

The Distributional Effects of Firm Demand Changes: Evidence from U.S. Worker-Owner Data*

Sean Wang[†] Samuel Young[‡]

October 24, 2025

Abstract

This paper analyzes which individuals benefit when firm demand increases and whether the same groups bear the costs of demand decreases. We develop a flexible framework to estimate the incidence of firm shocks with data on each firm’s workers and owners. Specifically, we use linked firm-worker-owner tax data from U.S. pass-through firms to track how each dollar that firms distribute as wages and business profits is allocated among individuals. These data allow us to conduct the first joint analysis of how changes in firm demand affect both workers and owners. We leverage export-demand variation and value-added fluctuations as firm-specific demand shocks. The incidence of these shocks is highly unequal. Individuals in the top 1% of the national income distribution receive up to 60% of the income changes, while those in the bottom 50% receive less than 15%. This unequal distribution arises because firm owners receive most of the income changes from the shocks and are disproportionately in the top of the income distribution. Moreover, the incidence is asymmetric for positive versus negative demand changes. Workers bear 26% of the losses from a negative shock that reduces firm value added by 14% but receive only 10% of the gains from a similarly sized positive shock. This asymmetry arises primarily from the additional costs of job loss for workers following negative demand shocks to the firm.

*We thank Emek Basker, Henry Hyatt, Martha Stinson, and audiences at the 2025 ASSA meetings, 2025 CESifo Area Conference on Labor Economics, 3rd Oslo Public-Labor Workshop, and the U.S. Census Bureau for their helpful comments. Young is grateful to the Washington Center for Equitable Growth and the Upjohn Institute for financial support. All errors are our own. *Disclaimer: Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 7508149: CBDRB-FY25-CES007-03, CBDRB-FY25-CES007-09). Samuel Young was also a part-time employee of the U.S. Census Bureau at the time of writing this paper.*

Contact information: sean.y.wang@census.gov and sgyoung@asu.edu.

[†]U.S. Census Bureau, Center for Economic Studies

[‡]Arizona State University, W.P. Carey School of Business

1 Introduction

Which individuals benefit from firm demand increases? For example, if a firm suddenly wins a major new customer, what share of the income gains would go to individuals in different parts of the income distribution or different demographic groups? Conversely, if the firm loses the customer, would the same groups that receive the gains also bear the losses? Answering these questions is crucial to understanding how policies that encourage firm growth or provide firms with downside protection affect income inequality. Moreover, the answers speak to theoretical debates about wage-setting (e.g., the extent to which owners insure workers against firm-specific earnings risk).

Prior research on these questions has *separately* studied how firm shocks affect workers and owners. On the one hand, workers receive a small share of the gains from positive firm productivity or demand shocks (Card et al., 2018), but they can experience large earnings declines from negative shocks (e.g., the cost of job loss (Jacobson et al., 1993)). On the other hand, private firm owners face substantial income risk (Moskowitz and Vissing-Jørgensen, 2002; Hall and Woodward, 2010; DeBacker et al., 2023), but there is little evidence on how their income responds to firm-specific shocks. However, no prior research has *jointly* analyzed how firm demand changes affect workers and owners. A joint analysis is necessary to compare the relative magnitudes of these channels and determine who ultimately receives most of the income gains and shoulders most of the losses.

This paper provides a comprehensive analysis of the incidence of firm demand changes using linked *firm-worker-owner* data that cover over a quarter of private employment in the U.S. We use administrative tax data to identify firms’ workers and owners and measure who receives each dollar of income the firm distributes in wages and business income. These disaggregated individual-level income measures allow us to estimate the pass-through of firm shocks to firms’ workers and owners. We use these estimates to characterize the incidence of firm-specific demand changes.

Our analysis uncovers three stylized facts. First, the incidence of firm demand changes is highly skewed toward high-income individuals. We find that 30–60% of income changes from the firm shocks accrue to individuals in the top 1% of the national income distribution. The shock incidence—that is, the distribution of income changes—is substantially more unequal than the distribution of baseline income that firms pay to workers and owners. Second, this unequal shock incidence is driven by firm owners. We show that the pass-through of demand shocks to owners’ total income is six to nine times larger than the pass-through of the same shocks to workers’ income. Because owners are concentrated in the top of the income distribution, these pass-through differences drive the unequal incidence. Third, the incidence is asymmetric for positive versus negative demand changes. Although workers receive a small share of the benefits of positive shocks, they bear a *larger* share of losses from negative shocks, especially after we account for income changes due to job loss.

For this analysis, we combine administrative Form W-2 and Form K-1 tax filings to identify all workers and owners of S corporations in the U.S. These firms provide an ideal context for studying the incidence of firm-level shocks because they are privately held by individual shareholders and must annually allocate all business profits (net income) to these individuals (i.e., they are pass-through businesses). Moreover, these firms employ 27% of U.S. workers, and their substantial growth has

contributed to the rise in top-income inequality (Cooper et al., 2016). Since we observe the identities of the firms’ workers and owners, we can calculate the share of income from each firm that goes to different groups. Furthermore, we can track workers and owners after they leave a firm, allowing us to measure income changes resulting from job loss or business closure.

We use these data to estimate the distributional effects of firm demand changes. We analyze two complementary firm-level shocks: export-demand shocks and fluctuations in firm value-added.¹ The export-demand shocks pass several tests supporting that they are exogenous, firm-specific, and non-skill-biased (i.e., *Hicks neutral*). However, we can only construct these shocks for exporters. In contrast, we can construct value-added shocks for almost every firm, allowing us to characterize the distribution of income changes for a much broader sample. Furthermore, we can separately analyze the incidence of positive versus negative value-added shocks. We find that the export-demand and value-added shocks yield similar results on an overlapping set of firms, suggesting that the properties of the two shocks are similar.

We construct the export-demand shocks using changes in countries’ demand for specific products as demand shifters for firms that previously exported to those countries (Hummels et al., 2014; Garin and Silvério, 2024).² We find that an increase in export demand increases workers’ and owners’ income from the firm, but the effect is much larger for owners. In response to the same-sized demand increase, worker income increases by 1% while owner income, including wages and business income, increases by 9%. The larger pass-through of shocks to owners than to workers holds whether we analyze (1) income received only from the treated firm (capturing the firm’s wage and profit distribution policy) or (2) total income from all firms (capturing earnings changes for individuals who leave the treated firm and spillover effects on income from other firms). We find that the effects on earnings for workers are also biased towards high earners: the effects for workers in the top 10% of the income distribution are two to three times as large as the pass-through for workers in the rest of the distribution.

We use these estimates to calculate the incidence of firm demand shocks across the income distribution. Specifically, we use the heterogeneous worker and owner treatment effects to construct the share of the total change in income from the shock that goes to each part of the income distribution. This analysis requires observing firm owners in order to place them in the income distribution. The incidence of these shocks is highly unequal; individuals in the top 1% of the national income distribution receive 58% of all income changes, while those in the bottom 50% receive 9%. The unequal incidence is driven by owners, who receive 83% of all income changes and are often at the top of the income distribution. Additionally, owners capture twice as large a share of income *changes* from the demand shocks as their *baseline* share of all the income paid to workers and owners. Moreover, we demonstrate that accounting for owners in wage data is important for accurately distributing the incidence, as owners receive a substantial portion of their compensation

¹While the export shocks capture changes in demand (i.e., output prices), the value-added shocks also include changes in productivity. For brevity, we refer to both as *demand shocks*, but the value-added shocks are closer to *revenue TFP shocks*.

²For example, consider two shoe producers, one of which historically exported to France and the other to Japan. If France experiences an increase in the demand for shoes, this would lead to a larger export-demand shock for the first firm.

in the form of wages. If we instead classified W-2 wage payments to owners as worker income, we would overstate the workers’ share of the incidence by 40%.

To analyze a broader set of firms, we also construct shocks based on firm-specific value-added innovations (Guiso et al., 2005; Lamadon et al., 2022). These shocks are the residual change in firm value-added, conditional on granular industry and location controls. We can define these shocks for any firm with consistent tax data, and as a result, we validate our results for exporters with estimates for almost all S corporations. The incidence of these shocks is also highly skewed toward top earners; individuals in the top 1% of the income distribution receive around 30% of the incidence, and individuals in the top 10% receive around 60%. These estimates are smaller than those for the export-demand shock primarily because firm owners in this broader sample are less likely to be in the top 1% of the income distribution. The wide coverage of these shocks also enables us to estimate the incidence by sex and race/ethnicity. We find that women and racial/ethnic minority groups receive a small share of the incidence of these shocks (29% and 24%, respectively). The relatively low share of income changes accruing to these minority groups mirrors their low share of baseline income from firms, largely due to these groups’ disproportionately low rates of business ownership.

Finally, we use the value-added shock to analyze whether positive versus negative demand changes have asymmetric effects on workers and owners. We find that the share of the incidence that workers receive is asymmetric, driven by nonlinear treatment effects for both groups. For positive demand shocks, workers receive a progressively *smaller* share of the income gains as the shock magnitude increases. Conversely, when firms experience negative demand shocks, workers bear a disproportionately *larger* share of the losses as the shock magnitude grows. We find a larger degree of asymmetry when we measure workers’ and owners’ total income, even if they leave the treated firm. For example, using this outcome, workers bear 26% of the losses from a *negative* shock that reduces firm value added by 14% but only receive 10% of the gains from a similarly sized *positive* value-added change. One reason the asymmetry is starker when we include individuals who leave the firm (rather than restricting the sample to job stayers) is that workers face greater extensive-margin exposure to negative shocks than owners (e.g., higher probability of job loss).

Our key contribution is showing that incorporating owner data enables us to answer a broader set of distributional questions than the existing rent-sharing literature (Card et al., 2018).³ First, by mapping all firm profits to individual owners, we can go beyond calculating only the overall share of *rents* that workers receive. Instead, we fully characterize the incidence of firm demand changes across the entire income distribution and by demographic groups. Second, we provide a unified framework to incorporate multiple income change margins into the incidence calculation. For example, we can calculate the incidence (1) only including income changes for workers and owners who stay at the firm and (2) incorporating the total income changes for workers and owners who

³A closely related paper is Kline et al. (2019), who find that the rents from successful patent applications disproportionately go to owners and high-earning workers. In addition to the above contributions, a key difference between our papers is the nature of the shocks. Specifically, we emphasize the importance of non-skill-biased (Hicks-neutral) shocks for our analysis and implement tests that support this assumption. In contrast, the patents Kline et al. (2019) study may be skill-biased, which could explain the disproportionate effect on high-earners. We also use value-added innovations to characterize the distribution of income for a much broader sample of firms and document the asymmetric responses to positive versus negative shocks.

leave the firm (i.e., including the costs of job loss and business closure). Our framework can also be applied to flexibly estimate the distributional incidence of other firm-level shocks and policies. It is particularly useful for studying shocks that may jointly affect a firm’s survival, hiring and firing choices, and wage and profit-sharing policies, because it consistently aggregates the income changes from all these margins.

Additionally, observing owners allows us to overcome methodological challenges in the rent-sharing literature. For example, we show the importance of correctly classifying wage and salary payments to owners. If we misclassify owner wage income as worker income, we would overstate workers’ share of the incidence by almost 40%. Furthermore, because owners are often at the top of the income distribution, the bias is even more severe when analyzing top-earning workers.

Our key finding is that the gains from firm-specific demand changes at private U.S. firms are concentrated among individuals at the top of the income distribution. This result demonstrates that even non-skill-biased firm shocks can substantially increase *income* inequality due to the concentrated distribution of business income. In contrast, the effect of these shocks on *wage* inequality is substantially less, though still positive. Thus, the effect on top-*income* inequality is primarily determined by the distribution of business ownership, rather than the canonical determinants of wage inequality, such as human capital, technology, or labor market conduct.

Our analysis also contributes to understanding the income dynamics of private business owners. The concentrated incidence of gains in our setting may be surprising given that most firms in our sample are considered small businesses (on average, they have 35–55 employees). Yet, this pattern aligns with [Cooper et al. \(2016\)](#) and [Smith et al. \(2019\)](#), who document that owners of relatively small businesses in the U.S. often have high incomes. We extend their insights by showing that the distribution of income changes from demand shocks is substantially more unequal than the distribution of all baseline income paid to workers and owners. Furthermore, although [Smith et al. \(2019\)](#) document that these owners’ income largely reflects human capital (labor income), we find that their income is an order of magnitude more responsive to firm demand shocks than that of other high-income workers. As a result, the allocation of business income is crucial for understanding why the income dynamics of business owners differ from those of high-income workers.

The results in this paper suggest that high earners (particularly owners) are more likely to benefit from policies targeted towards increasing firm demand. Specifically, firm-level policies, such as government procurement, likely have properties similar to our shocks. Consequently, our results suggest that while such policies benefit workers, most gains still accrue to owners and high-income individuals, even when targeted at relatively small firms. On the other hand, policies aimed at providing firms with downside protection, while still primarily benefiting owners, may generate relatively larger benefits for workers. This contribution builds on prior analyses of the distributional impact of policies using linked worker-owner data. For example, the K-1 and W-2 tax data have been recently used to similarly evaluate the distributional impacts of corporate tax changes ([Risch, 2024](#); [Kennedy et al., 2022](#)), minimum wage increases ([Risch and Rao, 2022](#)), and the Paycheck Protection Program ([Splinter et al., 2024](#)). However, to our knowledge, we are the first to use linked

worker-owner data to estimate the distributional impact of demand changes.⁴

Our finding that workers bear a greater share of the income losses from negative firm shocks than the gains from positive shocks is novel in the empirical literature. It builds on past research that has tested for asymmetric pass-through of positive versus negative shocks with only worker data.⁵ However, research using only worker data is unable to benchmark the magnitude of the change in workers’ income against the change in owners’ income resulting from the same shock. In contrast, since we directly estimate the change in owners’ income, we can calculate which group bears a greater share of the total income changes. Furthermore, similar to [Chan et al. \(2023\)](#) and [Friedrich et al. \(2024\)](#), this analysis highlights the importance of including the extensive margin when analyzing the relative firm-specific income risk that workers and owners face. Specifically, we find that the asymmetric incidence arises primarily from the additional costs of job loss for workers following negative demand changes.

Moreover, this evidence that the incidence of demand changes is asymmetric contradicts the predictions of several standard wage-setting theories. Specifically, these theories share the prediction that workers should receive the same share of the benefits from positive firm-specific demand shocks as the losses from negative shocks. For example, theories of *firm insurance* assume that owners insure workers against negative shocks in return for receiving the profits from positive shocks, implying symmetric incidence.⁶ Similarly, bargaining models imply that the share of surplus changes that workers receive should be equal for surplus increases and decreases. This implicit prediction of incidence symmetry affects the implications of these models for the sources of workers’ income risk and whether firm-specific shocks impact overall income inequality. In contrast, we document incidence asymmetry. While there are some caveats in directly mapping our empirical estimates to these model predictions, more precisely testing the symmetric incidence of demand changes is a promising area for future empirical and theoretical research.

The rest of the paper is structured as follows. Section 2 outlines our framework for estimating the incidence of firm shocks. Section 3 describes our linked firm-worker-owner data. Section 4 documents the unequal incidence of the export-demand shocks. Section 5 presents the incidence estimates for the value-added shocks. Section 6 discusses the results.

2 Firm-Level Shock Incidence Conceptual Framework

In this section, we outline our framework for estimating the incidence of firm-specific shocks. By *incidence* we mean the share of the total income gains or losses caused by the shock that accrues to different groups of individuals (e.g., workers versus owners, different parts of the income distribution,

⁴In standard models of firm production and wage-setting, the distributional effects of corporate taxes or minimum wages are not necessarily similar to those of demand shocks. Minimum wages change input prices, while corporate taxes change post-tax profits for owners. Additionally, both policies are market-wide shocks. In contrast, our demand changes affect firm-specific value added and, consequently, the pre-tax surplus that can be split between workers and owners.

⁵[Juhn et al. \(2018\)](#); [Cho and Krueger \(2022\)](#); [Chan et al. \(2023\)](#); [Grübener and Rozsypal \(2023\)](#); [Acemoglu et al. \(2023\)](#); [Friedrich et al. \(2024\)](#); [Mertens et al. \(2024\)](#); [Merkle \(2024\)](#) test for asymmetry with worker data and arrive at mixed conclusions.

⁶See [Baily \(1974\)](#); [Azariadis \(1975\)](#) for theoretical justifications and Section 6 and Appendix D for further discussion.

or demographic groups). The incidence depends on both group-specific treatment effect estimates and the baseline share of income from the firm that each group receives.

Firm Shock Assumptions Our goal is to analyze firm demand changes that are orthogonal to market-wide confounders or relative productivity shifts among different types of workers. More formally, we are interested in firm shocks that satisfy three assumptions.⁷ First, we want the shocks to be *exogenous* demand shifters that do not reflect other factors like labor supply shifts. Second, we want the shocks to be *firm-specific*, that is, we do not want firms in the same labor market to experience correlated shocks. Market-wide shocks could affect workers' wages through several channels that we want to exclude (e.g., they could increase the market-level wage by moving along an aggregate labor supply curve). Finally, we want the shocks to be *Hicks neutral* (i.e., non-skill-biased). Hicks neutrality ensures that relative productivity changes between types of workers or between labor and capital do not drive differential effects between groups of workers and owners.

To formalize the assumptions in a stylized framework, consider a firm j with revenue

$$R_j = P_j \cdot A_j \cdot F_j(K_j, L_j, M_j) \quad (1)$$

where P_j is the product's price, A_j is the firm's total factor productivity (TFP), and F_j is the firm's production function over capital, labor, and materials. We are interested in analyzing changes in $P_j \cdot A_j$. Although $P_j \cdot A_j$ includes changes in the firm's product demand via P_j and changes in the firm's Hicks-neutral productivity via A_j , we refer to the shocks as demand shocks for simplicity.⁸ We can decompose $P_j \cdot A_j$ into a firm-specific z_j and market-wide component $\psi_{m(j)}$

$$\ln(P_j \cdot A_j) = z_j + \psi_{m(j)}. \quad (2)$$

This framework helps illustrate the necessary assumptions for a candidate firm shock z_j . Shock exogeneity implies that z_j is orthogonal to other changes in the firm's production function F_j or the price of inputs. Firm-specificity implies that z_j is orthogonal to $\psi_{m(j)}$. Finally, since z_j scales up the entire production function, it will always be Hicks neutral because it is orthogonal to changes in the relative productivity of the different inputs (which are also captured in F_j). In Section 4, we evaluate several testable implications of these assumptions for the export-demand shock.

Estimating the Incidence of Firm Shocks We define the incidence of a firm shock z_j on a specific group as the share of the total income changes caused by the shock that goes to individuals in that group. For example, if a shock increased workers' total salaries by \$20,000 while boosting owner profits by \$80,000, workers would receive 20% of the incidence. Conceptually, this notion of

⁷While our incidence framework could be applied to shocks that do not meet all these assumptions, we impose these restrictions due to the economic and policy relevance of analyzing this type of shock.

⁸The first export-demand shock we analyze only includes variation in product demand, so it is a true demand shock. The second value-added shock we analyze plausibly includes variation in both demand and productivity. An alternative way to think about $P_j \cdot A_j$ is that it is the firm's total *revenue* factor productivity (TFPR). TFPR is a common proxy for a firm's TFP because any estimate of TFP based on revenue data will capture variation in TFP and output prices (Foster et al., 2017).

incidence is analogous to distributional tables that policy agencies construct for tax-policy changes, which show what share of the changes in tax burden is borne by households in each part of the income distribution (e.g., [Habib and Heller \(2022\)](#)). Indeed, our framework could be used in other contexts to produce distributional tables for government policies implemented at the firm level.

Next, we illustrate how to estimate the incidence of z_j on workers versus owners, but the framework can be flexibly applied to other groups. With our firm-worker-owner data, we can estimate how a firm shock affects workers' and owners' income. Specifically, we estimate the following specifications where w_i refers to worker income and π_i refers to owner income:

$$\ln(w_i) = \gamma^W + \alpha^W \times z_{J(i)} + \varepsilon_i \quad (3)$$

$$\ln(\pi_i) = \gamma^O + \alpha^O \times z_{J(i)} + \varepsilon_i. \quad (4)$$

The α^W and α^O coefficients measure the *proportional* effect of z_j on worker and owner income (e.g., a shock increases worker income by 5% and owner income by 10%). However, our notion of incidence depends on the *level* changes in income (e.g., a shock increases worker income by \$5,000 and owner income by \$50,000). Accordingly, we define the worker incidence for a specific firm as:

$$\beta_j = \left(\sum_{i:J(i)=j} \alpha^W \times w_i \right) / \left(\sum_{i:J(i)=j} \alpha^O \times \pi_i + \sum_{i:J(i)=j} \alpha^W \times w_i \right) \quad (5)$$

where the two summations are over all the firm's workers and owners. By multiplying the percentage treatment effects, α , by baseline income, w and π , we convert proportional changes into level changes. Finally, we define the share of the shock incidence borne by workers as the β_j for the average firm:

$$\bar{\beta} = \frac{\alpha^W \times \bar{s}^W}{\alpha^O \times \bar{s}^O + \alpha^W \times \bar{s}^W} \quad (6)$$

where $\bar{s}^W = E[s_j^W]$ is the average firm-level share of income at baseline that workers receive and $\bar{s}^O = E[s_j^O]$ is the average firm-level share of income that owners receive.⁹ Equation 6 highlights the key inputs for estimating incidence: the treatment effects of a firm-specific shock on workers and owners, and the average baseline share of income from each firm that goes to workers and owners.

Framework Flexibility and Limitations An advantage of our incidence framework is that it can easily be extended to consider other groups of individuals or income concepts. First, to extend the framework to groups beyond workers versus owners, we can implement Equation 6 with different group-specific treatment effects and baseline income shares. For example, to estimate the incidence across the income distribution, we would estimate separate α s for each income group and use income-group-specific baseline shares. Second, we can modify the sample and definition of w

⁹Because the incidence is a nonlinear function of the earnings share, the incidence for the average firm is different than the average β_j across all firms. Our approach is analogous to evaluating marginal effects at the mean rather than computing average marginal effects across all units in nonlinear models.

and π in equations 3 and 4 to calculate the incidence under alternative income change concepts. For example, in our analysis, we consider both workers’ and owners’ income only from the original treated firm and their total income from all firms, even if they leave the treated firm. These variations allow us to analyze the distinct questions of (1) who bears the incidence only incorporating changes in the firms’ wage and profit-sharing policy, and (2) who bears the incidence when also incorporating extensive margin changes (e.g., workers and owners leaving the firm due to layoffs or firm closure).

The framework also illustrates which effects of firm demand changes we exclude from our incidence concept. Specifically, we only include the effects of the shock on the firm’s original workers and owners. For example, when we estimate the incidence across income or demographic groups, the underlying sample includes only the firms’ incumbent workers and owners in these groups. Consequently, we exclude the effects of firm shocks on several relevant groups. First, we exclude spillover effects on the workers or owners of non-treated firms. Such spillovers could arise from input-output linkages or via equilibrium labor-market changes. Second, we exclude the effects on consumers. Finally, our incidence concept currently excludes the effects of a firm’s new hires. However, we are working on an extension that will allow us to incorporate new hires.

Our framework is related to papers that use the degree of *rent-sharing* as a measure of worker bargaining power (Card et al., 2018). In our case, if we calculate Equation 6 only including workers and owners who remain employed, we would also obtain an estimate of worker bargaining power.¹⁰ However, our framework’s flexibility allows us to estimate several other concepts beyond worker bargaining power. For example, we can calculate the incidence across other income and demographic groups, and we can incorporate additional income change margins into the incidence calculation. Finally, our incidence concept differs from analyses of how changes in firm demand and productivity affect the labor share (Kehrig and Vincent, 2021; Mertens and Schoefer, 2025).¹¹

3 Linked Firm-Worker-Owner Data in the U.S.

Our analysis is based on linked firm-worker-owner data for a substantial share of U.S. firms and workers. Specifically, we combine W-2 and K-1 tax data to identify all workers and owners of S corporations and measure who receives each dollar of income from these firms. We use our linked data to show that the distribution of total income from S corporations is substantially more skewed toward high-income, white, and male individuals than the distribution of wage and salary payments.

3.1 Taxation and Owner Compensation at S Corporations

We analyze firms structured as S corporations because administrative data provide detailed information about their workers and owners. These are private firms that are generally closely held with

¹⁰Specifically, in a “union” wage-bargaining model with the threat of shutdown (e.g., Abowd and Lemieux (1993)), our estimate of the workers’ share of the incidence would identify workers’ bargaining power.

¹¹Specifically, we measure changes in workers’ and owners’ incomes relative to their *counterfactual earnings* in the absence of the firm shock. For example, consider a worker who is laid off but immediately finds an equally paid job elsewhere, experiencing no income loss. In our framework, this worker bears none of the incidence of the negative shock. In contrast, absent other firm-level changes, such a layoff would reduce the firm’s labor share.

few shareholders (they are restricted to, at most, 100 individual owners). These firms are taxed as *pass-through entities*, meaning that they do not pay corporate taxes but, instead, their owners pay individual income taxes on their share of the firm’s profits. More precisely, the firm distributes net income as follows. First, the firm calculates its annual net income as total revenue minus total costs (costs include the cost of goods sold, wages and benefits, other expenses, taxes, interest, and depreciation). Then, the firm distributes total net income across the owners, proportional to their ownership shares. These distributions to each owner can be positive or negative (i.e., business income losses), but there are some restrictions on whether owners can claim the losses on their individual tax returns (Goodman et al., 2023). Due to their structure, S corporations provide an ideal setting to analyze the incidence of firm shocks. First, we can observe the universe of owners and their business income based on tax data.¹² Second, the distribution of profits at pass-throughs is straightforward (e.g., there are no retained earnings or financially motivated stock buybacks).

S corporations are an increasingly important segment of the U.S. economy. In 2022, these firms employed 27% of all workers and operated 46% of establishments (U.S. Census Bureau, 2022) (see Table A2). In comparison, Schlingemann and Stulz (2022) estimate that in 2018 publicly traded firms employed between 18–22% of domestic U.S. workers. Additionally, Cooper et al. (2016) show that the rise of pass-through firms over the past two decades explains a sizable share of rising top-income inequality. Moreover, despite the perception that S corporations are only prevalent in certain industries, Table A2 shows that they are prevalent across almost all industries.

Owners of S corporations can receive both business income and wage and salary compensation from their firm, further underscoring the importance of linked worker-owner data.¹³ In our sample, the average owner receives 45% of their total compensation as wage and salary payments.¹⁴ Consequently, relying solely on business income data to measure owner compensation would substantially understate their total income. Moreover, analyzing wage and salary data without identifying the owners would lead to misclassifying some owners as workers. This bias is especially important for top-earning “workers” (Hyatt et al. (2020) show that for 79% of S corporations, at least one of the top-three wage earners is an owner). In Section 6, we argue that these biases are also plausible for other types of U.S. firms. We use our linked data to address these biases by classifying owner wage and salary payments as owner compensation.¹⁵

¹²There is no comprehensive administrative data on C corporation owners (public or private). Additionally, although partnerships are also pass-throughs, they are often owned by other businesses, making it difficult to identify their ultimate owners (Love, 2021).

¹³The IRS requires S corporations to pay owners who are actively involved in the company a “reasonable salary” in addition to business income compensation (see IRS (2008) and IRS (2025) for details).

¹⁴Smith et al. (2019) find that the average S corporation owner receives 66% of their compensation from wages and salaries. Since we impose a 10-employee minimum for our sample, it is more skewed towards larger firms, where owners receive a smaller share of their compensation in the form of wages and salaries.

¹⁵There are several reasons why classifying income by ownership status (owners versus workers) is more appealing than trying to separate capital versus labor income. Specifically, the distinction between labor and capital income in this setting is ambiguous. First, tax incentives heavily influence whether firms pay owners in wages or business income (Smith et al., 2022). Second, it is unclear whether owner business income is actually *capital income*. For example, Smith et al. (2019) argue that 75% of pass-through business income should be considered labor income. Motivated by this evidence, Piketty et al. (2022) write that “In the real world, the frontier between capital and labour income flows is often fuzzy—or at least more difficult to draw than what is generally assumed in theoretical models.”

3.2 Firm-Worker-Owner Tax and Survey Data

We link administrative tax records and survey data to construct a comprehensive dataset that combines firms’ income statement data with detailed information on their workers and owners. See Appendix Section C for details on our data and sample construction. Our primary worker and owner data sources are K-1 and W-2 tax forms. All S corporations are required to file a Schedule K-1 form for each of their owners. These forms provide the identity of all owners and include the amount of annual business income each owner receives from the firm.¹⁶ All firms must also report wage and salary payments to individuals on Form W-2. The W-2 data provide worker-firm linkages and total worker earnings, including wages, salaries, bonuses, tips, and exercised stock options. Finally, we use Schedule C tax data to measure individuals’ self-employment income. We link together these tax records using the Census’s Protected Identification Keys (PIKs). This allows us to construct a comprehensive measure of each individual’s *annual earned income* from all firms and from self-employment. However, we are unable to analyze household-level *total taxable income*, which includes capital gains, dividends, and non-market income (e.g., Social Security and unemployment). We use our measure of earned income to rank individuals in the national income distribution.¹⁷ Additionally, we merge data to identify the age, sex, and ethnicity/race of each owner and worker (from the SSA NUMIDENT, Decennial Censuses, and ACS).

We use other Census datasets to build a firm-level panel with detailed firm characteristics and income statement data. Our starting point is the Longitudinal Business Database (LBD) (Chow et al., 2021). The LBD covers the universe of private, non-farm firms and provides us with firm employment, detailed industry, legal form of organization (e.g., S corporations), and location. To construct the export-demand shock, we merge data from the Longitudinal Firm Trade Transactions Database (LFTTD) that provides us with annual firm-by-product-by-country export values for the universe of exporters (Kamal and Ouyang, 2020). Finally, we measure firms’ revenue and value added using firm-level tax data on the Business Register.¹⁸ We link all firm-level datasets and the K-1s and W-2s using Employer Identification Number (EIN) linking and fuzzy name and address matching. However, our unit of analysis is at the firm level, which is more aggregated than EINs. Specifically, a firm corresponds to the set of EINs ultimately owned or controlled by the same parent legal entity.¹⁹

¹⁶Due to limitations in which K-1 boxes we have access to at the Census, we only observe owners’ income from the firm’s “ordinary business income” from Box 1 and not the firm’s other sources of net income (e.g., rental, interest, capital gain, or dividend income). This motivates our sample restriction to non-financial firms because the other income sources are more likely to be substantial for financial firms.

¹⁷Specifically, each year, we rank *individuals* based on their total W-2 wage and salary income, their K-1 business income, and their Schedule C self-employment income. Consequently, our income percentile measures are at the individual level rather than the household level. We use publicly released income percentile thresholds from Auten and Splinter (2024) to assign individuals to income percentiles. We adjust their reported thresholds to account for the fact that our income measure is at the individual level and only includes a subset of total income. In 2018, the nominal thresholds for the 50th, 90th, and 99th percentiles were \$38,067, \$117,760, and \$339,046, respectively.

¹⁸We define revenue as “Gross receipts or sales less returns and allowances” (Box 1C on IRS form 1120-S). We define value added as revenue minus “Cost of goods sold” (Box 2 on IRS form 1120-S).

¹⁹We use the LBDFID variable as our concept of a firm. We account for spurious changes in LBDFIDs over time (documented in Haltiwanger et al. (2013)) by tracking groups of establishments that simultaneously switch between the same LBDFIDs.

3.3 Sample Selection and the Baseline Firm Compensation Distribution

Sample Construction We use the previously described firm-worker-owner dataset to construct our analysis sample. We start with all non-financial and non-real-estate S corporations between 2013 and 2018 (the time range is restricted by our K-1 data, which spans from 2013 to 2022). We restrict our sample to firms with matched W-2, K-1, and Business Register data in the base year to ensure that we observe at least one owner and worker for each firm. To reduce the influence of part-time and part-year workers, we restrict our sample to workers who earned at least \$15,000 in real earnings from the firm in the base year. Finally, we require that the firm has at least 10 employees in the base year.²⁰

To analyze the export-demand and value-added shocks, we implement further shock-specific restrictions. We implement these restrictions and those in the previous paragraph based on the firm’s time $t - 1$ values relative to when the shock is defined. For the export-demand shock sample, we restrict to firms with at least \$5,000 in real exports per worker in the base year (Appendix Figure A2 shows estimates for different threshold values). This ensures that changes in the firms’ export demand substantially impact the firms’ overall outcomes. Across all six base years (2013–2018), this yields a sample of around 115,000 treated firm-year observations (all counts are rounded to comply with Census Bureau guidance on disclosure avoidance). On average, the firms have 59 W-2 workers and 2.9 owners.²¹ Since we can only define these shocks for exporting firms, the sample is concentrated in manufacturing and wholesale trade. For the value-added shock sample, we can include all S corporations with non-missing tax data. However, we restrict the analysis to a random 20% sample for computational simplicity. This yields around 478,000 treated firm-year observations. On average, the firms have 85 workers and 2.1 owners. Since this sample is representative of all S corporations, its industry distribution is more diverse, with a large share of firms in accommodation and food services, healthcare, professional services, retail trade, and construction (see Table A2).

For each shock-specific sample, we construct panels that follow firms, workers, and owners over time, whether or not they continue to meet the baseline sample restrictions. Specifically, we construct a firm-level panel and an individual-level panel that track firms and individuals for five years before and three years after the shocks. We construct two versions of the individual-level panels. The first version only includes individuals (workers and owners) who remain at the treated firm and only includes the income they receive from the treated firm. We refer to this as the *stayers sample*. The second sample includes all individuals who received income from the treated firm in the base year, regardless of whether they remain at the firm. It includes their income from all firms and their self-employment income. We refer to this as the *all-income sample*.

²⁰The vast majority of employer firms in the U.S. are very small (e.g., 1–2 employees). Consequently, if we did not impose a minimum employee threshold, our sample would be dominated by tiny firms.

²¹Our measure of workers is the annual number of W-2s issued by the firm. Since this includes part-year employees, it is an upper bound on the number of full-time equivalent employees at any point in the year. If we scale our employment counts by the average relationship between W-2 and LBD employees in [Stinson and Wang \(2025\)](#), the average numbers of employees are 35 and 55 for the export-demand and value-added samples, respectively. We use these rescaled point-in-time employment numbers in the introduction and discussion sections.

The Baseline Distribution of Firm Compensation We use our linked worker-owner data to calculate the share of baseline income that different groups receive from each firm. Specifically, for each firm j at baseline ($t - 1$), we calculate the income share for group g as

$$s_j^g = \frac{\sum_{i:i \in g, J(i)=j} (w_{i,j} + \pi_{i,j})}{\sum_{i:J(i)=j} (w_{i,j} + \pi_{i,j})} \quad (7)$$

where $w_{i,j}$ and $\pi_{i,j}$ are wages and business income respectively. We then average s_j^g across all firms. These estimates correspond to the necessary baseline shares in the incidence formula in Equation 6.

Figure 1 and Table A3 present our estimates of the baseline income shares and show that the distribution is skewed towards high-income, male, and white individuals. Figure 1 Panel A shows that for the export-demand sample, on average 29% of the income from the firm goes to the top 1% of the income distribution, while only 16% goes to the bottom 50%. The large share of income going to the top 1% is driven by the prevalence of owners in the top 1% (69% of income that owners receive goes to owners in the top 1% while the corresponding number for workers is only 3%). Panel B shows that the baseline distribution for the value-added shock sample is also highly skewed. However, a smaller share goes to the top 1% because a relatively larger share of owners in the broader sample are in the 90–99th percentile of the income distribution rather than the top 1%. Table A3 Panel B also shows the baseline income distribution by sex and race and ethnicity. We again find that business income is more concentrated than wage income, with 76% accruing to men, and 78% accruing to white individuals, although this concentration is less extreme than what we find by segments of the income distribution.

More broadly, these estimates show that unequal rates of business ownership tend to amplify existing *wage* inequality into even greater degrees of *income* inequality. These results are also consistent with research showing that business owners are concentrated in the top of the income distribution (Cooper et al., 2016) and are more likely to be male and white (Fairlie and Robb, 2008). As a result, we find that accounting for firm ownership and the distribution of business income is especially important for understanding income inequality at the top of the income distribution.

4 The Incidence of Export-Demand Shocks

We use our linked worker-firm-owner data to document the unequal incidence of firm-level demand changes. In this section, we analyze demand changes driven by product demand shifts for exporting firms. Increases in export demand raise worker and owner income from the firm, but the effects are substantially larger for owners and high-earning workers. Consequently, more than half of the income changes from these shocks go to the top 1% of the income distribution.

4.1 Export-Demand Shock Empirical Strategy

The export-demand shocks use country-level product demand changes (e.g., a decline in the demand for shoes in France) as demand shifters for firms that previously exported that product to that

country. Our shock construction follows [Hummels et al. \(2014\)](#) and [Garin and Silvério \(2024\)](#).

Shock Construction Consider a firm j that exports to country c by product p pairs. We define the firm’s baseline share of total exports to each country-product pair as

$$\omega_{jpct_0} = \frac{X_{jpct_0}}{X_{jt_0}}. \quad (8)$$

We define these shares using three years of lagged data from the Longitudinal Firm Trade Transactions Database (LFTTD) ([Kamal and Ouyang, 2020](#)). We define a country c ’s change in non-U.S. imports for product p as

$$\sigma_{pct} = 2 \times \frac{M_{pct_{t+1}} - M_{pct_t}}{M_{pct_{t+1}} + M_{pct_t}} \quad (9)$$

where M is the total imports of product p from all countries except the U.S.²² We measure country-by-product imports using the harmonized BACI trade flow data (see [Gaulier and Zignago \(2010\)](#) for details). For each exporting firm, we define their export-demand shock in a *shift-share* form as

$$z_{jt} = \sum_{p,c} \omega_{jpct_0} \times \sigma_{pct}. \quad (10)$$

The demand shock in Equation 10 includes variation from product-wide demand changes (e.g., a global decline in the demand for shoes) and product-country-specific changes. The variation from product-wide changes may pose a problem for our strategy because such variation is less likely to be firm-specific and exogenous.²³ To purge our shocks of the product-wide variation, we follow the recentering approach in [Borusyak and Hull \(2023\)](#) and control for the expected shock if the country-product growth rates were randomly assigned within HS-4 product groups.²⁴ Consequently, all identifying variation is at the country-product level (e.g., country-specific changes in product demand or exchange rate fluctuations). In Section 4.4, we show that these export-demand shocks satisfy several testable implications of the exogeneity, firm-specific, and non-skill-biased assumptions.

Empirical Specification and Outcomes We estimate the dynamic effects of these export-demand shocks on a series of firm- and individual-level outcomes. For each outcome, we estimate

²²We exclude U.S. imports to avoid U.S.-specific shocks (e.g., productivity changes) from affecting the country-specific demand change. The functional form in Equation 9 is [Davis et al. \(1996\)](#)’s symmetric growth rate, which approximates log differences while accommodating zeros.

²³This variation would not be firm-specific if exporters of the same product were clustered in the same labor market ([Garin and Silvério, 2024](#)). Regarding shock exogeneity, it is more plausible that a growing firm selects into producing a *product* that will experience an increase in demand than a country-product pair that will experience an increase in demand.

²⁴ The idea of separating country-product versus product-wide export demand variation is from [Garin and Silvério \(2024\)](#). In practice, we construct the recentered shock as $\bar{z}_{jt} = \sum_{p,c} \omega_{jpct_0} \times \bar{\sigma}_{pt}$ where $\bar{\sigma}_{pt}$ is the average product import growth rate across all countries weighted by the distribution of the country-product pairs U.S. firms export to. This approach also clarifies our identifying assumption relative to the recent shift-share frameworks ([Goldsmith-Pinkham et al., 2020](#); [Borusyak et al., 2022](#)). Specifically, we assume that the shocks are randomly assigned and rely on asymptotics as the number of firms and country-product groups goes to infinity.

the following long-difference specification for different time intervals n

$$\frac{Y_{j,t_0+n} - Y_{j,t_0}}{\text{abs}(Y_{j,t_0})} = \gamma + \alpha_n \times z_{jt_0} + X'_{jt_0}\beta_n + \varepsilon_{jt}. \quad (11)$$

The outcome functional form is an *absolute percent change* that has a percent-change interpretation but accommodates negatives.²⁵ DeBacker et al. (2023) use the same functional form to analyze the dynamics of owner income at pass-through firms. For all estimates, we set $t_0 = -1$ so that the α_n coefficients represent the percent change relative to one year before the shock. We also estimate a *pooled specification* that includes long differences defined until $t + 1$, $t + 2$, and $t + 3$ and estimate one coefficient for the pooled time horizons. X_{jt_0} includes firm- and worker-level controls. In our preferred specification, we only include year fixed effects and the expected product-level shock, $\bar{z}_{jt_0}$ (see footnote 24). We winsorize all outcomes at the 1st and 99th percentiles. We cluster standard errors at the firm level but also show that the standard errors are only slightly larger when we calculate shock-level standard errors following Borusyak et al. (2022). Finally, for all worker- and owner-level regressions, we weight the regression by the inverse number of workers or owners at the firm at time $t - 1$ so that our firm- and individual-level estimates give equal weight to each firm.

Since we define shocks for multiple base years (“cohorts”), we stack data across cohorts and estimate pooled treatment effects. Specifically, for each cohort from 2013 to 2018, we construct firm- and individual-level panels that track treated firms and individuals for five years before and three years after the shock year. We then stack these panels together and estimate the pooled treatment effects. For all specifications, we included cohort fixed effects and we interact all controls by cohort. Consequently, this approach is equivalent to running a separate regression for each cohort, except that we use the pooled OLS regression weights to combine the cohort-specific treatment effects. This methodology is analogous to Cengiz et al. (2019)’s stacked approach to difference-in-differences designs with staggered treatment adoption.

4.2 The Effect of Export-Demand Shocks on Firms, Workers, and Owners

Firm-Level Results Increases in export demand lead to an immediate and persistent increase in firm-level exports, revenue, and employment. Figure 2 plots the dynamic effects of the export-demand shocks. Specifically, it plots the estimated α_n coefficients from Equation 11 for five years before the shock and three years after the shock. To help interpret the coefficient magnitudes, note that the shock units themselves have no natural interpretation (i.e., it is an instrument for firm demand and we report the *reduced-form*). However, the relative effects on different firm outcomes or on workers versus owners are interpretable. A one-unit increase in the shock immediately leads to an 8% increase in total firm exports, which persists for the three years following the shock. Similarly, the shock leads to persistent increases in firm revenue and employment. This evidence shows that

²⁵Both business income and firm-level value added can be negative. This functional form is scale-invariant, can accommodate negatives at time t_0 and $t_0 + n$, and can accommodate zeros at time $t_0 + n$. For positive values, it measures conventional percent changes. For negative baseline values, the functional form implies that $\alpha = 10\%$ will increase the outcome by $.1 \times |Y_{jt_0}|$ (e.g., a 10% increase in business income from a baseline of -\$100 ends up at -\$90).

these are persistent rather than transitory demand shocks (in several wage-setting models, persistent and transitory shocks have different effects on wages (Guiso et al., 2005)). Finally, for all three outcomes, there is limited evidence of pre-trends (i.e., different growth rates between high and low-shock firms). The lack of pre-trends supports the shock exogeneity assumption. Moreover, the lack of pre-trends and stable treatment effects imply that we can use our worker and owner estimates over different time horizons to infer the dynamic treatment effects on both groups.

To summarize the firm-level effects, Figure A1 plots pooled three-year estimates. The shock increases exports by 7.9%, revenue by 3.7%, value added by 3.2%, employment by 2.0%, and total compensation (the sum of wages and business income) by 4.4%. The larger effect on exports than revenue is expected because firms have non-export revenue that is not directly affected by the shock.

Individual-Level Results Next, we estimate how export-demand shocks affect the incomes of workers and owners. For this analysis, we focus on the two samples of individuals described in Section 3. The first *stayers sample* is restricted to workers and owners who were at the treated firm at time $t - 1$ and remained employed afterward (we consider various post-shock tenure restrictions). For this sample, we only include individuals’ income from the treated firm. Estimates from this sample isolate changes in the firms’ wage-setting and profit-sharing policies. In other words, they capture intensive-margin earnings changes for workers and owners who remain at the treated firm while avoiding extensive-margin changes (e.g., earnings changes due to switching firms).

The second *all-income* sample captures a broader set of earnings changes for workers and owners. Specifically, this sample includes all individuals who received income from the treated firm at time $t - 1$, whether or not they remain. For this sample, we include their wage and business income from all firms and self-employment.²⁶ Consequently, this sample allows us to measure the total risk workers and owners face from firm shocks, including the risks of layoff and business closure.

We first analyze the dynamic effects of the export-demand shocks on workers and owners in the stayers sample. Figure 3 plots the α_n coefficients for workers in Panel A and for owners in Panel B. We estimate separate regressions for owners and workers. For both panels, we use a fully balanced panel of individuals who were continuously employed by the firm from time $t - 2$ to $t + 3$. This restriction ensures that composition differences or part-year employment do not drive the results. Panel A shows that a one-unit increase in the shock causes a persistent increase in workers’ income of around 1%. Panel B similarly shows a persistent increase in owners’ income, except the magnitude is almost an order of magnitude larger, around 8–9%. Currently, we do not find statistically significant pre-trends for workers or owners, although the owner estimates are noisy. We are in the process of expanding our sample, which will allow us to obtain more precise estimates.

To provide a summary measure of these results, Table 1 Panel A presents pooled three-year estimates. For these pooled estimates, we still restrict to incumbents, but relax the restriction that

²⁶For this sample, we assign an income of zero to individuals who are not observed in the data in a given year. Since we observe the universe of wage and self-employment income, this imputation would only be incorrect for workers in cases where the worker leaves the country or switches to “under-the-table” employment. For owners, since we do not observe business income from C corporation ownership (e.g., dividends), this imputation may understate owner income in some cases.

the individuals were employed for all three post-shock years.²⁷ The pooled estimates show that the shock increased owner income by 9.1% (SE = 2.4%) and worker income by 1.1% (SE = 0.2%).

Our longitudinal data allow us to expand the income concepts we analyze. First, we separately analyze the effect on owners' business income and W-2 income. Table 1 shows that owners' W-2 income increases by 1.9% (SE = 0.8%) in response to a one-unit shock increase. This shows that although owners' wages are affected by demand changes, their business income is much more sensitive. Second, we estimate the same effects for workers and owners using the all-income sample. The last two columns in Table 1 show that the worker and owner estimates from this sample are 1% (SE = 0.2%) and 7.4% (SE = 1.9%), respectively.²⁸ These estimates show that owners' income is much more sensitive to demand changes than workers' income across both samples and income concepts.

Additionally, we estimate the impact of the export-demand shock on the job transitions of workers and owners. Specifically, Panel C in Table 1 presents three-year pooled estimates on separating from the treated firm and becoming nonemployed.²⁹ We find significant extensive-margin effects for workers but not owners. Specifically, a one-unit increase in the export-demand shock decreases the probability that a worker separates from the firm by 0.7 pct. pts. (SE = 0.3) and decreases the probability of nonemployment by 0.2 pct. pts. (SE = 0.01). The corresponding estimates for owners are -0.2 and -0.1, both of which are insignificant. Overall, these estimates show that workers are more exposed to firm shocks on the extensive margin than owners. In Section 5.4, we demonstrate that this differential extensive margin exposure partially explains why workers bear a greater share of the incidence of negative than positive demand changes.

Because our ultimate goal is to estimate the incidence of these shocks across the income distribution, we next estimate heterogeneous effects for workers in different parts of the income distribution. Specifically, we estimate a modified version of Equation 11 that includes interactions between the shock and worker income percentile groups.³⁰ Table 1 Panel B includes these heteroge-

²⁷While the balanced panel eliminates composition effects, it also restricts the sample in a way that is not ideal for our incidence analysis (e.g., we would only calculate the incidence among workers and owners who remained employed for three years). For the pooled estimates, we only require that the individual was employed at the firm at time $t - 1$ and in one of the $t + 1$, $t + 2$, or $t + 3$ years. We use these estimates from our stayers sample incidence analysis. Reassuringly, the worker and owner estimates are very similar to the balanced panel.

²⁸There are several potential explanations for why the all-income estimates are smaller in magnitude than the stayers sample estimates. First, the all-income baseline income measure is broader (e.g., it includes income from all firms, not just the treated firm). Consequently, the same size income change from the treated firm would lead to a smaller percentage increase in total income. Second, this sample includes extensive margin changes, which could mitigate the effects of the shock (e.g., if workers left the firm rather than face large wage cuts). Finally, since we only require employment at the firm in year $t - 1$, the sample may be relatively less attached to the firm and therefore less affected by the shock.

²⁹We define a job separation as an indicator for not receiving any W-2 or K-1 income from the treated firm. We define nonemployment as an indicator for not receiving any W-2, K-1, or Schedule C income in a given year. Our definition of nonemployment is very strict because it excludes unemployment spells that last less than a full year.

³⁰Specifically, we classify each worker into an income group H_i based on their national income percentile in the year before the shock. We then estimate the following modified specification

$$\frac{Y_{i,t_0+n} - Y_{i,t_0}}{\text{abs}(Y_{i,t_0})} = \gamma + \sum_h \alpha_{n,h} \times z_{J(i),t_0} \times \mathbb{1}[H_i = h] + X'_{it_0} \beta_{n,h} + \varepsilon_{it} \quad (12)$$

where the coefficients of interest are the group-specific treatment effects $\alpha_{n,h}$. We estimate three-year pooled estimates for each group. We allow the year-fixed effects and the average product control to vary by heterogeneity group.

neous effects for workers in the 0–50th, 50–90th, and 90–100th percentiles of the income distribution. Across both the stayers and all-income samples, the treatment effects are around 2.5 times larger for workers in the top 10% of the income (e.g., for the stayers sample, the top 10% estimate is 2.7% (SE = 0.9%) versus 0.9% (SE = 0.2%) for workers in the 50–90th percentiles). This evidence shows that high-income workers are more exposed to export-demand shocks than low-income workers.

4.3 The Incidence of Export-Demand Shocks across the Income Distribution

We use the previous estimates of the export-demand shocks’ impacts on workers and owners to calculate the incidence of firm-specific demand changes across the income distribution. Specifically, we combine the worker and owner treatment effect estimates in Table 1 (including the heterogeneous worker estimates) with the baseline income shares in Figure 1 to calculate the group-specific incidence following Equation 6. Figure 4, Panel A shows the incidence implied by the stayers sample treatment effect estimates. Each bar represents the share of the total income change from the shock that goes to each income group. To calculate the standard errors for the share going to each group, we apply the delta method to Equation 6 and assume the baseline shares are constants. For example, the far-right bar shows that 58% of the incidence goes to individuals in the top 1% of the income distribution. Within each bar, the colors show the contribution from workers (in red) and owners (in blue). This worker-owner decomposition shows that the reason the incidence is so skewed is that owners receive a much larger share of the income change than workers and owners are concentrated in the top of the income distribution (e.g., overall owners receive 83% of the incidence and 69% of owners are in the top 1%). In contrast, even though we find larger treatment effects for high-earning workers, they contribute relatively little to the overall incidence because there are relatively few (income-weighted) high-earning workers compared to owners.

We can compare the distribution of income changes from export shocks in Figure 4 directly to the distribution of baseline income at the same firms in Figure 1, Panel A. This comparison highlights the central role of owners in concentrating the distribution of income gains at the top of the income distribution. The distribution of income changes is significantly more concentrated towards top earners than the distribution of baseline income. Individuals in the top 1% of the income distribution receive less than 30% of the overall wage and business income from the firm in the baseline year, but they receive nearly 60% of the overall income changes due to the export shock. In contrast, while individuals in the 50th–90th income percentiles collect the largest share of baseline income—nearly 40% for all wage and business income—they only account for around 10% of income gains from the export shock. Moreover, although the total income to workers accounts for 58% of all baseline income at the firm, they only account for 17% of all income changes. These differences between the baseline income distribution and our incidence distribution are driven by the significantly higher pass-through of the shock to owner earnings.

These results also highlight the importance of classifying owner wage and salary W-2 payments as owner income, rather than worker income, when assessing the distributional incidence of firm shocks. Specifically, 6% of the total worker and owner income gains or losses from the shock go to

owners via W-2 payments.³¹ Without the owner data, we would not know that these were actually W-2 payments to owners, and we would misclassify this income as worker income. Although the relative response of owner W-2 earnings to the export shock is similar in magnitude to the response of high-income workers, wage payments to owners constitute a meaningful 14% share of all baseline income at the firms in our sample. Since the share of income changes to workers is already low, misclassifying owner wages as worker income would overstate the worker share of the incidence by nearly 40% (from 17% to 24%). Moreover, the owner W-2 payments are concentrated in the top of the income distribution, so the upward bias would be even larger for high-earning workers.

Finally, the incidence of the export-demand shock is similar when we use the all-income sample. Panel B of Figure 4 shows the incidence implied by the all-income sample estimates. The top 1% share of the incidence is 59% versus 58% for the stayers sample, and the overall owner share is 82% versus 83%. This similarity shows that, regardless of whether we include extensive-margin earnings changes, most of the export-demand shock incidence goes to owners and the top 1% of the income distribution.

4.4 Export-Demand Shock Assumption Tests

The export-demand shocks satisfy several testable implications of the exogeneity, firm-specificity, and Hicks-neutrality assumptions.

Testing Shock Exogeneity We provide several pieces of institutional and empirical evidence consistent with the shock exogeneity assumption. Specifically, this assumption implies that the shocks are orthogonal to all other firm-level changes that could affect workers’ or owners’ income (e.g., changes in labor supply, other input prices, or firm productivity). Economically, the most plausible threat to this assumption is that some firms select into faster-growing country-product pairs because of the firm’s expectations about future firm-level changes. For example, firms anticipating a decline in the market wage might choose to export to a faster-growing country.

First, the absence of pre-trends in firm- or individual-level outcomes implies that firms are not selecting country-product pairs based on *prior* firm-level changes. Specifically, Figures 2 and 3 illustrate that firms experiencing larger export-demand shocks did not exhibit differential trends in prior export, revenue, employment, wages, or owner profit growth. Additionally, in unreported results, we show that the country-product level import growth rates, σ_{pct} , exhibit a very low autocorrelation. These facts rule out several types of selection based on anticipated firm-level changes that would violate our shock exogeneity assumption. Specifically, selection bias could only arise if firms select export destinations based on knowledge of future country-product demand changes and future firm-level changes that are unrelated to their respective past changes.

Our second piece of supporting evidence is that the treatment effects exhibit a *dose-response* relationship with firms’ baseline export intensity. Specifically, Figure A2 plots the effect of the

³¹The 6% number is based on using the owner W-2 treatment effect of 1.9% (Table 1) and the fact that 14% of the baseline firm income goes to owner W-2 payments in our incidence formula in Equation 5.

export-demand shock on firm sales and employment for different bins of firms’ baseline exports per worker. For firms with exports per worker less than \$1,000, we do not detect effects on either outcome. In contrast, for firms with exports per worker greater than \$1,000, we detect significant effects that increase with the baseline export intensity. This evidence is inconsistent with firms selecting their exporting destinations based on anticipated future firm-level changes; in that case, we would expect spurious treatment effects even for low-intensity exporters.³²

Testing Firm-Specificity To test whether the shocks are firm-specific, we show that our estimates are quantitatively unchanged when we add time-varying local labor market controls. Specifically, Figure A1 shows that the three-year pooled estimates for firm-level exports, revenue, value added, employment, and total income are virtually unchanged when we add detailed industry by location by year fixed effects.³³ If firms in the same labor market exported to the same country-product destinations, we would expect the labor market controls to attenuate the coefficients.

Testing for Skill-Biased Shocks Next, we show that the effects of the export-demand shocks on firms’ workforce composition are inconsistent with the shocks being skill-biased. This finding addresses the concern that the unequal incidence of export-demand changes is driven by the shocks increasing the productivity of high-skilled workers or owners relative to lower-skilled workers.

We implement the test proposed in Lindner et al. (2022) to identify the skill-bias of firm-specific shocks. To provide intuition, consider a firm in a competitive labor market with a CES production function over high- and low-skilled workers. In this case, a Hicks-neutral shock would not affect the relative *skill ratio* (i.e., the ratio of high- to low-skilled workers). Lindner et al. (2022) show that with imperfectly competitive labor markets, Hicks-neutral shocks can affect the skill ratio. However, the joint responses of the *skill ratio* and the *skill premium* (i.e., the ratio of high- to low-skilled workers’ wages) can be used to identify skill-biased shocks. Specifically, skill-biased shocks are the only type of shock that can simultaneously increase the skill ratio and the skill premium. Intuitively, if a Hicks-neutral shock increases the relative wages of high-skilled workers (e.g., the firm has more monopsony power over high-skilled workers), the firm should move down its high-skilled worker demand curve and hire relatively fewer of these workers. The only time a firm would hire relatively more high-skilled workers despite their higher relative wage is when the shock is skill-biased.

We implement this test by analyzing the effect of the export-demand shock on the share of firms’ workers in different skill groups. First, we split workers into three skill groups based on whether they are in the 0–50th, 50–90th, or 90–100th percentiles of the earnings distribution (the same groups we used to analyze heterogeneous earnings effects).³⁴ Then, for each firm j we calculate

³²Garin and Silvério (2024) provide another compelling test of export-demand shock exogeneity using Portuguese firms that export to multiple country-product pairs. Specifically, firms that export to one fast-growing country-product pair are not more likely to export to other fast-growing pairs, which is inconsistent with export destination selection biasing the results.

³³The industry and location fixed effects include four-digit NAICS industry FEs, commuting zone by two-digit NAICS sector FEs, and county fixed effects. We let all the labor market control coefficients differ by treatment cohort.

³⁴To ensure that the treatment effect of the shock on workers’ earnings does not mechanically affect a firm’s workforce composition, we calculate workers’ average earnings percentiles over the entire sample period to classify them into skill groups. In future analyses, we will use workers’ AKM fixed effects or fixed education levels to further mitigate this concern.

its share of workers in skill group T as s_{jt}^T (owners excluded). Finally, we estimate Equation 11 with the level difference in employment shares, $s_{j,t_0+n}^T - s_{j,t_0}^T$, as the outcome.

The results from the worker composition analysis in Table 2 indicate that increases in export demand cause firms to increase the relative share of low-skilled workers. The first two columns show that a one-unit increase in the export-demand shock leads to a 0.34 pct. pts. (SE = 0.13) increase in the share of low-skilled workers (0–50th percentile) and a –0.28 pct. pts. (SE = 0.13) decrease for middle-skilled workers (50–90th percentile). The share of workers in the highest 90+ percentile also decreases by 0.06 pct. pts., but is statistically insignificant. Interpreted through the skill-bias test, this evidence is inconsistent with skill-biased shocks because the skill ratio decreases while the skill premium increases (i.e., the worker heterogeneity estimates in Table 1). Hence, the shock must be Hicks neutral or low-skill biased. Finally, the fourth column in Table 2 reports that the shock decreased the share of owners relative to total firm headcount by –0.20 pct. pts. (SE = 0.1). This result shows that firm demand increases cause a larger proportional increase in the number of employees than owners.

Standard Error Robustness Our standard errors are similar when we calculate them with firm-level clustering or at the shock-level following Borusyak et al. (2022). Figure A1 shows both standard errors for our firm-level outcomes, and Table 1 Panel A shows the two types of standard errors for our individual-level outcomes. The shock-level standard errors are only slightly larger than the firm-level standard errors—for our main individual-level estimates, the largest proportional difference is 0.002 versus 0.003 for worker total income, and the largest absolute difference is 0.024 versus 0.027 for the stayers sample owner estimates.

5 The Incidence of Value-Added Shocks

To complement the exporter analysis, we construct firm demand shocks based on idiosyncratic value-added fluctuations. We can construct these shocks for almost every S corporation, expanding our sample to represent a quarter of U.S. private-sector employees. The incidence of these shocks is also very unequal across the income distribution. Next, because these shocks have more statistical power, we can expand our analysis in two ways. First, we show that women and minority race and ethnicity groups bear a small share of the incidence. Second, we show that incidence is asymmetric for positive versus negative demand changes; workers bear a larger share of the losses from negative shocks than the gains from positive shocks.

5.1 Value-Added Shock Construction

We define our value-added shocks as annual firm-specific value-added changes, controlling for market-wide trends. Our shock construction is similar to Guiso et al. (2005), Lamadon et al. (2022), and Friedrich et al. (2024), although we impose simpler time series assumptions on the value-added

process. Specifically, for each firm j , we assume that log value added follows a random walk

$$\ln(VA_{j,t+1}) = \ln(VA_{j,t}) + \gamma_{m(j),t} + z_{jt} \quad (13)$$

where $\gamma_{m(j),t}$ captures the market-wide value-added innovation and z_{jt} isolates the firm-specific component. In practice, we recover our value-added shocks z_{jt} as the residuals from regressing first differences in log value added on time-varying market-level controls (four-digit NAICS fixed effects, commuting zone $\times$ two-digit NAICS fixed effects, and county fixed effects). Following our construction of the export shock sample, we construct “stacked” analysis panels at the firm and individual level that include all cohorts of firms with value-added shocks from 2013 to 2018. For computational simplicity, we restrict to a 20% sample. We cluster standard errors at the firm level.

These value-added shocks complement the previous export-demand shocks in several ways. First, they can be defined for all S corporations with non-missing tax data, which expands our sample to include a more representative set of industries and increases the statistical power of the shocks. Second, these shocks are commonly used in the literature that tests for evidence of *firm insurance* (Guiso et al., 2005; Guiso and Pistaferri, 2020). Since we provide a novel test of whether firms insure workers in Section 5.4, it is helpful to study analogous shocks. Third, the value-added shocks could capture firm-specific demand changes (e.g., the firm receiving a new contract) or productivity changes (e.g., the firm adopting new technology), while the export-demand shocks only capture changes in product demand. However, one concern with these shocks is that they may also include some firm-specific changes that do not satisfy our shock assumptions outlined in Section 2. This concern motivates the following exploration and tests of the value-added shock properties.

5.2 Value-Added Shock Validation

Shock Dynamics and Firm-Level Effects The value-added shocks lead to sudden and persistent effects on firm-level outcomes. Figure 5 Panel A plots the dynamic effects of the value-added shocks on firm-level value added. Specifically, it plots the estimated α_n coefficients from estimating Equation 11 with firm value added as the outcome. For all our analyses of the value-added shocks, we only include cohort-by-year fixed effects as controls. The estimates show that the value-added shock is relatively orthogonal to lagged changes in value added (e.g., the pre-period estimates are never more than 5% in absolute value). However, the shock leads to a large and persistent increase in firm-level value added of around 70%. These estimates show that value-added shocks exhibit dynamics similar to those of export-demand shocks. Next, to summarize the shock’s effects on firm-level outcomes, Figure 5 Panel B plots the pooled three-year treatment effects on firm revenue, value added, employment, and total compensation. The shock increases firm revenue by 50%, value added by 68%, employment by 29%, and total compensation by 57%.

Value-Added and Export-Demand Shock Overlapping Sample Results We next show that the value-added and export-demand shocks yield similar results on an overlapping set of firms.

Specifically, we estimate the effect of the value-added shocks on the same set of firms and individuals that we used to analyze the export-demand shocks. We find that the relative effects of the different shocks on worker versus owner earnings are very similar.³⁵ The relative effect size determines the overall shock incidence. To provide a summary measure of the similarity between the estimates from the two shocks, Figure A3 plots the implied incidence across the income distribution from both shock estimates (Panel A shows that export-demand incidence, which is identical to Figure 4, and Panel B shows that value-added shock incidence). The incidence estimates are remarkably similar; for example, the overall owner share is 83% for the export-demand shock and 85% for the value-added shock. This similarity suggests that the properties of the value-added shocks are similar to those of the export-demand shocks. Since we previously presented several pieces of evidence that the export-demand shocks satisfy our shock assumptions, this similarity partially mitigates the concern that the value-added shocks severely violate our shock assumptions.

5.3 The Incidence of Value-Added Shocks

After assessing the properties of the value-added shocks, we estimate the incidence of these shocks across the income distribution and other demographic groups. To estimate the incidence across income groups, we follow the same procedure we used for the export-demand shocks. Table 3 presents analogous estimates of how the value-added shock affects workers' and owners' income and job transitions. The results confirm that owners' income is much more sensitive to value-added changes than workers' income. For the stayers sample, a one-unit value-added shock increases owner income by 100% but only increases worker income by 17%. Interestingly, however, we do not find that high-earning workers are more sensitive to value-added changes than low-earning workers.³⁶ One additional difference relative to the export-demand shock results is that workers and owners are exposed to the value-added shocks on the extensive margin. Specifically, we find that a positive value-added shock significantly reduces the probability of a job separation and of nonemployment by similar magnitudes for workers and owners. One explanation for why owners experience fewer "job separations" following positive shocks is that their firms may be less likely to go out of business.

Next, we use the worker and owner value-added shock estimates to calculate the incidence of the shock across the income distribution. Figure 6 Panel A shows the incidence implied by the stayers sample estimates, and Panel B shows the incidence implied by the all-income sample estimates. These estimates show that the incidence of value-added shocks is also skewed towards the top of the distribution. For the stayers and the all-income samples, the top 10% of the income distribution receives 59% and 67% of the incidence, respectively.

However, individuals in the top 1% of the income distribution receive a smaller share of the

³⁵For example, for the stayers sample, we find that owner income increased by 102% and worker income increased by 12% in response to a one-unit value-added shock (the ratio is very similar to the estimates from the export-demand shock). We also confirm that owners' W-2 income significantly increases and that higher-earning workers experience larger effects.

³⁶For the stayers sample, the coefficient estimates for the three income groups are 20.1%, 15.0%, and 20.1%, respectively. The reason that we find income heterogeneity for the export-demand shocks but not the value-added shocks is due to the different samples used to estimate the shocks, rather than the different shocks. Specifically, we still find significant differences between the income groups when we estimated the effects of the value-added shocks on the export-demand shock sample in Section 5.2.

incidence of the value-added shocks than the export-demand shocks. Comparing Figure 4 with Figure 6, we find that individuals in the top 1% of the income distribution accrue twice as large a share of the changes from export-demand shocks as from value-added shocks. The primary reason for this difference is that owners in the broader value-added sample are less likely to be in the top 1% of the income distribution. Figure 1 shows that although owners collect approximately 40% of baseline income in both samples, nearly 69% of owner income at the export-demand shock firms goes to owners in the top 1%, whereas only 38% of owner income at the value-added shock firms goes to owners in the top 1%. However, in both cases, the vast majority of income changes—83% for the export-demand shocks and 77% for the value-added shocks—accrue to owners.

Finally, the broader representativeness of the value-added shock is well suited for characterizing the distributional incidence across other demographic groups. Specifically, we calculate the share of income changes by sex and race and ethnicity. To construct these incidence estimates, we first estimate worker by demographic-group specific treatment effects using the same specification we use to estimate the income group heterogeneity (see Footnote 30).³⁷ We then use these coefficients and the baseline income shares by demographic group in Table A3 to calculate the incidence.

Panels C and D of Figure 6 show the incidence of the value-added shocks for men and women and for each racial or ethnic group. We also find that the distribution of income gains is more unequal than the distribution of baseline income, although the differences by sex or race and ethnicity are generally smaller than the differences by income group. We find that men receive 75% of the total income changes. For comparison, Table A3 reports that men receive 66% of the baseline income from these firms. Similarly, we find that white workers receive 77% of the income changes, while they receive 71% of the baseline income. In both cases, differences in business ownership rates across the groups lead to greater differences in the share of income changes received by those same groups.

5.4 The Asymmetric Effects of Value-Added Shocks

The previous incidence estimates combine who bears the costs of *negative* firm demand changes with who receives the benefits of *positive* changes. However, the individuals who bear the costs of negative shocks may not be the same as those who receive the benefits of positive shocks. Additionally, it is plausible that incidence could exhibit asymmetry in either direction. For example, if wages are downward rigid, workers could benefit more from positive shocks than they suffer from negative shocks. Alternatively, if firms lay off workers in response to negative shocks, the large cost of job loss could lead to the opposite asymmetry. Next, we empirically test for incidence asymmetry.

We test for asymmetric effects by estimating separate *nonlinear* shock treatment effects for workers and owners. Intuitively, the linear treatment effects we previously estimated impose that a constant share of incidence always goes to workers versus owners, independent of the shock’s magnitude or sign. In contrast, nonlinear treatment effects allow the sensitivity of workers’ and

³⁷Across the sex and race/ethnicity groups, the heterogeneous treatment effects are relatively similar, and this degree of heterogeneity plays a small role in shaping the overall incidence. Thus, we do not report these estimates.

owners' income to the shock to depend on the shock's magnitude.

Specifically, we estimate a linear spline extension of Equation 11:

$$\frac{w_{i,t_0+n} - w_{i,t_0}}{\text{abs}(w_{i,t_0})} = \gamma + \alpha_n z_{J(i),t_0} + \sum_g \alpha_{n,g} (z_{J(i),t_0} - k_g)_+ + X'_{i,t_0} \beta_n + \varepsilon_{it}, \quad (14)$$

where k_g are spline knots (points at which we allow the slope of the shock's effect to change), and $(z_{J(i),t_0} - k_g)_+ = \max(z_{J(i),t_0} - k_g, 0)$. We select four spline knots at -0.2 , -0.05 , 0.05 , and 0.2 , dividing the shock distribution into five distinct regions, each with potentially different slopes. For example, the estimated slope for shocks in the interval $[-0.2, -0.05)$ is $\hat{\alpha}_n + \hat{\alpha}_{n,-0.2}$. We separately estimate pooled three-year effects for workers and owners using Equation 14.

Figure 7 plots the nonlinear treatment effect estimates for workers and owners. Panels A and B present estimates from the all-income sample, and Panels C and D present the stayers sample results. Specifically, for different potential realizations of the value-added shock, z , we plot the predicted proportional change implied by the estimates from Equation 14.³⁸ For example, Panel A shows that workers at firms with a negative -0.2 value-added shock would, on average, experience six percentage points lower total earnings growth than workers at firms with a positive 0.2 value-added shock. The estimates in Panel B show that the owners of these two firms would experience a 56 percentage point average difference in total earnings growth. These estimates confirm our prior finding that owners' income is substantially more sensitive to firm shocks than workers' income.

The estimates in Figure 7 also illustrate that workers and owners exhibit differential nonlinear responses to the value-added shocks. First, for both workers and owners, shocks that are larger in magnitude (both positive and negative) yield diminishing marginal changes in individuals' income growth (i.e., the slopes flatten at the extremes of the shock distribution). One reason for this nonlinearity is that more extreme shock value realizations also have small effects on real firm-level outcomes like sales.³⁹ However, workers and owners also exhibit different degrees of nonlinear treatment effects, which leads to an asymmetric shock incidence. Consider, for example, how workers and owners in the all-income sample are differentially affected by progressively more negative firm-level shocks. For workers (Panel A), as the value-added shock decreases from -0.05 to -0.4 , the slope flattens slightly, but they continue to experience larger total earnings losses from more negative shocks. In contrast, for owners (Panel B), as the value-added shock decreases from -0.05 to -0.4 , the average decline in owners' total income is relatively constant. Moreover, workers and owners exhibit the opposite asymmetry for progressively larger positive shocks—owners experience larger

³⁸The predicted proportional change is $E[Y_{it}|z_{J(i),t_0}]$ which is equal to $\hat{\alpha}_n \times z_{J(i),t_0} + \sum_g \hat{\alpha}_{n,g} \times (z_{J(i),t_0} - k_g)_+$. We normalize the predicted effects to be zero when $z_{J(i),t_0} = 0$. Due to this normalization, the predicted change for any given z does not have an absolute interpretation; however, relative comparisons between the predicted effects for different values of z remain meaningful. Specifically, the *relative* effects between different z s are interpretable as the difference in income growth for individuals at firms with the different shock realizations. We calculate the standard errors for each predicted proportional change using the Delta method. For now, we assume all covariances between the estimated spline coefficients are zero when calculating the standard errors.

³⁹Figure A4 presents plots of analogous nonlinear treatment effect estimates for firm-level revenue, value added, and employment growth rates. It shows that more extreme shock realizations have smaller marginal effects on all three outcomes. One explanation for this nonlinearity is that larger-in-magnitude shock realizations contain more measurement error, making their impact on firm-level outcomes more muted.

income gains for more positive shock realizations. In contrast, the effect on workers' income flattens out. Finally, to address the imprecision of these estimates, we will expand our analysis to the full 100% sample in the future and formally test for differential nonlinear effects between workers and owners.

Together, these nonlinear treatment effects for the all-income sample imply that workers bear a larger share of the total losses from negative shocks than their share of the gains from positive shocks. Figure 8 plots how the share of the incidence borne by workers varies for different-sized value-added shocks. Specifically, we use the implied nonlinear treatment effect estimates from Figure 7 and the baseline income shares to calculate Equation 6 for different realizations of the value-added shock z . The red line plots the incidence for the all-income sample and shows that while workers bear 26% of the losses from a negative -0.2 value-added shock, they only receive 9.5% of the gains for a similarly sized positive 0.2 value-added shock.⁴⁰

For the stayers sample, we find some asymmetry in the share of incidence borne by workers, but the degree of the asymmetry is less than that of the all-income sample. Specifically, the blue line in Figure 8 shows asymmetry estimates for this sample. While workers bear a progressively smaller share of the gains from positive shocks, the asymmetry for negative shocks is more mixed. The major difference between the two samples is that the stayers sample does not include extensive-margin income changes (e.g., the cost of job loss or gains from switching jobs). We explore whether this difference drives the asymmetry differences by analyzing nonlinear extensive margin effects on workers and owners in Figure 9. Specifically, the figure plots estimates from Equation 14 for two outcomes. The outcome in Panel A is an indicator for leaving the treated firm (i.e., no longer receiving any income). The outcome in Panel B is an indicator for nonemployment (we impose a strict definition of nonemployment: receiving no annual income from wages, salaries, business income, or self-employment).

We find that workers are more exposed on the extensive margin to negative firm-level shocks than owners (e.g., they are more likely to experience job loss). Panel A shows that in response to a negative -0.20 value-added shock, workers are almost three times as likely to leave the firm as owners (6.3 percentage points versus 2.2 percentage points). Consequently, this differential extensive-margin sensitivity, paired with the large earnings declines from job loss for workers (Jacobson et al., 1993), is one explanation for why workers in the all-income sample bear a much larger share of the losses from negative shocks than gains from positive shocks. In Panel B of Figure 9, we test whether workers are also more exposed to nonemployment after negative shocks and do not find significant differences between workers and owners. Negative firm shocks increase the probability of nonemployment for workers and owners, but the difference is much smaller (1.2 percentage points for workers versus 0.6 percentage points for owners). Overall, this evidence shows that only considering workers who remain employed misses that job loss is a major component of how workers are exposed to negative firm shocks.

⁴⁰To put these numbers in perspective, the firm-level first stage estimates in Figure A4 show that a negative -0.20 value-added shock led to a relative decrease in firm-level value added of 14%, while a $.20$ value-added shock led to a relative increase in firm-level value added of 26%.

6 Discussion: Wage-Setting and Policy Implications

In this section, we describe how our results contribute to (1) our understanding of how firm-specific shocks affect workers and owners, (2) debates about which wage-setting models best match the empirical evidence, and (3) discussions about the distributional consequences of firm-specific policies.

How are workers and owners affected by firm shocks? Our key contribution is to provide a unified framework for calculating the incidence of firm-specific shocks using linked worker-owner data. This framework enables us to make conceptual and methodological advances to the existing *rent-sharing literature* (see [Card et al. \(2018\)](#) for a summary).

Conceptually, we can answer a broader set of distributional questions than the existing literature. Rent-sharing studies measure what share of changes in *firm surplus*—that is, the remaining income after subtracting firm costs available to be distributed between workers and owners—goes to workers who remain at the treated firm. These studies assume that the residual surplus that does not go to workers goes to a representative owner, whose characteristics are unknown. In contrast, by explicitly mapping firm profits to individual owners, we can go beyond estimating the aggregate share of surplus going to workers and instead distribute the surplus across different parts of the income distribution, as well as demographic or geographic groups. For example, our framework allows us to construct distributional tables for firm shocks or firm policies analogous to those commonly used in tax policy analysis. Because some workers are at the top of the income distribution and some owners may be in the middle (e.g., small business owners or public company shareholders), these distributional tables are more informative than knowing only what share of surplus goes to workers.

Additionally, our framework allows us to incorporate multiple income-change margins into the incidence calculation. For example, we can calculate the incidence (1) only including income changes for workers and owners who remain at the firm and (2) incorporating the total income changes for workers or owners who leave the firm (i.e., including the costs of job loss and business closure). This broader approach is valuable for analyzing policy interventions or firm shocks that simultaneously affect firm survival, employment, wage negotiations, and profit-sharing.

Methodologically, we address multiple sources of measurement error that could bias existing rent-sharing estimates. First, we highlight a new issue in analyzing how firm shocks affect workers: misclassifying wage and salary payments to owners as worker income can bias these estimates. Specifically, we find that such a misclassification would overstate the worker share of the incidence by 40%. Although we only illustrate this bias for S corporations, it is also likely present for other U.S. firm structures.⁴¹ Moreover, this misclassification bias may explain why our estimates of heterogeneous pass-through by worker income differ from those in prior studies. While many papers find that the pass-through of firm shocks to wages is larger for high-earning workers ([Juhn et al.,](#)

⁴¹First, the bias is likely present for C corporations (43% of U.S. employment), especially private C corporations. Specifically, C corporations face a stronger tax incentive than S corporations to pay owners with wages and salaries rather than business income ([Smith et al., 2022](#)). Owners of Sole Proprietorships (4% of employment) and Partnerships (12% of private employment) are not legally supposed to pay owners via W-2s. However, [Hyatt et al. \(2020\)](#) surprisingly find that 24% of partnerships with employees pay at least one partner in wages, which suggests that this measurement error issue may affect partnerships as well.

2018; Kline et al., 2019; Cho and Krueger, 2022; Chan et al., 2023; Wallskog et al., 2024), our results are more mixed. We find evidence of such heterogeneity for export-demand shocks but not for value-added shocks. This discrepancy could be due to the owner W-2 classification bias since high-earning W-2 recipients are more likely than low-earning recipients to be owners.⁴²

Second, our framework can circumvent issues due to measurement error in firm surplus. Existing rent-sharing estimates are sensitive to measurement error in firm surplus, a difficult-to-measure concept that depends on capital prices and the value of workers’ outside options (see Card et al. (2018); Kline et al. (2019) for evidence that the bias can be substantial). Our framework does not require measuring the surplus from firm-level data. Instead, we directly define the total change in surplus as the sum of the changes in worker and owner income.

Finally, we contribute to understanding how demand changes impact a firm’s owners, a topic less studied than their effects on workers.⁴³ Specifically, we demonstrate that firm owners are highly vulnerable to idiosyncratic fluctuations in their firms’ performance. Moreover, owners remain highly exposed, even when we analyze their total income, which includes both their salary and business income from other firms. This finding suggests that owners cannot fully self-insure against firm-specific risk by switching jobs or diversifying their ownership in private firms.⁴⁴ This evidence shows that idiosyncratic firm-specific shocks are a contributor to the earnings volatility of private firm owners (Moskowitz and Vissing-Jørgensen, 2002; DeBacker et al., 2023). Furthermore, our finding complements Smith et al. (2019)’s evidence that these owners’ incomes largely reflect human capital (labor income). Although their analysis highlights similarities between firm owners and high-income employees, we find that owners’ incomes are substantially more sensitive to firm-level shocks. This result emphasizes that ownership significantly amplifies exposure to firm-specific risk compared to high-earning employees without ownership stakes.

Asymmetric incidence and wage-setting models Our finding that owners bear more of the incidence of positive firm shocks than negative shocks contradicts the predictions of several prominent theories of wage-setting. These theories imply that owners receive the same share of the benefits from positive shocks as the losses from negative shocks. For example, wage bargaining models with the canonical *Nash bargaining solution* imply that owners always receive a fixed share of any changes in the firm surplus. Consequently, the share of total income gains or losses that owners receive from firm shocks is independent of the sign and size of the shock (see Appendix D).⁴⁵

Similarly, theories of firm insurance or risk-sharing between workers and firms imply a symmetric

⁴²Another difference between our estimates and many of the worker-heterogeneity estimates cited above is that we rank workers based on the national income distribution, while many of those papers use the within-firm earnings distribution. However, we would still expect the owner W-2 bias to be present for within-firm earnings rankings.

⁴³While there is a large literature on how firm shocks affect public firm CEOs, there is less research on how shocks affect the owners of private firms. Since these owners are often owner-managers, these are distinct questions. One exception is Grindaker et al. (2024) who study the effects of bankruptcy on the CEOs of smaller Norwegian firms, who are likely also firm owners.

⁴⁴Note, however, that our comprehensive income measure in Footnote 17 does not include all forms of income that owners could use to self-insure. Specifically, we cannot include capital gains, dividends, or other forms of investment income. In particular, we miss business income from owning a private C corporation.

⁴⁵Appendix Section D shows that a “union” wage-bargaining model with multi-worker firms also yields a symmetric incidence, although the incidence notion is slightly different than what we currently calculate in this paper.

worker incidence (Baily, 1974; Azariadis, 1975). In these models, firms offer workers contracts that “insure” workers by promising not to lower their wages in response to negative shocks, in return for allowing the firm owners to receive the gains from positive shocks. Consequently, the owners’ share of the incidence is generally symmetric and determined by the relative risk aversion of owners and workers.⁴⁶ Moreover, wage-setting models that predict asymmetric responses to positive and negative firm shocks often predict the opposite asymmetry to what we find. In particular, several extensions to the baseline firm-insurance model predict that workers’ wages respond more to positive shocks than negative shocks (see Appendix Section D for details). Similarly, *downward wage rigidity* implies that workers’ wages are more responsive to positive than negative shocks.⁴⁷

In contrast to these theoretical predictions, we document the opposite type of asymmetry: workers bear a larger share of income losses from negative shocks than gains from positive shocks. This asymmetric incidence is particularly pronounced when we incorporate income changes from extensive-margin job transitions. Many of the simple wage-setting theories mentioned above abstract from costly job separations, so our findings highlight the importance of incorporating this margin for understanding who is affected by firm growth and contraction. Notably, even when we focus only on intensive-margin earnings—closer to the wage concepts traditionally modeled—the asymmetry persists. For example, Figure 8 illustrates that as positive value-added shocks increase in magnitude from 0.05 to 0.2, the workers’ share of total income gains declines from 16% to 9%. In practical terms, firms appear to initially raise workers’ wages in response to moderately sized positive shocks (e.g., give workers a 5% wage increase in a good year), but for larger positive shocks, they stop raising workers’ wages, and owners receive all additional surplus. This behavior contradicts bargaining and risk-sharing models, which predict continuous proportional wage growth as the magnitude of shocks increases. Our results thus suggest a need for theories that explicitly incorporate asymmetric responses across both intensive and extensive margins.

There are several caveats to the conclusion that our findings contradict these theoretical predictions. First, some of our treatment effect estimates for extreme shock realizations in Figure 7 are imprecise, implying that there is some uncertainty in our incidence asymmetry estimates. Second, we do not currently include new hires in our notion of incidence. This exclusion could drive our asymmetry result because new hires are a key way that firms respond to positive shocks (Davis et al., 2013). In future versions, we will address these issues by expanding the sample from our current 20% random sample and incorporating new hires into our framework. Finally, our intensive-margin

⁴⁶In Appendix Section D, we show symmetric incidence in a simple risk-sharing model when workers and owners have CARA utility. Beyond that special case, the models may predict nonlinear responses, but they still require a degree of symmetry (if workers are exposed to downside risk but experience no upside gains, the firm is not providing insurance to the workers).

⁴⁷Due to differences in the earnings concepts used, our estimates could be consistent with previous empirical evidence documenting downward nominal hourly wage rigidity. First, Grigsby et al. (2021) find that only 2.5% of workers annually receive a nominal *base wage* cut, but 17% of workers receive a nominal *total compensation* cut, where total compensation includes commissions, bonuses, and overtime. Our earnings measure includes these additional forms of compensation, so nominal compensation reductions are quite plausible. Second, our wage decline measures are relative to workers at control firms. For example, if workers at the treated firm received a 5% annual wage increase, but workers at the control firm received a 10% increase, we would define this as a *negative* 5% earnings treatment effect despite the nominal wage increase for the treated workers. Third, our earnings responses could include hours adjustments.

earnings outcome captures annual compensation rather than hourly wages, potentially conflating wage changes with adjustments in hours or days worked. However, we find similar short- and long-run treatment effects, while we would expect hours adjustments to be more temporary.

By jointly analyzing worker and owner outcomes, we introduce a new test of whether firms provide workers insurance against firm fluctuations (see [Beaudry and DiNardo \(1991\)](#); [Guiso et al. \(2005\)](#) and [Sockin and Sockin \(2025\)](#) for other tests). Specifically, these theories not only predict a small overall pass-through of firm shocks to workers but also that the pass-through should be symmetric. This test builds on [Guiso et al. \(2005\)](#), who use the small overall pass-through of firm shocks to workers as evidence of firm insurance. For example, a relatively small overall pass-through may not be consistent with firm insurance if the wage changes are concentrated in response to negative shocks or if workers still face considerable extensive-margin earnings risk.

Moreover, by simultaneously estimating the effects on workers and owners, we build on past papers that have tested for asymmetric pass-through using only worker data ([Juhn et al., 2018](#); [Cho and Krueger, 2022](#); [Grübener and Rozsypal, 2023](#); [Chan et al., 2023](#); [Acemoglu et al., 2023](#); [Friedrich et al., 2024](#); [Mertens et al., 2024](#); [Merkle, 2024](#)). A challenge in asymmetry analyses using only worker-level data is distinguishing whether asymmetric wage responses to positive versus negative shocks reflect asymmetry in the shocks’ impact on firm-level surplus or asymmetric wage pass-through.⁴⁸ This challenge is related to the difficulties in accurately measuring firm surplus in the rent-sharing literature. In contrast, since we estimate the change in owners’ income, we can benchmark the magnitude of the change in workers’ income against the change in owners’ income to directly calculate which group bears a greater share of the total income changes.

The Distributional Incidence of Firm-Specific Government Policies Our results imply that several government policies targeted at firms may have an unequal distributional incidence. Specifically, firm-level policies like government procurement or firm-specific subsidies could satisfy the same shock assumptions as our shocks (e.g., firm specific and Hicks neutrality). Consequently, the distributional incidence of these policies may be quite unequal. Additionally, our shocks have a very unequal incidence, even at relatively small firms (on average, they have 35–55 employees). This result implies that targeting policies at “small businesses” may not result in a more equitable incidence across the income distribution. Moreover, our shock asymmetry result suggests that policies that provide firms with downside protection (e.g., preventing firm closures) may have a more equitable incidence than policies that encourage firm growth. Finally, another benefit of our incidence framework is that it can be used to directly analyze the incidence of firm-level policies. We view applying the framework to construct distributional tables for specific firm-level policies, similar to those currently constructed for tax policies, as a policy-relevant area for future research.

⁴⁸For example, Figure [A4](#) provides evidence of asymmetric effects of our value-added shocks on firm-level outcomes.

7 Conclusion

This paper analyzes which individuals bear the incidence of firm demand changes. Relative to previous research, we use firm-worker-owner data to jointly analyze how firm demand changes affect workers and owners. We find that the incidence of firm demand changes is highly skewed toward high-income individuals. Specifically, 30–60% of income changes from the firm shocks accrue to individuals in the top 1% of the national income distribution. This unequal distribution arises because firm owners are more exposed to firm shocks than workers and are disproportionately in the top of the distribution. However, the incidence is asymmetric for positive versus negative demand changes. Although workers receive a small share of the benefits of positive shocks, they bear a larger share of losses from negative shocks, especially after we account for income changes due to job loss.

Our analysis makes several contributions to understanding the distributional consequences of firm demand changes. By incorporating owner data and mapping all firm profits to individuals, we can fully characterize the incidence of these shocks across the income distribution and demographic groups. This approach enables us to construct distributional tables for firm shocks analogous to those used in tax policy analysis, while flexibly incorporating multiple margins of income adjustment. Methodologically, we demonstrate the importance of correctly classifying owner wage payments, showing that misclassification would substantially overstate workers’ share of the incidence. Our findings also have implications for wage-setting theory and public policy. First, the asymmetric incidence we document contradicts an implicit prediction of several standard wage-setting models that the relative worker and owner responses to shocks should be symmetric. Second, our results suggest that some firm-targeted policies may have very unequal distributional consequences.

Although our analysis covers more than a quarter of U.S. employment, our estimates are specific to the effects of recent firm-specific shocks at S corporations in the U.S. Extending our analysis along all three dimensions is a promising area for future research. First, a key open question is whether the incidence of demand changes is more equitable at other types of U.S. firms or during other periods. In particular, contrary to some conventional wisdom, existing evidence suggests that the incidence of firm shocks may be more equitable among large publicly traded firms than among the smaller privately owned firms we study.⁴⁹ Second, our estimates for firm-specific, skill-neutral shocks may not be informative about the effects of other types of firm shocks. In particular, applying our framework to study skill-biased and market-wide shocks could provide micro-estimates that inform how these changes have shaped the evolution of income inequality over recent decades. Finally, extending this analysis internationally could reveal whether the unequal incidence we document reflects specific features of the U.S. labor market. While we know that the baseline distribution of income in European countries is more equal than in the U.S. (Blanchet et al., 2022), our framework shows that shock incidence also depends on differential pass-through rates to workers versus owners, not just baseline income shares. These open questions illustrate the benefits of using worker-owner data to uncover the full distributional consequences of firm shocks.

⁴⁹Cooper et al. (2016) show that C corporation dividends are more equitably distributed than pass-through firms’ profits.

References


- Abowd, J. A. and T. Lemieux (1993) “The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada,” *The Quarterly Journal of Economics*, Vol. 108, pp. 983–1014.
- Acemoglu, Daron, Alex Xi He, and Daniel le Maire (2023) “Eclipse of Rent-Sharing: The Effects of Managers’ Business Education on Wages and the Labor Share in the US and Denmark.”
- Auten, Gerald and David Splinter (2024) “Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends,” *Journal of Political Economy*, Vol. 132, pp. 2179–2227.
- Azariadis, Costas (1975) “Implicit Contracts and Underemployment Equilibria,” *Journal of Political Economy*, Vol. 83, pp. 1183–1202, Publisher: University of Chicago Press.
- Baily, Martin Neil (1974) “Wages and Employment under Uncertain Demand,” *The Review of Economic Studies*, Vol. 41, pp. 37–50, Publisher: [Oxford University Press, Review of Economic Studies, Ltd.].
- Beaudry, Paul and John DiNardo (1991) “The Effect of Implicit Contracts on the Movement of Wages Over the Business Cycle: Evidence from Micro Data,” *Journal of Political Economy*, pp. 665–688.
- Blanchet, Thomas, Lucas Chancel, and Amory Gethin (2022) “Why Is Europe More Equal than the United States?” *American Economic Journal: Applied Economics*, Vol. 14, pp. 480–518.
- Borusyak, Kirill and Peter Hull (2023) “Nonrandom Exposure to Exogenous Shocks,” *Econometrica*, Vol. 91, pp. 2155–2185.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2022) “Quasi-Experimental Shift-Share Research Designs,” *The Review of Economic Studies*, Vol. 89, pp. 181–213.
- Card, David, Ana Rute Cardoso, Jorg Heining, and Patrick Kline (2018) “Firms and labor market inequality: Evidence and some theory,” *Journal of Labor Economics*, Vol. 36, pp. S13–S70.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019) “The Effect of Minimum Wages on Low-Wage Jobs,” *The Quarterly Journal of Economics*, Vol. 134, pp. 1405–1454.
- Chan, Mons, Sergio Salgado, and Ming Xu (2023) “Heterogeneous Passthrough from TFP to Wages,” *Working Paper*.
- Cho, David and Alan B. Krueger (2022) “Rent Sharing within Firms,” *Journal of Labor Economics*, Vol. 40, pp. S17–S38, Publisher: The University of Chicago Press.
- Chow, Melissa, Teresa Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White (2021) “Redesigning the Longitudinal Business Database,” *NBER Working Paper*, p. w28839, Place: Cambridge, MA.
- Cooper, Michael, John McClelland, James Pearce, Richard Prisinzano, Joseph Sullivan, Danny Yagan, Owen Zidar, and Eric Zwick (2016) *Business in the United States: Who Owns It, and How Much Tax Do They Pay?*, Vol. 30 of Tax Policy and the Economy, Cambridge: MIT Press.
- Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger (2013) “The Establishment-Level Behavior of Vacancies and Hiring*,” *The Quarterly Journal of Economics*, Vol. 128, pp. 581–622.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh (1996) *Job Creation and Destruction*, Cambridge, MA, USA: MIT Press.
- DeBacker, Jason, Vasia Panousi, and Shanthi Ramnath (2023) “A Risky Venture: Income Dynamics among Pass-Through Business Owners,” *American Economic Journal: Macroeconomics*, Vol. 15, pp. 444–474.
- Fairlie, Robert W. and Alicia M. Robb (2008) *Race and Entrepreneurial Success: Black-, Asian-, and White-Owned Businesses in the United States*: The MIT Press.
- Foster, Lucia S., Cheryl A. Grim, John Haltiwanger, and Zoltan Wolf (2017) “Macro and Micro Dynamics of Productivity: From Devilish Details to Insights,” *NBER Working Paper*.

- Friedrich, Ben, Lisa Luan, Costas Meghir, and Luigi Pistaferri (2024) “Earnings Dynamics and Firm-Level Shocks,” *Working Paper*.
- Garin, Andrew and Filipe Silvério (2024) “How Responsive Are Wages to Firm-Specific Changes in Labour Demand? Evidence from Idiosyncratic Export Demand Shocks,” *The Review of Economic Studies*, Vol. 91, pp. 1671–1710.
- Gaulier, Guillaume and Soledad Zignago (2010) “BACI: International Trade Database at the Product-Level. The 1994-2007 Version,” *Working Papers*, Number: 2010-23 Publisher: CEPII research center.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift (2020) “Bartik Instruments: What, When, Why, and How,” *American Economic Review*, Vol. 110, pp. 2586–2624.
- Goodman, Lucas, Elena Patel, and Molly Saunders-Scott (2023) “Implications of Tax Loss Asymmetry for Owners of S Corporations,” *American Economic Journal: Economic Policy*, Vol. 15, pp. 342–369.
- Grigsby, John, Erik Hurst, and Ahu Yildirmaz (2021) “Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data,” *American Economic Review*, Vol. 111, pp. 428–471.
- Grindaker, Morten, Andreas R Kostøl, and Kasper Roszbach (2024) “Make it or Break it: Corporate Bankruptcy and Management Careers.”
- Grübener, Philipp and Filip Rozsypal (2023) “Firm Dynamics and Earnings Risk,” *Working Paper*.
- Guiso, Luigi and Luigi Pistaferri (2020) “The insurance role of the firm,” *The Geneva Risk and Insurance Review*, Vol. 45, pp. 1–23.
- Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi (2005) “Insurance within the firm,” *Journal of Political Economy*, Vol. 113, pp. 1054–1087.
- Habib, Bilal and Rebecca Heller (2022) “Current Work on the Distributional Analysis of Household Income Resulting From Policy Changes,” *Congressional Budget Office Working Paper 2022-09*.
- Hall, Robert E. and Susan E. Woodward (2010) “The Burden of the Nondiversifiable Risk of Entrepreneurship,” *The American Economic Review*, Vol. 100, pp. 1163–1194.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda (2013) “Who Creates Jobs? Small versus Large versus Young,” *Review of Economics and Statistics*, Vol. 95, pp. 347–361.
- Hummels, David, Rasmus Jorgensen, Jakob Munch, and Chong Xiang (2014) “The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data,” *American Economic Review*.
- Hyatt, Henry, Seth Murray, and L. Kristin Sandusky (2020) “Business Income Dynamics and Labor Market Fluidity,” *IZA Journal of Labor Economics*, Vol. 10.
- IRS (2008) “Wage Compensation for S Corporation Officers,” *IRS Fact Sheet FS-2008-25*.
- (2025) “S corporation employees, shareholders and corporate officers.”
- Jacobson, Louis, Robert LaLonde, and Daniel Sullivan (1993) “Earnings Losses of Displaced Workers,” *American Economic Review*, Vol. 83, pp. 685–709.
- Juhn, Chinhui, Kristin McCue, Holly Monti, and Brooks Pierce (2018) “Firm Performance and the Volatility of Worker Earnings,” *Journal of Labor Economics*, Vol. 36, pp. S99–S131, Publisher: The University of Chicago Press.
- Kamal, Fariha and Wei Ouyang (2020) “Identifying U.S. Merchandise Traders: Integrating Customs Transactions with Business Administrative Data,” *CES Working Paper 20-28*.
- Kehrig, Matthias and Nicolas Vincent (2021) “The Micro-Level Anatomy of the Labor Share Decline,” *The Quarterly Journal of Economics*, Vol. 136, pp. 1031–1087.
- Kennedy, Patrick J, Christine Dobridge, Paul Landefeld, and Paul Mortenson (2022) “The Efficiency-Equity Tradeoff of the Corporate Income Tax: Evidence from the Tax Cuts and Jobs Act,” *Working Paper*.

- Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar (2019) “Who profits from patents? rent-sharing at innovative firms,” *Quarterly Journal of Economics*, Vol. 134, pp. 1343–1404.
- Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler (2022) “Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market,” *American Economic Review*, Vol. 112, pp. 169–212.
- Lindner, Attila, Balazs Murakozy, and Balazs Reizer (2022) “Firm-level Technological Change and Skill Demand,” *Working Paper*, p. 50.
- Love, Michael (2021) “Where in the World Does Partnership Income Go? Evidence of a Growing Use of Tax Havens,” *SSRN Electronic Journal*.
- Merkle, Matthew (2024) “Job Tenure and Rent Sharing,” *Working Paper*.
- Mertens, Matthias, Steffen Mueller, and Georg Neuschaeffer (2024) “Identifying rent-sharing using firms’ energy input mix,” *Working Paper*.
- Mertens, Matthias and Benjamin Schoefer (2025) “From Labor to Intermediates: Firm Growth, Input Substitution, and Monopsony,” *Working Paper*.
- Moskowitz, Tobias J. and Annette Vissing-Jørgensen (2002) “The Returns to Entrepreneurial Investment: A Private Equity Premium Puzzle?” *American Economic Review*, Vol. 92, pp. 745–778.
- Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2022) “Rethinking Capital and Wealth Taxation,” *World Inequality Lab Working Papers*.
- Risch, Max (2024) “Does Taxing Business Owners Affect Employees? Evidence From A Change in the Top Marginal Tax Rate,” *Quarterly Journal of Economics*, Vol. 139, pp. 637–692.
- Risch, Max and Nirupama Rao (2022) “How Do Small Firms Accommodate Minimum Wage Increases? Evidence from Matched Employer-Employee Tax Returns,” *Working Paper*.
- Schlingemann, Frederik P. and René M. Stulz (2022) “Have exchange-listed firms become less important for the economy?” *Journal of Financial Economics*, Vol. 143, pp. 927–958.
- Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick (2019) “Capitalists in the Twenty-First Century,” *The Quarterly Journal of Economics*, Vol. 134, pp. 1675–1745.
- (2022) “The Rise of Pass-Throughs and the Decline of the Labor Share,” *American Economic Review: Insights*, Vol. 4, pp. 323–340.
- Sockin, Jason and Michael Sockin (2025) “Variable Pay and Risk Sharing Between Firms and Workers.”
- Splinter, David, Eric Heiser, Michael Love, and Jacob Mortenson (2024) “The Paycheck Protection Program: Progressivity and Tax Effects.”
- Stinson, Martha and Sean Wang (2025) “The Composition of Firm Workforces from 2006–2022: Findings from the Business Dynamics Statistics of Human Capital Experimental Product,” *CES Working Paper 25-20*.
- U.S. Census Bureau (2022) “County Business Patterns National Statistics.”
- Wallskog, Melanie, Nicholas Bloom, Scott W Ohlmacher, and Cristina Tello-Trillo (2024) “Within-Firm Pay Inequality and Productivity,” *NBER Working Paper*.


8 Figures

Figure 1: Baseline Distribution of Firm Compensation Across the Income Distribution




Note: This figure reports the average share of each firm's total wage and business income allocated to different income groups, where total income for each individual combines wages from W-2 forms and business income from K-1 forms. For each firm j , we define the share of baseline income going to group g as $s_j^g = \sum_{i \in g} w_i / \sum_i w_i$, where w_i is the income that individual i receives from the firm. Each bar reports the average of s_j^g across all firms in that sample for a given group. The shares are calculated using all individuals in our stayers sample, but they are very similar if we use the all-income sample. To construct the income groups, we rank individuals each year based on their combined W-2 wages, K-1 business income, and self-employment income (Schedule C). Because this measure is at the individual level and excludes certain forms of income, we adjust the publicly available percentile thresholds from [Auten and Splinter \(2024\)](#) accordingly. The 2018 thresholds for the 50th, 90th, and 99th percentiles were approximately \$38,067, \$117,760, and \$339,046, respectively. Appendix Table A3 presents the raw numbers behind this figure.

Figure 2: Dynamic Effects of Export-Demand Shocks on Firm-Level Outcomes



Note: This figure plots the dynamic firm-level effects (α_n) of the export-demand shock constructed in Section 4. The sample includes S corporations that satisfy our baseline restrictions (non-finance, at least 10 employees, and at least \$5,000 in real exports per worker in the base year). We estimate the absolute value percent-change specification in Equation 11, measuring outcomes as long differences relative to the year before the shock ($t - 1$). We control for the expected product-level shock (see Section 4 for details) and include year and cohort fixed effects. Standard errors are clustered at the firm level, and 95% confidence intervals are indicated by the vertical lines around each point estimate.


Figure 3: Dynamic Effects of Export-Demand Shocks on Individual-Level Outcomes




Note: This figure plots the dynamic individual-level effects of the export-demand shock constructed in Section 4. Panel A shows the effects on workers; Panel B shows the effects on owners. Worker income includes all wage and salary payments reported on W-2 forms. Owner income includes both business income (from K-1 forms) and any wage/salary compensation the owner receives from the firm. We measure both outcomes using the *absolute percent change* functional form $\frac{Y_{j,t+n} - Y_{j,t}}{\text{abs}(Y_{j,t_0})}$ as in Equation 11. The sample in each panel is restricted to a balanced panel of individuals—those who remain at the firm from $t - 2$ through $t + 3$. We control for the expected product-level shock as well as year and cohort fixed effects. We weight the regressions by the inverse number of workers or owners at the firm at time $t - 1$ so that our firm- and individual-level results weight each firm the same. Standard errors are clustered at the firm level, and the vertical lines around each point denote 95% confidence intervals.

Figure 4: The Incidence of Export-Demand Shocks Across the Income Distribution

Panel A. Stayers Sample




Panel B. All-Income Sample




Note: This figure plots the incidence of the export-demand shocks across the income distribution. Incidence means the total income change from the shock that goes to individuals in each group. The estimates are based on Equation 6, which combines group-specific treatment effects and baseline income shares. We calculate standard errors using the delta method and assume that the baseline income shares are constants. We partition the sample into owners and workers in three income bins (bottom 50%, 50–90%, and the top 10%) for estimating the coefficients. These estimates are presented in Table 1. The baseline income shares are presented in Table A3, Panel A. See Footnote 17 for details on the income percentile construction. Panel A of this figure restricts the sample to incumbent individuals who are employed at the firm at least one year after the shock. Panel B includes all individuals who received income from the treated firm at time $t - 1$, regardless of whether they remain. We change our baseline income shares to match the different samples.

Figure 5: Effects of Value-Added Shocks on Firm-Level Outcomes

Panel A. Dynamic Value-Added Estimates



Panel B. Pooled Firm Estimates



Note: This figure illustrates how the value-added shock defined in Section 5 affects firm-level outcomes. The sample for both panels includes non-finance S corporations with at least 10 employees in the base year (see Section 3 for details).


Panel A plots the *dynamic* event-study coefficients for firms' value added, measured as the *absolute percent change* long difference relative to one year before the shock ($t = -1$). The estimates are from estimating Equation 11 with firm-level value added as the outcome. We stack cohorts from 2013 to 2018 and include cohort and year-fixed effects, as described in Section 5.

Panel B presents the *three-year pooled* long-difference estimates for four firm-level outcomes: revenue, value added, employment, and total compensation (the sum of payroll plus business income). Each bar indicates the average pooled effect from $t + 1$ to $t + 3$. Standard errors are clustered at the firm level, and vertical lines indicate 95% confidence intervals.


Note for Figure 6 (next page): This figure plots the incidence of the value-added shocks across income, sex, and race/ethnicity groups. Incidence means the total income change from the shock that goes to individuals in each group. The estimates are based on Equation 6, which combines group-specific treatment effects and baseline income shares. We calculate standard errors using the delta method and assume that the baseline income shares are constants. For the income-group estimates, we partition the sample into owners and workers in three income bins (bottom 50%, 50–90%, and the top 10%) for estimating the coefficients. For the demographic group estimates, we allow the worker coefficients to differ by demographic group. These estimates are presented in Table 3. The baseline income shares are presented in Table A3, Panel A. See Footnote 17 for details on the income percentile construction. The stayers sample restricts the sample to incumbent individuals who are employed at the firm for at least one year after the shock. The all-income sample includes all individuals who received income from the treated firm at time $t - 1$, regardless of whether they remain. We change our baseline income shares to match the different samples. The other/missing race and ethnicity category includes Native Hawaiian and Pacific Islander, American Indian and Alaska Native, and mixed-race individuals. The white, black, Asian, and other/missing categories only include non-Hispanic individuals.

Figure 6: The Incidence of Value-Added Shocks Across Income, Sex, and Racial Groups


Panel A. Income Groups – Stayers Sample




Panel B. Income Groups – All-Income Sample



Panel C. Sex – All-Income Sample




Panel D. Race/Ethnicity – All-Income Sample




Note: See previous page.

Figure 7: Nonlinear Worker and Owner Responses to Value-Added Shocks


Panel A. Worker Coefficients – All-Income Sample




Panel B. Owner Coefficients – All-Income Sample



Panel C. Worker Coefficients – Stayers Sample




Panel D. Owner Coefficients – Stayers Sample



Note: See next page.

Note for Figure 7 (previous page): This figure plots the estimated nonlinear effects of firm-level value-added shocks on worker and owner income. The panels plot the implied proportional change in income for different-sized value-added shocks that are implied by estimating Equation 14. For example, the panels show that a .1 increase in the firm value-added shock would increase owner income by 14.7% and increase worker income by 1.7%. The nonlinear treatment effects are reported separately for the all-income sample (Panels A and B) and the stayers sample (Panels C and D). We also report separate estimates for workers (Panels A and C) and owners (Panels B and D). The vertical dashed lines indicate the cutoffs in the value-added shock distribution where we allow the treatment effect slope to change. The dashed lines represent 95% confidence intervals.


Figure 8: Asymmetric Worker Incidence Share by Value-Added Shock Magnitude




Note: This figure plots how the workers' share of the incidence varies across the value-added shock distribution. Specifically, using the implied treatment effect estimates from Figure 7, we calculate the workers' and owners' share of the value-added shock incidence for each part of the shock distribution using Equation 6 and the baseline income shares in Table A3. The x-axis represents different values of the support of the value-added shock distribution. The figure plots workers' overall share of the incidence for the all-income sample in red and for the stayers sample in blue.

Figure 9: Nonlinear Worker and Owner Extensive-Margin Responses to Value-Added Shocks

Panel A. Effect on Remaining at the Treated Firm



Panel B. Effect on Nonemployment (Zero Annual Income)



Note: This figure plots the estimated nonlinear effects of firm-level value-added shocks on extensive-margin outcomes for workers and owners. The panels plot the implied probability of leaving the treated firm or becoming nonemployed for different-sized value-added shocks, as implied by estimating Equation 14. We define the indicator for leaving the treated firm as not receiving any K-1 or W-2 income from the firm in a given year. We define the indicator for nonemployment as not receiving any K-1, W-2, or Schedule C income from any firms in a given year. The vertical dashed lines indicate the cutoffs in the value-added shock distribution where we allow the treatment effect slope to change. The horizontal dashed lines represent 95% confidence intervals.

9 Tables

Table 1: The Effects of Export-Demand Shocks on Individual-Level Outcomes

Panel A: Worker and Owner Earnings Effect Estimates					
Sample:	Stayers Sample			All-Income Sample	
Income Measure:	Owner Total	Owner W-2	Worker Total	Owner Total	Worker Total
Pooled 3-Year % Income Change	0.091*** (0.024) [0.027]	0.019* (0.008) [0.009]	0.011*** (0.002) [0.003]	0.074*** (0.019) [0.022]	0.010*** (0.002) [0.003]

Panel B: Worker Earnings Effect Heterogeneity by Baseline Income		
Sample:	Stayers Sample	All-Income Sample
Income Measure:	Worker Total	Worker Total
Pooled 3-Year Estimate		
0–50th Pctl. Workers:	0.011** (0.004)	0.010 (0.004)
50th–90th Pctl. Workers:	0.009*** (0.002)	0.010*** (0.002)
90th+ Pctl. Workers:	0.027** (0.009)	0.025*** (0.007)

Panel C: Worker and Owner Extensive Margin Separation Effect Estimates				
Sample:	Owners		Workers	
Separation Measure:	Job Separation	Nonemployment	Job Separation	Nonemployment
Pooled 3-Year Separation Rate	−0.002 (0.003)	−0.001 (0.002)	−0.007** (0.003)	−0.002** (0.001)

Note: This table includes worker- and owner-level estimates of the pooled three-year effect of the export-demand shock. In panels A and C, all coefficients are from the pooled version of Equation 11. In panel B, the estimates are from the heterogeneity specification in Equation 12. The stayers sample columns restrict the sample to incumbent individuals who are employed at the firm for at least one year after the shock. The all-income sample columns include all individuals who received income from the treated firm at time $t - 1$, regardless of whether they remain. The “Owner Total” outcome includes owners’ business income and wage and salary income. The “Owner W-2” outcome only includes owners’ wage and salary income. In Panel C, the Job Separation outcome is an indicator for no longer being employed at the treated firm, where employment is defined as receiving any W-2 or K-1 income from the firm. The Nonemployment outcome is defined as not receiving any income (W-2, K-1, or Schedule C) from any firm in a given year. In Panel A, the standard errors in parentheses cluster at the firm level, while the standard errors in brackets are calculated following [Borusyak et al. \(2022\)](#). The earnings percentiles used in Panel B are described in Footnote 17. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2: Shock Skill Bias Test: The Effects on Workforce Composition

	Firm-Level Share of Workers			Owner Share of Workers + Owners
	0–50th Earnings Pct.	50–90th Earnings Pct.	90+ Earnings Pct.	
Pooled 3-Year Composition Change	0.0034** (0.0013)	−0.0028* (0.0013)	−0.0006 (0.0034)	−0.0020 (0.0010)

Note: This table reports the pooled three-year effect of the export-demand shock on firms' workforce composition across worker skill groups and owners. The first three columns present estimates from Equation 11, where the dependent variable is the change in the firm's employment share for each skill group relative to the baseline year. Specifically, we calculate the share of workers in skill group T as s_{jt}^T (owners excluded) and define the firm-level outcome variable as $s_{jt}^T - s_{jt_0}^T$. Workers are assigned to skill groups based on their average earnings percentiles over the entire sample period: 0–50th, 50–90th, and 90+ percentiles. The last column shows the change in the share of owners relative to total firm headcount (workers and owners combined). Standard errors clustered at the firm level are reported in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: The Effects of Value-Added Shocks on Individual-Level Outcomes

Panel A: Worker and Owner Earnings Effect Estimates					
Sample:	Stayers Sample			All-Income Sample	
Income Measure:	Owner Total	Owner W-2	Worker Total	Owner Total	Worker Total
Pooled 3-Year % Income Change	1.000*** (0.013)	0.260*** (0.004)	0.174*** (0.001)	0.604*** (0.009)	0.068*** (0.001)

Panel B: Worker Earnings Effect Heterogeneity by Baseline Income		
Sample:	Stayers Sample	All-Income Sample
Income Measure:	Worker Total	Worker Total
Pooled 3-Year Estimate		
0–50th Pctl. Workers:	0.201*** (0.002)	0.071*** (0.002)
50th–90th Pctl. Workers:	0.150*** (0.002)	0.065*** (0.001)
90th+ Pctl. Workers:	0.201*** (0.005)	0.077*** (0.003)

Panel C: Worker and Owner Extensive Margin Separation Effect Estimates				
Sample:	Owners		Workers	
Separation Measure:	Job Separation	Nonemployment	Job Separation	Nonemployment
Pooled 3-Year Separation Rate	−0.130*** (0.001)	−0.027*** (0.001)	−0.130*** (0.001)	−0.013*** (0.000)


Note: This table includes worker- and owner-level estimates of the pooled three-year effect of the value-added shock. In panels A and C, all coefficients are from the pooled version of Equation 11. In panel B, the estimates are from the heterogeneity specification in Equation 12. The stayers sample columns restrict the sample to incumbent individuals who are employed at the firm for at least one year after the shock. The all-income columns include all individuals who received income from the treated firm at time $t - 1$, regardless of whether they remain. The “Owner Total” outcome includes owners’ business income and wage and salary income. The “Owner W-2” outcome only includes owners’ wage and salary income. In Panel C, the Job Separation outcome is an indicator for no longer being employed at the treated firm, where employment is defined as receiving any W-2 or K-1 income from the firm. The Nonemployment outcome is defined as not receiving any income (W-2, K-1, or Schedule C) from any firm in a given year. The earnings percentiles used in Panel B are described in Footnote 17.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Appendix for
The Distributional Effects of Firm Demand Changes:
Evidence from U.S. Linked Worker-Owner Data


A Appendix Figures

Figure A1: Pooled Effects of Export-Demand Shocks on Firm-Level Outcomes




Note: This figure reports three-year pooled long-difference estimates of the export-demand shock on firm-level outcomes, following Equation 11. These pooled estimates are analogous to the dynamic worker-level estimates in Figure 2. The outcomes are measured relative to one year before the shock. The red bars (Baseline) include year, cohort, and expected product-level shock controls, with standard errors clustered by firm. The blue bars (Shift-Share SEs) use the same Baseline specification but compute standard errors following the shock-level procedure in [Borusyak et al. \(2022\)](#). The orange bars (Ind. + Loc. Controls) add four-digit NAICS industry fixed effects, commuting zone $\times$ two-digit NAICS sector fixed effects, and county fixed effects, each interacted with the treatment cohort. The sample includes non-finance S corporations meeting our baseline restrictions (at least 10 employees and at least \$5,000 in real exports per worker). The vertical lines around each bar represent 95% confidence intervals.

Figure A2: Export-Demand *Dosage Response* Effects on Firm-Level Outcomes




Note: This figure reports the three-year long-difference effects of the export-demand shock by bins of firms' baseline exports per worker, illustrating a “dosage response” relationship. We measure the outcomes (sales and employment) as the *absolute percent change* following Equation 11 and plot separate estimates for firms in each bin of baseline export intensity. The sample includes non-finance S corporations that meet all our baseline restrictions except the exports per worker restriction. We control for the expected product-level shock as well as year and cohort fixed effects. We interact treatment and all controls with the baseline exports/worker bins. The plotted estimates are from the interactions between treatment and the export-heterogeneity bins. The vertical lines represent 95% confidence intervals, with standard errors clustered by firm.

Figure A3: Comparing the Export-Demand and Value-Added Shocks on a Common Sample



Note: This figure compares the distributional incidence of export-demand shocks and value-added shocks for the same subset of S corporations that have both measures available. Panel A replicates the main stayers sample export shock incidence estimates (see Figure 4 and Section 4), showing how much of the total income change accrues to workers (red bars) versus owners (blue bars) in each part of the national income distribution. Panel B reports the analogous incidence estimates when using the value-added shocks constructed in Section 5, applied to the identical set of firms and individuals. The total owner share reported in each panel is the fraction of overall income changes captured by owners.

Figure A4: Asymmetric Effects of Value-Added Shocks on Firm-Level Outcomes



Note: This figure plots the estimated nonlinear effects of firm-level value-added shocks on firm-level outcomes. It plots the implied proportional change for the firm outcomes for different-sized value-added shocks from estimating Equation 14. The dashed horizontal lines represent 95% confidence intervals. The vertical dashed lines indicate the cutoffs in the value-added shock distribution where we allow the treatment effect slope to change.

B Appendix Tables

Table A1: Individual Income Group Thresholds

Year	Percentile		
	50th	90th	99th
2013	32,358	100,273	285,778
2014	33,177	103,773	305,176
2015	33,787	107,312	305,064
2016	34,689	109,302	307,558
2017	35,912	113,786	322,250
2018	38,068	117,761	339,049
2019	39,844	122,640	354,108

Note: This table reports the reference estimates of the 50th, 90th, and 99th percentiles of the national individual income distribution for our measure of total income. These income percentiles are calculated from the individual pre-tax national income estimates from [Auten and Splinter \(2024\)](#), adjusted to reflect each income group's share of pre-tax national income from wage and pass-through income. We use these thresholds to assign individuals to income groups based on their total wage, pass-through, and self-employment income in the same corresponding year. All amounts are in nominal dollars. For details, see Appendix Section [C](#)

Table A2: Employment and Establishment Shares by Legal Form of Organization and Industry

Industry	Legal Form of Organization				
	C Corps	S Corps	Partnerships	Sole Props	Govt., Nonprofits, & Other
All industries					
<i>Employee share</i>	43.3	27.1	12.2	3.9	13.5
<i>Establishment share</i>	24.5	46.0	12.0	9.9	7.6
Construction					
<i>Employee share</i>	25.4	59.3	9.9	5.2	0.1
<i>Establishment share</i>	13.8	63.7	8.6	13.7	0.1
Manufacturing					
<i>Employee share</i>	63.3	24.4	9.2	2.5	0.6
<i>Establishment share</i>	31.9	46.4	13.4	7.9	0.4
Wholesale trade					
<i>Employee share</i>	55.1	31.5	9.6	2.5	1.3
<i>Establishment share</i>	39.5	46.5	8.9	3.9	1.2
Retail trade					
<i>Employee share</i>	62.0	24.9	8.3	3.5	1.3
<i>Establishment share</i>	38.8	40.1	10.6	9.1	1.4
Transportation and warehousing					
<i>Employee share</i>	65.7	23.1	7.3	3.3	0.6
<i>Establishment share</i>	28.2	50.0	9.9	11.3	0.6
Finance, insurance, and real estate					
<i>Employee share</i>	61.8	16.8	12.3	3.6	5.6
<i>Establishment share</i>	31.3	41.3	17.0	7.2	3.2
Professional, scientific & technical					
<i>Employee share</i>	43.8	31.8	16.9	4.3	3.2
<i>Establishment share</i>	18.7	61.7	9.7	8.8	1.1
Administrative and support					
<i>Employee share</i>	46.1	35.1	13.6	3.8	1.4
<i>Establishment share</i>	21.5	50.5	11.3	14.7	2.0
Education					
<i>Employee share</i>	7.7	9.2	3.6	2.4	77.2
<i>Establishment share</i>	12.0	37.7	9.5	9.8	31.0
Health care and social assistance					
<i>Employee share</i>	17.7	20.8	11.2	3.5	46.8
<i>Establishment share</i>	16.7	44.5	9.4	10.5	18.9
Accommodation and food services					
<i>Employee share</i>	28.7	37.4	25.8	6.8	1.3
<i>Establishment share</i>	22.0	40.6	24.8	11.9	0.7
All other industries					
<i>Employee share</i>	46.3	16.4	10.2	3.6	23.5
<i>Establishment share</i>	21.2	33.7	9.5	9.2	26.5

Note: This table reports the distribution of employment and establishments by legal form of organization (LFO) and industry for the year 2022. The data are based on the Census Bureau's publicly released County Business Patterns (CBP) data ([U.S. Census Bureau, 2022](#)).

Table A3: Average Share of Total Firm Wages and Business Income to Individual Groups

Panel A: Export-Demand Shock Sample			
	Share of Income to Each Group (%)		
	Workers	Owners	Total
<u>Income Groups</u>			
0–50th Percentiles	13.2	2.5	15.7
50–90th Percentiles	37.1	2.3	39.4
90–99th Percentiles	8.5	7.6	16.1
Top 1%	1.5	27.3	28.8
<u>Total</u>	60.3	39.7	
Panel B: Value-Added Shock Sample			
	Share of Income to Each Group (%)		
	Workers	Owners	Total
<u>Income Groups</u>			
0–50th Percentiles	24.0	2.7	26.7
50–90th Percentiles	33.0	7.1	40.1
90–99th Percentiles	4.8	13.2	18.0
Top 1%	0.6	13.9	14.5
<u>Sex</u>			
Male	37.7	28.1	65.8
Female	25.4	8.8	34.2
<u>Race/Ethnicity</u>			
White	42.0	29.0	71.0
Black	3.3	0.6	3.9
Asian	2.8	3.6	6.4
Hispanic	6.8	1.8	8.6
Other/Missing	8.0	2.0	10.0
<u>Total</u>	62.9	37.0	

Note: This table reports the average share of each firm’s total wage and business income allocated to different income groups, where total income for each individual combines wages from W-2 forms and business income from K-1 forms. For each firm j , we define the share of baseline income going to group g as $s_j^g = \sum_{i \in g} w_i / \sum_i w_i$, where w_i is the income that individual i receives from the firm. Each row reports the average of s_j^g across all firms in that sample for a given group. Panel A displays results for the export-demand shock sample and Panel B for the Value-Added Shock sample. Within each panel, the columns show the portion of total income going to workers, owners, and both combined (“Total”). The shares are calculated using all individuals in our stayers sample, but they are very similar if we use the all-income samples. To construct the income groups, we rank individuals each year based on their combined W-2 wages, K-1 business income, and self-employment income (Schedule C). Because this measure is at the individual level and excludes certain forms of income, we adjust the publicly available percentile thresholds from [Auten and Splinter \(2024\)](#) accordingly. For reference, the 2018 thresholds for the 50th, 90th, and 99th percentiles were approximately \$38,067, \$117,760, and \$339,046, respectively. The other/missing race and ethnicity category includes Native Hawaiian and Pacific Islander, American Indian and Alaska Native, and mixed-race individuals. The white, black, Asian, and other/missing categories only include non-Hispanic individuals.

C Census and IRS Data Construction and Linking

C.1 Worker-Level Wage, Business Income, and Self-Employment Data

We combine multiple sources of administrative tax data including individuals’ wages and salaries (W-2s), business income (K-1s), and self-employment income (Schedule C) to construct a comprehensive measure of individual-level income. We link these datasets together using the Census Bureau’s unique Protected Identification Key (PIK), which provides an anonymized identifier for each individual.

W-2 Wage and Salary Data We measure individuals’ wage and salary earnings using the universe of IRS W-2 records. Employers file these records annually for each employee (i.e., they do not include payments to contractors). We define W-2 earnings as the sum of Box 1 income (wages, tips, other compensation) and a subset of Box 12 income (deferred compensation subject to FICA taxes but not income taxes).⁵⁰ Consequently, our W-2 compensation measure includes all taxable wage and salary income (e.g., wages, bonuses, commissions, tips, and certain non-cash benefits), as well as income placed in deferred retirement accounts that is not yet subject to income taxes. Additionally, W-2 earnings include some forms of equity compensation (Eisfeldt et al., 2023), although such compensation is less common at the S corporations we study. However, our W-2 earnings concept does not include non-taxable benefits, such as health insurance premiums (including employee and employer contributions) or employer contributions to retirement plans.

K-1 Business Ownership and Income Data We identify the owners of pass-through businesses and their business income using IRS K-1 records. These are individual-level records that are filed annually by S corporations and partnerships for each of their owners. Although we focus on S corporations for our analysis, we also include business income from partnerships when we measure individuals’ total income from all firms.

We first provide a brief overview of business income distributions for S corporations.⁵¹ Each year, firms calculate their firm-level net income. Net income is defined as the firm’s total revenue from “ordinary business activities” minus expenses (e.g., cost of goods sold, wages and benefits, and other operating expenses, depreciation, and interest expenses). The firm then allocates this net income to each owner in proportion to their ownership share of the firm. These allocations are reported on the K-1 tax forms. The individual owners then report these business income distributions on their personal tax returns. Form K-1s separate business income into different categories (reported in different boxes) based on the type of income (e.g., ordinary income, real-estate income, interest income, and dividend income). Due to data availability at the Census, we define owners’ business income as K-1 Box 1 (ordinary business income or losses). Ordinary business income includes all income or losses from the business’s primary operating activity (e.g., sales of goods and services) but excludes several types of financial income. Consequently, we exclude firms in finance and insurance (NAICS 52) and real estate and rental and leasing (NAICS 53) from our analysis because we cannot accurately measure their business income.⁵²

It’s important to distinguish between being taxed on business income and taking distributions from the firm. Owners of S corporations are taxed annually on their share of the firm’s ordinary business income, regardless of whether this income is distributed to their personal accounts. Conversely, distributions represent transfers of money from the company to the owner. They are typically tax-free, provided the distributions do not exceed the owner’s basis in the corporation (see Nitti (2014) for details about owners’ basis). Additionally, an owner’s business income can include losses, which can be used to offset their other

⁵⁰The restriction to a subset of Box 12 income is due to the Census Bureau not receiving all W-2 boxes. According to Bee et al. (2023), our subset of Box 12 income includes “elective deferrals to (retirement) plans under Box 12 codes D: 401(k), E: 403(b), F:408(k)(6), G: 457(b), and H: 501(c)(18)(D). These boxes cover 96.3 percent of all elective retirement contributions on W-2s.”

⁵¹Business income distributions at partnerships are largely similar, except that they are not required to allocate total business income to owners based on their ownership shares.

⁵²Cooper et al. (2016) show that for non-finance pass-through firms, K-1 Box 1 includes a substantial amount of total business income distributed to owners. The limited availability of K-1 boxes is also why we exclude partnerships in our analysis. Specifically, we do not observe Schedule 1065 K-1 boxes 4a (guaranteed payments for services) or 4b (guaranteed payments for capital). Since partnerships cannot pay owner/managers W-2 wages, they instead provide these owners with guaranteed payments. Since we do not observe these payments, we cannot accurately measure the total income of owners in partnerships and exclude them from our analysis.

taxable income; however, these deductible losses are limited by the owner’s basis. In our measure of business income, we include the full amount of these losses as reported on the Schedule K-1 forms, even though some owners may not immediately claim the full loss due to basis limitations. Including these losses prior to basis adjustments does not pose a significant problem since unclaimed losses can typically be carried forward and realized in future periods when sufficient basis is available, reflecting their eventual economic impact on the owner’s taxable income.

A known difficulty with allocating pass-through income to individuals is that the proximate recipient of K-1 income may be another business entity, which then ultimately distributes the income to individuals (i.e., the “owner” on the K-1 is an EIN that corresponds to another pass-through firm rather than an SSN that corresponds to an individual) (Cooper et al., 2016; Love, 2021). We partially address this issue by resolving simple tiered K-1 structures.⁵³ Although tiered K-1s are much less common for S corporations, as their owners are typically individuals, this process helps us accurately measure individuals’ total income, including business income from partnerships.

Schedule C Self-Employment Data We measure individuals’ self-employment status and income using IRS Schedule C and Schedule SE records. Specifically, we define each individual’s self-employment income as their total Schedule SE self-employment income if they also filed a Schedule C. We allocate self-employment to individuals rather than households based on which PIK is associated with the Schedule SE filing (e.g., if a sole proprietorship firm is co-owned by a couple and they both report Schedule SE income, we would separately measure their individual self-employment income).

Constructing Individual-Level Earnings Percentile Measures Our measure of income is the total *individual* receipt of income from wages and salaries and income from pass-through businesses and self-employment. Correspondingly, we use the data from Auten and Splinter (2024) to create reference estimates of the 50th, 90th, and 99th percentiles of the national individual-level distribution of a comparable income measure. This ensures that while the individuals in our sample are drawn from a specific population of individuals who receive wage or business income from an in-sample S corporation, they are assigned to income groups using nationally representative and publicly available definitions. In addition, by using both individual-level definitions of income and individual-level distributional estimates, we avoid the issues that arise from aligning the definitions of tax units with individuals.

Specifically, we use the nominal income group thresholds for the individual-weighted distribution of pre-tax national income from Table C3a in Auten and Splinter (2024), and we adjust those income group thresholds using each income group’s share of pre-tax national income coming from wage and pass-through income from Table C5a in Auten and Splinter (2024). Since the income shares are only available by income group, we use the published income shares for the income group immediately below each income threshold to help ensure that the adjustments are more representative of individuals who are at the cutoff.⁵⁴ As a result, our distributional thresholds reflect the adjustments to fiscal income & income groups made by Auten and Splinter (2024), including sample corrections and adjustments to the number of individuals in a tax unit, but it excludes the expansions to fiscal income (e.g., retained earnings and taxes at C corporations, employer payroll taxes and benefits, retirement account income, or imputed rents) that are also not captured in our income measures.

We calculate the nominal income thresholds for each year and assign individuals to income groups using thresholds contemporaneous with the year of income. As a result, our income group definitions are relative to the national distribution of income in the same year, and we do not need to make any additional

⁵³Specifically, we resolve all *tiered entities* where the chain of business income payments is of length one. For example, consider the following chain of K-1s: $EIN1 \rightarrow EIN2 \rightarrow SSN1$. In this case, we would allocate the business income from EIN 1 to SSN 1. However, we do not resolve more complex chains of K-1 payments.

⁵⁴We use the income share estimates for individuals in the 0–50th percentiles of the income distribution to adjust the 50th percentile income threshold, the income share estimates for individuals in the 50–90th percentiles to adjust the 90th percentile income threshold, and the income share estimates for individuals in the 95–99th percentiles to adjust the 99th percentile income threshold. Our key assumption for this approach is that since the shares are weighted by income, they will be more representative of the income mix for individuals near the top of any income interval.

adjustments for changes in prices or incomes over time. Table A1 reports the nominal income thresholds that we use to determine income groups in each year.

Individual-Level Demographics We use survey and administrative data from the SSA Numident, the 2000 and 2010 Decennial Censuses, and all available waves of the American Community Surveys to identify the demographics and educational attainment of individuals in our data. Whenever multiple sources of data are available, we use a file assembled by the LEHD program that uses the most recent available data for all individuals who ever appear in one of the component datasets. Since we draw from several population-level demographic datasets, the expected coverage for individuals in our sample is also high. From [Graham et al. \(2022\)](#), the expected share of workers missing data on sex and age is approximately 5%, and the expected share of workers missing data on race and ethnicity is approximately 20%. The share of workers missing data on educational attainment is much higher (85%), primarily because data on educational attainment is only available from the 2000 Census long form and the American Community Surveys. We only use education responses collected after age 25, so coverage rates are as high as 25% for individuals born before 1976, and decline steadily for younger cohorts.⁵⁵

We calculate demographic group-specific income shares and pass-through estimates based on only valid observed data, and we do not impute any demographic characteristics. All individuals with missing sex or race and ethnicity data are tabulated in a separate group. Correspondingly, group-specific baseline income shares can be considered lower bounds for the true income share of each group.

C.2 Firm-Level Employment, Income, and Exporting Data

Firm Identifiers: EINs versus Census Firm IDs The firm-level datasets we combine for our analysis are sometimes at the Employer Identification Number (EIN)-level and sometimes at the more aggregated Census firm ID-level datasets. An EIN is an administrative taxpayer ID assigned by the IRS. An EIN could correspond to one or more establishments (i.e., physical locations where business is conducted). However, a single firm may file different types of taxes under different EINs (e.g., a firm might file payroll and corporate taxes under different EINs or use separate EINs for its East and West Coast operations).⁵⁶ The Census Bureau aggregates multiple EINs corresponding to the same firm under a single Census firm ID (or “Enterprise Unit”). Specifically, the Census defines firms as groups of EINs that are under the same ownership or control (the survey question specifically asks whether another company “owns more than 50 percent of the voting stock of your company.”) For our analysis, we aggregate all EIN-level data to the Census firm ID level. For example, our firm-level revenue measure is the sum of all EIN-level revenue associated with the same firm ID, and we aggregate individuals’ business income and wages from firms in a similar manner.

Census Longitudinal Business Database (LBD) Our baseline dataset of U.S. firms is the Longitudinal Business Database ([Chow et al., 2021](#)). It is an annual establishment-level dataset covering the universe of non-farm, private-sector firms in the U.S. It includes measures of establishment-level employment, payroll, industry, and location. We use [Fort and Klimek \(2016\)](#)’s 2012 NAICS codes to classify each establishment into time-consistent industry codes. The LBD also contains a crosswalk between establishments, EINs, and Census firm IDs. We define firm-level industry and location measures based on which industry or location has the largest share of employment at the firm. Crucially, for our analysis, the LBD also contains the ‘legal form of organization’ (LFO) for each EIN, which is derived from firm-level IRS tax filings (e.g., whether the EIN files taxes as a C corporation, S corporation, etc.). Finally, we construct a *Longitudinal Firm ID* based on the LBDFID variable in the LBD. This variable enables us to consistently track firms across

⁵⁵For details on the differential coverage of educational data by age, see the discussions in Section 2.2 of ([Stinson and Wang, 2025](#)).

⁵⁶[Song et al. \(2019\)](#) write that: “An EIN is not always the same, however, as the ultimate parent firm. Typically, this is because large firms file taxes at a slightly lower level than the ultimate parent firm. The 4,233 New York Stock Exchange publicly listed firms in the Dun and Bradstreet database reported operating 13,377 EINs in 2015, or an average of 3.2 EINs each. For example, according to Dun and Bradstreet, Walmart operates an EIN called “Walmart Stores”, which operates the domestic retail stores, with different EINs for the Supercenter, Neighborhood Market, Sam’s Club, and online divisions. As another example, Stanford University has four EINs: the university, the bookstore, the main hospital, and the children’s hospitals.”

time and corrects for known issues with the LBD’s firm ID breaking across time (Haltiwanger et al., 2013). Specifically, we utilize the flow of establishments across firm IDs to correct for breaks in the LBD’s firm IDs. For example, if all establishments at firm ID A move to firm ID B in a given year, we assume that firm IDs A and B belong to the same Longitudinal Firm ID. We use these longitudinal Firm IDs as the firm-level unit of observation for all our analyses (e.g., a worker separation is defined as a worker leaving a Longitudinal Firm ID).

Firm-Level Sales and Cost of Goods Sold Data We measure firm-level sales/revenue and the cost of goods sold (COGS) using extracts of firm-level IRS tax filings available in the Census Bureau’s Business Register (BR). We follow Haltiwanger et al. (2019) to extract these data from the BR, although we have modified the process to include more matches and to extract firm-level COGS from the Business Register. We use these firm-level measures as outcome measures and to construct the value-added shocks. For S corporations, we define the following measures:

- **Sales/Revenue** = Form 1120-S Box 1c (“gross receipts or sales less returns and allowances from your core trade or business”).
- **Cost of Goods Sold** = Form 1120-S Box 2 (“cost of goods sold”).
- **Value Added** = Sales - Cost of Goods Sold.

We calculate the firm-level measures of sales and COGS as the sum of sales and COGS across all EINs under the same Census firm ID.

Firm-Level Export Data The Longitudinal Firm Trade Transactions Database (LFTTD) provides us with annual firm-by-product-by-country export and import values for the universe of exporters in the U.S. (Kamal and Ouyang, 2020). The data include the value of all “merchandise exports” (i.e., excluding services) from U.S. firms to foreign countries. The data are derived from U.S. Customs and Border Protection customs declaration forms. Export values are separated into different products based on their HS-6 product code (see [here](#) for details about HS codes).

C.3 Sample Construction and Restrictions

Sample Construction and Matching For our analysis, we focus on six “cohorts” of firms that are treated in the years 2013–2018. We then follow these cohorts for five years before treatment and three years afterward. The year restrictions are necessitated by the availability of K-1 data, which we only have for the years 2013–2022.⁵⁷

For each cohort, we start with the universe of LBD firms and select all firms with (1) 10+ employees, (2) at least 75% of employment at S corporation EINs, and (3) not in the finance and insurance (NAICS 52) or real estate and rental and leasing (NAICS 53) industries.⁵⁸ For each cohort, we track these treated firms across time, regardless of whether they continue to meet the sample restrictions.⁵⁹

For each in-sample firm, we merge on the firm’s ownership (from K-1s), workers (from W-2s), sales and COGs (from Form 1120-S firm-level tax filings), and export data (from the LFTTD). We link together these datasets using different EIN-matching methods:

- **LFTTD**: we link each LFTTD observation to an LBD Firm using the LFTTD’s EIN to FIRMID crosswalk described in Kamal and Ouyang (2020).

⁵⁷Prior to 2013, we observe the owners of each S corporation but do not observe the business income they received from their firm. For these years, we impute individual-level business income from each firm based on the number of firm owners and the total firm-level net income on the firm’s 1120-S. We use these imputed business income measures for the pre-trend coefficients for some cohorts, but they never influence any of our treatment-effect estimates.

⁵⁸For each cohort t , we make all sample restrictions based on year $t - 1$ because year t employment could be affected by treatment.

⁵⁹For the export-demand and value-added shocks, a single firm can be treated multiple times if they are in multiple cohorts. We stack data across cohorts and estimate treatment effects pooled across cohorts.

- *W-2s*: we link W-2s to the LBD following the EIN matching process in [Stinson and Wang \(2025\)](#). They show that this yields W-2 matches for 97% of all LBD firms because firms typically use the same EIN to file their Form 941 payroll taxes and their W-2s.
- *K-1s and Form 1120-S*: The K-1 ownership information and firm-level income statement data come from different parts of the same IRS tax filings (i.e., the S corp K-1s are from Schedule K-1 of Form 1120-S). Consequently, we match these data to the LBD observations in tandem based on EIN matches. This matching process is more complicated than the W-2 matching process because firms often use different EINs to file their income taxes (e.g., forms 1120-S and Schedule K-1s) and their payroll taxes (e.g., Form 941). For example, [Haltiwanger et al. \(2017\)](#) estimate that 20% of firms use different EINs to file these taxes. To match these firms, we search for EIN matches using various types of matching strategies, including:
 - *Direct EIN Matches*: We first check for direct matches where the firm uses the same EIN in the same year to file income and payroll taxes.
 - *EIN Matches using Lagged and Leading EINs*: For income EINs that we cannot match directly, we check for matches using the lagged and leading EINs associated with the same LBD firms. For example, if LBD firm 1 used EIN A in year $t - 1$ and we observe a K-1 for firm EIN A in time t that doesn't directly match any of the LBD firms, we would match firm 1 to EIN A. The rationale is that firms may change which EINs they use to file taxes over time, so the lagged and leading EIN crosswalks help link the different types of EINs together
 - *Name and Address Matches*: We match previously unmatched EINs based on direct firm name and location (e.g., city/state or ZIP code) matches.

Sample Restrictions: After the previously described matching process, we construct separate analysis samples for the export-demand shocks and value-added shocks. We first implement the following sample restrictions on the individual and firm-level data.

- *Firm-Level Sample Restrictions*:
 1. The firm has an export-demand (value-added) shock defined at time t .
 2. For the export-demand shock sample, we further restrict to firms that have at least \$5,000 in exports per worker at time $t - 1$.
 3. The firm has non-zero sales at time $t - 1$. This requires a match to the BR-based firm-level tax data. It is necessary because one of our main outcome variables is firm-level sales, so we need to observe sales at time $t - 1$ to calculate the absolute percent change.
- *Individual-Level Sample Restrictions*: for each firm, we define the set of treated individuals as all individuals who received non-zero income from the firm at time $t - 1$ (either W-2 income or K-1 business income). For workers (i.e., individuals who only receive W-2 income), we require that they receive at least \$15,000 in real annual income in year $t - 1$. This threshold is roughly equivalent to earning the federal minimum wage while working full-time, thereby mitigating the effects of transitions from part-time to full-time work.

Finally, for our final analysis sample, we require that each firm has at least one matched W-2 employee and one matched K-1 owner at time $t - 1$ (i.e., we require a K-1 and W-2 match at time $t - 1$). If we included firms with incomplete W-2 and K-1 matches, we could not accurately measure the share of income going to workers or owners because data on one of the groups would be missing.

D Incidence Symmetry in Wage-Setting Models

Our notion of the worker incidence, defined in Equation 5, is the share of total income changes going to workers, divided by the total income changes going to both workers and owners. Consequently, our incidence measure is a *levels* concept rather than an *elasticity*. In Section 5.4, we show that the worker incidence is *asymmetric* both when we only consider income to workers and owners who stay at the treated firm and when we include extensive-margin income changes.

In this section, we sketch simple wage-bargaining and firm-insurance models and show that both models yield a *symmetric* incidence where the incidence notion matches the levels incidence measure we use in our paper. Although more complicated versions of the models can yield asymmetric responses, these simple versions illustrate that incidence symmetry is an implicit prediction of several prominent models.

D.1 Wage-Bargaining Models

In the simplest version of Nash bargaining, a worker's wage is $w = b_w + \gamma \cdot (z - b_w)$ where b_w is the value of the worker's outside option, z is the worker's productivity, and γ is the worker's bargaining power. Consequently, the worker always gets a share γ of the surplus changes, and owners get $1 - \gamma$, so the incidence is symmetric.

Union Wage-Bargaining Model

Next, we show that the symmetric incidence also holds when firms employ multiple workers and can adjust employment and wages in response to productivity shocks. Specifically, we adopt the two-stage “right-to-manage” union bargaining framework of Abowd and Lemieux (1993).

All workers jointly bargain with the firm and then split the surplus equally. The firm first chooses the optimal amount of labor, and then workers and the firm bargain over the wage, conditional on the optimal labor. Revenue is $R = z \cdot F(L)$, where z is a Hicks-neutral productivity shifter and L is employment.

Stage 2: wage setting (given L). The optimal wage maximizes a Nash product of the firm and workers' total surplus with worker bargaining power γ :

$$\max_w \left[L(w - b_w) \right]^\gamma \left[R - wL - b_f \right]^{1-\gamma}$$

where b_w and b_f are the worker and firm outside options. Setting $b_f = 0$, the negotiated wage is

$$w(L) = (1 - \gamma) b_w + \gamma \frac{R}{L}, \tag{A1}$$

a weighted average of the worker's outside option b_w and average revenue per worker.

Stage 1: employment choice. The firm chooses L anticipating A1. The resulting first-order condition equates marginal revenue and the outside wage, $z F'(L) = b_w$.⁶⁰

Surplus split. The total surplus to be split between workers and owners is $S = R - L b_w$. From A1 owner profits π and total worker surplus $L(w - b_w)$ are equal to

$$\pi = (1 - \gamma) S, \quad L(w - b_w) = \gamma S.$$

So owners and workers always receive fixed shares $1 - \gamma$ and γ .

⁶⁰Because we set $b_f = 0$, all terms involving γ cancel in the employment FOC, so the equilibrium L and the ratio R/L are *independent of bargaining power*. Because our incidence measure divides changes in wages and profits by total surplus, we omit the algebra behind this result.

Incidence of productivity shocks. For any shock dz the change in surplus can be decomposed into the portions going to owners, incumbent workers, and new hires

$$\frac{dS}{dz} = \underbrace{\frac{d\pi}{dz}}_{\text{Owners}} + \underbrace{L \times \frac{dw}{dz}}_{\text{Incumbent Workers}} + \underbrace{\frac{dL}{dz} \times \gamma \times \left(\frac{S}{L}\right)}_{\text{New Hires}}. \quad (\text{A2})$$

Equation A2 shows that the owners' share of surplus changes (the owners' incidence) is constant and equal to

$$\frac{d\pi}{dS} = (1 - \gamma). \quad (\text{A3})$$

Consequently, this extended bargaining model with endogenous new hires preserves the symmetric incidence prediction. However, the worker surplus here is different from what we currently measure in the paper (i.e., our empirical estimates of the total change in workers' income plus owners' income do not perfectly map to dS). Specifically, Equation A2 includes surplus accruing to new hires, while we do not include that term in our incidence definition (we only include the first two terms in Equation A2). This is one motivation for our ongoing effort to incorporate new hires into our incidence framework, allowing us to calculate a notion of incidence that includes new hires.

D.2 Simple Firm-Insurance Model

We sketch a simple risk-sharing model between a worker and a firm and show that it predicts a symmetric incidence of firm productivity shocks. The below formulation follows Section 2.1 in Chapter 6 of [Cahuc and Zylberberg \(2004\)](#) but assumes away endogenous labor supply.

Environment. A single risk-averse worker and a possibly risk-averse owner contract before an idiosyncratic productivity shock z is realized. Wages are $w(z)$ and profits are $\pi(z) = z - w(z)$. The worker and owner utility functions are $u(w)$ and $v(\pi)$, respectively.

Contract. Given the worker's outside utility $\bar{u}$, the firm chooses a wage schedule $A = \{w(z)\}$ that satisfies

$$\max_A \mathbb{E}[v(\pi(z))] \quad \text{s.t.} \quad \mathbb{E}[u(w(z))] \geq \bar{u}.$$

The Lagrangian $\mathbb{E}[v] + \lambda(\mathbb{E}[u] - \bar{u})$ yields the point-wise first-order condition

$$v'(\pi(z)) = \lambda u'(w(z)) \quad \forall z. \quad (\text{A4})$$

Pass-through. Differentiating A4 with respect to z gives

$$v''(\pi) \left(1 - \frac{dw}{dz}\right) = \lambda u''(w) \frac{dw}{dz},$$

so

$$\frac{dw}{dz} = \frac{v''(\pi)}{v''(\pi) + \lambda u''(w)}. \quad (\text{A5})$$

Incidence. Worker incidence as defined by Equation 5 corresponds to

$$\beta = \frac{\Delta w}{\Delta w + \Delta \pi} = \frac{dw/dz}{(dw/dz) + (1 - dw/dz)} = \frac{dw}{dz}.$$

Symmetric pass-through with CARA utility. With constant-absolute-risk-aversion utilities $u(w) = -e^{-\gamma_w w}$ and $v(\pi) = -e^{-\gamma_o \pi}$, A4 implies $\gamma_o e^{-\gamma_o \pi} = \lambda \gamma_w e^{-\gamma_w w}$. Substituting into A5 yields the constant

$$\beta = \frac{\gamma_o}{\gamma_o + \gamma_w}. \quad (\text{A6})$$

Equation A6 highlights that the wage pass-through rises when the firm owner is *more* risk-averse ($|\gamma_o|$ increases) because owners require greater insurance, and it falls when workers are more risk-averse ($|\gamma_w|$ increases). In the polar case of a risk-neutral firm $\gamma_o = 0$, pass-through is zero, and the owner receives all of the gains from positive productivity shocks, while bearing all of the losses from negative productivity shocks. Crucially, however, the ratio dw/dz is *constant*, so the division of gains and losses is identical for positive and negative shocks and does not vary with their magnitude. The model therefore predicts a symmetric, size-invariant incidence.

D.3 Extended Firm-Insurance Models that Predict Asymmetry

Several extensions to the baseline firm-insurance model predict that wages respond more to positive shocks than negative shocks. However, this is the opposite degree of asymmetry that we find. For example, Harris and Holmstrom (1982) assume one-sided firm commitment where firms can commit to not renege on long-term contracts, but workers will renege when advantageous. In this case, wages increase in response to *positive* productivity signals but are unaffected by *negative* signals. Balke and Lamadon (2022) show similar asymmetry in response to worker productivity shocks without imperfect monitoring considerations. Furthermore, Thomas and Worrall (1988) assume that both workers and firms can renege on the long-term contract (i.e., both sides feature limited commitment). In this case, the wage response to shocks is nonlinear because wages respond more to large shocks (where there is a larger incentive to renege) than to small shocks. However, the two-sided limited commitment model does not feature the positive versus negative asymmetry we document. However, none of the papers mentioned in this paragraph consider *firm-specific* productivity shocks, which we study in this paper. Instead, they consider market-wide or worker-specific productivity shocks.

Appendix References

- Abowd, J. A. and T. Lemieux (1993) “The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada,” *The Quarterly Journal of Economics*, Vol. 108, pp. 983–1014.
- Auten, Gerald and David Splinter (2024) “Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends,” *Journal of Political Economy*, Vol. 132, pp. 2179–2227.
- Balke, Neele and Thibaut Lamadon (2022) “Productivity Shocks, Long-Term Contracts, and Earnings Dynamics,” *American Economic Review*, Vol. 112, pp. 2139–2177.
- Bee, Adam, Joshua Mitchell, Nikolas Mittag, Jonathan Rothbaum, Carl Sanders, Lawrence Schmidt, and Matthew Unrath (2023) “National Experimental Wellbeing Statistics,” *CES Working Paper Number 23-04*.
- Cahuc, Pierre and André Zylberberg (2004) *Labor Economics*, Cambridge, MA: MIT press, second edition edition.
- Chow, Melissa, Teresa Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White (2021) “Redesigning the Longitudinal Business Database,” *NBER Working Paper*, p. w28839, Place: Cambridge, MA.
- Cooper, Michael, John McClelland, James Pearce, Richard Prisinzano, Joseph Sullivan, Danny Yagan, Owen Zidar, and Eric Zwick (2016) *Business in the United States: Who Owns It, and How Much Tax Do They Pay?*, Vol. 30 of Tax Policy and the Economy, Cambridge: MIT Press.
- Eisfeldt, Andrea L., Antonio Falato, and Mindy Z. Xiaolan (2023) “Human Capitalists,” *NBER Macroeconomics Annual*, Vol. 37, pp. 1–61, Publisher: The University of Chicago Press.
- Fort, Teresa C and Shawn Klimek (2016) “The Effects Of Industry Classification Changes On Us Employment Composition,” *CES Working Paper 18-28*, p. 24.
- Graham, Matthew, Erika McEntarfer, Kevin McKinney, Stephen Tibbets, and Lee Tucker (2022) “LEHD Snapshot Documentation, Release S2021_R2022Q4,” *Working Papers*, Number: 22-51 Publisher: Center for Economic Studies, U.S. Census Bureau.
- Haltiwanger, John, Ron Jarmin, Robert Kulick, Javier Miranda, and Veronika Penciakova (2019) “Augmenting the LBD with Firm-Level Revenue,” *CES Technical Notes CES-TN-2019-02*.
- Haltiwanger, John, Ron S Jarmin, Robert Kulick, and Javier Miranda (2017) “High Growth Young Firms: Contribution to Job, Output and Productivity Growth,” *CARRA Working Paper Series*, p. 75.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda (2013) “Who Creates Jobs? Small versus Large versus Young,” *Review of Economics and Statistics*, Vol. 95, pp. 347–361.
- Harris, Milton and Bengt Holmstrom (1982) “A Theory of Wage Dynamics,” *The Review of Economic Studies*, Vol. 49, p. 315.
- Kamal, Fariha and Wei Ouyang (2020) “Identifying U.S. Merchandise Traders: Integrating Customs Transactions with Business Administrative Data,” *CES Working Paper 20-28*.
- Love, Michael (2021) “Where in the World Does Partnership Income Go? Evidence of a Growing Use of Tax Havens,” *SSRN Electronic Journal*.
- Nitti, Tony (2014) “Determining the Taxability of S Corporation Distributions: Part I.”
- Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter (2019) “Firming Up Inequality,” *The Quarterly Journal of Economics*, Vol. 134, pp. 1–50.
- Stinson, Martha and Sean Wang (2025) “The Composition of Firm Workforces from 2006–2022: Findings from the Business Dynamics Statistics of Human Capital Experimental Product,” *CES Working Paper 25-20*.
- Thomas, Jonathan and Tim Worrall (1988) “Self-Enforcing Wage Contracts,” *The Review of Economic Studies*, Vol. 55, pp. 541–554.