The Incidence of Distortions*

David Atkin
MIT and NBER

Baptiste Bernadac
UCLA

Dave Donaldson
MIT and NBER

Tishara Garg
MIT

Federico Huneeus
Duke and Central Bank of Chile

PRELIMINARY AND INCOMPLETE

Abstract

Economic distortions—such as market power, taxes, credit constraints, etc.— are fundamental in understanding income differences across countries. Recent work has documented the pervasive extent of economic distortions and how they lead to substantial aggregate productivity loss. Far less well understood is how these phenomena affect members of society differently. In this paper we combine unique datasets from Chile, linking workers and owners to firms, firms to each other, firms to consumers, and firms and consumers to the government, in order to quantify the incidence of distortions for the first time.

*This version: November 2024. Author contacts: atkin@mit.edu, bbernadac@ucla.edu, ddonald@mit.edu, tishara@mit.edu, and federico.huneeus@duke.edu. We are grateful to Joaquin Galeno and Qirui Wang for outstanding research assistance, to STEG for funding, and to numerous audience members for comments that have improved this work. The findings, interpretations, and conclusions are those of the authors and do not necessarily represent those of the Central Bank of Chile nor its board members.

1 Introduction

Economic distortions are pervasive and fundamental to understanding why countries sit at different levels of development. These distortions arise, for example, when firms enjoy market power, face borrowing constraints, or pay taxes or tariffs on outputs and inputs. Market failures such as these reduce the economy's overall efficiency, and hence its aggregate amount of production. A large literature has produced a deep understanding of the consequences that these distortions have for aggregate productivity (Restuccia and Rogerson (2008), Hsieh and Klenow (2009)). What is far less well understood, however, is the extent to which the burden of these distortions is shared equally across households—the question of "Who bears the incidence of distortions?" remains open. If the poor bear a greater incidence then distortions have clear equity costs as well as efficiency ones.

In this paper we build and analyze a new dataset from Chile that is designed to illuminate the incidence of economic distortions across households for the first time. In particular, we merge data on consumer-to-firm transactions and firm-to-firm transactions—both at the product level—with firm-to-employee wage payments, firm-to-individual ownership registries, debtor-to-lender financial loan transactions, and individual- and firm-level tax and transfer payments. The result is a micro-level system of national income and product accounts depicting the flow of goods and services, from each household to each firm in the supply of factors, from each firm to every other firm, and finally from each firm back to each household in the form of final consumption.

Central to this exercise is the fact that large stores and chains in Chile now transmit to the tax authorities not only the quantities and prices of each product sold but also the tax ID of the consumer they are sold to. This granularity allows us to link individuals to both the products they purchase and the stores at which they shop, with the firm-to-firm transaction data further providing visibility on where that product was sourced from, and at what price, all the way up the supply chain. In contrast to widely-used retail scanner data, the coverage goes beyond supermarket retail—with household consumption surveys at the store-brand level filling in the remaining gaps via statistical matching—and the consumer and retailer identifiers allow us to merge these data with firm-level datasets (on their employees, owners, suppliers, and tax payments) and household-level datasets (on their employment, firm ownership, pension claims, tax payments, and transfer receipts).

We then embed this new dataset inside a general equilibrium model of the Chilean economy. This model follows the framework laid out in Baqaee and Farhi (2020). Households supply heterogeneous amounts of factors of production (capital and various types of labor) to firms, firms use arbitrary technologies to make output with factors and in-

termediates, households consume firms' goods with heterogeneous preferences, and the government makes net tax/transfer payments from every household and firm in heterogeneous ways.

On top of these flows of goods and services, we then add an arbitrary set of distortions on every bilateral exchange. For example, when firms sell to any given firm or consumer, they may charge a markup as a result of their output market power. Analogously, when they buy from any factor or firm they may enjoy input market power and hence charge a markdown. Similarly, taxes, bribes, and credit constraints drive a wedge between the price that the seller effectively receives and the price that the buyer effectively pays, and this is the essence of any economic distortion.

While it is well known that such distortions lead to Pareto-inefficient outcomes, less focus has been paid to their heterogeneous consequences across the household distribution. Clearly, households may be differentially exposed to such distortions through their consumption (e.g. when they buy from suppliers with high product market power), their supply of factors of production (e.g. when they work for employers with high labor market power), their ownership of firms (and hence their capture of the rents that result from market power), and their participation in tax and transfer schemes. Further, in all of these cases, a household can be both directly and indirectly exposed due to their position in supply chains—for example, when a consumer buys a final good, they will exposed both to any markup charged by the final seller as well as to those charged by sellers further up the supply chain.

The data architecture that we pair with the Baqaee and Farhi (2020) model provides visibility on all of these interconnected phenomena for the first time. Our approach goes beyond the analysis of specific distortions in specific sectors that characterizes much of the limited existing work on this topic. Such comprehensiveness is important. When a second-best economy is simultaneously affected by multiple distortions the welfare effect of eliminating any specific distortion will be a function of other distortions in the economy. For example, what may appear as a harmful distortion when focusing within a sector may actually be beneficial if its presence happens to mitigate the effects of other distortions across sectors. Thus, to fully grasp both the equity and efficiency consequences of distortions one has to go beyond specific distortions in specific sectors. Our wide-reaching approach quantifies the relevance of this issue of overlapping distortions.

A central challenge when studying the impact of distortions arises in simply arriving at measures of the distortions themselves. We employ standard techniques from the misallocation literature to do so, noting both that these methods only uncover products of input and output distortions and that these same products are all that are required to

estimate incidence.

Armed with such estimates of distortions on exchanges throughout the economy, as well as data on the network of such linkages between individuals, households, firms and the government, our final step is to conduct a series of counterfactual simulations that illuminate the incidence of distortions in our model economy. Our main exercise reduces all distortions proportionately to their size.[1] While this across-the-board reform is an extreme scenario, it relates to the typical goal of work in the misallocation literature, which is to assess the aggregate productivity gains that result from reducing distortions. We go beyond such aggregate impacts and answer the question of who is relatively harmed and helped by the presence of the distortions that are in place throughout the economy, uncovering differential incidence across groups based on income levels, age, and gender. We find that the burden of distortions falls relatively more on the shoulders of the poor, the young and women. The channels though are very different across groups. The poor gain relatively more from reducing distortions because of larger reductions in the prices of the types of good they buy, larger increases in labor income and smaller reductions in profit income. The young and women gain more due to bigger gains in labor income.

Related Literature

This paper relates to several strands of the literature. First, we draw on theoretical and empirical tools in the literature that quantifies aggregate efficiency losses from distortions. This includes the seminal work of Restuccia and Rogerson (2008) and Hsieh and Klenow (2009), extended to a networked economy by Jones (2013) and Baqaee and Farhi (2019).

However, as discussed above, our primary interest lies not in the aggregate extent of misallocation—the effect on a hypothetical representative agent who buys all goods and owns all factors of production—but in its heterogeneous consequences—its incidence—across the many different types of agents throughout any modern economy. This interest is inspired by examples such as those stressed by Schmitz (2020), who has argued that the particularly odious consequences of monopoly power may derive more from its unequal incidence than from its effects on aggregate efficiency. This would be the case if, for example, poor households are more likely (than rich households) to consume goods with high markups, or to derive income from sources that are less likely to be connected (directly or indirectly) to firms that charge markups. We are unaware of attempts to quan-

[1] Our approach also allows us to remove certain wedges or sets of wedges and asks which particular distortions are most responsible for the unequal incidence—both because some wedges may be larger than others and because their incidence differs conditional on size—and where different sets of policies lie on the efficiency-equity frontier. Such an analysis can also shed light on the question posed above about whether certain distortions are countervailing or reinforcing.

tify these broad inequality consequences of output market power. Nor are we aware of efforts to expand the scope of this point by seeing how it interacts with other distortions, beyond output market power, such as monopsony (vis-a-vis labor markets and intermediates markets), credit constraints, and taxes. This is surprising given the potential for countervailing incidence effects across different types of distortions.

By contrast, recent work has made large advances in documenting the incidence of distortions within particular sectors where previous data has allowed progress. For example, Faber and Fally (2022), Gupta (2022), and Sangani (2023) examine how retail markups differ across the income distribution. Similarly, Burstein, Cravino, and Rojas (2024), using a database from Chile that comprises one component of ours, document variation in markups on intermediates across types of buyer firms. Turning to other types of distortions, Faber (2014) and Acosta and Cox (2024) study the extent to which tariffs affect consumers across the income distribution heterogeneously, and Sharma (2023) evaluates the extent to which labor markdowns differ by gender of the employee. The study of the incidence of tax policies has, of course, a long tradition, ranging from partial equilibrium (seller vs. buyer incidence) to general equilibrium (often capital vs. labor). In this vein, an evolving modern approach, exemplified by Conlon, Rao, and Wang (2022), uses microdata to ask questions such as "Who pays sin taxes?" (on alcohol, cigarettes, and sugary drinks) across demographic categories. Our work aims to build on the lessons of studies like these in order to arrive at a bigger-picture understanding of how multiple types of distortions (markups, markdowns, credit constraints, corruption, taxes, etc.), with multiple types of exposure (firms, workers, and consumers, each in both their direct and indirect forms) overlap and interact to determine the overall incidence of distortions.

Finally, our work draws connections between recent contributions in general equilibrium measurement. For example, Adão et al. (2022) and Andersen et al. (2022) build micro-level versions of national accounts by connecting workers and owners to firms, firms to one another, and, in the latter case, firms to consumers. But these studies have considered neoclassical environments with no market distortions, and connecting consumers merely to the firms at which they shop may offer an incomplete picture if distortions vary across products as well as across firms. On the other hand, studies such as Atkin and Donaldson (2022) and Manelici et al. (2024) have built detailed pictures of market distortions in order to understand certain aggregate effects of trade and FDI shocks, respectively, but they have not emphasized the connections of individuals to the economy's distortions in order to quantify the incidence of misallocation.

2 Theory

This section describes the background theory of a general equilibrium economy with arbitrary distortions. We draw heavily on the presentation in Baqaee and Farhi (2020).

2.1 Set-up

The economy consists of C consumers, F primary factors, and N firms. We incorporate international trade and trade imbalances by modeling a "rest of the world" firm and consumer.

Consumers and factors. Each consumer $c \in \mathcal{C}$ has preferences

$$U_c(x_{c1}, \ldots, x_{cN}),$$

over the consumption amounts x_{ci} for each good $i \in \mathcal{N}$, and where each $U_c(\cdot)$ is homothetic (although each consumer can have arbitrarily different homothetic preferences). Each primary factor $f \in \mathcal{F}$ is in fixed supply, with each consumer c owning a share Φ_{cf} of the aggregate factor supply L_f. Similarly, consumer c owns a share Φ_{ci} of firm i. Denoting the prices of good i and factor f as p_i and w_f, respectively, each consumer faces the budget constraint

$$\sum_{i \in \mathcal{N}} p_i x_{ci} = \sum_{f \in \mathcal{F}} \Phi_{cf} w_f L_f + \sum_{i \in \mathcal{N}} \Phi_{ci} \pi_i + T_c \equiv \chi_c,$$

where we let π_i denote any (after-tax) firm profits earned by firm i and T_c denote the net personal taxes paid by household c. Finally, we take the total value of consumption in this economy as the numeraire (i.e. $\sum_i \sum_c p_i x_{ci} = 1$) and let $b_{ci} \equiv p_i x_{ci} / \chi_c$ denote the share of consumer c's expenditure devoted to good i.

Firms. Each good i is produced by a unique and single-product firm that we also denote by i.[2] Each firm has access to its own arbitrary, but constant returns-to-scale, production function that uses potentially all factors and other goods as inputs. Denoting the firm's total output as y_i, its Domar weight (i.e. the share of its sales in total economy-wide consumption) is denoted by $\lambda_i \equiv p_i y_i$. We let p_j^I denote the price of any input (which could be either a good selling for p_j or a factor selling for w_f), such that when firm i purchases

[2] We map multi-product firms (including firms that sell a single product to multiple buyers at different prices) into single-product firms by assigning inputs to each product proportionally. We make an exception for retailers given their central role in mediating trade between consumers and firms and the nature of their production function. We break multi-product retailers into single-product retailers and assign any input that is later sold as output to the relevant single-product retailer. All remaining inputs are then shared proportionately across the single-product retailers.

the quantity of inputs x_{ij}, then its after-tax profits are $\pi_i = p_i y_i - \sum_{j \in \mathcal{F},\mathcal{N}} p_j^I x_{ij} - T_i$ when it rebates a total amount T_i to the government in net taxes. Finally, we let $\Omega_{ij} \equiv \frac{p_j^I x_{ij}}{p_i y_i}$ denote the share of input j in total revenues of firm i and let Ω denote the square matrix composed by the elements Ω_{ij}. Similarly, we define the Leontief inverse of this matrix by $\Psi \equiv (I - \Omega)^{-1}$.

Distortions. We allow for a general treatment of distortions. When a firm i buys an input (a good or a factor) at price p_j^I, it may nevertheless be the case that the as-if marginal cost to the buying firm, from buying one more unit of input j, is $\tau_{ij} p_j^I$ rather than the price p_j^I that the input supplier receives. When this is the case, the seller's and buyer's marginal incentives are not aligned via linear prices, which results in what we refer to as an input distortion (or wedge). The source of such distortions is potentially manifold. One example is monopsony power, in which the buyer recognizes that purchasing an extra unit costs not just p_j^I but also an additional amount (i.e. an extra $p_j^I(\tau_{ij} - 1)$) that comes from increasing the price of infra-marginal units.[3] Another example is simply an input tax collected by the firm, such as a payroll tax, or a value-added tax. Regardless of their underlying source, the combined effect of input distortions determines the value of τ_{ij} and this value is sufficient for what follows.[4] Given such distortions, we let $C_i \equiv \sum_{j \in \mathcal{F},\mathcal{N}} x_{ij} \tau_{ij} p_j^I$ denote the firm's total costs, inclusive of the distortions, and $c_i \equiv C_i / y_i$ the corresponding unit (and marginal) cost.

In addition to these arbitrary distortions τ_{ij} on each input, we also allow for scenarios in which the price p_i that the customers of firm i pay is different from firm i's marginal cost c_i, where any such output wedge is denoted by $\mu_i \equiv p_i / c_i$. This allows for firm i to enjoy market power in its output market and hence charge a markup, for example, or for it to collect sales taxes (rebated to the government in T_i).[5] We then define $\widetilde{\Omega}_{ij} \equiv \tau_{ij} \mu_i \Omega_{ij}$ as the share of the distortion-inclusive cost of input j in firm i's total costs, with the matrices $\widetilde{\Omega}$ and $\widetilde{\Psi} \equiv (I - \widetilde{\Omega})^{-1}$ defined analogously.[6]

Government. Finally, we set $\sum_c T_c = \sum_i T_i$ to close the government's budget constraint.

[3] Similarly, if firm i faces a binding credit constraint, the price it pays the lender of a marginal unit of capital (i.e. p_j^I) is lower than the shadow price of capital within firm i (i.e. $\tau_{ij} p_j^I$ for some $\tau_{ij} > 1$).

[4] Among multiple sources of τ_{ij}, one important distinction is between those due to taxes and those not. The revenues from the former accrue to the government and are included in T_i whereas those from the latter accrue to the firm and are included in π_i. We therefore keep these distinct when we turn to the data.

[5] Recalling that our treatment of single-product firms implicitly encompasses multi-product firms as well, this setup allows for a firm that charges different wedges on each customer.

[6] The expressions below hinge on the difference between $\widetilde{\Omega}_{ij}$ and Ω_{ij}, and hence only on the combined wedge $\tau_{ij}\mu_i$ rather than either of the components τ_{ij} or μ_i alone. We nevertheless keep the distinction between input and output wedges fully general for expositional purposes.

2.2 The Incidence of Distortions

Our strategy for quantifying the general equilibrium incidence of distortions proceeds by solving for the changes in individuals' real incomes that would occur if distortions were to change. For example, one such set of changes we consider in Section 5 below considers the complete removal of existing wedges.

To solve for the effects of such hypotheticals, begin by noting that the change in real income $\mathcal{Y}_c$ for consumer c is composed of two terms: the change in their nominal income χ_c and the change in the price index that is appropriate for their particular utility function. Using the envelope theorem to simplify the latter effect, the change in real income due to any vector of small price changes $d\ln p$ is

$$d\ln \mathcal{Y}_c = d\ln \chi_c - \sum_{i\in N} b_{ci} d\ln p_i. \tag{1}$$

In turn, the price changes can themselves be written, using the envelope theorem applied to each firm's costs, as

$$d\ln p_i = \sum_{j\in\mathcal{N}} \widetilde{\Psi}_{ij} d\ln \mu_j + \sum_{k\in\mathcal{N},\mathcal{F}} \widetilde{\Psi}_{ij}\widetilde{\Omega}_{jk} d\ln \tau_{jk} + \sum_{f\in\mathcal{F}} \widetilde{\Psi}_{if} d\ln w_f. \tag{2}$$

For example, if the markup on good j changes, this effect will propagate forward along the supply-chain to the price of good i in accordance with the Leontief inverse weight $\widetilde{\Psi}_{ij}$—this appropriately sums all senses in which the cost of good i depends, both directly and indirectly, on the price of input j.

In addition, the change in distortions under consideration will affect the prices w_f of each factor f, and these factor price changes will affect both good and factor prices (due to the last term of Equation (2), weighted by the Leontief-inverse exposure elements $\widetilde{\Psi}_{if}$) as well as the incomes $d\ln \chi_c$ in Equation (1). These changes in income themselves satisfy

$$d\chi_c = \sum_{f\in\mathcal{F}} \Phi_{cf} L_f dw_f + \sum_{i\in\mathcal{N}} \Phi_{ci} d\pi_i + dT_c, \tag{3}$$

where dT_c allows for any change in net government transfers that is desired to be simulated as part of the counterfactual. Similarly, the change in after-tax profits $d\pi_i$ satisfies

$$d\pi_i = \left(\frac{\pi_i - T_i}{\lambda_i}\right) d\lambda_i + \lambda_i \sum_{j\in\mathcal{F},\mathcal{N}} \Omega_{ij}(d\ln \mu_i + d\ln \tau_{ij} - d\ln \widetilde{\Omega}_{ij}) - dT_i, \tag{4}$$

where dT_i is the change in net taxes rebated by the firm.[7]

[7]When solving for the effects of a change in wedges below, we keep track of the component that arises

Finally, the change in firm i's Domar weight is given by

$$d\lambda_i = -\sum_{l,m\in\mathcal{F},\mathcal{N}} \lambda_l \Omega_{lm}(d\ln\mu_l + d\ln\tau_{lm})\Psi_{mi} + \sum_{k\in\mathcal{F},\mathcal{N}} \mu_k^{-1}\lambda_k \text{Cov}_{\widetilde{\Omega}^{(k)}}(d\ln\widetilde{\Omega}^{(k)}, diag(\tau^{(k)})^{-1}\Psi_{(i)}) + \sum_{c\in\mathcal{C}} d\chi_c \sum_{k\in\mathcal{N}} b_{ck}\Psi_{ki}, \tag{5}$$

where for any matrix A, the notation $A^{(k)}$ denotes the vector formed from row k (and analogously, $A_{(k)}$ that for column k), the weighted covariance operator $Cov_a(b,c)$ denotes the covariance of vectors b and c weighted by the vector a, and $diag(\tau^{(k)})$ is the diagonal matrix with the ith diagonal element equal to τ_{ki}. Because one can always think of a factor as a firm that uses no inputs, changes in factor prices satisfy an analogous expression to (5) but for dw_f instead of $d\lambda_i$.

The intuition behind Equation (5) is as follows. The first term captures the direct effect of wedges on demand for i: for example, an increase in the input wedge τ_{lm} will be a negative demand shock for supplier m, and this demand shock will propagate backward to firm i via Ψ_{mi}. The second term captures substitution within each buyer k: a relative change in input prices might cause this buyer to substitute input shares away from relatively costlier inputs and towards others. The consequences of these relative demand shocks for inputs for the sales of firm i depends on the direct plus indirect importance of these inputs in i's supply chain (i.e. the entries of $\Psi_{(i)}$) and the size of buyer k (i.e. λ_k). And the final term incorporates the fact that an income shock to consumer c can lead to more total spending on firm i even if the expenditure share spent on that good (the change in which is captured in the second term) is constant.

The final step is to solve for the change in input shares $d\ln\widetilde{\Omega}_{ij}$ that appear in Equation (5) above. Unlike all previous expressions that rely only on optimizing behavior (via the application of the envelope theorem), this final step requires an understanding of how buyers substitute across suppliers in response to changes in the prices that they pay. As in Baqaee and Farhi (2019), a unified treatment of such substitution is as follows. First, let any buyer i have a single CES utility/production function with elasticity θ_i among all inputs.[8] Then the change in input cost shares for such a buying entity i is given by

$$d\ln\widetilde{\Omega}_{ij} = (1-\theta_i)\left(d\ln p_j^I - \sum_{k\in\mathcal{F},\mathcal{N}} \widetilde{\Omega}_{ik} d\ln p_k^I\right). \tag{6}$$

Put together, the system of equations in (2)-(6) constitutes a linear system that can be

from any change in tax policy (which, as discussed above, matters only for the determination of dT_i).

[8]To allow for asymmetric substitution (as in the case of nested CES preferences or technologies) we create an additional fictitious good and a buyer that plays the role of a CES bundle in any lower nest.

solved for any given values of elasticities θ_i, wedges μ_i and τ_{ij}, values of the input share $\widetilde{\Omega}$ and Leontief-inverse Ψ matrices, and values of the consumption shares b_{ci}. A solution to this system can then be used as a first-order approximation to the question of interest (valid to the extent that the exogenous changes in wedges fed into the system is small), or used in each step (updating the allocation each time) of a simple iterative algorithm that solves for arbitrary changes exactly.

Stepping back, these expressions make clear that the incidence of distortions in any economy will hinge on two considerations. The first is the analog of what tax analysts call "statutory" incidence: who actually pays the tax. Here, the broader question is: who pays the wedge and to whom do the rents from that wedge accrue? Combining Equations (1) and (2) demonstrates, for example, the simple sense in which consumers c with high values of $\sum_i b_{ci}\widetilde{\Psi}_{ij}$ are those who are, in the "statutory" sense, paying the wedge μ_j. For example, those consumers who buy from supermarkets with large markups or from supermarkets that themselves source from high-markup producers. The flip-side of these statements is clear from Equations (3) and (4): the owners (as embodied in the ownership matrix Φ_{ci}) of high-markup firms are earning income from large markups. Similarly, marked-down wages are being paid for by workers (in Φ_{cf}) and the owners of the firms marking down wages are receiving these payments.

However, as tax analysts recognize well, statutory incidence is not the end of the story for the "economic" incidence that ultimately matters. Indeed, in the simplest possible case of a competitive model of a single market, for example, the elasticities of supply and demand in that market uniquely determine the effect of a tax on consumer and producer prices, and hence the "economic" incidence among the two agents within that one market, and statutory incidence is irrelevant.

While this same force is at work above—in Equations (2) and (3) via the dependence on factor price responses $d\ln w_f$—the broader sense of "statutory" incidence of distortions does matter when we look across markets. The actors in any realistic economy participate in the bilateral exchange of each good and factor to vastly heterogeneous extents. That is, who buys and sells products that are directly or indirectly marked up, who works for and owns firms that buy inputs that are directly or indirectly marked down, etc., will vary enormously. The formulae above combine heterogeneous extents of statutory incidence—as measured in the spending patterns of consumers, the supply chain patterns of firms, and the ownership matrices of which consumers own which firms and factors—with general equilibrium elasticities of supply and demand in order to measure economic incidence correctly. The next two sections illustrate how we will measure these phenomena in order to put the theory of this section into practice.

3 Data

Our analysis draws on eight administrative datasets from Chile's Servicio de Impuestos Internos, their Internal Revenue Service equivalent (henceforth, IRS), and two survey datasets from the statistical agency (INE for its acronym in Spanish). These datasets cover the entire formal private sector in Chile. Below, we provide an overview of the main data sources and key variables, as well as how we deal with informality. We use data from 2022 in all cases, the most recent year where all sources are jointly available.[9]

First, we use a firm-to-firm electronic receipt dataset that is based on value-added tax (VAT) records. This dataset electronically records all transactions between formal firms in the economy. Thus, for each firm, we know the complete list of buyers and suppliers that the firm trades with (including those in the public sector). There are no reporting thresholds involved. The dataset includes the value of the transaction, the price and the product involved. Products are codified using around 2,500 standardized product categories via a string-matching procedure. Second, we use a firm-to-firm dataset similar to the first one but for international transactions that reports all imports and exports between domestic firms and foreign firms originating from customs records.

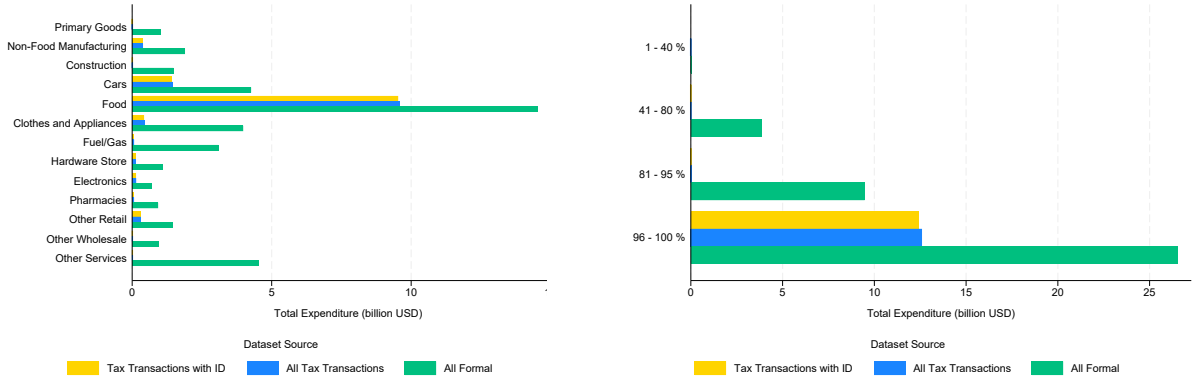
Third, we use a similar dataset to the two above but covering transactions of individual products between firms and households based on electronic receipts collected as part of the VAT system. This dataset electronically records all transactions between formal firms and individuals in the economy. Administrative data mapping firms to consumers are very uncommon and come about through the government's electronic filing system, which requests that every purchase be associated with a customer tax ID. Consumers routinely comply with this request, and do so almost always when shopping at larger retailers in part because these stores use this tax ID to link customers to their loyalty or rewards programs.[10] Thus, for each firm, we have a list of individuals the firm sells to and the products that they purchased. There are no reporting thresholds although coverage is incomplete, both for informal retailers and smaller stores that typically do not report tax IDs

[9]This paper was developed via an agreement within the scope of the research agenda conducted by the Central Bank of Chile (BCCh for its Spanish acronym) in economic and financial affairs in its purview. The BCCh has access to anonymized information from various public and private entities, by collaboration agreements signed with these institutions. To secure the privacy of workers and firms, the BCCh mandates that the development, extraction, and publication of the results should not allow the identification, directly or indirectly, of natural or legal persons. Officials of the BCCh processed the disaggregated data on our behalf. We implemented all analysis and neither involved nor compromised the IRS in doing so. The information contained in the databases of the IRS is of a tax nature originating in self-declarations of taxpayers presented to the Service; therefore, the veracity of the data is not the responsibility of the Service.

[10]This database provides customer-level itemized sales information for 2,124 of the country's largest firms in terms of their total final consumption sales.

in their submissions (an issue we address by apportioning this consumption using consumption surveys described below). Beyond having customer tax IDs, the variables are the same as those in the firm-to-firm dataset (value, price, detailed product classifications). For the remaining formal firms, only total sales to all customers are available.

Figure 1: Expenditures Across Retailer Type and Firm Size By Data Source



(a) Across Retailer Types

(b) Across Firm Size

Notes: Panel (a) displays total consumer expenditures across retailer types broken down by the source and level of detail in the consumption data: all formal sales to final consumers, formal sales recorded at the transaction level, and formal sales both recorded at the transaction level and attached to consumer tax IDs. Panel (b) reports the same for groups based on firm size.

Figure 1 displays the composition of expenditures captured by three different types of administrative consumption records described above. The green bars plot all formal sales data to final consumers, the blue bars plot all expenditures for which we have transaction-level data, and the yellow bars plot all expenditures for which we have both transaction-level data and tax IDs. Panel (a) shows the allocations of these three types of data across different types of retailer, for example food retailers, construction, or service firms, while Panel (b) shows the allocations across different firm size bins. The fact that the blue and yellow bars are almost identical shows that, conditional on shopping at a store for which we have transaction-level records, almost all customers provide their tax IDs. In terms of coverage by retailer type and firm size, stores reporting transaction-level data are under-represented in services and over represented in food retail and cars, and are much larger than the average store—consistent with the tax authorities sharing transaction-level data only for the largest retailers, primarily supermarkets, department stores and chains.[11]

Fourth, we use a matched employer-employee dataset (from IRS tax affidavits 1887

[11] Anecdotally, these larger stores are more likely to record tax IDs than smaller ones, so were the tax authorities to share transaction-level data for smaller stores there would likely be a smaller share of expenditures attached to tax IDs.

and 1879) that reports annual earnings from each job that a worker has. Earnings include wages, salaries, bonuses, tips, and other sources of labor income deemed taxable by the IRS. As earnings are reported net of social security payments, we adjust the earnings measure to include these payments. Fifth, we use the ownership linkages of firms from tax records (IRS tax affidavits 4415 and 4416). This dataset includes, for each firm, the complete list of owners of the firm (which can themselves be both firms and individuals) as well as the share of the firm that each owner possesses. Sixth, we use government-to-household data that combines a dataset of transfers and another on income tax payments that allows us to build direct net transfers. The transfers data records the main direct transfers that the government makes to households every year. For this dataset, we know, for each type of transfer, the total amount of the transfer and the policy the transfer corresponds to. The income tax dataset, on the other hand, records the income tax payments from households to the government (which apply only to the top income deciles of the country).

The six aforementioned datasets record all the relevant transactions and relationships that firms and individuals have with different agents in the economy: other firms (both domestic and foreign), households, government, workers, banks, and capital owners. Furthermore, all the aforementioned datasets that record transactions include a measure of the relevant price involved in the transaction (with the exception of government transfers and ownership links, where such information is not needed for our analysis).

Two other administrative datasets serve as complements to these bilateral administrative datasets. We use a civil registry database to provide year of birth, gender, marriages, place of birth and the father and mother of each individual. These data provide both demographic information useful for grouping individuals into groups when exploring differences in incidence and allow us to combine individuals into households, which is crucial since many consumption and income choices are made in part at the household level (for example, supermarket purchases that are recorded at the individual level). Eighth and finally, we use an administrative dataset (IRS tax forms 29 and 22) that contains each firm's balance sheets to measure total sales, material costs, investments, and fixed assets for each firm.

All individuals and all formal firms in Chile are assigned a unique tax ID that is consistently recorded across the datasets above, which enables all of the merges we require. In what follows, we define a firm as a tax ID.[12] Given the centrality of these bilateral administrative datasets to our analysis, Table 1 presents statistics about the scale and other attributes of these datasets.

[12] As all tax forms are reported at the headquarters level, plant-level information is not available.

Table 1: Descriptive statistics on the scale of the bilateral administrative datasets

1. Firm-to-Firm Domestic Trade	# Buyers 1,354,408	# Suppliers 624,073	# Pairs 35,993,564	# Transactions 2.1 Billion
2. Firm-to-Firm International Trade	# Buyers 93,423	# Suppliers 155,283	# Pairs 273,110	# Transactions 5,298,769
3. Firm-to-Individual Consumption	# Consumers 13,453,311	# Suppliers 2,124	# Pairs 43,626,887	# Transactions 6 Billion
4. Firm-to-Workers Employment/Wages	# Firms 702,729	# Workers 8,242,191	# Pairs 13,138,247	# Jobs per Worker 1.6
5. Firm-to-Individual Ownership	# Owners 1,781,539	# Owned 1,445,504	# Pairs 3,172,853	Median Ownership Share 34%
6. Gov-to-Individuals Net Transfers	# Individuals 8,021,862	# Policies 10	# Policy Pairs 16,262,917	# Policy Transactions 16,495,680

Notes: This table presents statistics on the size of each transactional dataset described in Section 3. Each statistic is conditional on non-zero flows, which leads to differences in the number of firms and individuals across datasets. For example, not all firms engage in international trade.

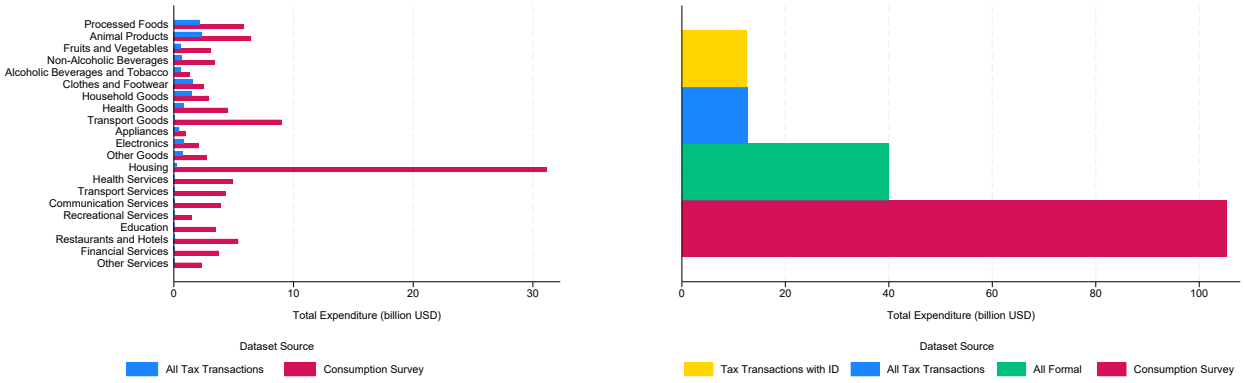
While the administrative datasets cover close to the universe of formal economic transactions in the economy, they miss informal economic activity that is an important feature of Chile's economy, as it is of many other countries. Thus, in order to complement the administrative data, we use three large-scale government surveys that capture informal transactions. The first of these is a detailed consumption survey conducted by the Chilean government as a key input into inflation and poverty measurement (the IX Encuesta de Presupuestos Familiares fielded between October 2021 and September 2022 and surveying around 44,000 individuals). A lengthy questionnaire is administered to a large and representative sample of households. Households report the full list of consumption expenses they had in a year, including the associated prices, quantities, product description, and the store it was purchased from.[13] Surveyors both ask to see receipts of all purchases to increase accuracy and separately interview every member of the family to ensure no consumption is missed.

Panel (a) of Figure 2 compares product-group by product-group the coverage in the consumption survey (applying the survey weights) to the consumption captured at the transaction-level in the firm-to-individual administrative data that covers expenditure at the largest retailers. As would be expected, the consumption survey total expenditures are much larger but the gap also differs across sectors, with the largest shortfalls in the administrative data found in services, particularly housing where the consumption survey includes imputed rent paid by owner-occupiers. Panel (b) displays the aggregates across all four types of data we use to measure consumption; all formal sales to final consumers,

[13] Recall periods vary depending on the durability of products, with shorter recall periods for less durable products that we convert to annual expenditures.

formal sales recorded at the transaction level, formal sales both recorded at the transaction level and attached to consumer tax IDs, and all sales in the consumption survey (expanded via survey weights).

Figure 2: Consumption Data Coverage: Across Products and Aggregates



(a) Across Products (b) Aggregates

Notes: Panel (a) displays total expenditure in consumption across product categories. The blue bars represent consumption documented in administrative tax data. The red bars represent consumption documented in the consumption survey (expanded via survey weights). Panel (b) presents aggregates across all four sets of sales data; all formal sales to final consumers, formal sales recorded at the transaction level, formal sales both recorded at the transaction level and attached to consumer tax IDs, and all sales in the consumption survey (expanded via survey weights).

The consumption survey also includes detailed demographics (including income data) that allow us to match surveyed individuals to similar groups of individuals in the administrative data for whom informal spending is missing, and for whom formal spending at small stores may not be attached to their tax IDs. Specifically, we follow Blanchet, Saez, and Zucman (2022) and draw on statistical matching methods that provide the matches that minimize the distances between common variables in two datasets via recent advances in optimal transport algorithms. This approach preserves the multivariate distribution of the consumption survey data that is being brought into the administrative data and avoids extrapolation that would occur using predictive models that may generate extreme outliers. We match not only on detailed demographics (e.g. age, gender, marriage status, location) but also on the full vector of expenditures at supermarkets and department stores (expenditures that are well-measured in both the firm-to-individual administrative data and in the consumption surveys). This match provides measures of expenditure on each individual-product-firm type triplet, with expenditures not covered by the firm-to-individual administrative data allocated proportionately to the corresponding stores of that type, e.g. a specialized shoe retailer, in the location that the individual

resides in.[14]

The income module of the household survey also contains multiple questions capturing income from government programs as well as from both informal and formal employment. We complement these data with a rich labor force survey that reports formal and informal labor market activity for a sample of representative workers in the economy (the 2020–2022 Encuesta Suplementaria de Ingresos surveying approximately 300,000 individuals). Again this data is matched to demographics that allow us to fill in missing data in the administrative data using a similar statistical matching process to how we deal with consumption above (in this case matching on demographics and the vector of formal income by sector). Finally, we bring in a survey of small firms designed by the Chilean government to measure the informal sector (the 2022 Encuesta de Microemprendimiento 7 surveying 8,500 small business owners). Using the included survey weights, we use this survey to populate the informal firms in the economy. These informal firms are then statistically matched to the informal firm owners and workers reported in the labor force surveys prior to performing the match between the labor force survey and the administrative data described above.

Taken together, we believe these datasets may provide the most complete mapping of interactions between economic agents in an economy that have been assembled to date. As we have emphasized above, this level of detail is necessary for understanding incidence in the presence of overlapping distortions that may mitigate or exacerbate each other.

4 Measurement

Our implementation of the incidence analysis presented in Section 2 requires three empirical inputs: (i) matrices of household-level exposure to distortions; (ii) elasticities of substitution within firms' technologies and consumers' preferences; and (iii) the size of the various distortions themselves. We discuss each of these in turn.

4.1 Exposure Matrices

To implement our incidence formulae from Section 2, we require knowledge of three fundamental matrices that together capture household exposure to distortions. Using the definitions from Section 2, these are the matrices of households' consumption shares b,

[14] We also impose that total expenditure at each formal retailer equals their final consumption sales in the administrative records that cover all formal firms but are neither matched to individual tax IDs not broken down at the item-transaction level. The residual between the consumption survey predicted expenditures and the total formal sales is informal expenditures.

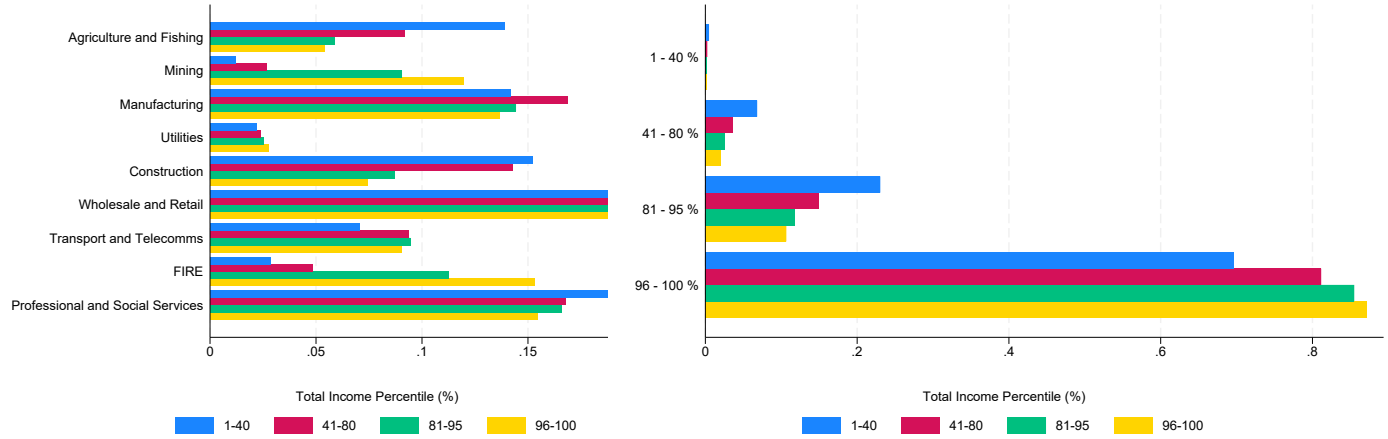
input-output revenue shares Ω, and household-to-firm and -factor ownership shares Φ. Household consumption shares for each household i (b_i) are measured directly from the firm-to-household electronic VAT receipt data combined with the consumption survey for non-captured shopping trips and are at the store-product level. Input-output revenue shares between a buyer i and a supplier j (Ω_{ij}) come directly from the spending by buyers on particular suppliers, relative to total sales of buyers, in the firm-to-firm electronic transaction data and employee-employer data. Recall that we break retailers and wholesalers—multiproduct firms that sells inputs with minimal transformation—into multiple fictitious single-product firms, one for each final product in the original retailer/wholesaler, each of which uses as its input the corresponding product found in the original firm's input purchases. Any retailer/wholesaler inputs not sold as output are attributed to overhead costs that are shared across the single-product firms. For all other multi-product firms, we allocate all inputs proportionately to revenues. Firms' total sales—used also to measure the Domar weights λ_i—are recorded in firms' tax filings.

Finally, household-to-firm and -factor ownership shares between a household c and a firm i (Φ_{ci}) or factor f (Φ_{cf}) come directly from the firm ownership (in case of capital and profits) and employer-employee records (in the case of labor). Thus, except for the non-captured shopping trips that we impute using detailed store-product-level consumption surveys matched to household characteristics and informal labor matched from the same surveys, all the elements in the key exposure matrices are directly observed in Chile's expansive administrative datasets.

One example that illustrates the richness of the exposure matrices we observe is presented in Figure 3. In panel (a) we display elements of the matrix Φ_{cf}, where the set of individuals c is broken into four income categories and the set of factors f is taken to be sector-specific labor for each of nine sectors (the merge of individuals into households is still forthcoming). Large distinctions across income categories are apparent. For example, the lowest-income individuals are almost three times as dependent on the agricultural and fishing sector for their (labor) income as the richest individuals are, whereas the richest individuals are about ten times more dependent on the mining sector and six times more dependent on finance, real estate and insurance than the poorest are. Panel (b) presents an analogous depiction of Φ_{cf}, but now where the factor groups f are based on the size (four bins) of the firm size distribution at which individuals work. The poor are relatively more likely to work in smaller firms than the rich are. Both of these figures highlight how different the exposure to labor distortions in various sectors and firm sizes is depending on where in the income distribution an individual lies.

Other types of exposure matrices are also easy to visualize (in aggregated forms). For

Figure 3: Exposure of Individuals to Labor Income



(a) By sector

(b) By firm size

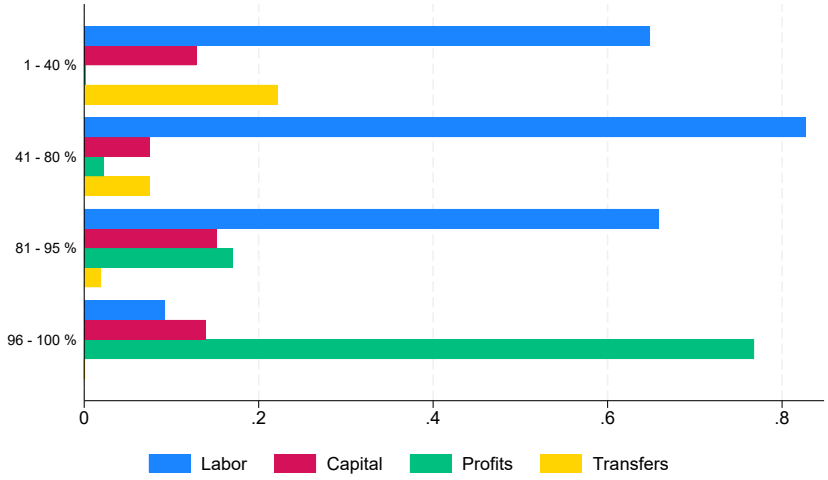
Notes: Panel (a) displays elements of the matrix Φ_{cf} for four types of individual c and nine types of sector-specific labor (i.e. f). By construction, $\sum_f \Phi_{cf} = 1$ for each c. Panel (b) does the same but for groups of f based on the size of the firm at which individuals are employed.

example, a natural next question is how total income divides into labor relative to capital and the role of government transfers. Figure 4 displays such a breakdown. In so doing, we disaggregate capital earnings into two components: (i) an amount that corresponds to the "fair" return on each firm's capital (based on Chile's interbank lending rate); and (ii) the additional profits that each firm earns, beyond labor and intermediate costs and the aforementioned fair return on their capital. Evidently, labor and transfer income comprise virtually all of the earnings for individuals in the bottom 80 percentiles of the income distribution. By contrast, for individuals in the top five percentiles of the income distribution we see that the vast majority of their income derives from firm ownership, and of that the bulk is profits rather than the fair return on capital.[15]

As with labor income above, Figure 5 displays features of the ownership matrix Φ_{ci}, again for individuals c based on their position in the income distribution. In panel (a) the firms i are broken down by sector and in (b) they are broken down by firm size bins. As shown above, while rich individuals are of course more likely to own firms, the entries of Φ_{ci} sum to one within any c. Thus, these figures portray how a given type of income group's profit income is distributed across types of firms (sectors in (a), and sizes in (b)). Again it is apparent that firm ownership type is heterogeneous across the income distri-

[15]Note that this result is in part driven by the fact that we are distributing all profits to income. In reality, a fraction of profits remain inside the firm with the value of the ownership stake rising (and so it would only be recorded as income to the IRS when the firm distributes it as dividends or is sold).

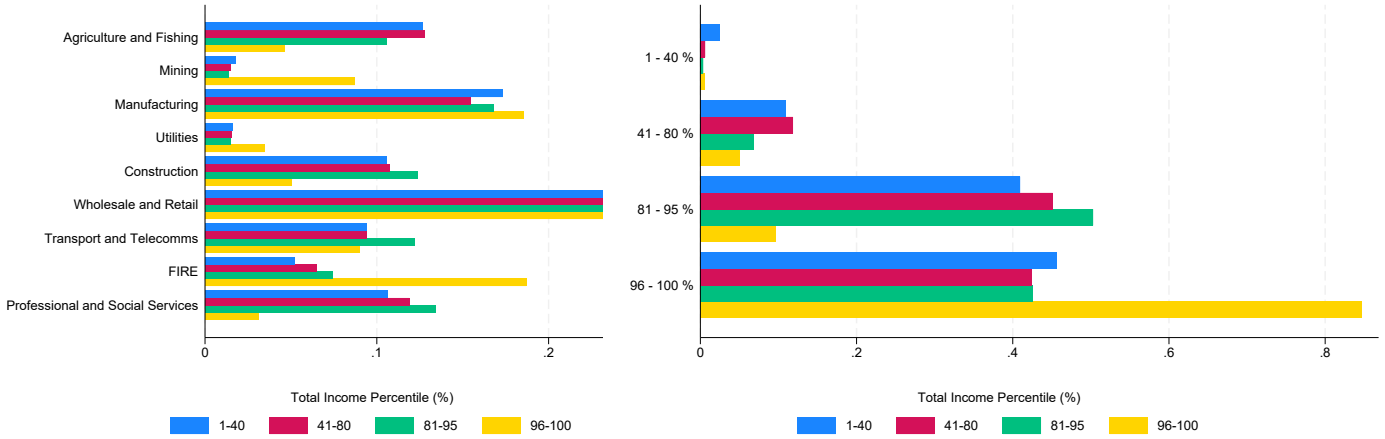
Figure 4: Income by Source



Notes: This figure displays the breakdown, for any type of individual (based on their position in the income distribution) c, of their earnings by source. Four types of source are indicated: that from all types of labor, that from all types of capital (based on a fair return), that from the profits that are earned in all types of firms, and that from government transfers.

bution; for example, poorer individuals are relatively more likely to own agriculture and fishing firms, and small- and medium-sized firms more generally.

Figure 5: Exposure of Individuals to Profit Income



(a) By sector

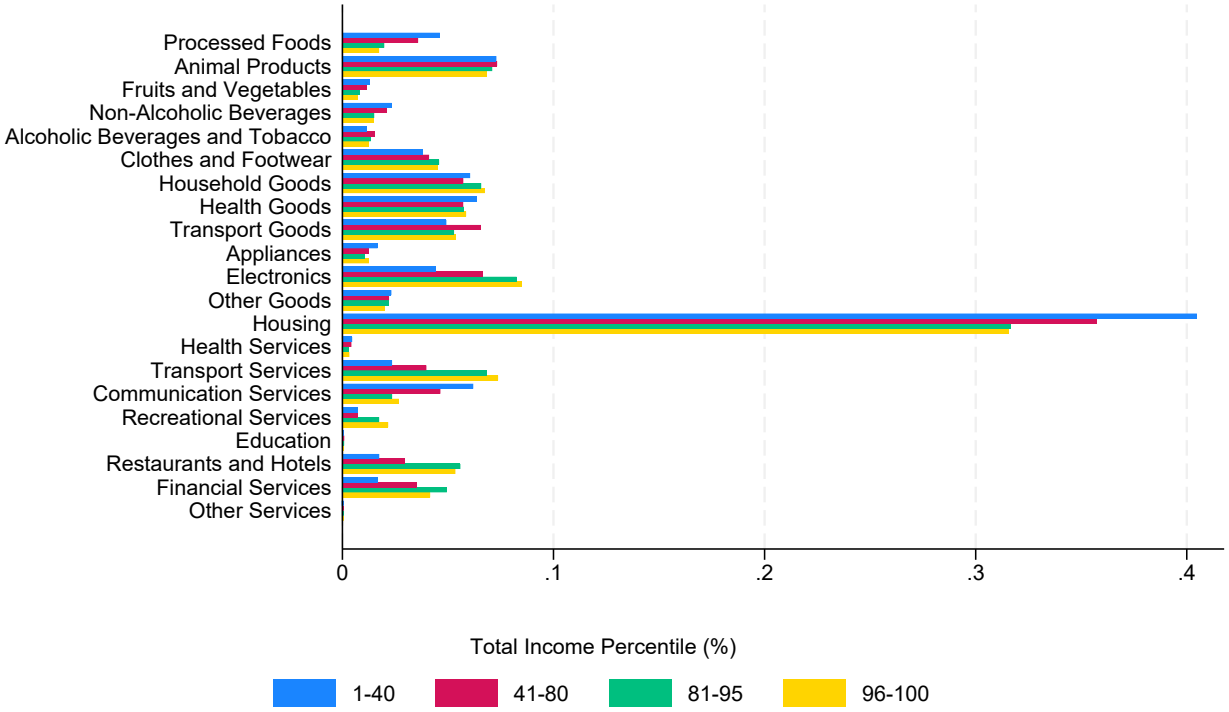
(b) By firm size

Notes: Panel (a) displays elements of the matrix Φ_{ci} for four types of individuals c and nine types of firms (i.e. i) owned, by sector. By construction, $\sum_i \Phi_{ci} = 1$ for each c. Panel (b) does the same but for groups of i based on the size of the firm owned.

Similarly, Figure 6 turns to the case of the matrix of consumer expenditure shares,

b_{ci}. Using the information encoded in the linked firm-to-individuals transaction records as well as the statistically-matched expenditures from the consumption survey, we again summarize these expenditure shares by individuals' income groups. We break down each groups' expenditure shares by product category, revealing substantial cross-group heterogeneity in product consumption patterns across broad products. Of course, such summary measures hide the cross-group heterogeneity that our data uncovers within these broad categories.

Figure 6: Consumer Expenditure Shares From Firm-to-Individual Transactions and Consumption Surveys



Notes: Figure displays elements of the matrix b_{ci}, as revealed in the transaction-level individual-firm matched consumption data as well as the statistically-matched expenditures from the consumption survey, for four types of individuals c (based on income groups) and 19 types of firms (i.e. i) owned, by product category. By construction, $\sum_i b_{ci} = 1$ for each c.

Finally, Figure 7 displays elements of the firm-to-firm cost share matrix $\widetilde{\Omega}_{ij}$ (corresponding to a disaggregated version of standard input-output tables). Panel (a) reports aggregates based on narrow (6-digit) categories that are available in the data. Panel (b) shows how this matrix changes through supply chain linkages by displaying the indirect linkages constructed from elements of the Lenotief inverse of this matrix, $\widetilde{\Psi}_{ij} - \widetilde{\Omega}_{ij} - I$.

Figure 7: Firm-to-Firm Expenditure Shares



(a) Direct Linkages: $\widetilde{\Omega}_{ij}$ (b) Indirect Linkages: $\widetilde{\Psi}_{ij} - \tilde{\Omega}_{ij} - I$

Notes: Panel (a) displays elements of the matrix of firm-to-firm input cost shares $\widetilde{\Omega}_{ij}$ for groups of buying and selling firms i and j using 6-digit sectors. Panel (b) does the same but for the indirect linkages through supply chains captured through the Lenotief inverse of this matrix, $\widetilde{\Psi}_{ij} - \tilde{\Omega}_{ij} - I$.

4.2 Elasticities of Substitution

The second requirement of our simulation procedure is the elasticities of substitution θ_i that appear in Equation (6), in both technology and preferences. Technological elasticities also play a role in our estimation of distortions, as discussed further below. In both cases we use functional form choices that have been common in prior work.

Beginning with preferences, we specify the nested demand system from Equation (6) as a two-tier system across products and sectors. In particular, at the lower level, the elasticity θ_s denotes substitution across products within sector s, and the upper-level elasticity across sectors is denoted by θ_U. We obtain estimates of the parameters θ_s from Gervais and Jensen (2019) and the parameter θ_U from Redding and Weinstein (2024), which consider products and sectors using similar levels of aggregation as we do.

On the production side, we assume that firms' production functions take the Cobb-Douglas form across groups of inputs: capital, labor, and material inputs coming from each sector. This amounts to setting $\theta_i = 1$ in Equation (6) for every firm when it concerns substitution across such groups. Then, as with common specifications of Cobb-Douglas technologies, we assume that inputs within these groups (e.g. different versions of capital) are perfect substitutes.

4.3 Distortions

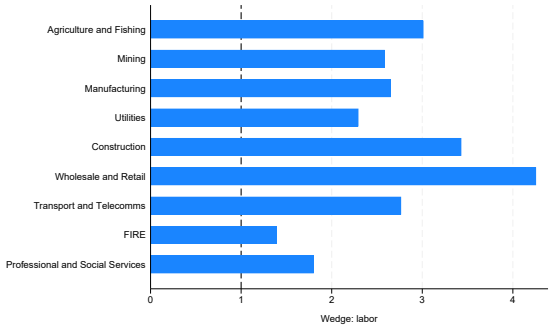
The final input into our analysis is the size of the wedges μ_i and τ_{ij} themselves. As emphasized by De Loecker and Warzynski (2012), for example, this is straightforward once estimates of firms' technologies (as discussed above) are known. A key identity in Hall (1986) expresses a firm's markup (or more generally, its output wedge) μ_i as the ratio of the firm's output elasticity, for any distortion-free input, to the firm's expenditure on that input as a share of the firm's total revenues. In our case, with potential bilateral distortions τ_{ij} on each input j, the analog of this result is

$$\frac{\eta_{ij}}{\Omega_{ij}} = \mu_i \tau_{ij}, \tag{7}$$

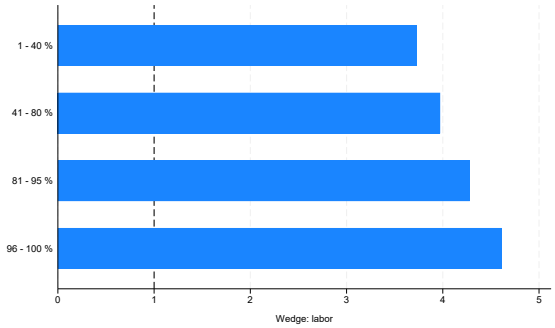
where η_{ij} is the elasticity of firm i's output with respect to input j. Equation (7) highlights how the combination $\mu_i \tau_{ij}$ is identified but the separate components μ_i and τ_{ij} are not. However, this is inconsequential for our study of incidence because, as discussed above, our formulae depend only on $\mu_i \tau_{ij}$.

In the case of Cobb-Douglas technologies, as assumed above, we follow the approach developed in Hsieh and Klenow (2009) for measuring the output elasticities η_{ij}. This method assumes that US firms use (on average within each sector) inputs in a non-misallocated fashion, and that their technologies are Cobb-Douglas over capital, labor, and each sector's type of materials. In this case the output elasticities η_{ij} in Equation (7) are obtained from the average shares (across firms, within each sector) of each type of input in the costs of US firms. The resulting wedge estimates from applying this procedure, $\mu_i \tau_{ij}$ for j equals labor, capital, and materials show considerable dispersion, as found by Hsieh and Klenow (2009) for the cases of China and India. There is also considerable dispersion of average wedges across sectors and across firm size bins, with Figure 8 displaying these averages by input type. Particularly notable are the large wedges on wholesale and retail labor, and the finding that wedges tend to be bigger in larger firms—patterns that will interact with the cross-group heterogeneity in exposure detailed above when calculating incidence.

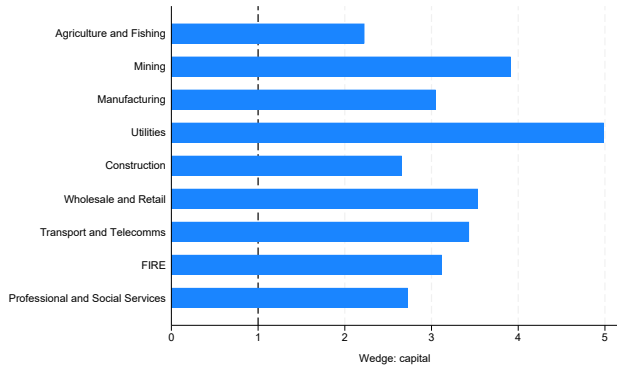
Figure 8: Estimates of Average Wedges By Sector and Firm Size



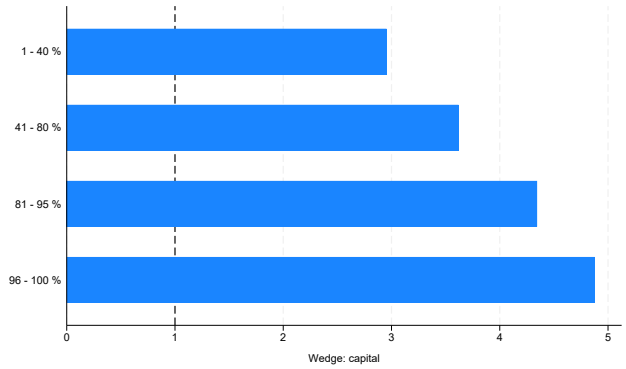
(a) Labor Wedges by Sector



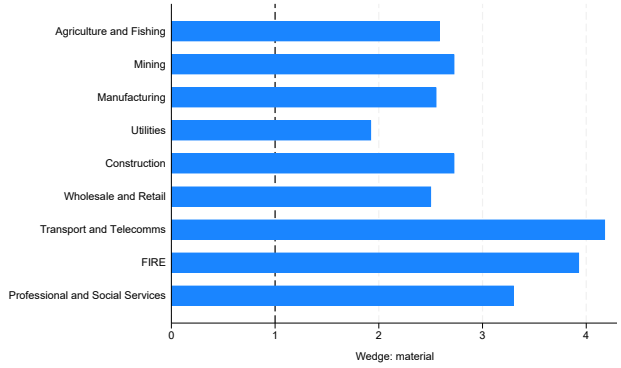
(b) Labor Wedges by Firm Size



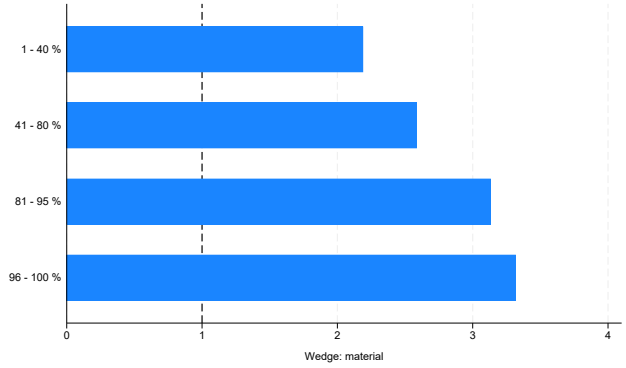
(c) Capital Wedges by Sector



(d) Capital Wedges by Firm Size



(e) Materials Wedges by Sector



(f) Materials Wedges by Firm Size

Notes: Panel (a) displays average values of the estimated wedge $\mu_i \tau_{ij}$, where j is labor, obtained by using the Hsieh and Klenow (2009) method, for groups of firms i arranged by sector. Panel (c) and (e) do the same for the wedge $\mu_i \tau_{ij}$, where j is capital amd materials, respectively. Panel (b), (d) and (f) repeats these same wedges for labor,capital, and materials respectively, but for groups of firms i arranged by firm size rather than sector.

5 Counterfactuals

Having populated the formulae in Section 2 with the data and estimates described in Sections 3 and 4, we can perform a number of counterfactuals designed to address the key question at the heart of this paper: what is the incidence of distortions?

The first set of counterfactuals asks how the burden of distortions is shared across various groups in society. To do so, we can use Equation (1) to calculate the welfare changes experienced by different groups if all distortions in the economy were removed. Specifically, we eliminate all capital, labor and materials-input wedges such that consumers and downstream firms pay only the marginal costs of production, and firms pay labor and capital their value marginal products. We focus on differences in incidence not only across the income distribution, but also by gender and age. This allows us to explore hypotheses such as do the poor bear more of the incidence of distortions, or do certain age groups, or do women? And are these differences quantitatively important?

To better understand where any differences come from, we decompose the total welfare effects for each demographic group in a number of ways. First, it is straightforward to decompose the total effect into that coming from the consumption side, from factor services, and from changes in non-factor services such as rents and transfers. Furthermore, we can break down the consumption-side effects in two additional ways. First, we can decompose the impacts from reducing different combinations of distortion—labor, capital and material wedges, in all three cases multiplied by output wedges. Second, we further decompose each of these into those coming from changing goods and services prices (the first two terms in Equation (2)) and those coming from changing factor prices (the last terms in Equation (2)).

Our second set of counterfactuals asks which particular distortions are most responsible for the unequal burden of distortions. By removing certain wedges or sets of wedges we explore which distortions matter most for welfare and to which groups, both because some wedges are larger in magnitude and because their incidence differs conditional on size. In addition, this set of counterfactuals asks which wedges reinforce each other and which are countervailing, in the sense that interacting distortions may partially correct the harmful effects that each would cause in isolation. For example, distortions a firm faces on the output side may counteract those they face on the input side, or greater heterogeneity in wedges across firms within a sector may mitigate misallocation due to that sector having a relatively low average level of wedges.

Finally, our third set of counterfactuals explore a related question: what are the tradeoffs between equity and efficiency of specific policy changes, and what is the equity-

efficiency frontier that policymakers face? For example, how much would policies that reduce wedges in credit markets or the monopsony power of firms raise aggregate output and would this come at the expense of increases in certain forms of inequality? Are there other policy mixes that achieve the same increase in output but in a more equitable manner?

5.1 The Distributional Impact of Eliminating Wedges

As a first counterfactual calculation, we begin with the exercise described above in which all wedges are removed. To be consistent with the first-order approximation underlying the results of Section 2, we implement a reduction of 1% of all wedges, thereby reducing each wedge in proportion to its size.

Figure 9 displays the results of this exercise. The hypothetical change in wedges in the economy induces considerable improvements in the real income of individuals who start in three of the four income bands that we have introduced above. Members of the poorer groups gain the most (in terms of percentage growth in their real income) with real income declining for the richest group. The magnitudes range from 0.6% in the poorer groups to a loss of 0.4% in the richest.[16] The fact that there are broad gains for most groups is not surprising as the removal of distortions raises efficiency and hence the total size of the pie.

Following Section 2, we break down these results into four sets of channels. First, we examine a partial-equilibrium consumption channel, which ignores factor price changes. We show this channel separately for labor, capital and materials wedges. Figure 9 shows that all income groups gain from this channel. The biggest gains come from shrinking labor wedges, then capital, and finally the smallest gains arise from shrinking material input wedges. Poorer consumers gain the most from the reductions in labor wedges through this channel, in part because labor wedges were particularly large in retail and wholesale and poorer consumers disproportionately source from this sector. Gains from the capital and material channels are fairly similar across income groups.

Second, we consider the general equilibrium forces of changes in factor income of labor and capital.[17] As a response to the reduction in wedges, both factor prices increase. However, gains are unevenly distributed between income groups. The poorest groups gain relatively more from increases in labor income, whereas the richest group gains significantly less. On the other hand, the poorest groups gain almost nothing from capital income changes, whereas the richest group gains a lot. This is to be expected given the

[16] Given the 1% reduction in the wedges, these numbers can be interpreted as elasticities.

[17] Given that factor supply is fixed and that the numeraire is GDP, one can interpret this outcome as changes in factor shares of GDP.

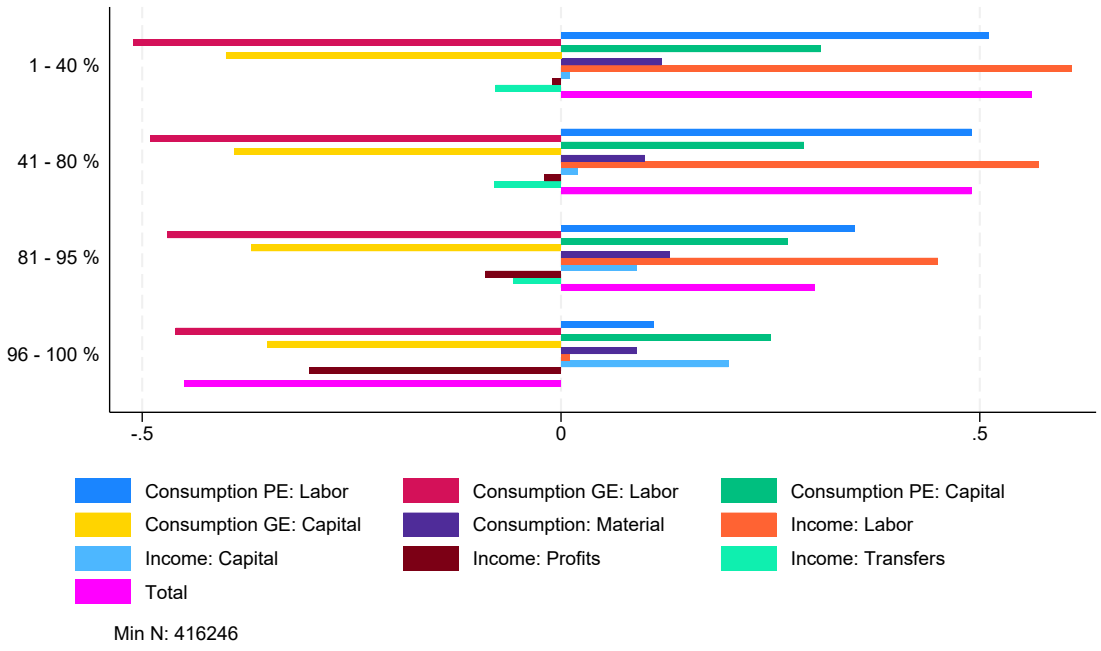


Figure 9: Welfare Impacts of Reducing All Distortions by 1%: Across Income

Notes: This figure displays the impacts on the real income $\mathcal{Y}_c$ of individuals c across income (grouped by their percentile in the initial income distribution) of a counterfactual exercise that reduces all distortions in the economy by 1%. Impacts are broken down by 4 channels. First, the partial-equilibrium consumption exposure to labor, capital and material wedges. Second, the general equilibrium changes in factor income coming from changes in factor prices. Third, the impact that factor price changes have on consumption prices (labelled GE). Finally, the impact on non-factor income, which involves profits and transfers.

evidence in Figure 4, namely that richer households rely relatively more on capital income than poorer households.

The third channel we investigate corresponds to the role that general equilibrium forces due to factor price changes have on final consumption prices. Figure 9 also show that individuals lose from this channel. This is natural given that we showed increases in factor income, which imply that the cost of production goes up. These losses are relatively large but fairly evenly distributed across the income distribution.[18]

Finally, we turn to the channel of non-factor income—that is, profits and (net) transfers. We find that these are unevenly distributed across the income distribution. Given that wedges go down, tax revenues also go down and thus transfers go down because the government keeps a balanced budget. We distribute these reductions in transfers proportional to the initial distribution of transfers across the income distribution. Thus, the poorest households lose the most out of this whereas the richest households are close to

[18]We note that such conclusions may change once labor factors are defined by skill group and location rather than just location as in the current draft.

unaffected. On the other hand, the reduction in wedges reduces overall profits. These losses are largest for the richest group of individuals, since they rely the most on profit income, and are very sizable. As the transfer loss to the poor is smaller than the profit loss to the rich, non-factor income channels make the impact of distortions more regressive and lie behind the absolute losses we find for the rich.

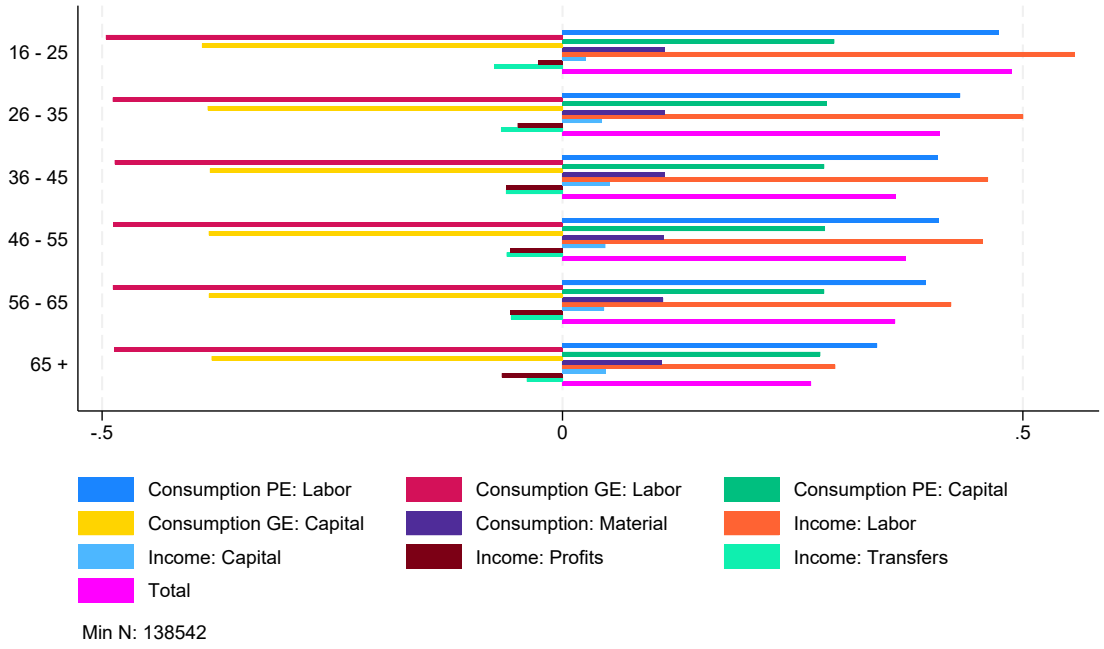


Figure 10: Welfare Impacts of Reducing All Distortions by 1%: Across Age

Notes: This figure displays the impacts on the real income $\mathcal{Y}_c$ of individuals c across age groups of a counterfactual exercise that reduces all distortions in the economy by 1%. Impacts are broken down by four channels. First, the partial-equilibrium consumption exposure to labor, capital and material wedges. Second, the general equilibrium changes in factor income coming from changes in factor prices. Third, the impact that factor price changes have on consumption prices (labeled GE). Finally, the impact on non-factor income, which involves profits and transfers.

Having examined incidence through the lens of heterogeneity across individuals grouped by income, we turn now to two additional differences. Figure 10 explores incidence by age, and Figure 11 by gender. The gains from reducing wedges are positive for all age groups, though older individuals gain relatively less from reduction in wedges. This result comes from the fact that older individuals lose more from profit reductions and gain less from labor income increases. These individuals also gain more from capital income increases, but this is not enough to counteract the relatively lower gains on the two other channels. Turning to gender differences, Figure 11 shows that women gain relatively more than men from reducing distortions. This difference comes mostly from differences in changes in labor income, which women are more exposed to than men are. This is con-

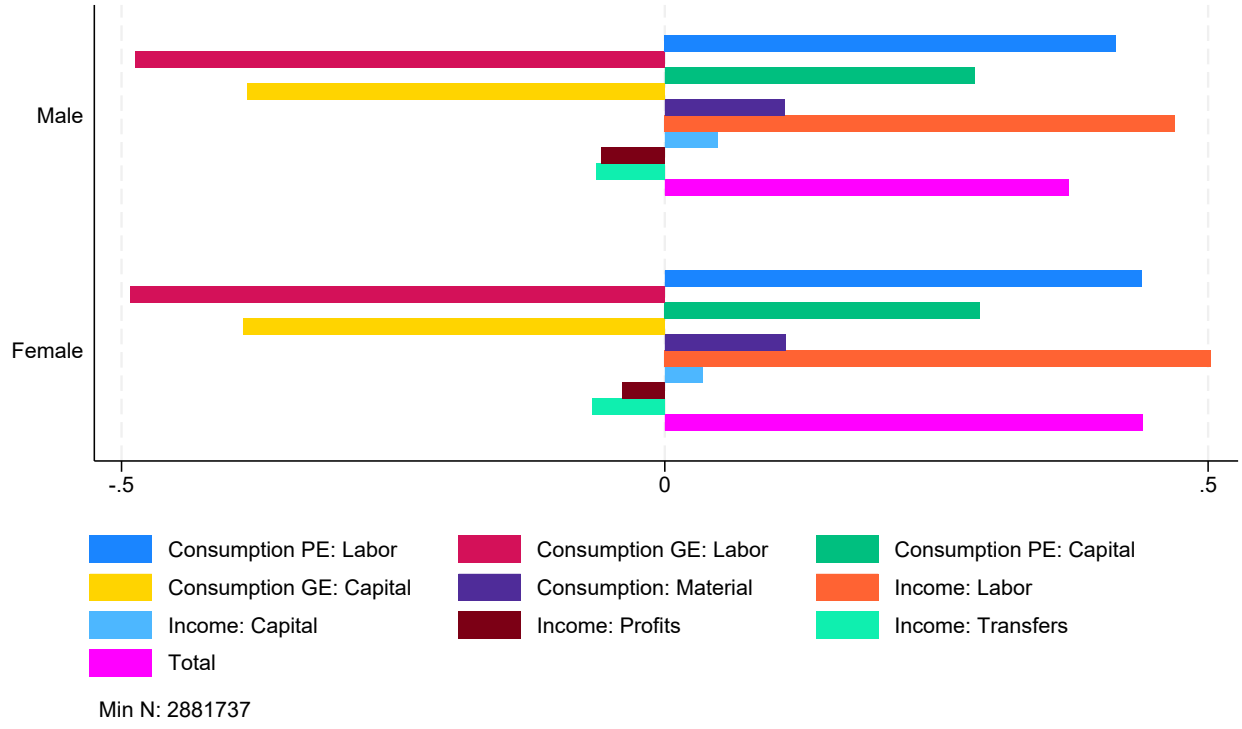


Figure 11: Welfare Impacts of Reducing All Distortions by 1%: Across Gender

Notes: This figure displays the impacts on the real income $\mathcal{Y}_c$ of individuals c across gender of a counterfactual exercise that reduces all distortions in the economy by 1%. Impacts are broken down by 4 channels. First, the partial-equilibrium consumption exposure to labor, capital and material wedges. Second, the general equilibrium changes in factor income coming from changes in factor prices. Third, the impact that factor price changes have on consumption prices (labeled GE). Finally, the impact on non-factor income, which involves profits and transfers.

sistent with the well-established feature of the gender wage gap. Through the lens of our model, a fraction of this gap is due to distortions, which are partially reduced in this counterfactual, benefiting women relatively more.

6 Conclusion

Recent work has documented the pervasive extent of economic distortions—such as market power, taxes, tariffs, credit constraints, etc.—and how they lead to substantial misallocation, or aggregate productivity loss. Far less well understood is how these phenomena affect members of society differently. In this paper we combine unique datasets from Chile, linking workers and owners to firms, firms to each other, firms to consumers, and firms and consumers to the government, in order to quantify the full incidence of distortions for the first time.

We find that the burden of distortions falls relatively more on the shoulders of the poor, the young and women. The channels, though, are very different across groups. The poor gain relatively more from reducing distortions because of larger reductions in the prices of the types of good they buy, larger increases in labor income, and smaller reductions in profit income. The young and women gain more due to bigger gains in labor income.

Future versions of this analysis will incorporate three additional features into the modeled economy. First, we will merge in data on individuals' education status so as to differentiate labor by skill. Second, we will incorporate a survey of informal firms in order to have such firms provide income to those who work at or own them. Finally, we will use data on households' state pension claims in order to broaden the picture of how households heterogeneously own shares of private-sector firms.

References

Acosta, M. and L. Cox (2024). *The Regressive Nature of the US tariff code: Origins and Implications*. Tech. rep. Working Paper.

Adão, R., P. Carrillo, A. Costinot, D. Donaldson, and D. Pomeranz (2022). "Imports, Exports, and Earnings Inequality: Measures of Exposure and Estimates of Incidence". *The Quarterly Journal of Economics* 137, pp. 1553–1614.

Andersen, A. L., K. Huber, N. Johannesen, L. Straub, and E. T. Vestergaard (Nov. 2022). *Disaggregated Economic Accounts*. Working Paper 30630. National Bureau of Economic Research.

Atkin, D. and D. Donaldson (2022). "The Role of Trade in Economic Development". *Handbook of International Economics: International Trade, Volume 5*. Ed. by G. Gopinath, E. Helpman, and K. Rogoff. Vol. 5. Handbook of International Economics. Elsevier, pp. 1–59.

Baqaee, D. and E. Farhi (2019). "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem". *Econometrica* 87, pp. 1155–1203.

Baqaee, D. and E. Farhi (2020). "Productivity and Misallocation in General Equilibrium". *The Quarterly Journal of Economics* 135, pp. 105–163.

Blanchet, T., E. Saez, and G. Zucman (July 2022). *Real-Time Inequality*. Working Paper 30229. National Bureau of Economic Research.

Burstein, A., J. Cravino, and M. Rojas (2024). *Input Price Dispersion Across Buyers and Misallocation*. Tech. rep. Central Bank of Chile.

Conlon, C., N. Rao, and Y. Wang (2022). "Who Pays Sin Taxes? Understanding the Overlapping Burdens of Corrective Taxes". *Review of Economics and Statistics*, pp. 1–27.

De Loecker, J. and F. Warzynski (2012). "Markups and Firm-Level Export Status". *American Economic Review* 102, pp. 2437–71.

Faber, B. (2014). "Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices". UC Berkeley Department of Economics.

Faber, B. and T. Fally (May 2022). "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data". *The Review of Economic Studies* 89.3, pp. 1420–1459. eprint: https://academic.oup.com/restud/article-pdf/89/3/1420/43615976/rdab061.pdf.

Gervais, A. and J. B. Jensen (2019). "The tradability of services: Geographic concentration and trade costs". *Journal of International Economics* 118, pp. 331–350.

Gupta, A. (2022). "Demand for Quality, Variable Markups and Misallocation: Evidence from India". Dartmouth working paper.

Hall, R. E. (1986). "Market Structure and Macroeconomic Fluctuations". *Brookings Papers on Economic Activity* 2, pp. 285–338.

Hsieh, C.-T. and P. J. Klenow (2009). "Misallocation and Manufacturing TFP in China and India". *Quarterly Journal of Economics* 124.4, pp. 1403–1448.

Jones, C. I. (2013). "Misallocation, Economic Growth, and Input–Output Economics". *Advances in Economics and Econometrics*, pp. 419–456.

Manelici, I., J. P. Vasquez, M. Ulate, and R. D. Zarate (2024). *The Gains from Foreign Multinationals in an Economy with Distortions*. Tech. rep.

Redding, S. J. and D. E. Weinstein (2024). "Accounting for trade patterns". *Journal of International Economics* 150, p. 103910.

Restuccia, D. and R. Rogerson (2008). "Policy distortions and aggregate productivity with heterogeneous establishments". *Review of Economic Dynamics* 11.4, pp. 707–720.

Sangani, K. (2023). *Markups Across the Income Distribution: Measurement and Implications*. Tech. rep.

Schmitz, J. A. (2020). *Monopolies Inflict Great Harm on Low-and Middle-Income Americans*. Tech. rep. Federal Reserve Bank of Minneapolis.

Sharma, G. (2023). *Monopsony and Gender*. Tech. rep.