism more generally.

Taken together, Table 3 and Figures 6-8 provide a useful picture of the over-time behavior of prices. While the average innovation to prices is not trending over time, the average disguises two disparate trends. Price innovations are smaller for exporters with larger market shares and are falling for older exporters relative to younger exporters (Table 3). This leads to increased dispersion in price innovations, and a growing skewness in prices caused by an increasing incidence of large and positive

price innovations among the younger exporters (Figure 8). The result is steadily rising prices for young exporters (Figure 6), that is plausibly connected to falling market shares (Figure 5) and the reduced skewness in growth rates for young exporters (Figure 7).

6. A Simple Model of Import Market Dynamism

In this section, we show how a simple extension of Melitz (2003) allows the age of a market to affect the dynamics of productivity and market entry and can generate all the business dynamism facts described in Sections 3, 4, and 5. There are two key features of this model.

One, while productivity shocks are identically distributed across firms and over time, they are additive rather than multiplicative. Consistent with the evidence in Table 3, the impact of a given additive shock is smaller the more productive a firm is and the larger its market share. Two, import markets are growing over time, which combined with initial differences in firm productivity, leads to an entry order effect. Early entrants have high productivity while growing market size enables later entrants to succeed despite lower productivity. The relative strength of these two effects determines whether market dynamism, as captured by standard measures in the literature, is rising or falling.

In Melitz (2003), firms differ according to (randomly drawn) productivity, and face fixed and marginal costs of entering a single foreign market. The canonical model is solved for a general equilibrium characterized by entry, prices and revenues for each firm and wages that clear the market. There is no time dependence. One can think about the model as a one-time comparative static where the equilibrium adjusts in the face of changing trade costs. Or, if one were to use this model to understand changes over a long period of time, it would be a multi-period sequence in which firms' productivity draws are independent over time.

We depart from this baseline in two ways. First, we allow productivity to evolve according to an MA(1) process so that the time path of firm productivity, prices, revenue, and entry depend on past conditions. This additive process also differs from papers in the business dynamism literature that use an AR(1) process for productivity. Two, we examine these productivity dynamics in the context of the imports of a small open economy, sales into which do not affect wages in the exporting countries. We use this approach because our focus is on import market dynamics rather than the general equilibrium

of firm supply in the domestic market. It also allows us to abstract from any mechanism related to labor reallocation or constrained aggregate labor supply.

6.1 The Model

Suppressing market *i* and product *h* subscripts, consumers have CES utility over varieties of a product with elasticity of substitution σ, which generates demand for firm *j*

$$q_{jt} = \frac{p_{jt}^{-\sigma}}{\sum_{j \in J} p_{jt}^{1-\sigma}} E_t$$

where the summation in the denominator runs over all firms $j \in J$ that are active in the market, and E_t is total expenditures on this good. Goods are produced using labor paid a wage *w* so that $p_{jt} = \frac{\mu w}{\varphi_{jt}}$, where μ is the usual constant markup over marginal costs and productivity φ_{jt} is firm-time specific.

Initial productivity φ_{j0} is drawn from a Pareto distribution with positive support. Subsequent productivity depends on lagged productivity and a normally distributed innovation term that is mean zero and has the same variance across firms and over time $\varepsilon_{jt} \sim N(0, \gamma^2)$.

$$\varphi_{jt} = \varphi_{j,t-1} + \varepsilon_{jt}$$

Solving recursively through earlier periods, this can be written as

$$\varphi_{jt} = \varphi_{j0} + \sum_t \varepsilon_{jt}$$

A few things are notable about this productivity process. First, a given shock has a larger percentage impact on productivity (and prices) the lower is the initial productivity draw. [18] Assume that shocks are small relative to the initial productivity level, we can then write the change in prices as

$$ln\left(\frac{p_{gt}}{p_{gt-1}}\right) = ln\left(\frac{\varphi_{gt-1}}{\varphi_{gt}}\right) = ln\left(\frac{1}{\left(1 + \frac{\varepsilon_{gt}}{\varphi_{gt-1}}\right)}\right) \approx -\frac{\varepsilon_{gt}}{\varphi_{gt-1}}$$

[18] Applying an AR(1) regression model to our simulated data, one finds an AR(1) coefficient less than one, and residuals of that AR(1) are decreasing in market shares, consistent with Table 4.

Second, while the variance of the shock is the same for all time periods, the variance of productivity and prices across firms in a given time period depends on an accumulation of independent shocks. That is,

$$var(\varepsilon_{gt} + \varepsilon_{gt-1}) = var(\varepsilon_{gt}) + var(\varepsilon_{gt-1}) + 2cov(\varepsilon_{gt}, \varepsilon_{gt-1}) = 2\gamma^2$$

And so on as more shocks accumulate. This means that the variance of the productivity distribution grows over time, consistent with our price evidence in Figure 8, and the evidence from the US-focused business dynamism as reported in Decker et al (2020).

The effect of these shocks on mean (and higher moments) productivity and prices, as well as other measures of business dynamism, depends not only on the distribution of shocks but on equilibrium entry and exit. To understand these effects, we solve for entry in the manner of Melitz (2003). When selling into an import market, firm *j* faces a fixed cost of entry F, and will enter this market if revenues are sufficient to cover fixed entry costs, $\frac{R_{jt}}{\sigma} > F$. Substituting and arranging,

$$\frac{\varphi_{jt}^{\sigma-1}}{\sum_{j\in J} \varphi_{jt}^{\sigma-1}} > \frac{F\sigma}{E_t} \qquad (6)$$

Firms will enter if their productivity in this period, relative to a weighted average of productivity for competitors active in the market, exceeds fixed costs of entry relative to market expenditures. Similarly, firms will exit if, having previously entered, shocks to productivity changes cause them to fall below the entry threshold in subsequent periods. (Unlike Melitz, we do not assume a constant exogenous firm death rate.) This sets up the main features of the model.

Because productivity has memory, each period exhibits a mix of persistent incumbents, newly entering firms, and exiting firms.[19] The rate at which this takes place and how that entry, along with prices and market shares, change over time depends on a tension between two forces. First, increased expenditures E_t allow firms with lower productivity to enter. Since time *t* productivity depends on initial conditions, this creates an "entry order" effect. Firms with high initial productivity enter early; firms with low initial productivity enter late.

Second, the tendency of increased expenditures to allow lower productivity firms to enter is partially counteracted by the cumulative effect of shocks to productivity, which affects the denominator

[19] In a version of this model where firms redraw productivity independently every period, such as the canonical Melitz (2003), entry and exit rates are nearly 100% in every period.

of Eq (6). Recall, the accumulation of mean-zero normal shocks increases the variance of the productivity distribution. Were the summation $\sum_{j\in J}\varphi_{jt}^{\sigma-1}$ to run over all firms, so that $E(\sum_j \varepsilon_{jt}) = 0$, then the effect of rising variance would be to increase $\sum_{j\in J}\varphi_{jt}^{\sigma-1}$, since that summation is minimized (for a given mean) when φ_{jt} is a constant. However, the summation runs only over firms active in the market, and that is determined by productivity levels.

This creates a tendency for positive, not negative, productivity shocks to be selected in the market. An accumulation of positively selected shocks over time increases the mass of active firms in the right tail of the productivity distribution. (There is also more mass in the left tail of the overall productivity distribution, but these left tail firms do not enter the market and so do not appear in the data.) In short, holding expenditures fixed, an accumulation of mean-zero additive productivity shocks steadily increases the average productivity of active firms, raising the competitive bar to enter the market. That effect is magnified when sigma is high, because high sigma means that high productivity, low priced firms get a larger market share and have a bigger impact on the price index. It is weakened if expenditures are growing, because growing expenditures lower the competitive bar to enter the market.

6.2 Simulation Results

We rely on simulations to show these effects in combination affect measures of dynamism, as the mix of distributions employed and the time dependence of the productivity processes make analytical solutions difficult. This should not be thought of as a calibration per se, but rather, an illustrative simulation to show how additive productivity shocks and growing market size can generate business dynamism facts consistent with establishment and trade data.

We assume there are 1000 potential entrant firms, divided between 100 exporting countries, in proportion to the size of that exporter.[20] We generate an initial productivity distribution using draws from a Pareto distribution, and solve for the initial equilibrium characterized by the price index, and price and revenue for each firm that enters the market. We then redraw values of ε_{jt} as above, change aggregate expenditures in line with an assumption on market growth, and re-solve for period 2.

Iterating for 30 periods (as in our import market data) completes one run of the simulation. We then restart a new run of the simulation with a different initial set of productivity draws and repeat for

[20] We use the world share of GDP for the largest 100 countries in 2015.

30 periods, eventually completing 50 runs. Finally, we create new sets of simulations (50 runs of 30 years each) by using the same underlying productivity draws for a given sequence but experimenting with changing parameter values for σ and using different growth rates for expenditures. Changing these parameters generates different equilibrium entry, price index, and revenue values from the same sequence of productivity draws.

To explore similarities between the model simulation and the earlier empirics, we recover from this process the business dynamism measures captured in Sections 3-5 and run the regression in Eq (1) to assess, e.g. how entry rates and other variables change over time. We can think of the distinct model runs as essentially creating a separate import market (*ih*) 50 times, corresponding to the 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.

We begin using a baseline case of 5% growth in expenditures (corresponding to annualized average growth for world trade from 1992-2021), $\sigma = 4$. Initial productivity is distributed Pareto with a scale parameter 3, and shape parameter 10 (mean=3.33, sd =.37), and annual shocks are distributed normally $\mathcal{N}(0,0.2)$. Figure 9 displays the results of the simulation.

Panel A shows the distribution of productivity, conditional on entry, for various years. As discussed above, a series of mean zero additive shocks generates a distribution with an increasing mass in the right tail. Panel B shows the productivity of the marginal entrant, comparing no expenditure growth to 3% or 5% annualized growth. With no growth in expenditure, the productivity of marginal entrants has to steadily rise to counteract the increasing mass of high productivity firms. Faster growth counteracts this mechanism, and ultimately results in falling productivity for marginal entrants in the high growth scenario. Panel C shows the impact on the price index, which falls in all three growth scenarios, reflecting the increasing mass of firms in the right tail of the productivity distribution. However, it falls fastest when growth is greatest because in this scenario greater entry lowers the price index through variety effects.

The remaining panels of Figure 9 capture the patterns of firm level entry rates, market share volatility, and prices for young and old firms. Focusing on the 5% growth scenario, all match the findings

from the US and World data reported in Sections 3-5: entry rates and market share reallocation falling, prices for young firms rising and market shares falling relative to old firms.[21]

It is helpful to explore differences across growth rates in the simulations to better understand the mechanisms at work. Entry rates are falling in all growth scenarios, but more quickly the more rapid is growth. Without growth, repeated additive shocks increase the mass of high productivity firms, steadily raising the bar to entry and slowing entry rates.[22] With growth, marginal entrants can be less and less productive. However, two forces lead faster growth to slow the entry rate. First, faster growth in previous periods means that more firms are already in the market, and the price index is lower (variety effects). Second, faster growth in previous periods means the market has worked its way further down the productivity distribution, so that a marginal entrant in a given period is worse than in a slower growing market.

These entry dynamics also directly lead to the rising relative price and falling relative market share of old and young firms in the 5% growth scenario. With no growth, the price of young/old firms shows relatively little change: the growing mass of old firms in the right tail of the productivity distribution is counteracted by the rising productivity necessary for young firms to enter the market. However, with growth, young firms can be progressively less productive than old firms and still enter, causing a rising relative price. In all scenarios we find a falling relative market share for young firms since the main driver of this effect (increasing incumbent productivity) holds with or without expenditure growth.

Market share volatility is relatively flat in the no growth scenario, but is falling over time for rapidly growing markets. The key mechanism here is that rapidly growing markets experience entry by more high-priced and low market share firms. Even though these young firms face shocks to productivity that are larger in percentage terms than the older firms, the absolute movement in their market shares is very low.

[21] We have also experimented with aggregating the firm level data to exporting countries and running similar regressions at the country rather than the firm level. Results are qualitatively similar, and are often quantitatively similar as well. One exception is that exit rates are much flatter in the country level regressions, similar to that seen in Figure 1.

[22] Not shown in the figure: choosing higher σ values enables high productivity firms to capture larger market shares, and magnifies the effect of the increasing mass of highly productive firms.

6.3 Some Checks on Model Implications

Recall that the model in Section 6 was built to match particular facts from the data, notably the use of additive shocks to capture the negative relationship between price residuals and market share. In this section we examine two additional implications of the model: the connection between entry and expenditure growth rates, and the entry order effect.

In the model, the accumulation of mean-zero additive shocks generates a rising variance in productivity, and an increasing mass of firms in the right tail of the productivity distribution. This slows both entry rates and market reallocation, but the extent to which this happens depends on growth rates.

To look for this effect in the World trade data, we calculate the annualized growth rate for each *ih*, then separate markets into high- (above 75th percentile), medium- (25th-75th percentile), and low-growth (below 25th percentile) markets. We then re-estimate Eq (1) for entry rates and market reallocation for each sub-sample. Results are in Figure 10: the fastest growing markets see the most rapidly declining entry rates and market share reallocation.

An additional implication of the model with expenditure growth is an "entry order" effect. Persistent productivity means that exporters can be ordered from highest to lowest productivity, and that exporters who start more productive will remain, on average, more productive over time. As the size of the import market grows, the zero-profit productivity hurdle falls, allowing in less productive firms at the margin. Since entrants' productivity is progressively worse over time, the prices they charge rise, both absolutely and relative to older firms.

A challenge for taking the entry order hypothesis to the data is that we cannot typically see productivity or prices for firms who have not entered a market. But using the trade data we can make use of price data from other markets. Recall that an importer *i* typically purchases from only a small fraction of the available exporters of product *h* in a given year *t*. Our dependent variable is entry timing between 1991-2005: the number of years after 1991 that exporter *j* enters that market.[23] Our explanatory variable is the average price of exporter *j* selling product *h* on all markets other than *i,* in

[23] In some *ih* markets, the market "starts", i.e. has the first reported data, sometime after 1991. In these cases, we calculate entry relative to the first year that import market *ih* starts.

years 2006-2015. We rank order these exporter prices from low to high within a given *iht*. A positive correlation means that relatively high-priced exporters on other markets enter market *ih* late.

Table 4 reports US and world entry order regressions with various fixed effects. All regressions show a strong positive correlation between an exporter's rank 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.

7. Conclusion

We use bilateral trade data for many countries, products, and years to measure business dynamism in import markets. We show that dynamism is slowing in a way that strongly resembles falling dynamism as demonstrated in the literature using establishment data. Entry rates for new exporters fall over time, reducing market churn, and the reallocation of market shares declines significantly. Market shares for young exporters fall, and prices rise, relative to longer-tenured incumbents, and skewness in growth rates is falling for young exporters. These patterns occur in over 90% of products and 90% of countries.

A central idea in the establishment-level business dynamism literature is that a lack of dynamism could result either from a reduction in the variance of productivity shocks hitting firms, or from rising adjustment costs that prevent reallocation of inputs toward more productive firms. Since the variance of productivity is rising in the establishment data, authors point to an inability to reallocate labor across firms as the culprit behind declining dynamism.

However, this is unlikely to explain declining dynamism in import markets – while labor moves relatively little across countries, the extent of this movement has been increasing over time. This means that international migration is unlikely to be a main factor slowing market reallocation across exporting countries. Instead, we show that an extension of Melitz (2003) with additive productivity shocks and growing expenditures can explain all these facts.

Consistent with our evidence, additive shocks mean that the impact of a given shock is smaller the more productive is a firm and the larger its market share. While the shocks are identically

distributed in each period, the accumulation of normally distributed shocks over time leads to a rising variance in productivity. Entry selects on positive shocks, so the average productivity of incumbents rises over time. When we combine this with market growth, the model generates an entry order effect in which early entrants have high productivity while growing market size enables later entrants to succeed despite progressively lower productivity. We confirm that entry order effects are an important feature of world trade data, and that a number of our business dynamism measures in imports data drop off most quickly in the fastest growing markets.

We think it is interesting that market dynamism using the imports data strongly correlates with market dynamism in the establishment data. We are hesitant to push this interpretation too far because it may be that the mechanisms we explore here are different from those that cause business dynamism to slow at the establishment level within the US. Still, it appears that declining dynamism is a robust feature of market aging that holds across a broad range of markets and products, and our mechanism – growth plus additive productivity shocks – could plausibly hold in a wide range of contexts. In any event, the richness of trade data may enable future research that isolates underlying causes and consequences of declining dynamism in broader contexts.

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: 103490.

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.

Amiti, Mary and Heise, Sebastian, 2025. "U.S. Market Concentration and Import Competition" *The Review of Economic Studies*, 92(2): 737–771.

Amiti, Mary, Stephen J. Redding, and David E. Weinstein. 2019. “The Impact of the 2018 Tariffs on Prices and Welfare.” *Journal of Economic Perspectives*, 33 (4): 187-210.

Arkolakis, Costas, Arnaud Costinot, Dave Donaldson and Andrés Rodríguez-Clare, 2019. "The Elusive Pro-Competitive Effects of Trade", *The Review of Economic Studies*, 86(1): 46–80,

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2017. “Concentrating on the Fall of the Labor Share.” *American Economic Review*, 107 (5): 180-85.

Aw, Bee Yan, Mark J. Roberts, and Daniel Yi Xu. 2011. "R&D Investment, Exporting, and Productivity Dynamics." *American Economic Review* 101 (4): 1312–44.

Baqaee, David, Burstein Ariel, Duprez Cedri, Farhi Emmanuel. 2023. “Supplier Churn and Growth: A Micro-to-Macro Analysis.” NBER Working Paper 31231.

Besedes, Tibor and Prusa, Thomas J. 2006. “Ins, outs, and the duration of trade.” *Canadian Journal of Economics*, 39(1): 266-295.

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

Blonigen, Bruce and Soderbery, Anson. 2010. “Measuring the Benefits of Product Variety with an Accurate Variety Set.” *Journal of International Economics*, 82(2): 168-180.

Cooper, Russell, Haltiwanger, John, and Willis, Jonathan 2024. “Declining Responsiveness at the Establishment Level: Sources and Productivity Implications.” NBER Working Paper 32130.

Covarrubias, Matias, Gutierrez German, Philippon Thomas. 2020. “From Good to Bad Concentration? US Industries over the Past 30 Years.” *NBER Macroecon. Annu*., 34: 1-46.

Davis, Steven J. and Haltiwanger, John. 2014. “Labor Market Fluidity and Economic Performance.“NBER working paper 20479.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2014. “The Role of Entrepreneurship in US Job Creation and Economic Dynamism. *Journal of Economic Perspectives*, 28 (3): 3-24.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2016. “Where Has All the Skewness Gone? The Decline in High-Growth (Young) Firms in the US.” *European Economic Review,* 86: 4-23.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2017. “Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown.” *American Economic Review*, 107 (5): 322-26.

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

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal. 2020. “The Return to Protectionism.” *The Quarterly Journal of Economics*, 135(1): 1–55.

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

Hopenhayn, Hugo, Julian Neira, Rish Singhania, 2022 “From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship, and the Labor Share” *Econometrica* 90(4), 1879-1914.

Hummels, David, and Peter J. Klenow. 2005. “The Variety and Quality of a Nation's Exports*.” American Economic Review*, 95 (3): 704-723.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang. 2014. "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data." American Economic Review, 104 (6): 1597-1629.

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

Philipon, Thomas, 2022, “Additive Growth” NBER Working Paper 29950

Pierce, Justin R. and Schott, Peter K. 2012. “Concording U.S. Harmonized system codes over time.” *Journal of Official Statistics*, 28(1): 53-68.

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

Sanghamitra Das, Mark J. Roberts, James R. Tybout, (2007), “Market Entry Costs, Producer Heterogeneity, and Export Dynamics”, *Econometrica*, 75(3), 837-873

Table 1: Summary Statistics on Trade Dynamism Measures

	(1) mean	(2) p25	(3) p50	(4) p75	(5) min	(6) max	(7) N
World							
Gross churning	0.670	0.052	0.233	1.384	0	2	14963257
Net entry	0.068	-0.095	0.001	0.181	-2	2	14963257
Entry rate	0.369	0.003	0.047	0.320	0	2	14963257
Exit rate	0.301	0.0004	0.032	0.219	0	2	14963257
US							
Gross churning	0.326	0.030	0.099	0.318	0	2	335785
Net entry	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 for the weighted measures of entry rate, exit rate and gross and net churning, calculated for each of the two samples described in Section 3.

Table 2: Sign Tests on Trade Dynamism Trends

		Positive		Negative		Sign test matches aggregate trend	
	Number of Tests	p≤0.05	p>0.05	p≤0.05	p>0.05	Sign match	Sign match p≤0.05
By product							
Entry rate	5038	4992	46	0	0		
Exit rate	5038	1168	1967	588	1315		
Gross churning	5038	4701	204	77	56		
Market share reallocation	5037	3780	688	167	402		
Young partner's market share	5010	4992	14	1	3		
Rel price of young partners	4918	71	612	2783	1452		
By country						**World Trade Share**	
Entry rate	100	95	1	4	0	0.913	0.894
Exit rate	100	50	2	44	4	0.603	0.598
Gross churning	100	97	0	2	0	0.908	0.908
Market share reallocation	80	78	0	2	0	0.907	0.907
Young partner's market share	82	78	1	0	3	0.902	0.901
Rel price of young partners	78	0	0	76	2	0.901	0.897

Note: We replace the vector of calendar year dummies in Eq (1) with a simple indicator variable that takes a value of 1 for years 2-6 for which we have data for that importer-product market, and value 0 for all subsequent years and estimate that coefficient for each product, pooling over importers; and each importer, pooling over products. This table provides a count of the number of coefficients, by sign and significance.

Table 3: Price Residuals and Exporter Age

Dependent variable: residuals from regression of price on lagged price

Exporter characteristics	United States			World		
	(1)	(2)	(3)	(4)	(5)	(6)
market share	-0.5734***		-0.5647***	-0.1053***		-0.1052***
	(0.0039)		(0.0039)	(0.0006)		(0.0006)
age 6-10		-0.0048*	0.0128***		-0.0399***	-0.0403***
		(0.0026)	(0.0026)		(0.0003)	(0.0003)
age 11-15		-0.0133***	0.0166***		-0.0687***	-0.0692***
		(0.0028)	(0.0028)		(0.0004)	(0.0004)
age 16-20		-0.0500***	-0.0106***		-0.0937***	-0.0939***
		(0.0030)	(0.0030)		(0.0004)	(0.0004)
age 21-25		-0.0787***	-0.0301***		-0.1295***	-0.1296***
		(0.0032)	(0.0032)		(0.0004)	(0.0004)
age 26-30		-0.1277***	-0.0668***		-0.1573***	-0.1575***
		(0.0036)	(0.0036)		(0.0006)	(0.0006)
Constant		0.0389***	0.0433***	0.0119***	0.0727***	0.0827***
		(0.0021)	(0.0021)	(0.0001)	(0.0001)	(0.0001)
N	4834312	4834312	4834312	107655102	107655102	107655102
R2	0.0865	0.0825	0.0866	0.1145	0.1151	0.1154
Product-year FE	YES	YES	YES			
Exporter-product-year FE				YES	YES	YES

Notes: Dependent variable is residuals from Eq (5) for US and World samples. US sample has exporter-product-time variation. World sample has importer-exporter-product-time variation.

Table 4: Predicting Exporters' Order of Entry

	United States		World	
	(1)	(2)	(3)	(4)
Price percentile	0.3301***	1.1219***	1.1263***	0.6542***
	(0.0283)	(0.0265)	(0.0044)	(0.0047)
N	289591	288842	8957499	8913353
R2	0.0892	0.3612	0.3315	0.3728
Product FE	YES	YES		
Partner FE		YES		
Importer-product FE			YES	
Partner-product FE				YES
Partner FE			YES	

Note: The dependent variable is the order of entry (between 1991 to 2005), defined as the difference between the first year an exporter *j* shows up on import market *ih* and the first year the importer *i* starts importing product *h*. The explanatory variable is exporter *j's* average (from 2006-2015) rank in the price distribution for product *h* in the US market (for US sample), or in the rest of world markets (for the World sample).

Figure 1: Entry Rate, Exit Rate, Gross Churning over Calendar Year

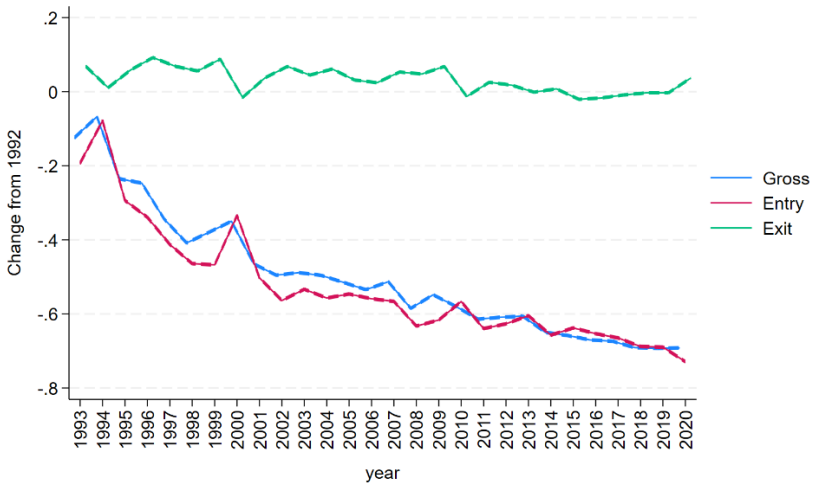

World

Note: This figure presents the change in entry rate, exit rate and gross churning (defined in Eqs 2 and 3) over calendar year. Dashed lines indicate the 95% confidence intervals for the estimates of coefficients of year dummies.

Figure 2a: Growth in Export Suppliers

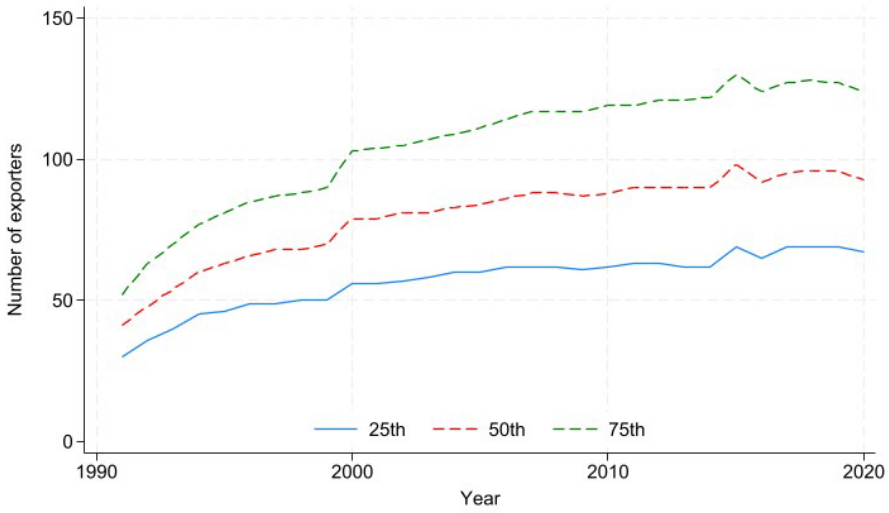

Figure 2b: The Fraction of Available Exporters an Importer Buys From

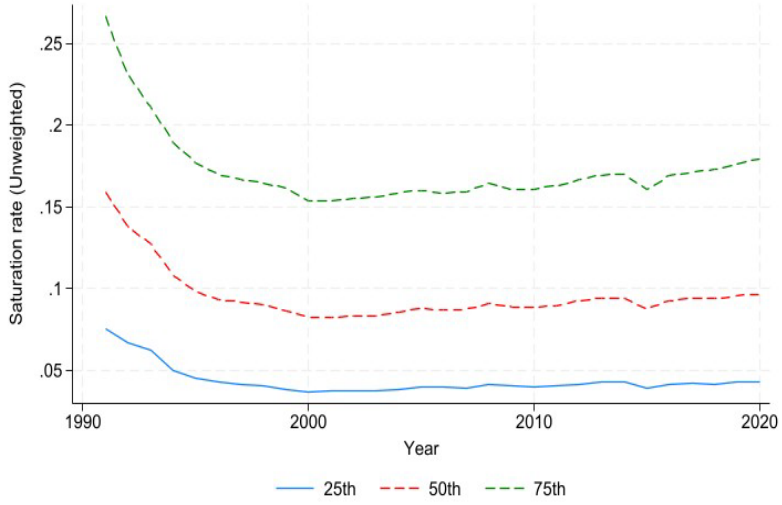

Note: In Figure 2a, we calculate the number of exporters who sell a product *h* to any importer market in each year, and plot percentiles of this count for each year. In Figure 2b, we calculate the number of exporters who sell to a particular importer-product market as a fraction of the number of exporters who sell that product to any market, and plot percentiles of this count for each year.

Figure 3: Entry Rates in Markets with Varying Saturation

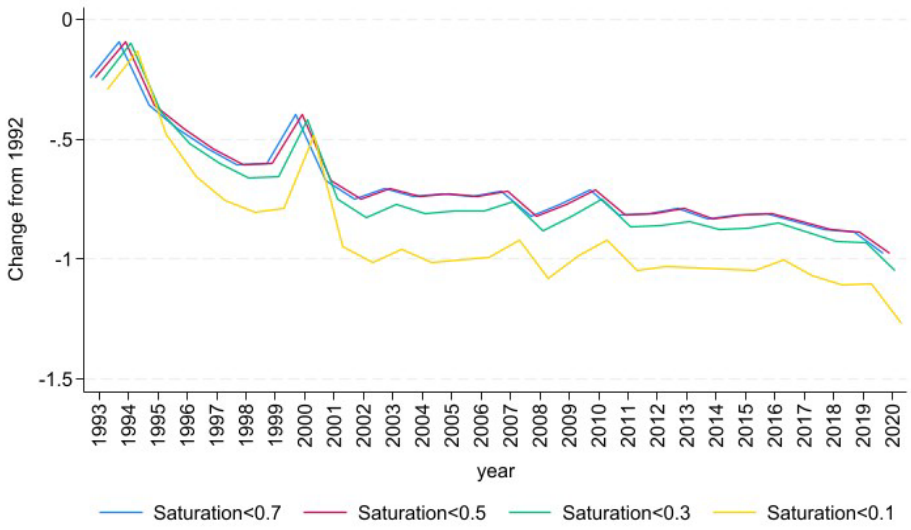

Note: We use a subsample of import markets (ih) where saturation rate never exceeds some threshold (70%, 50%, 30%, 10%).

Figure 4: Market Share Reallocation

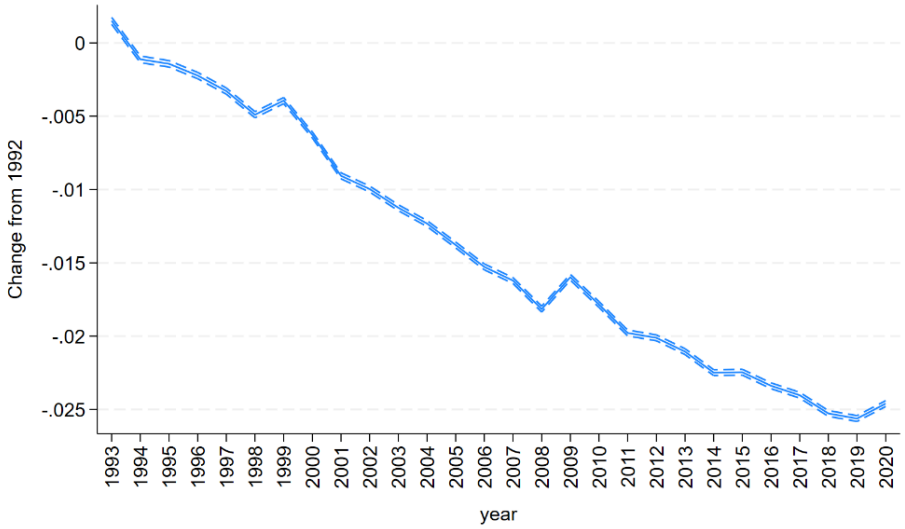

Note: This figure presents the change in market share volatility, defined as the absolute value of the change in market share as in Eq (4), over calendar years. Dashed lines indicate the 95% confidence intervals for the estimates of coefficients of year or age dummies.

Figure 5: Market Shares of Young and Old Exporters

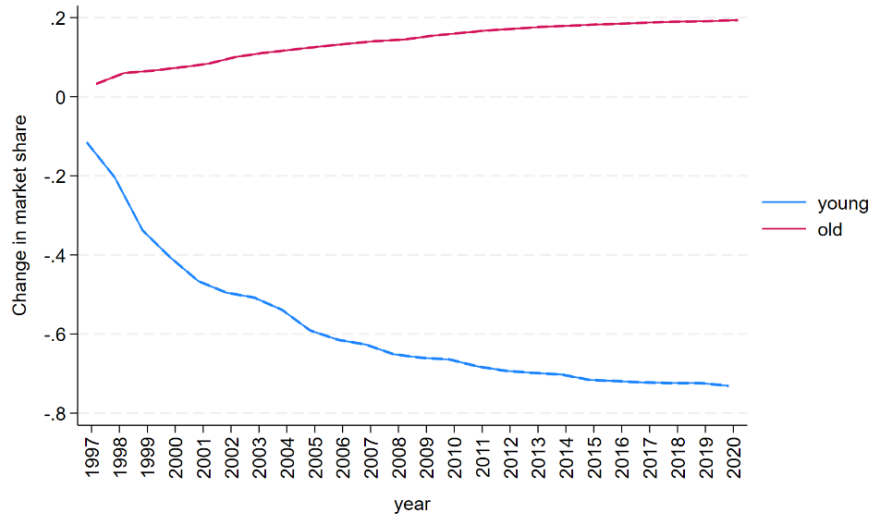

Notes: This figure depicts the trend of total market young exporters vs. old exporters on each import market (*ih*) over calendar year. Young includes varieties that are less than 6 years old on the *ih* market. Old includes varieties that have been on the import market for more than 5 years. Dashed lines indicate the 95% confidence intervals for the estimates of coefficients of year dummies.

Figure 6: Weighted Average Prices by Exporter Age

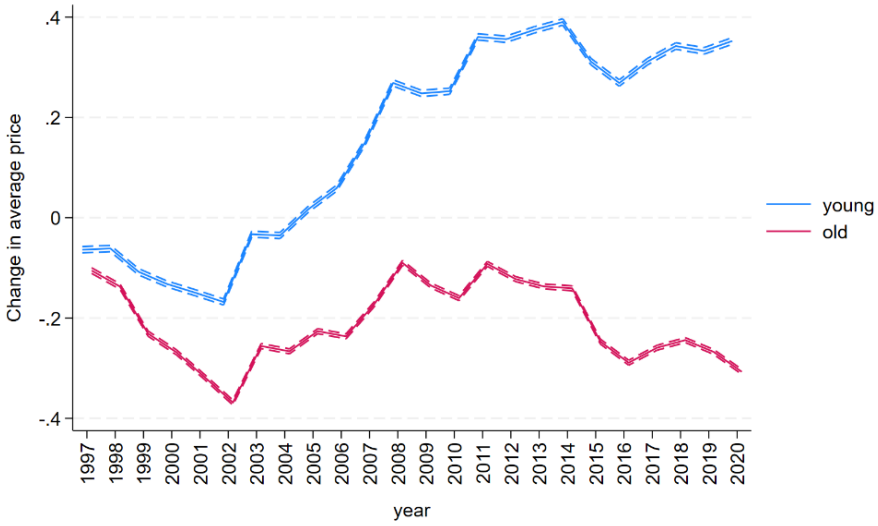

Notes: This figure depicts the trend of total market share and weighted average prices of young exporters vs. old exporters on each import market (*ih*) over calendar year. Young and Old are defined as in Figure .5 Dashed lines indicate the 95% confidence intervals for the estimates of coefficients of year dummies.

Figure 7 Declining Skewness in Young Exporter Growth

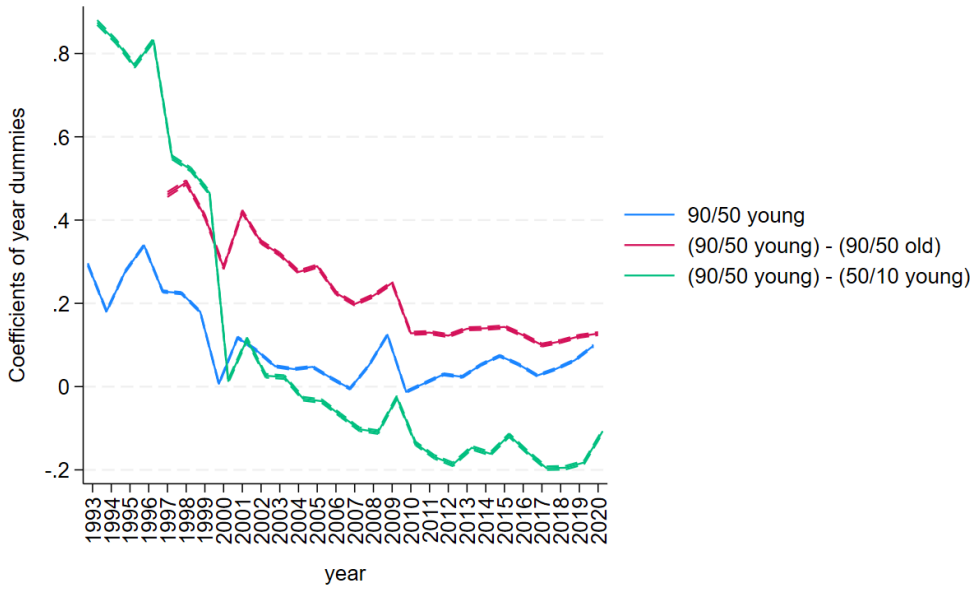

Notes: We calculate year-on-year growth rates for each importer-exporter-product-year flow, group these into old and young exporters for each importer-product. This figure presents the trend of 90-50 growth rates for young exporters over time, both absolutely and relative to 50-10 growth rates for young exporters and 90-50 growth rates for old exporters. Dashed lines indicate the 95% confidence intervals for the estimates of coefficients of year dummies.

Figure 8 Moments of Price Residuals

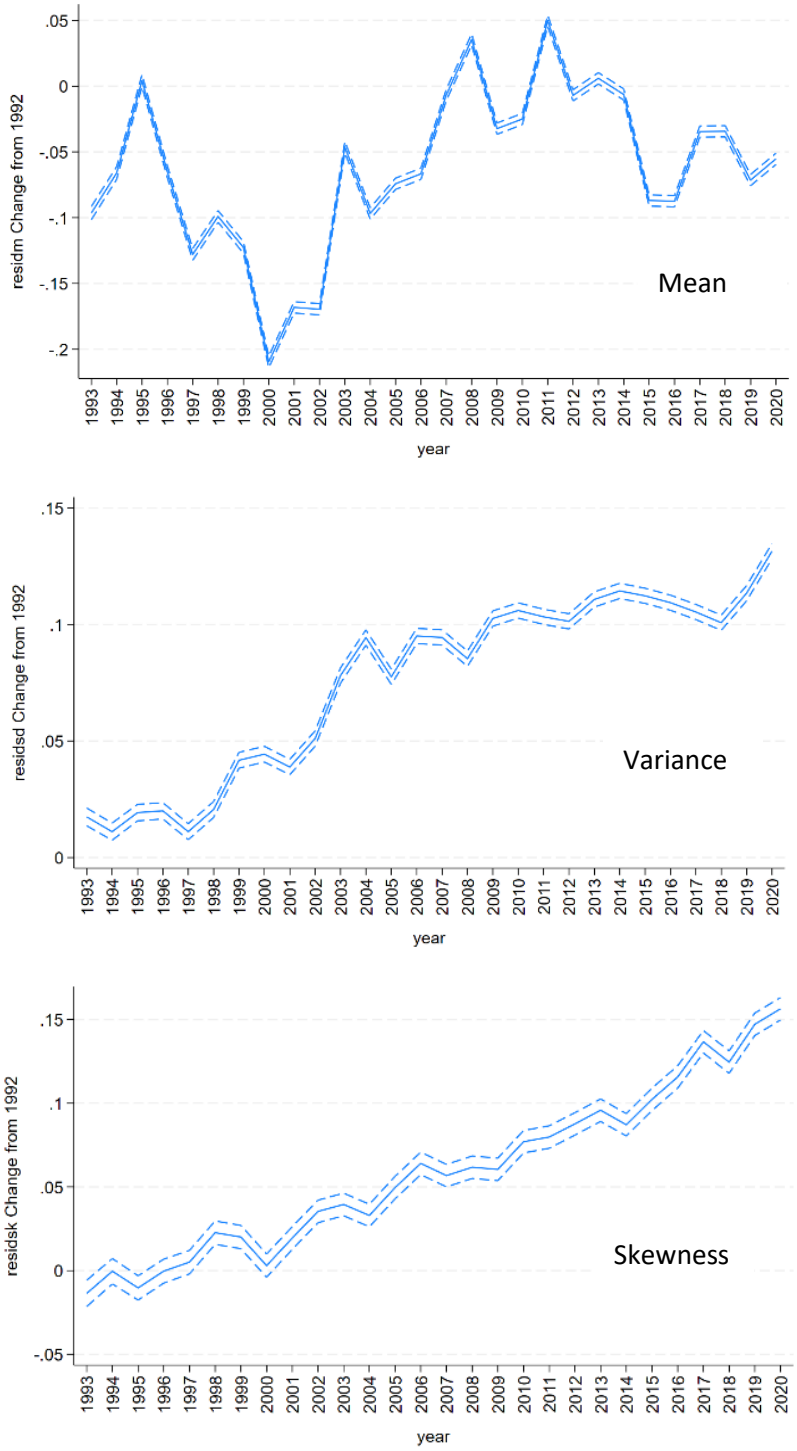

Note: This figure presents the time trend of the mean, standard deviation and skewness of the residual from running AR(1) model of delivered price using World import data. Dashed lines indicate the 95% confidence intervals for the estimates of coefficients of year dummies.

Figure 9 Simulation – Measures of Productivity and Dynamism with Varying Expenditure Growth

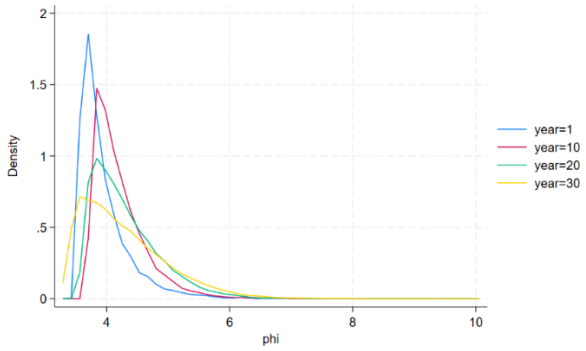

a. Productivity distribution

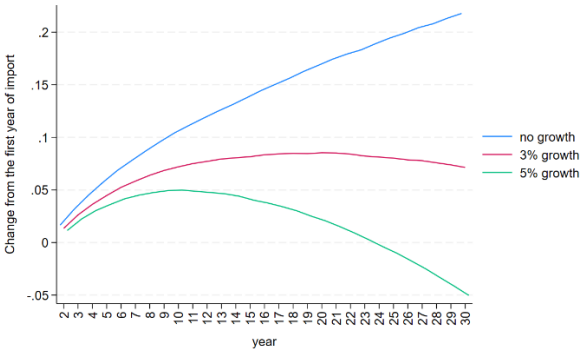

b. Productivity of marginal entrant

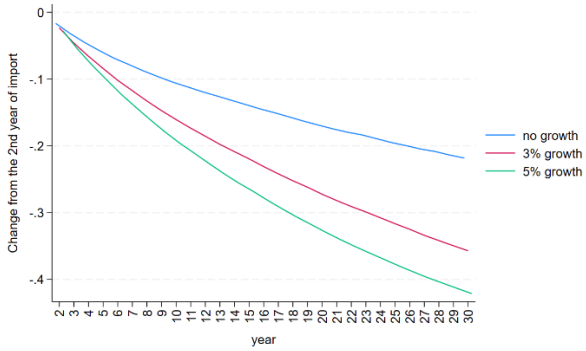

c. Price index

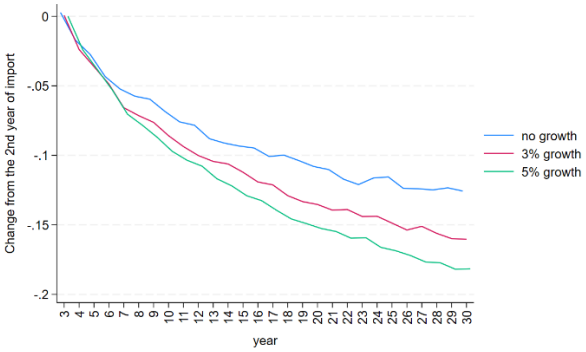

d. Entry rate

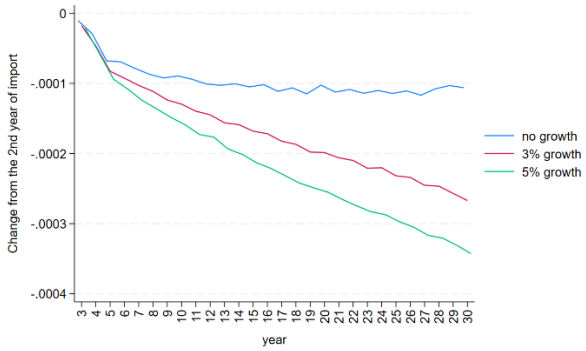

e. Market Share Reallocation

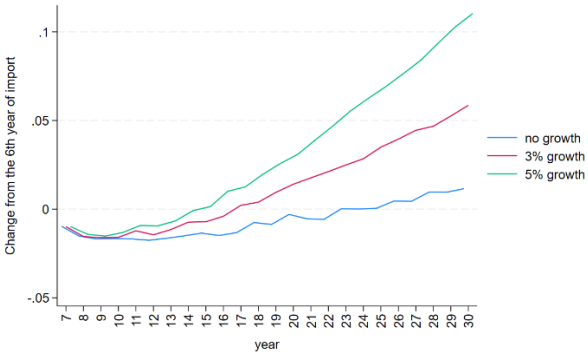

f. Relative Prices of Young Firms

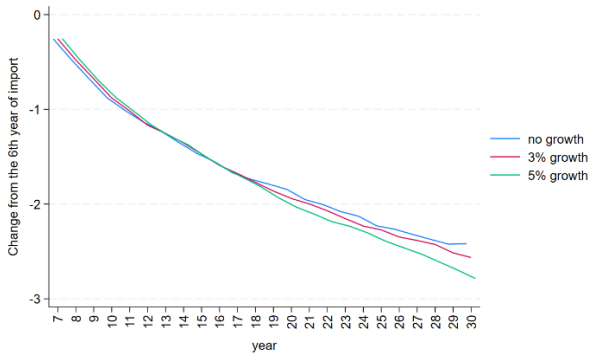

g. Relative Market Shares of Young Firms

Figure 10 Market Dynamism by Importer-product Growth Rate

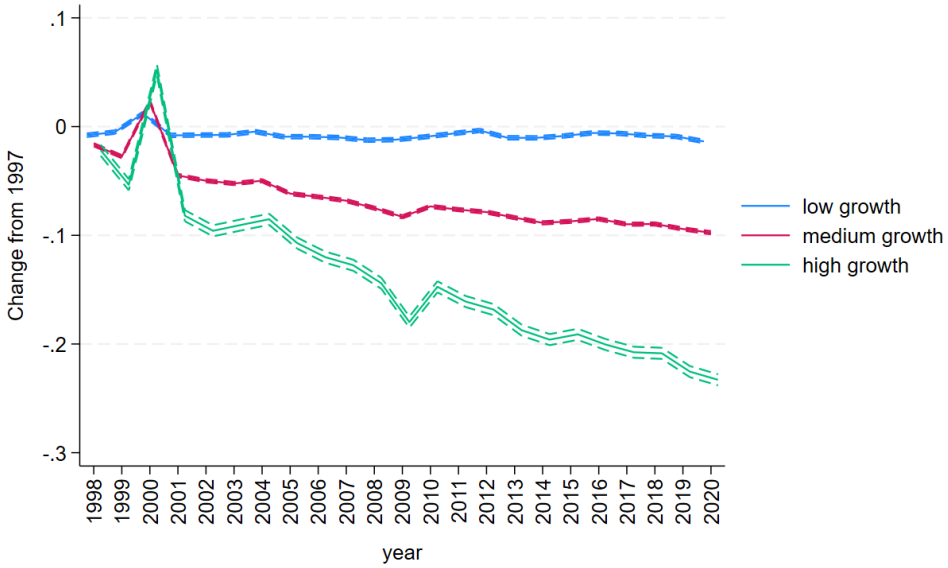

a. Entry rate

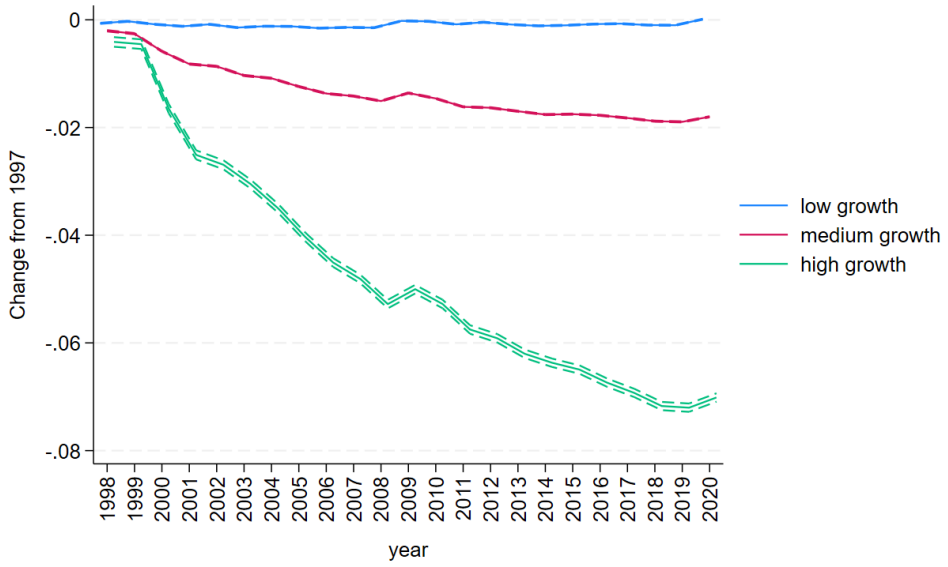

b. Market share reallocation

Note: This figure presents the change in entry rate and market share reallocation (defined in Eqs 2 and 4) over calendar year by growth rate. We calculate the annualized growth rate for each *ih*, then separate markets into high- (above 75th percentile), medium- (25th-75th percentile), and low- growth (below 25th percentile) markets. Dashed lines indicate the 95% confidence intervals for the estimates of coefficients of year dummies.

Appendix: Comparing Market Dynamism in the US: Imports and Establishment Data

The US data allows us to check whether measures of dynamism in imports data correlates (across industries and over time) with the closely related measure of dynamism in the establishment data from the Business Dynamics Statistics (BDS). The BDS provides statistics on the universe of establishments in each NAICS (2017 version) 4-digit industry. In this exercise, we use two tables BDS, "4-digit NAICS" and "4-digit NAICS by age".

"4-digit NAICS" provides statistics on establishment entry rate and exit rate at NAICS 4-digit level that are analogous to the exporter entry rate and exit rate in the import data. It also provides data on the job creation rate (JCR) and job destruction rate (JDR). The job creation is the rate of employment gains from expanding establishments from year *t–1* to year *t* including establishment startups. The job destruction rate is the rate of employment losses from contracting establishments from year *t–1* to year *t* including establishments shutting down. The absolute value of the difference between JCR and JDR is analogous to our measure of market share volatility using import data.

"4-digit NAIC by age" provides information on total employment of establishments with different lengths of tenure on the market. We calculate the employment share of establishments that are younger than 6 years old and older than 5 years old, by NAICS 4-digit industry. Note that in the calculation of total employment, we include employment in the establishments that are left-censored. They account for 65% of the total establishments in 1991, 51% in 2005 and 44% in 2020.

To make direct comparison, dynamism measures calculated using US import data at the HTS 10-digit level are concorded to NAICS 4-digit. For each NAICS 4-digit industry for a given year, we take the average of each measure over all the HTS10 digit product that are matched. We regress the average value from the imports data on the corresponding measure from the establishment data, and report the results in Table A1.

We find a strong positive correlation for each measure. The NAICS industries where entry rates are slowing the fastest over time in the BDS are also the NAICS industries where entry rates into individual import markets are slowing the fastest. Exit rates, market share volatility, and the market shares of young and old exporters similarly show a strong correlation between import markets dynamism and establishment level dynamism.

Figure A1: Market Dynamism over Time (United States)

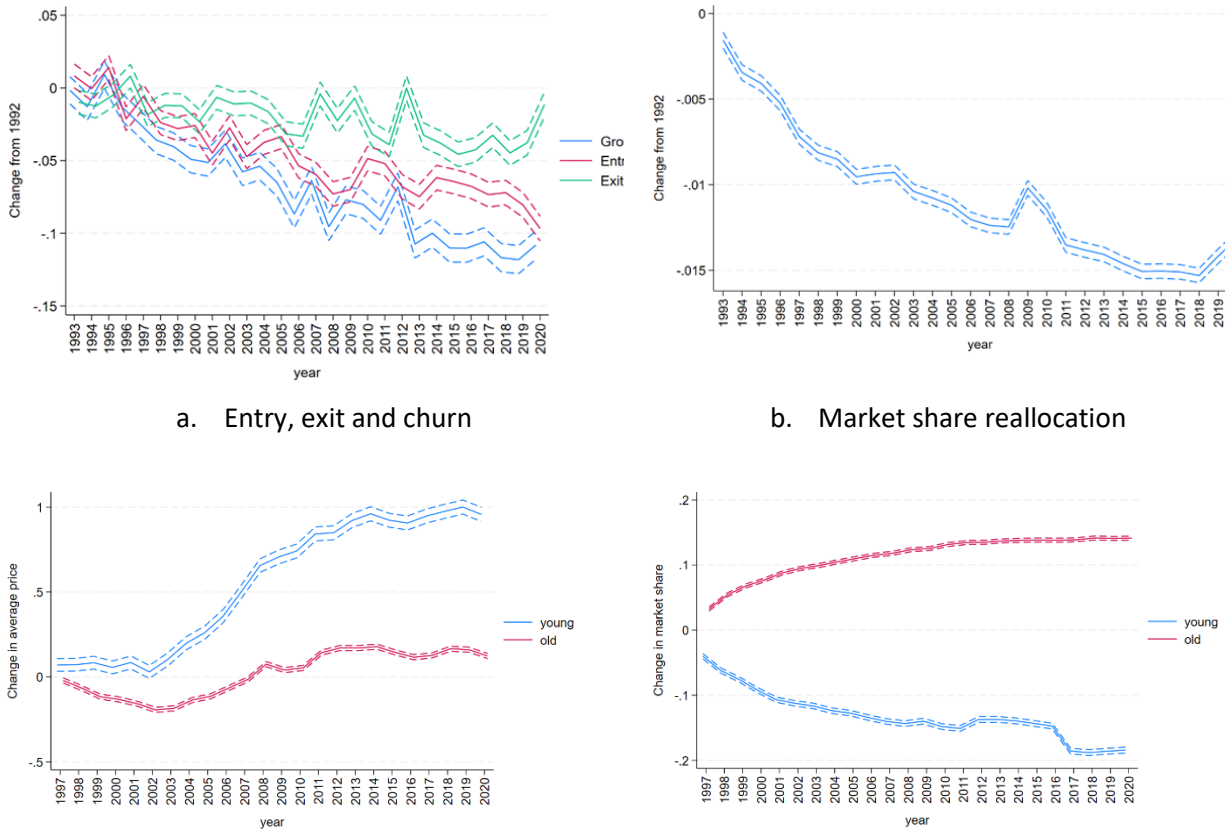

a. Entry, exit and churn

b. Market share reallocation

c. Weighted average price

d. Market share

Note: This figure presents the trend of different market dynamism measures using US trade data over 1996-2020.

Table A1: Correlating Dynamism - US Trade and Establishment Data

US Imports	Entry rate	Exit rate	Market share volatility	Market share of young exporters	Market share of old exporters
US Establishment	(1)	(2)	(3)	(4)	(5)
Entry rate	0.2977***				
	(0.0660)				
Exit rate		0.3582***			
		(0.0610)			
Job churning			0.0419***		
			(0.0040)		
Employment share of young establishments				4.7168***	
				(0.2804)	
Employment share of old establishments					0.3490***
					(0.0106)
constant	0.1192***	0.1052***	0.0189***	-0.1093***	0.8534***
	(0.0051)	(0.005)	(0.0009)	(0.0201)	(0.0033)
N	2655	2655	2760	2760	2300
R2	0.6182	0.6565	0.7943	0.1008	0.4921
NAICS FE	YES	YES	YES	YES	YES

Notes: The dependent variable in each column is the trade dynamism measure for the US market, aggregated to the NAICS 4-digit level in each year. The explanatory variable is the corresponding dynamism measure constructed from US establishment level data. “Old” exporters/establishments are older than 5 years, “Young” exporters/establishments are those less than or equal to 5 years old.