The Incidence of Distortions*

David Atkin
MIT and NBER

Baptiste Bernadac
UCLA

Dave Donaldson
MIT and NBER

Tishara Garg
Princeton

Federico Huneeus
Duke and Central Bank of Chile

PRELIMINARY DRAFT

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: December 2025. Author contacts: atkin@mit.edu, bbernadac@ucla.edu, ddonald@mit.edu, tgarg@princeton.edu, and federico.huneeus@duke.edu. We are grateful to Joaquin Galeno, Pedro Pessoa, Melody Ying, Vincent (Chengrui) Wang, Qirui Wang, and Jeremy Zhou 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, 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 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 electronic transaction records further providing visibility on where that product was sourced from, 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 (both capital and labor) to firms, firms use arbitrary technologies to make output with factors and intermediates, 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 allow for 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 are 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. We adapt modern computational tools for solving linear systems of equations with dimensionality in the millions—as is necessary, given the need to do general equilibrium analysis with individual- and firm-level microdata on even a relatively small country such as Chile—in order to do so. Our main exercise reduces all distortions proportionately to their size. 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.

Our second exercise removes 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. Such an analysis also sheds light on the question posed above about whether certain distortions are countervailing or reinforcing. Finally, we study the trade-off between equity and efficiency that exists in Chile—as a result of the tax and non-tax distortions that appear to be in place throughout its economy—by comparing aggregate impacts to distributional ones.

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 quantify 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 Setup

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, and the firm's profit (after corporate taxes, described below) is π_i. We denote the price of good i that the consumer pays by p_i and the payment for factor f that the owner receives as w_f. Finally, we let T_c denote the the net personal taxes (described below) paid by consumer c and s_c the consumer's net savings.[1] The consumer's budget constraint is therefore

$$\sum_{i \in \mathcal{N}} p_i x_{ci} = \sum_{f \in \mathcal{F}} \Phi_{cf} w_f L_f + \sum_{i \in \mathcal{N}} \Phi_{ci} \pi_i - s_c + T_c \equiv \chi_c. \tag{1}$$

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 on 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

[1] Because our model is static we treat s_c as constant in what follows.

[2] We map multi-product firms (including firms that sell a single product to multiple buyers) 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. In particular, 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.

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 $\lambda_i \equiv p_i y_i$. The firm purchases a quantity x_{ij} of inputs from firm j at price p_j and quantity x_{if} of factor services from factor f. Letting T_i denote all taxes (described below) that firm i rebates to the government, its after-tax profits are

$$\pi_i = p_i y_i - \sum_{j \in \mathcal{N}} p_j x_{ij} - \sum_{f \in \mathcal{F}} w_f x_{if} - T_i.$$

Finally, we define the $(N+F) \times (N+F)$ input-output matrix Ω such that firm-to-firm sales elements are $\Omega_{ij} \equiv \frac{p_j x_{ij}}{p_i y_i}$, factor-to-firm elements are $\Omega_{if} \equiv \frac{w_f x_{if}}{p_i y_i}$, and all other elements are zero by convention. Using this definition, the Leontief inverse of Ω is denoted by $\Psi \equiv (I - \Omega)^{-1}$.

Distortions. We allow for a general treatment of distortions in firm production. Following standard conventions, we define the wedge μ_{if} on firm i's use of factor f as the ratio of its (value-adjusted) marginal product for that factor divided by the factor's price. This can be written as

$$\mu_{if} \equiv \frac{\eta_{if}}{\Omega_{if}}, \tag{2}$$

where η_{if} is the elasticity of firm i's output with respect to factor f.[3] Similarly, the wedge on firm i's use of the input purchased from firm j is defined as

$$\mu_{ij} \equiv \frac{\eta_{ij}}{\Omega_{ij}}, \tag{3}$$

where η_{ij} is the elasticity of firm i's output with respect to input j.

We refer to μ_{if} and μ_{ij} as wedges because the values $\mu_{if} = \mu_{ij} = 1$ correspond to an economy that features production efficiency (and hence also Pareto efficiency, given the absence of distortions in consumption). Other values of wedges generically lead to inefficient aggregate production.

The wedges μ_{if} and μ_{ij} can arise due to a host of underlying market failures. One important source is market power. For example, consider a firm whose only input is factor f. If this firm has market power in its input market then it will hold back on hiring the factor (because hiring more will raise the price of inframarginal units) and push the value of its marginal product above the marked down price paid to the factor w_f, resulting in a wedge $\mu_{if} > 1$. If instead the firm has output market power then it will charge its buyers a price that is marked up above the firm's marginal cost, and this situation again results in

[3]Denoting firm i's production function by $G_i(\cdot)$, this expression obtains by combining the definitions of the wedge, $\mu_{if} \equiv \frac{p_i}{w_f}\frac{\partial G_i}{\partial x_{if}}$, the output elasticity $\eta_{if} \equiv \frac{x_{if}}{y_i}\frac{\partial G_i}{\partial x_{if}}$, and the input share $\Omega_{if} \equiv \frac{w_f x_{if}}{p_i y_i}$.

$\mu_{if} > 1$. Even if the firm simultaneously has market power in both markets then the value of μ_{if} summarizes the total effect of both underlying sources of distortion into a measure of how much factor f is underutilized (relative to the $\mu_{if} = 1$ benchmark) by this firm.[4]

Beyond market power, distortions can also arise due to taxation. While we introduce taxes explicitly next, it is important to note that since the wedges μ_{if} and μ_{ij} include all sources of distortion, they already incorporate the effects of any distortionary taxes. Continuing the example above, if the factor f in question is labor and the firm is charged a payroll tax then this would further contribute to a gap between the firm's value marginal product of labor and the wage w_f that its workers take home since the firm is paying a marginal price for labor that includes both w_f and the tax; that is, even a firm with no market power would have $\mu_{if} > 1$.

Finally, given wedges, we define cost-based input shares as $\widetilde{\Omega}_{if} \equiv \mu_{if}\Omega_{if}$ (and analogously $\widetilde{\Omega}_{ij}$). Similarly the cost-based Leontief inverse matrix is given by $\widetilde{\Psi} \equiv (I - \widetilde{\Omega})^{-1}$.

Government. As anticipated above, both consumers and firms interact with the government via the tax and transfer system. Beginning with consumers, we model the net transfer payment received by consumer c in equation (1) as

$$T_c = \Phi_{cg}T - \sum_{f \in \mathcal{F}} \frac{t_{cf} - 1}{t_{cf}} \Phi_{cf} w_f L_f - \sum_{i \in \mathcal{N}} \frac{t_{ci} - 1}{t_{ci}} \Phi_{ci}\pi_i. \tag{4}$$

Here, T denotes total net government tax revenue, of which consumer c is allocated the fixed share $\Phi_{cg} \geq 0$. However, this consumer pays income taxes t_{cf} on their earnings $(\Phi_{cf} w_f L_f)$ of each factor f and t_{ci} on their earnings from the net-of-corporate-taxes profits $(\Phi_{ci}\pi_i)$ of each firm i.[5] The assumption that income taxation is linear is without loss in our application below because we set the consumer-specific rates t_{cf} and t_{ci} below to reflect each individual's actual average tax rates at baseline (and adjust it during counterfactuals, as discussed further below) and because consumer income taxes are not distortionary in our model.[6]

Turning to firm taxation, we incorporate three key features of actual taxes in the Chilean context. First, we allow for a value-added tax (VAT), which is modeled by incorporating a

[4]A separate type of distortion arises (for example, in models of financial frictions) when the input seller j will not sell more than a given input quantity to firm i at the price p_j. If this constraint is binding then the shadow price of input j within firm i will exceed p_j—that is, the shadow price is given by $\mu_{ij}p_j$ for some $\mu_{ij} > 1$. Appendix A.8 illustrates the case of financial frictions using a simple example.

[5]Throughout, we define tax rates according to the convention that $t = 1$ corresponds to no tax at all, and $t > 1$ corresponds to a positive amount of taxation.

[6]This is because factors are in fixed supply to each consumer and there is no margin of entrepreneurial effort that makes profit taxation distortionary.

tax t_i on sales and a tax rebate t_{ij} on the purchase of any eligible inputs purchased by firm i from another firm j. However, these tax rates are firm pair-specific in order to allow for the fact that if firm i is informal it does not pay VAT at all (i.e $t_i = t_{ij} = 1$ for all j) and if firm i is formal but it buys from an informal firm j then firm i pays the sales tax but receives no rebate for inputs (i.e. $t_i > 1$ and $t_{ij} = 1$ for all informal j). The second type of firm tax that we allow for is a payroll tax t_{if} that firms must pay for all factors f that are a type of labor. Finally, the third type of firm tax in Chile is a corporate income tax on profits t_i^p. Putting these three types of taxes together, the total tax payments by firm i will be

$$T_i = \frac{t_i - 1}{t_i} p_i y_i - \sum_{j \in \mathcal{N}} \frac{t_{ij} - 1}{t_{ij}} p_j x_{ij} + \sum_{f \in \mathcal{F}} (t_{if} - 1) w_f x_{if} + (t_i^p - 1)\pi_i, \tag{5}$$

where the four terms respond to the tax payments corresponding to sales (inside VAT), rebates on non-factor inputs (also inside VAT), payroll taxes on labor factor inputs, and corporate income tax.[7]

Here, firm VAT and payroll taxation is generically distortionary because the contributions of these types of tax to T_i depend on the firm's decisions about output y_i and inputs (x_{ij} and x_{if}). As discussed above, the distortionary consequences of these taxes are already embodied in the composite wedges μ_{if} and μ_{ij}. However, tax- and non-tax contributions to these wedges need to be distinguished, as we do here, for two reasons: (i) the revenues from tax distortions accrue to the government, whereas the revenues from non-tax distortions accrue to the firm as profits; and (ii) below we perform counterfactual simulations that hold tax rates (and hence tax-related distortions) constant but reduce non-tax distortions.

Finally, the value of total net transfers T is set such that the government's budget balances:

$$\sum_{c \in \mathcal{C}} T_c = \sum_{i \in \mathcal{N}s} T_i.$$

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; another removes the non-tax components of all wedges while holding their tax components constant.

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.

[7] Recall that π_i denotes after-tax profits.

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

$$d\ln \mathcal{Y}_c = d\ln \chi_c - \sum_{i\in\mathcal{N}} b_{ci} d\ln p_i. \tag{6}$$

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

$$d\ln p_i = \sum_{k\in\mathcal{N},\mathcal{F}} \widetilde{\Psi}_{ij}\widetilde{\Omega}_{jk} d\ln \mu_{jk} + \sum_{f\in\mathcal{F}} \widetilde{\Psi}_{if} d\ln w_f. \tag{7}$$

For example, if the wedge on good j's use of input k shrinks (i.e. $d\ln\mu_{jk} < 0$), this effect will propagate forward along the supply-chain to the price of good i in accordance with the weight $\widetilde{\Psi}_{ij}\widetilde{\Omega}_{jk}$—this appropriately sums all senses in which the cost of good i depends, both directly and indirectly, on the price of input j (which in turn depends on μ_{jk}).

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 (7), weighted by the Leontief-inverse exposure elements $\widetilde{\Psi}_{if}$) as well as the incomes $d\ln\chi_c$ in Equation (6). These changes in income themselves satisfy

$$d\chi_c = \sum_{f\in\mathcal{F}} \Phi_{cf} L_f dw_f + \sum_{i\in\mathcal{N}} \Phi_{ci} d\pi_i + dT_c, \tag{8}$$

where dT_c follows from differentiating Equation (4), allowing endogenous variables (such as income) to change as well as for any change to the tax and transfer schemes that is desired to be simulated as part of the counterfactual. Similarly, the change in after-tax profits $d\pi_i$ satisfies

$$d\pi_i = \left(\frac{\pi_i + T_i}{\lambda_i}\right) d\lambda_i + \lambda_i \sum_{j\in,\mathcal{N},\mathcal{F}} \Omega_{ij}(d\ln\mu_{ij} - d\ln\widetilde{\Omega}_{ij}) - dT_i, \tag{9}$$

where dT_i is the change in net taxes rebated by the firm, which again follows from differentiating equation (5) with respect to all endogenous and exogenous elements.

The change in firm i's Domar weight is given by

$$\begin{aligned} d\lambda_i = &- \sum_{l\in\mathcal{N};m\in\mathcal{N}} \lambda_l \Omega_{lm} d\ln\mu_{lm} \Psi_{mi} \\ &+ \sum_{\mathcal{N}} \lambda_k \text{Cov}_{\widetilde{\Omega}^{(k)}}(d\ln\widetilde{\Omega}^{(k)}, diag(\mu^{(k)})^{-1}\Psi_{(i)}) + \sum_{c\in\mathcal{C}} d\chi_c \sum_{k\in\mathcal{N}} b_{ck}\Psi_{ki}, \end{aligned} \tag{10}$$

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 (10) but for dw_f instead of $d\lambda_i$.

The intuition behind Equation (10) is as follows. The first term captures the direct effect of wedges on demand for i: for example, an increase in the wedge μ_{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 Equations (9) and (10) 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. We return to the specification of these substitution forces below.

Put together, the system of equations in (7)-(10) constitutes a linear system that can be solved for any given values of elasticities θ_i, wedges μ_{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 (6) and (7) 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 (8) and (9): 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 (7) and (8) 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 nine 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.[8]

[8]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

First, we use a firm-to-firm electronic transaction records 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.[9] 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 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 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-

presented to the Service; therefore, the veracity of the data is not the responsibility of the Service.

[9] 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.

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

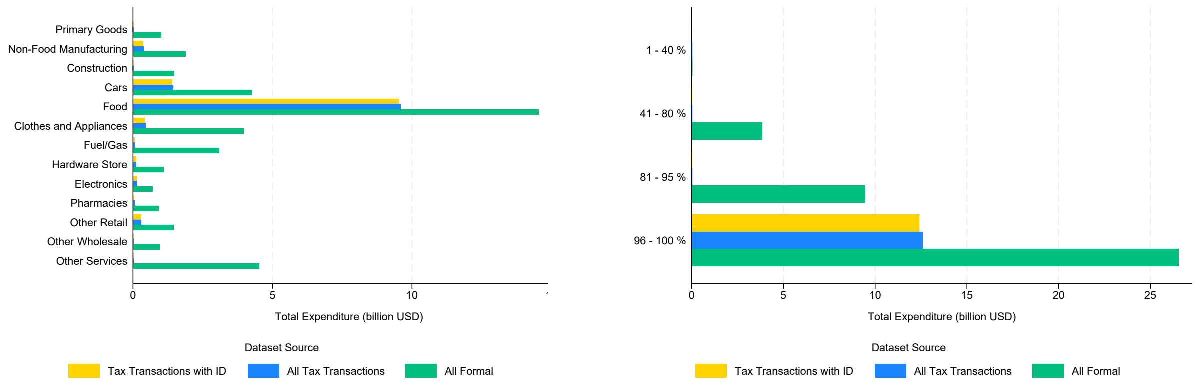

(a) Across Retailer Types

(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.

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.[10]

Fourth, we use matched employer-employee records (from IRS tax affidavits 1887 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 pension records to expand these ownership linkages. Chile has a mandatory private pension system with individuals able to choose between seven pension providers with each provider offering five funds of varying risk levels. The investments of each fund are public, allowing us to obtain each individual's shareholdings of any companies their pension fund invests in.[11]

[10] 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.

[11] Due to server confidentiality rules, the administrative pension data can only be merged in once other parts of the analysis are complete. Thus, we currently substitute actual pension holdings with predicted pension holdings based on the consumption survey (described below) that records pension payments and receipts.

Seventh, we use government-to-household linkages that combine 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 type of policy to which the transfer corresponds. 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 seven 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.

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.

While the administrative datasets cover close to the universe of formal economic transactions in the economy, they miss informal economic activity that is typically estimated to comprise about one quarter of employment (though considerably less of total output) in Chile's economy. 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.

[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.

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, 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 ad-

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

Figure 2: Consumption Data Coverage: Across Products and Aggregates

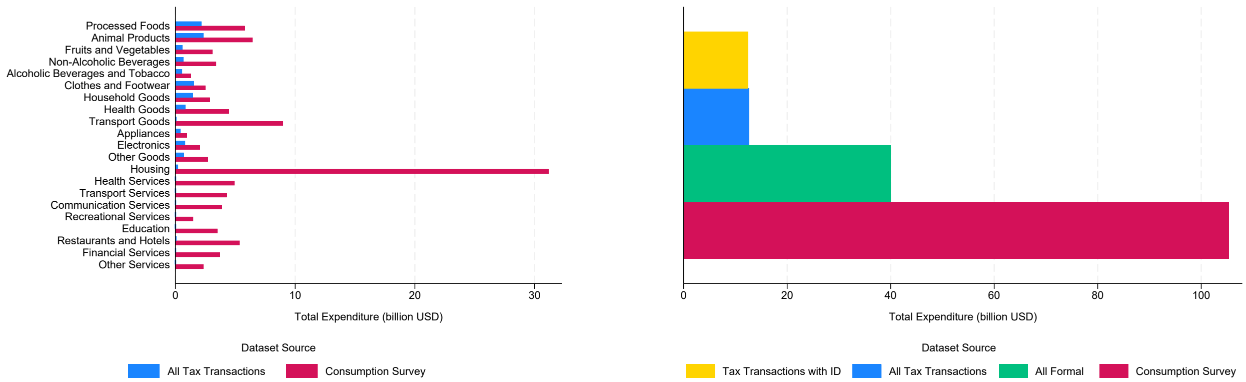

(a) Across Products

(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).

vances in optimal transport algorithms. This approach preserves the joint distribution of demographics and expenditures in the consumption survey data that is being brought into the administrative data and avoids extrapolation (and possibility of extreme outliers) that would occur using polynomial-based prediction models. We first form broad bins based on region, age and income. Within each group, survey respondents are matched (one-to-many) to individuals in the administrative data based on gender, income, labor and profit income shares, and 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 for each individual-product-firm type triplet. However, we still must allocate all these expenditures to specific retailers. We start by forming region-product-firm type groups. Excess expenditures in the consumption data within a group are allocated proportionately to the corresponding stores of that type, e.g. a specialized shoe retailer, in the location that the individual resides. Consumers are randomly matched to specific formal stores for whom we do not have linked tax-id data but whose total sales we still see in the administrative data.[14] The remaining consumer expenditure predicted from the survey match is allocated to synthetic informal firms of the

[14]Recall that we have total final goods sales for all formal firms but only firm-to-individual matched sales for the 2,124 largest sellers.

corresponding type (e.g. an informal specialized retailer).[15] These additional expenditures enable us to extend the detailed consumption patterns captured in administrative records to a more complete set of consumption transactions in the economy.

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 2021–2023 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 at the individual level on the vector of formal income by sector, again separately for each demographic group). Finally, we bring in a survey of small firms designed by the Chilean government to measure the informal sector (the Encuesta de Microemprendimiento surveys conducted in 2017, 2019 and 2022 surveying over 20,000 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 reported in the labor force surveys prior to performing the match between the labor force survey and the administrative data described above. This match allows us to supplement the administrative income data with informal income as well as ownership and profits of informal firms. Additionally, we augment the administrative firms' dataset with the informal firms created through this match.

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.

[15] To prevent very small or noisy predictions, we truncate extremely small consumers in these cases. We allow consumption to exceed the predicted consumption from the survey when expenditure in the firm-to-individual administrative data is greater than the matched total consumption from the survey.

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, 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 records 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 and pension records (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 the 1-digit 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.

Figure 3: Exposure of Individuals to Labor Income

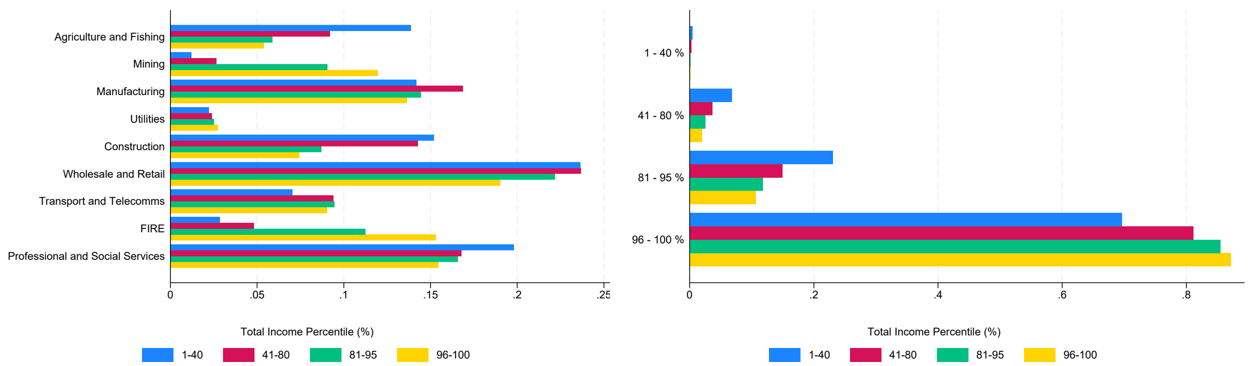

(a) By sector

(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.

Other types of exposure matrices are also easy to visualize (in aggregated forms). For 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.[16]

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

[16]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). An extreme case here are profits accruing to pension funds which we assign to the owners of those pension funds even though they only receive the returns to those claims upon retirement.

Figure 4: Income by Source

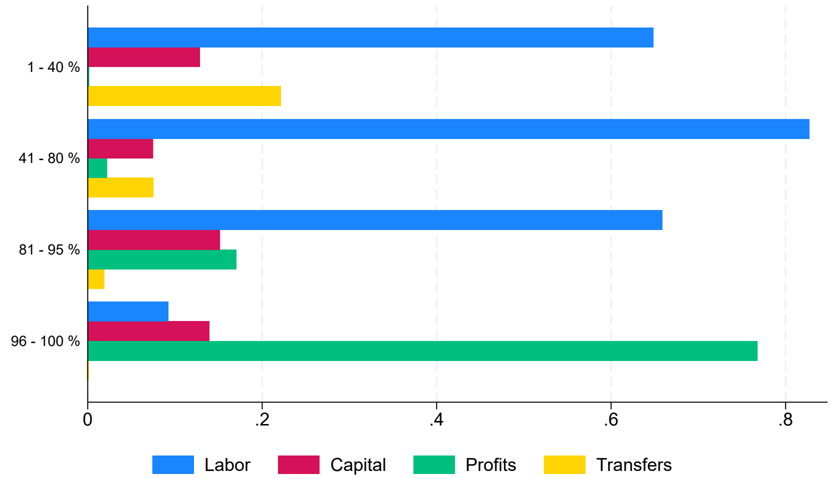

Notes: This figure displays the breakdown, for any type of individual (based on their position in the income distribution) c, of their earnings by source. Four types of source are indicated: that from all types of labor, that from all types of capital (based on a fair return), that from the profits that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

firms i are broken down by 1-digit 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 distribution; for example, poorer individuals are relatively more likely to own agriculture and fishing firms, and small- and medium-sized firms more generally.

Similarly, Figure 6 turns to the case of the matrix of consumer expenditure shares, beginning with a summary of direct shares (i.e. b_{ci}) in Panel (a). 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. Panel (b) continues by examining the indirect consumption exposure embodied in $\sum_i b_{ci}\tilde{\Psi}_{ij}$. As expected, a number of products are bought little in a "direct" fashion by consumers (low b_{ci} in Panel a) but are "indirectly" purchased to a far greater extent (high $\sum_i b_{ci}\tilde{\Psi}_{ij}$ in Panel b). This is the case, for example, for any product that is a widespread production input

Figure 5: Exposure of Individuals to Profit Income

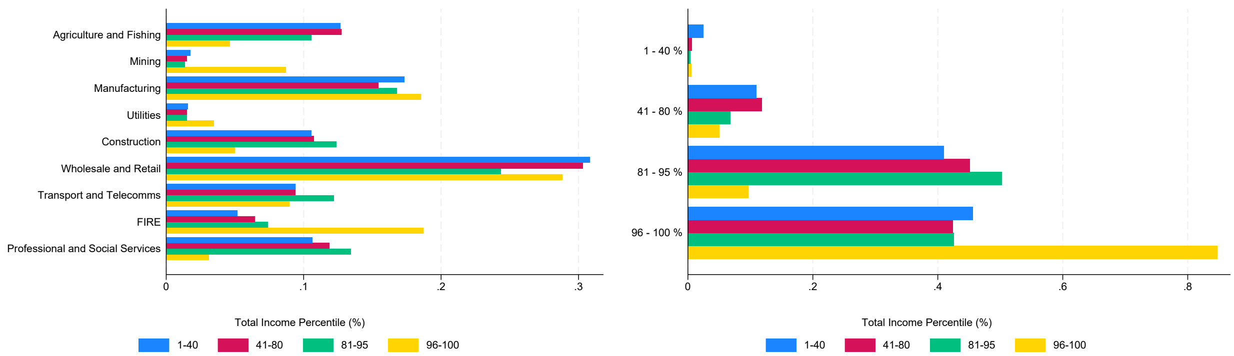

(a) By sector

(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.

but never features in final consumption.

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

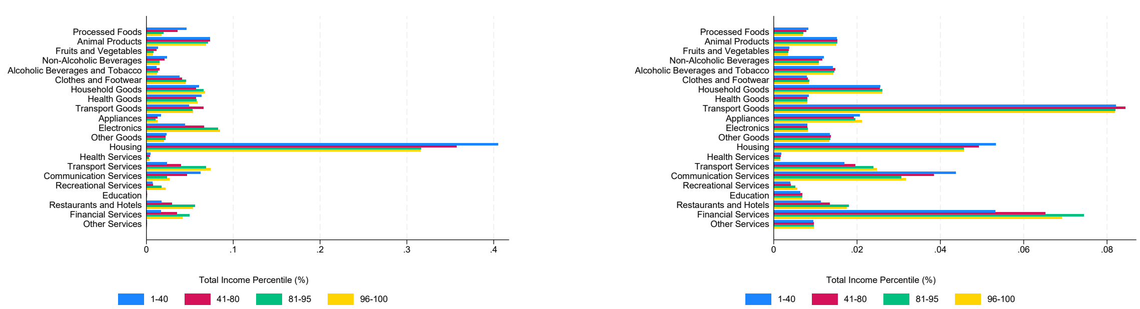

(a) Direct consumption: b_{ci}

(b) Indirect consumption: $\sum_i b_{ci} \tilde{\Psi}_{ij}$

Notes: Panel (a) 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) by product category i (aggregated across firms in each category. By construction, $\sum_i b_{ci} = 1$ for each c. Panel (b) displays the analog but for indirect consumption, defined as $\sum_i b_{ci} \tilde{\Psi}_{ij}$.

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, $\widetilde{\Psi}_{ij} - \tilde{\Omega}_{ij} - I$.

Figure 7: Firm-to-Firm Expenditure Shares

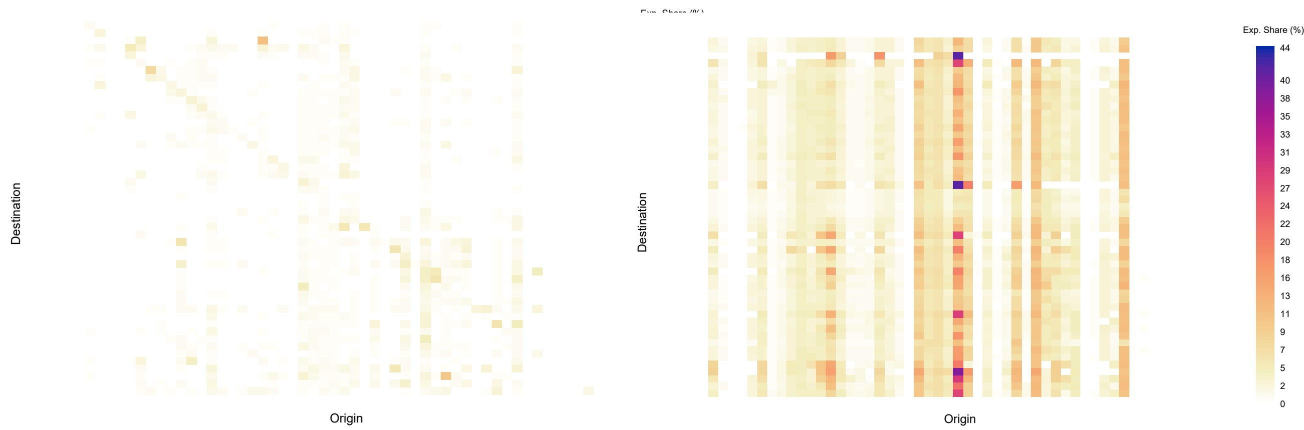

(a) Direct Linkages: $\widetilde{\Omega}_{ij}$

(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

As discussed in Section 2, the solution of Equations (7)-(10) requires the specification of the extent to which consumers' and firms' expenditure shares adjust in the face of price changes. In both cases we use functional form choices that have been common in prior work.

Beginning with preferences, we specify consumer demand as a nested demand system from across products and sectors.[17] In particular, at the lower level, we let the elasticity θ_s denote 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. The former study is based on US data and the latter is based on Chilean data.

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. 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.

[17]Despite the terminology here, sectors in production are not necessarily the same as sectors in consumption.

4.3 Distortions

The final input into our analysis is the size of the wedges μ_{if} and μ_{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. Indeed, as Equations (2) and (3) make clear, the key requirement is the output elasticities, η_{if} and η_{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_{if} and 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 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, μ_{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 substantial 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

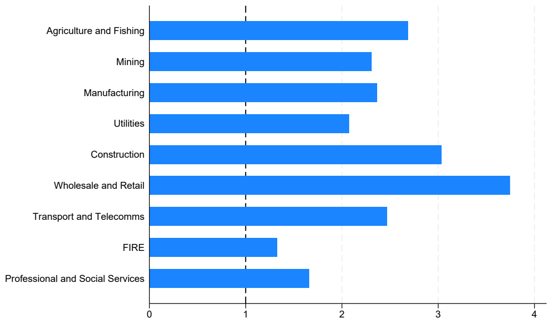

(a) Labor Wedges by Sector

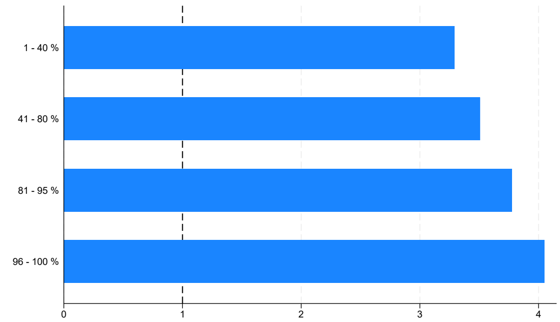

(b) Labor Wedges by Firm Size

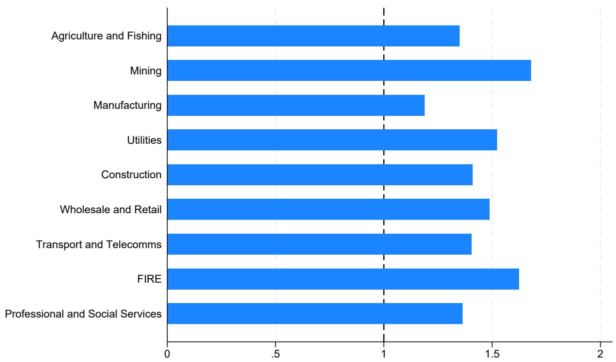

(c) Capital Wedges by Sector

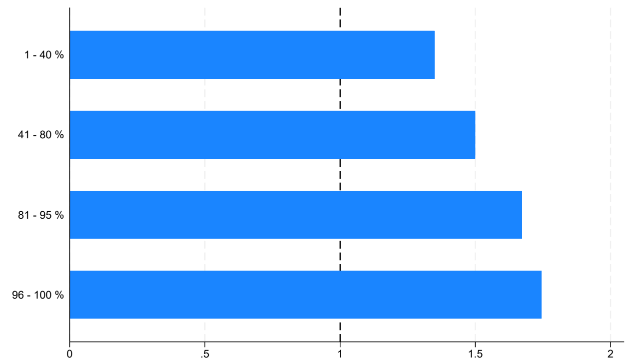

(d) Capital Wedges by Firm Size

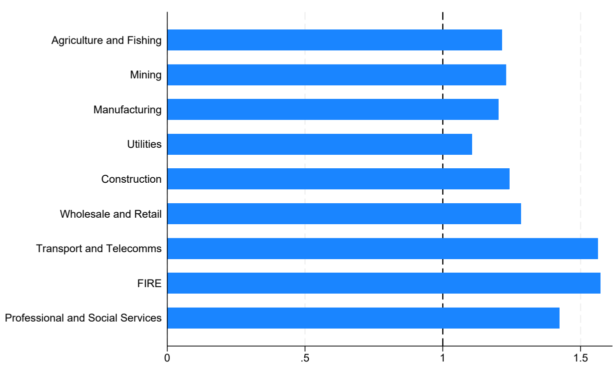

(e) Materials Wedges by Sector

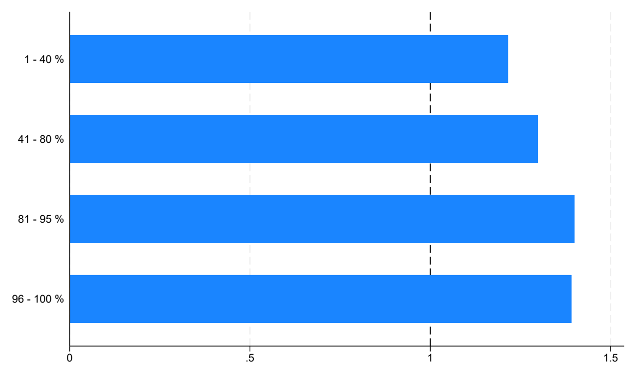

(f) Materials Wedges by Firm Size

Notes: Panel (a) displays average values of the estimated wedge μ_{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 μ_{ij}, where j is capital and 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 Measuring the Incidence of Distortions

Having populated the formulae in Section 2 with the data and estimates described in Sections 3 and 4, we now perform a number of counterfactual simulations 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 (6) 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? While executing any such counterfactuals, we hold constant features of the Chilean tax and transfers system such as the corporate profit tax rate and the value-added tax system.[18]

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 (21)) and those coming from changing factor prices (the last terms in Equation (21)).

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 heterogene-

[18] In particular, we assign to each firm the profit tax rate that they were paying in 2022, and to each firm-product the relevant statutory value-added tax rate (which varies by product).

ity 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 trade-offs 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?

In this preliminary draft, the calculations that follow are based on a version that is simplified in a number of dimensions:

- When reducing wedges we solve only for the linearized effect of a 1% change in any wedge. This calculates the local elasticity of responses to such small shocks, which amounts to the first step in an iterative algorithm for calculating the effects of large shocks (such as the complete elimination of wedges).
- Individuals are not grouped into households via the civil registry.
- Indirect firm ownership via pensions is not included. (But all firms' remaining ownership shares are allocated proportionally across direct owners so that ownership shares always sum to one.)
- We set $\theta_U = 1$ and $\theta_s = 1.3$ for all s.
- Types of labor are broken down by region but not further by, for example, education group.

We also focus on impacts across the income distribution, leaving the study of wider demographic impacts to future versions.

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 all of the four income bands that we have introduced above. These effects range from

0.4% to 0.55%.[19] Members of the upper middle-class group gain the most (in terms of percentage growth in their real income), though the difference across groups is relatively small. The fact that there are broad gains for all groups is not surprising as the removal of distortions raises efficiency and hence the total size of the pie.

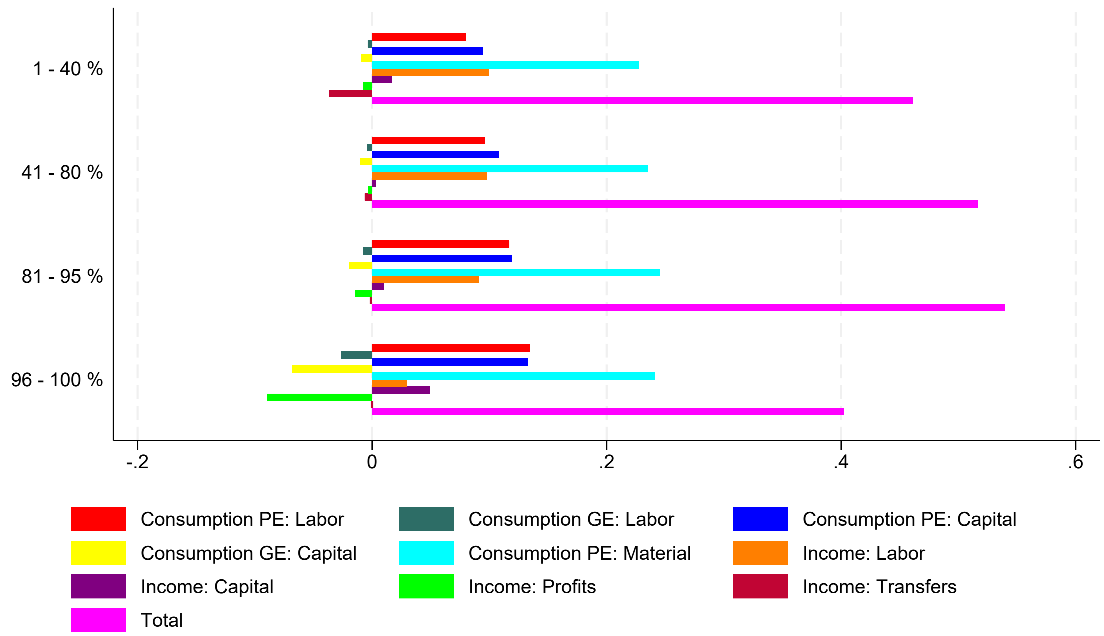

Figure 9: Welfare Impacts of Reducing All Distortions by 1%: Across Income

Notes: This figure displays the impacts on the real income $\mathcal{Y}_c$ (expressed as percent changes) 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.

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 report 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 material input wedges, then typically the gains from capital and labor wedges are similar but smaller than those from input wedges. These effects are again similar across income groups, though capital and labor wedge effects are slightly pro-rich and material wedge effects are slightly pro-middle class.

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

Second, we consider the general equilibrium forces that arise due to changes in factor income of labor and capital.[20] 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, gains from capital income changes tend to be pro-rich. This is to be expected given the 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.[21]

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 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. Overall, the heterogeneity of this profit channel drives most of the heterogeneity we see in total welfare impacts across income groups—at least in these preliminary results, which omit a number of potentially important sources of heterogeneity.

Having examined incidence through the lens of heterogeneity across individuals grouped by income, we turn now to three additional differences. Figure 10 explores incidence by age, Figure 11 by gender, and Figure 12 by region. The gains from reducing wedges are positive for all age groups, though older individuals gain relatively less from reduction in wedges. This result comes primarily from the fact that older individuals are more exposed to profit losses than younger individuals are. Turning to gender differences, Figure 11 shows that women gain relatively more than men from reducing distortions. This difference comes mostly from consumption-based heterogeneity, with women more exposed (on the consumption side) to labor wedges, capital wedges, and material wedges. Finally,

[20]Given that factor supply is fixed and that the numeraire is GDP, these effects can be interpreted as changes in factor shares of GDP.

[21]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.

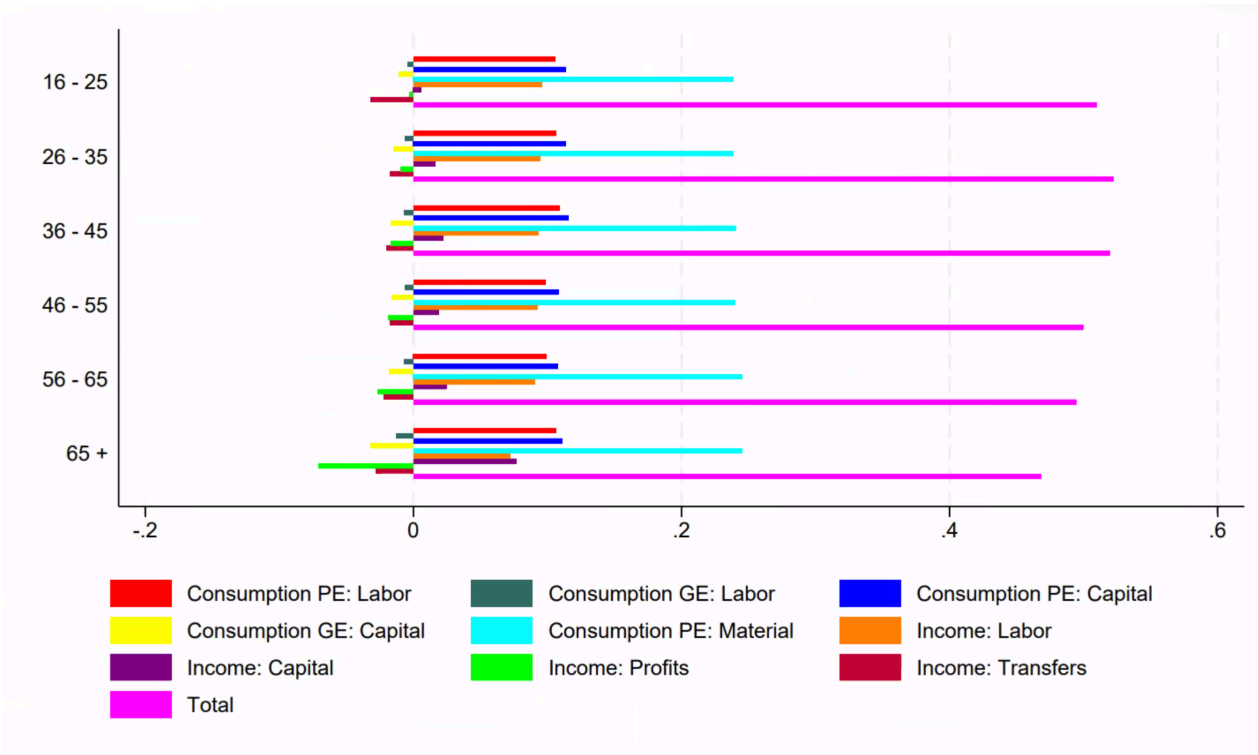

Figure 10: Welfare Impacts of Reducing All Distortions by 1%: Across Age

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.

with respect to geographic differences, in Figure 12 we split up Chile into three regions: Santiago (the country's political and commercial capital), other urban, and all rural. Evidently, residents of Santiago benefit least from reducing distortions, and this effect derives primarily from the exposure there of consumers to wedges (on all of labor, capital, and materials).

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.

5.2 Which Distortions Matter More?

The previous counterfactual exercise examined the impact of reducing all distortions throughout the economy. But it is natural to ask whether the incidence of certain types of distortions are borne more or less unequally than others. Towards this end, Figure 13 reports results from six separate simulations: the first is the baseline exercise of Section 5.1 (for purposes of comparison); the next three reduce (by 1%) only the labor, capital, and materials wedges in the economy, respectively; the fifth reduces all wedges, but only those in large firms; and the sixth reduces all wedges, but only those in manufacturing firms. Our

Figure 12: Welfare Impacts of Reducing All Distortions by 1%: Across Geography

Notes: This figure displays the impacts on the real income $\mathcal{Y}_c$ of individuals c across the location of individuals 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.

baseline findings about the incidence of distortions remain strikingly similar across these different types of distortions.

5.3 Equity-Efficiency Tradeoffs

As we have seen in the counterfactual simulations above, reductions in distortions have implications for inequality, especially at the very top. But they also have consequences for the economy's overall efficiency level. Our final analysis compares these two impacts for every simulation discussed thus far. Figure 14 reports the results. Each dot represents a separate counterfactual simulation, with the y-axis value referring to the impact of the counterfactual change on equality (as measured by the negative of the Gini coefficient across individual-level nominal income changes) and the x-axis value referring to the impact on aggregate efficiency (as measured by the Baqaee and Burstein 2022 value). Broadly, there is no evidence for an equity-efficiency tradeoff across these exercises—those that offer larger efficiency gains (such as that which reduces wedges on large firms only)

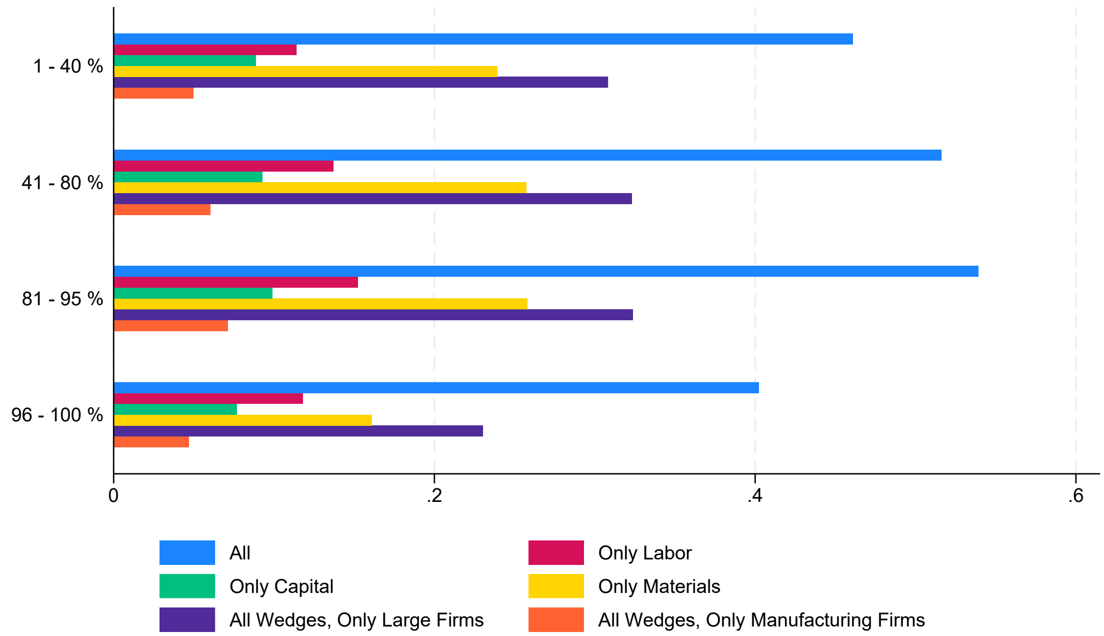

Figure 13: Welfare Impacts of Reducing Different Types of Distortions by 1%

Notes: This figure displays the impacts on the real income $\mathcal{Y}_c$ of individuals c across income groups of counterfactual exercise that reduce distortions in the economy by 1%. In the first, all distortions in the economy are reduced. Subsequent exercises reduce distortions only on labor, capital, materials, for large firms, and for manufacturing firms, respectively.

tend to also give rise to the greatest increases in (Gini-based) equality.

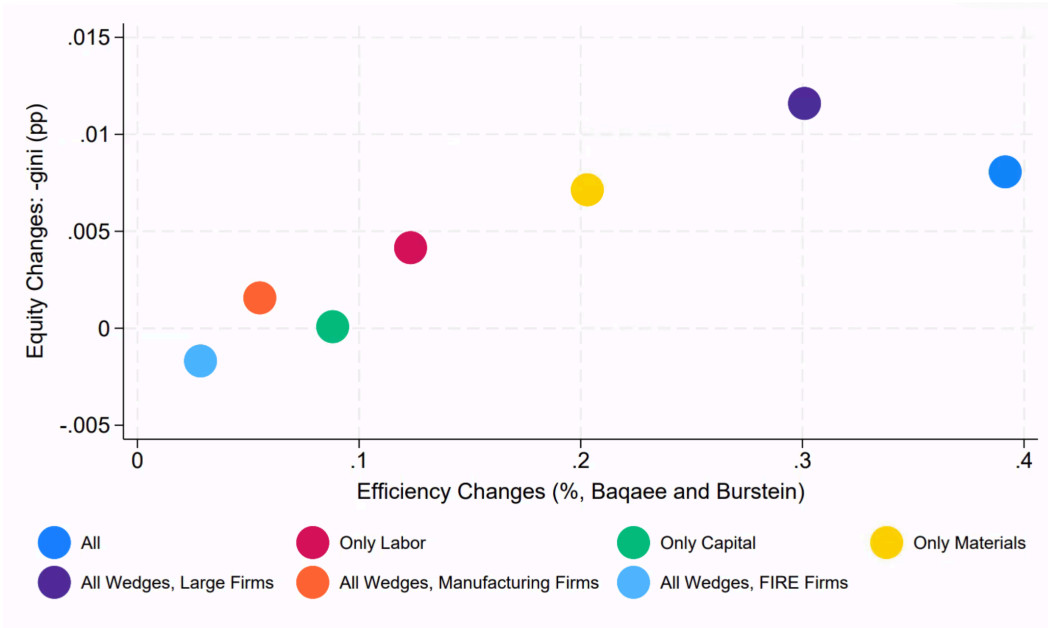

Figure 14: Equity-Efficiency Comparisons from Reducing Different Types of Distortions by 1%

Notes: Each dot in this figure displays the impacts of a separate counterfactual exercise. The y-axis reports the impact on the level of the negative of the Gini coefficient for nominal income (taken across all individuals). And the x-axis reports the impact on the level of the Baqaee and Burstein 2022 measure of aggregate efficiency.

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.

References

Acevedo, P., E. Luttini, M. Pizarro, D. Quevedo, and M. Rojas (2025). "Invoices rather than surveys: Using ML to monitor the economy". mimeo.

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. Elsevier, pp. 1–59.

Baqaee, D. and A. Burstein (2022). *Aggregate Welfare and Output with Heterogeneous Agents*. Tech. rep.

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.

BEA (2017). *US Input-Output table*. https://www.bea.gov/itable/input-output.

BEA/BLS (2017). *US capital expenses*. https://www.bea.gov/data/special-topics/integrated-industry-level-production-account-klems.

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.

Chile, B. C. de (2022). *Price Deflators*. https://si3.bcentral.cl/siete.

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.

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

Groningen, U. of (2014a). *WIOD National IO Table*. https://www.rug.nl/ggdc/valuechain/wiod/wiod-2016-release, in the section ”Input-Output tables of the WIOD 2016 release”.

Groningen, U. of (2014b). *WIOD Socio Economic Account*. https://www.rug.nl/ggdc/valuechain/wiod/wiod-2016-release, in the section ”Additional Data and Satellite Accounts”.

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

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.

INE (2012). *Industry classification*. https://www.ine.gob.cl/calidad-estadistica/clasificaciones, in the section ”CIIU”.

INE (2022a). *Consumption survey*. https://www.ine.gob.cl/estadisticas/sociales/ingresos-y-gastos/encuesta-de-presupuestos-familiares.

INE (2022b). *Employment survey*. https://www.ine.gob.cl/estadisticas/sociales/ingresos-y-gastos/encuesta-suplementaria-de-ingresos.

INE (2022c). *List of Municipalities*. https://www.ine.gob.cl/herramientas/portal-de-mapas/geodatos-abiertos.

INE (2022d). *Micro-enterprise survey*. https://www.ine.gob.cl/estadisticas/sociales/mercado-laboral/microemprendimiento.

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

KLEMS, W. (n.d.). https://www.worldklems.net/wkhome.

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.

SII (2022a). *Employer-Employee Administrative Data from Declaraciones Juradas N°1879 and N°1887*. https://www.sii.cl/destacados/renta/2024/propuesta_dj_1887.html and https://www.sii.cl/destacados/renta/2022/ayuda_1879.html.

SII (2022b). *Firm Ownership Administrative Data from Forms 4415 and 4416*. https://www.sii.cl/documentos/circulares/2007/circu31.htm.

SII (2022c). *Firm-to-firm electronic transaction Administrative Data from Factura Electronica*. https://www.sii.cl/servicios_online/1039-.html.

SII (2022d). *Firm-to-individual electronic transaction Administrative Data from Boleta Electronica*. https://www.sii.cl/servicios_online/3532-.html.

SII (2022e). *Forms 22 and 29*. https://www.sii.cl/servicios_online/1044-.html for Form 22 and https://www.sii.cl/servicios_online/1042-3264.html for Form 29.

SII (2022f). *Individual Income Declaration Forms*. https://www.sii.cl/servicios_online/1043-1518.html.

SP (2022). *AFP portfolio data*. https://www.spensiones.cl/apps/bdp/index.php, in the section "Carteras históricas de Inversión de los Fondos de Pensiones".

SRCI (2022). *Civil Registry Information*. https://www.registrocivil.cl/.

University of Chile, T. M. C. of the (2022). *National Socioeconomic Characterization Survey*. https://observatorio.ministeriodesarrollosocial.gob.cl/encuesta-casen-2022.

A Theory Appendix

In this appendix, we first provide detailed derivations of the key equations from Section 2. Then, we express the system of equations in matrix form, which we use for computational implementation.

A.1 Change in Welfare

Proposition 1.

$$d \ln \mathcal{Y}_c = d \ln \chi_c - \sum_{i \in \mathcal{N}} b_{ci} d \ln p_i. \tag{11}$$

Proof. Under homothetic preferences, $\frac{de_c(p,u)}{du} = e(p,1) \equiv P$. Moreover, by Shephard's lemma we have that $\frac{de_c(p,u)}{dp_i} = x_{ci}$, where x_{ci} is the consumption of good i by consumer c. Then, with a slight abuse of notation, by totally differentiating the expenditure function we have

$$\begin{aligned}
& e_c(p, \mathcal{Y}_c) = \chi_c \\
& \implies \frac{de_c(p,u)}{dp} dp + \frac{de_c(p,u)}{du} d\mathcal{Y}_c = d\chi_c \\
& \implies x_c dp + P d\mathcal{Y}_c = d\chi_c \\
& \implies x_c p d \ln p + P\mathcal{Y}_c d \ln \mathcal{Y}_c = \chi_c d \ln \chi_c \\
& \implies \sum_i x_{ci} p_i d \ln p_i + \chi_c d \ln \mathcal{Y}_c = \chi_c d \ln \chi_c \\
& \implies \sum_i b_{ci} d \ln p_i + d \ln \mathcal{Y}_c = d \ln \chi_c,
\end{aligned}$$

where the last line uses the definition $b_{ci} \equiv \frac{p_i x_{ci}}{\chi_c}$. The result follows.

□

A.2 Change in Prices

Proposition 2.

$$d \ln p_i = \sum_{j \in \mathcal{N};\ k \in \mathcal{N}, \mathcal{F}} \widetilde{\Psi}_{ij} \widetilde{\Omega}_{jk} d \ln \mu_{jk} + \sum_{f \in \mathcal{F}} \widetilde{\Psi}_{if} d \ln w_f. \tag{12}$$

Proof. Let's separate the mark-up and markdown component of the wedges as $\mu_{ij} = \mu_i \tau_{ij}$. By totally differentiating the CRS unit cost function and using Shephard's lemma, we can write

$$
\begin{aligned}
dc_i &= \sum_j \frac{dc_i}{d(\tau_{ij}p_j)} d(\tau_{ij}p_j) + \sum_f \frac{dc_i}{d(\tau_{if}w_f)} d(\tau_{if}w_f) \\
\implies dc_i &= \sum_j x_{ij} d(\tau_{ij}p_j) + \sum_f x_{if} d(\tau_{if}w_f) \\
\implies c_i d\ln c_i &= \sum_j x_{ij}\tau_{ij}p_j d(\ln \tau_{ij}p_j) + \sum_f x_{if}\tau_{if}w_f d(\ln \tau_{if}w_f) \\
\implies d\ln c_i &= \sum_j \tilde{\Omega}^p_{ij}(d\ln \tau_{ij} + d\ln p_j) + \sum_f \tilde{\Omega}^w_{if}(d\ln \tau_{if} + d\ln w_f),
\end{aligned}
$$

where $\tilde{\Omega}^p \equiv (\tilde{\Omega}_{ij})_{i\in\mathcal{N},j\in\mathcal{N}}$ and $\tilde{\Omega}^w \equiv (\tilde{\Omega}_{ij})_{i\in\mathcal{N},j\in\mathcal{F}}$. We also have $d\ln p_i = d\ln \mu_i + d\ln c_i$, and hence

$$
d\ln p_i = d\ln \mu_i + \sum_{j\in\mathcal{N}} \tilde{\Omega}^p_{ij}\left(d\ln \tau_{ij} + d\ln p_j\right) + \sum_{f\in F} \tilde{\Omega}^w_{if}(d\ln \tau_{if} + d\ln w_f),
$$

where $\tilde{\Omega}^p = (\tilde{\Omega}_{ij})_{i\in\mathcal{N},j\in\mathcal{N}}$ and $\tilde{\Omega}^w = (\tilde{\Omega}_{ij})_{i\in\mathcal{N},j\in\mathcal{F}}$

The first term captures the direct change in firm i's price due to the change in its markup, the second term captures the (direct and indirect) change induced by a change in goods prices faced by firm i, and the third term captures the (direct and indirect) change due to changes in factor prices faced by firm i.

Define $d\ln \tilde{U}_{ij} \equiv \tilde{\Omega}_{ij} d\ln \mu_i \tau_{ij}$ and $d\ln \tilde{U}_i \equiv \sum d\ln \tilde{U}_{ij}$. Using $\sum_{j\in\mathcal{N},\mathcal{F}} \tilde{\Omega}_{ij} = 1$ for all $i \in \mathcal{N}$, we can rewrite the above equation in the following matrix notation as

$$
d\ln p = d\ln \tilde{U} + \tilde{\Omega}^p d\ln p + \tilde{\Omega}^w d\ln w.
$$

Simplifying we have the desired result:

$$
d\ln p = \tilde{\Psi}^p d\ln \tilde{U} + \tilde{\Psi}^w d\ln w. \tag{13}
$$

□

A.3 Changes in Profits

Proposition 3.

$$
d\pi_i = \left(\frac{\pi_i + T_i}{\lambda_i}\right) d\lambda_i + \lambda_i \sum_{j\in\mathcal{F},\mathcal{N}} \Omega_{ij}(d\ln \mu_{ij} - d\ln \widetilde{\Omega}_{ij}) - dT_i, \tag{14}
$$

Proof. Firm profits are given by

$$\pi_i = p_i y_i - \sum_{j\in\mathcal{N},\mathcal{F}} p_{ij}x_{ij} - T_i = p_i y_i\left(1 - \sum_{j\in\mathcal{N},\mathcal{F}} \frac{p_{ij}x_{ij}}{p_i y_i}\right) - T_i.$$

Since total expenditure is the numeraire, we can express the above equation in terms of sales shares λ_i. Using $\Omega_{ij} = \frac{p_{ij}x_{ij}}{p_i y_i}$, we have

$$\pi_i = \lambda_i \sum_{j\in\mathcal{N},\mathcal{F}} (1-\Omega_{ij}) - T_i.$$

Totally differentiating this expression we obtain

$$d\pi_i = d\left(\lambda_i \sum_{j\in\mathcal{N},\mathcal{F}} (1-\Omega_{ij})\right) - dT_i = \lambda_i d\sum_{j\in\mathcal{N},\mathcal{F}} (1-\Omega_{ij}) + d\lambda_i \sum_{j\in\mathcal{N},\mathcal{F}} (1-\Omega_{ij}) - dT_i.$$

Finally, we arrive at the desired equation by substituting $d\sum_j(1-\Omega_{ij}) = -\sum_j d\Omega_{ij} = -\sum_j \Omega_{ij} d\ln\Omega_{ij} = \sum_j \Omega_{ij}(d\ln\tilde{\Omega}_{ij} - d\ln(\mu_{ij}))$ and $\pi_i = \lambda_i\sum_{j\in\mathcal{N},\mathcal{F}}(1-\Omega_{ij}) - T_i$, or hence $\sum_{j\in\mathcal{N},\mathcal{F}}(1-\Omega_{ij}) = \frac{\pi_i + T_i}{\lambda_i}$ in the expression above.

□

A.4 Changes in Sales and Factor Income Shares

Proposition 4. *For $i \in \mathcal{N},\mathcal{F}$*

$$\begin{aligned} d\lambda_i = &- \sum_{l\in\mathcal{N};m\in\mathcal{N}} \lambda_l \Omega_{lm} d\ln\mu_{lm}\Psi_{mi} \\ &+ \sum_{k\in\mathcal{C},\mathcal{N}} \lambda_k Cov_{\widetilde{\Omega}^{(k)}}(d\ln\widetilde{\Omega}^{(k)}, diag(\mu^{(k)})^{-1}\Psi_{(i)}) + \sum_{c\in\mathcal{C}} d\chi_c \sum_{k\in\mathcal{N}} b_{ck}\Psi_{ki}. \end{aligned} \tag{15}$$

Proof. First note that, for $i \in \mathcal{N}$, we have

$$\lambda_i = \sum_{c\in\mathcal{C}} \chi_c b_{ci} + \sum_{j\in\mathcal{N}} \lambda_j \Omega_{ji}.$$

Using this equation, we have

$$\lambda_i = \sum_{c\in\mathcal{C},j\in\mathcal{N}} \chi_c b_{cj} \Psi^p_{ji}$$

where $\Psi^p \equiv (I-\Omega^p)^{-1}$. For $f\in\mathcal{F}$, we have

$$\lambda_f = \sum_{i\in\mathcal{N}} \lambda_i\Omega_{if} = \sum_{i\in\mathcal{N}}\sum_{c\in\mathcal{C},k\in\mathcal{N}} \chi_c b_{ck}\Psi^p_{ki}\Omega_{if} = \sum_{i\in\mathcal{N}}\sum_{c\in\mathcal{C}} \chi_c b_{ck}\Psi_{kf}$$

where the last equality uses

$$\Psi = \begin{bmatrix} \Psi^p_{N\times N} & \Psi^p_{N\times N}\Omega^w_{N\times F} \\ \mathbf{0} & I \end{bmatrix}$$

where, recall, $\Omega^w = (\Omega_{ij})_{i\in\mathcal{N};\, j\in\mathcal{F}}$. Therefore, for $i \in \mathcal{N}, \mathcal{F}$, we have

$$\lambda_i = \sum_{c\in\mathcal{C}, j\in\mathcal{N}} \chi_c b_{cj} \Psi_{ji}, \tag{16}$$

which implies

$$d\lambda_i = \sum_{c\in\mathcal{C}, j\in\mathcal{N}} \chi_c b_{cj} (d\Psi_{ji}) + \chi_c (db_{cj}) \Psi_{ji} + (d\chi_c) b_{cj} \Psi_{ji}. \tag{17}$$

The third term of this equation coincides with the third term of (15).

To further characterize the first term of (17), note that

$$\begin{aligned} & \Psi = (I-\Omega)^{-1} \\ \implies & \Psi(I-\Omega) = I \\ \implies & \Psi = I + \Psi\Omega \\ \implies & d\Psi = d\Psi\Omega + \Psi d\Omega \\ \implies & d\Psi(I-\Omega) = \Psi d\Omega \\ \implies & d\Psi = \Psi d\Omega\Psi. \end{aligned}$$

This implies that we can substitute the expression for $d\Psi_{ji}$ in (17) to obtain

$$\begin{aligned} \sum_{c\in\mathcal{C}, j\in\mathcal{N}} \chi_c b_{cj}(d\Psi_{ji}) &= \sum_{c\in\mathcal{C}, j\in\mathcal{N}} \sum_{l\in\mathcal{N}, m\in\mathcal{N}} \chi_c b_{cj} \Psi_{jl} d\Omega_{lm} \Psi_{mi} \\ &= \sum_{l\in\mathcal{N}, m\in\mathcal{N}} \lambda_l d\Omega_{lm} \Psi_{mi} \\ &= \sum_{l\in\mathcal{N}, m\in\mathcal{N}} \lambda_l \Omega_{lm} d\ln\Omega_{lm} \Psi_{mi} \\ &= \sum_{l\in\mathcal{N}, m\in\mathcal{N}} \lambda_l \Omega_{lm} \left(d\ln\tilde{\Omega}_{lm} - d\ln\mu_{lm}\right) \Psi_{mi} \\ &= -\sum_{l\in\mathcal{N}, m\in\mathcal{N}} \lambda_l \Omega_{lm} d\ln\mu_{lm} \Psi_{mi} + \sum_{l\in\mathcal{N}, m\in\mathcal{N}} \lambda_l \Omega_{lm} d\ln\tilde{\Omega}_{lm} \Psi_{mi}, \end{aligned} \tag{18}$$

where the second equality uses (16). Note that the first term of Equation (18) coincides with the first term of Equation (15).

Also note that

$$\sum_{c\in\mathcal{C},j\in\mathcal{N}} \chi_c(db_{cj})\Psi_{ji} = \sum_{c\in\mathcal{C},j\in\mathcal{N}} \chi_c b_{cj} d\ln b_{cj}\Psi_{ji} \tag{19}$$

It remains to demonstrate that

$$\sum_{l\in\mathcal{N},m\in\mathcal{N}} \lambda_l\Omega_{lm} d\ln\tilde{\Omega}_{lm}\Psi_{mi} + \sum_{c\in\mathcal{C},j\in\mathcal{N}} \chi_c b_{ci} d\ln b_{cj}\Psi_{ji} = \sum_{l\in\mathcal{N},\mathcal{C};\ m\in\mathcal{N}} \lambda_l\Omega_{lm} d\ln\tilde{\Omega}_{lm}\Psi_{mi}$$

coincides with the second term of (15). To see this, note that

$$\begin{aligned}
&\sum_{l\in\mathcal{N},\mathcal{C};\ m\in\mathcal{N}} \lambda_l\Omega_{lm} d\ln\tilde{\Omega}_{lm}\Psi_{mi} \\
&\quad= \sum_{k\in\mathcal{N},\mathcal{C};\ m\in\mathcal{N}} \lambda_k\tilde{\Omega}_{km} d\ln\tilde{\Omega}_{km}(\mu_{km})^{-1}\Psi_{mi} \\
&\quad= \sum_{k\in\mathcal{N},\mathcal{C}} \lambda_k(\sum_{m\in\mathcal{N}} \tilde{\Omega}_{km} d\ln\tilde{\Omega}_{km}(\mu_{km})^{-1}\Psi_{mi} \\
&\quad= \sum_{k\in\mathcal{N},\mathcal{C}} \lambda_k\left(\sum_{m\in\mathcal{N}} \tilde{\Omega}_{km} d\ln\tilde{\Omega}_{km}(\mu_{km})^{-1}\Psi_{mi} - \sum_{m\in\mathcal{N}} \tilde{\Omega}_{km} d\ln\tilde{\Omega}_{km} \sum_{m\in\mathcal{N}} \tilde{\Omega}_{km}(\mu_{km})^{-1}\Psi_{mi}\right) \\
&\quad= \sum_{k\in\mathcal{N},\mathcal{C}} \lambda_k Cov_{\tilde{\Omega}^{(k)}}(d\ln\tilde{\Omega}^{(k)}, diag(\mu^{(k)})^{-1}\Psi_{(i)})
\end{aligned}$$

where the third equality uses the fact that $\sum_m \tilde{\Omega}_{km} = 1$ implies $\sum_m d\tilde{\Omega}_{km} = \sum_m \tilde{\Omega}_{km} d\ln\tilde{\Omega}_{km} = 0$. □

A.5 Change in Cost Shares

Proposition 5.

$$d\ln\widetilde{\Omega}_{ij} = (1-\theta_i)\left(d\ln p_j - \sum_{k\in\mathcal{F},\mathcal{N}} \widetilde{\Omega}_{ik} d\ln p_k\right). \tag{20}$$

Proof. Under CES demand, the consumption share for buyer $i\in\mathcal{C}$ of good $j\in\mathcal{N},\mathcal{F}$ is given by

$$\tilde{\Omega}_{ij} = \frac{\beta_{ij}^{\theta_i} p_j^{1-\theta_i}}{\sum_{l\in\mathcal{N},\mathcal{F}} \beta_{il}^{\theta_i} p_l^{1-\theta_i}},$$

where β_{ij} is an arbitrary preference parameter. Taking logs and differentiating both sides we have

$$
\begin{aligned}
d\ln\tilde{\Omega}_{ij} &= (1-\theta_i)d\ln p_j - \frac{\sum_{k\in\mathcal{N},\mathcal{F}}\beta_{ik}^{\theta_i}(1-\theta_i)p_k^{1-\theta_i}d\ln p_k}{\sum_{l\in\mathcal{N},\mathcal{F}}\beta_{il}^{\theta_i}p_l^{1-\theta_i}} = \\
&= (1-\theta_i)(d\ln p_j - \sum_{k\in\mathcal{N},\mathcal{F}}\tilde{\Omega}_{ik}d\ln p_k),
\end{aligned}
$$

where in the second equality we have used the definition of $\tilde{\Omega}_{ik}$. □

A.6 System of Equations in Matrix Form

We first reduce the system of equations further in three sets of unknowns: $d\ln\lambda_i$, $d\ln\lambda_f$, and $d\ln\chi_c$. Substituting the expression for $d\ln\tilde{\Omega}$ into the expressions for $d\ln\lambda_i$ we obtain

$$
\begin{aligned}
d\ln\lambda_i = &- \sum_{\mathcal{N},\ m\in\mathcal{N}}\frac{\lambda_l\Omega_{lm}}{\lambda_i}d\ln\mu_{lm}\Psi_{mi} \\
&+ \sum_{\mathcal{N}}\frac{\lambda_k}{\lambda_i}(1-\theta_k)\mathrm{Cov}_{\tilde{\Omega}^{(k)}}(d\ln\mathbf{p}, diag(\mu^{(k)})^{-1}\Psi_{(i)}) + \sum_{c\in\mathcal{C}}\frac{\sum_{k\in\mathcal{N}}\chi_c b_{ck}\Psi_{ki}}{\lambda_i}d\ln\chi_c,
\end{aligned}
$$

Substituting the expression for T_c in χ_c, defining the deflated $\tilde{\Phi}_{ci} = \frac{\Phi_{ci}}{t_{ci}}$ for $i \in N+F$, and differentiating, we get

$$
d\chi_c = \sum_f \tilde{\Phi}_{cf}\ln\lambda_f + \sum_i \tilde{\Phi}_{ci}\pi_i d\ln\pi_i + \Phi_{cg}dT
$$

Define $\tilde{t}_{ij} = t_{ij}^{-1}$ for $j \in N$ and $\tilde{t}_{if} = t_{if}$ for $f \in F$. Substituting T_i in the expression for π_i and differentiating, we have

$$
d\ln\pi_i = \left(d\ln\lambda_i + \frac{\lambda_i}{\pi_i t_i^p}\sum_{j\in N+F}\Omega_{ij}\tilde{t}_{ij}(d\ln\mu_{ij} - d\ln\tilde{t}_{ij})\right)
$$

where we have used that $d\tilde{\Omega}_{ij} = 0$ for all $i \in \mathcal{N}$, which itself comes from having assumed Cobb-Douglas technology. The expression for dT comes from differentiating the following expression for T which comes from the governemnt budget balance:

$$T = \sum_{i \in \mathcal{N}} \underbrace{\lambda_i \left(1 - \frac{1}{t_i} + \sum_{j \in \mathcal{N},\mathcal{F}} (\tilde{t}_{ij} - 1)\Omega_{ij}\right) + \frac{(t_i^p - 1)}{t_i^p}\pi_i}_{T_i}$$

$$+ \sum_{c \in \mathcal{C}} (\Phi_{cf} - \tilde{\Phi}_{cf})\lambda_f + (\Phi_{ci} - \tilde{\Phi}_{ci})\pi_i$$

We denote the change in sales shares of firms, income shares of factors, and income shares of individuals with a $(N + F + C) \times 1$ dimensional vector $d \ln \boldsymbol{\lambda}$. We solve for the changes by finding a solution to the system of equations

$$d \ln \boldsymbol{\lambda} = A d \ln \boldsymbol{\lambda} + B$$

or

$$d \ln \boldsymbol{\lambda} = (I - A)^{-1} B,$$

where A is a square matrix of dimension $(N + F + C)^2$ and B is a vector of dimension $(N + F + C) \times 1$. Finally, we partition these matrices such that

$$\begin{bmatrix} d \ln \lambda_{N+F \times 1} \\ d \ln \chi_{C \times 1} \end{bmatrix} = \begin{bmatrix} \mathbf{0}_{N+F \times N} & A^{12}_{N+F \times N} & A^{13}_{N+F \times C} \\ A^{21}_{C \times N} & A^{22}_{C \times F} & \mathbf{0}_{C \times C} \end{bmatrix} \begin{bmatrix} d \ln \lambda^p_{N \times 1} \\ d \ln \lambda^w_{F \times 1} \\ d \ln \chi_{C \times 1} \end{bmatrix} + \begin{bmatrix} B^1_{N+F \times 1} \\ B^2_{C \times 1} \end{bmatrix}$$

where

$$A^{12}_{if} = \sum_{c \in \mathcal{C}} \frac{\chi_c}{\lambda_i} (1 - \theta_c) \mathrm{Cov}_{\tilde{\Omega}^{(c)}}(\tilde{\Psi}_{(f)}, \Psi_{(i)})$$

$$A^{13}_{ic} = \frac{\sum_{k \in \mathcal{N}} \chi_c b_{ck} \Psi_{ki}}{\lambda_i}$$

$$A^{21}_{ci} = \frac{\tilde{\Phi}_{ci}\pi_i + \Phi_{cg}\left(\lambda_i \sum_{j \in \mathcal{N},\mathcal{F}} (\tilde{t}_{ij} - 1)\Omega_{ij} + \pi_i((t_i^p - 1) + \sum_c (\Phi_{ci} - \tilde{\Phi}_{ci}))\right)}{\chi_c}$$

$$A^{22}_{cf} = \frac{\tilde{\Phi}_{cf}\lambda_f + \Phi_{cg}\left(\lambda_f \sum_c (\Phi_{cf} - \tilde{\Phi}_{cf})\right)}{\chi_c}$$

and

$$
\begin{aligned}
B_i^1 =& - \sum_{l \in \mathcal{C},\mathcal{N},\, m \in \mathcal{N}} \frac{\lambda_l \Omega_{lm}}{\lambda_i} d\ln \mu_{lm} \Psi_{mi} \\
&+ \sum_{c \in \mathcal{C}} \frac{\chi_c}{\lambda_i}(1-\theta_c)\mathrm{Cov}_{\tilde{\Omega}^{(c)}}(\tilde{\Psi}^p d\ln \tilde{U}, \Psi_{(i)}) \\
B_c^2 =& \sum_{i \in \mathcal{N}} \frac{\tilde{\Phi}_{ci}}{\chi_c} \frac{\lambda_i}{t_i^p} \sum_{j \in \mathcal{F},\mathcal{N}} \Omega_{ij} \tilde{t}_{ij} (d\ln \mu_{ij} - d\ln \tilde{t}_{ij}) \\
&+ \frac{\Phi_{cg} T^g}{\chi_c} d\ln T^{g,direct}
\end{aligned}
$$

where

$$
\begin{aligned}
dT^{g,direct} =& \sum_{i \in \mathcal{N}} \lambda_i \Big(\sum_{j \in N+F} \Omega_{ij}(\tilde{t}_{ij} d\ln \tilde{t}_{ij} - (\tilde{t}_{ij}-1) d\ln \mu_{ij}) \\
&+ \frac{(t_i^p - 1) + \sum_{c \in \mathcal{C}}(\phi_{ci} - \tilde{\phi}_{ci})}{t_i^p} \sum_{j \in N+F} \Omega_{ij} \tilde{t}_{ij} (d\ln \mu_{ij} - d\ln \tilde{t}_{ij}) \Big)
\end{aligned}
$$

A.7 Special Case: No Consumer Heterogeneity

When there is no heterogeneity across consumers—that is, when the economy is populated by a single representative household—the welfare effects of marginal changes in wedges admit a particularly simple characterization. In this case, the change in aggregate welfare is given by

$$
d\ln \mathcal{Y} = - \sum_{i \in \mathcal{N}} b_{ci} d\ln p_i. \tag{21}
$$

Because the household's total income, which equals total expenditure, is normalized to unity, the welfare change depends solely on changes in the prices of goods produced by firms. The expression for $d\ln p_i$ is the same as in the general case: in addition to changes in wedges, it depends on changes in factor prices, which are determined by the following system of equations:

$$
\begin{aligned}
d\ln \lambda^w =& - \sum_{l \in \mathcal{C},\mathcal{N},\, m \in \mathcal{N}} \frac{\lambda_l \Omega_{lm}}{\lambda_f} d\ln \mu_{lm} \Psi_{mf} \\
&+ \frac{1}{\lambda_i}(1-\theta_c)\, \mathrm{Cov}_{\tilde{\Omega}^{(c)}}\Big(\tilde{\Psi}^p \, d\ln \tilde{U} + \tilde{\Psi}^w \, d\ln \lambda^w ,\ \Psi_{(f)} \Big) .
\end{aligned} \tag{22}
$$

In the representative-agent setting, there is no demand-side feedback operating through changes in the distribution of total factor income across households, simplifying the mapping from change in wedges to change in welfare.

A.8 Financial Frictions

Consider an economy with two firms and a single representative consumer. Suppose financial frictions are the sole source of misallocation. Both firms combine labor and capital via a Cobb–Douglas production technology to produce output for final consumption:

$$y_i = l_i^{\alpha} k_i^{1-\alpha}, \quad i \in \{1, 2\}.$$

Firm i maximizes

$$\pi_i = p_i l_i^{\alpha} k_i^{1-\alpha} - w l_i - r k_i,$$

where p_i is the output price faced by firm i. Firm 2 is subject to a financial constraint $k_2 \leq \bar{k}$. Profit maximization for firm 1 yields

$$l_1 = \alpha \frac{p_1 y_1}{w}, \quad k_1 = (1-\alpha) \frac{p_1 y_1}{r}. \tag{23}$$

Here, firm 1's scale is demand-determined. Assume $\bar{k}$ binds for firm 2, so that

$$k_2 = \bar{k}, \tag{24}$$

$$l_2 = \alpha \frac{p_2 y_2}{w} = \left(\frac{\alpha p_2}{w} \right)^{\frac{1}{1-\alpha}} \bar{k}. \tag{25}$$

The representative household is endowed with one unit each of labor and capital. Preferences are CES with elasticity of substitution θ. The household's income $w + r + \pi_1 + \pi_2$ is normalized to one. Under this normalization, total demand for each good is

$$p_1 y_1 = \lambda_1 = \frac{\beta p_1^{1-\theta}}{\beta p_1^{1-\theta} + (1-\beta) p_2^{1-\theta}}, \tag{26}$$

$$p_2 y_2 = \lambda_2 = \frac{(1-\beta) p_2^{1-\theta}}{\beta p_1^{1-\theta} + (1-\beta) p_2^{1-\theta}}, \tag{27}$$

where β and $1-\beta$ are CES preference weights. The equilibrium for the constrained economy is then

$$l_1 = \frac{\beta}{\beta + (1-\beta) x^{-(1-\theta)}}, \quad l_2 = \frac{1-\beta}{1-\beta+\beta x^{(1-\theta)}}, \quad w = \alpha,$$

$$k_1 = 1 - \bar{k}, \quad k_2 = \bar{k}, \quad r = \frac{1-\alpha}{1-\bar{k}} \cdot \frac{\beta}{\beta + (1-\beta) x^{-(1-\theta)}},$$

where $x \equiv p_1/p_2$ satisfies

$$x = \left(\frac{\beta \bar{k}}{(1-\beta)(1-\bar{k})} \right)^{\frac{1-\alpha}{1-(1-\alpha)(1-\theta)}}.$$

Letting $(\cdot)^*$ denote allocations in the unconstrained economy, firm profits are

$$\pi_1 = 0, \tag{28}$$

$$\pi_2 = \underbrace{\frac{1-\beta}{1-\beta+\beta x^{(1-\theta)}}}_{\lambda_2<\lambda_2^*=1-\beta} \left[1-\alpha-(1-\alpha)\underbrace{\left(\frac{\beta\bar{k}}{(1-\beta)(1-\bar{k})}\right)^{\frac{1}{1-(1-\alpha)(1-\theta)}}}_{1/\mu_{2k}<1}\right]. \tag{29}$$

In the unconstrained economy:

$$l_1^* = k_1^* = \beta, \quad l_2^* = k_2^* = 1-\beta, \quad w^* = \alpha, \quad r^* = 1-\alpha.$$

The same allocative effects would obtain if we only observed $(\Omega_{if})_{i=1,2;\, f=l,k}$, $(b_i)_{i=1,2}$, and the parameters α and θ, and removed the implied wedges—without knowing their source.

Complete removal of the financial constraint can *reduce* firm 2's profits. While financial frictions directly enter its profit maximization problem, the firm optimizes given prices and demand, not internalizing the general equilibrium feedback on those variables. Figure 15 , solid line, plots π_2 against $\frac{\bar{k}}{1-\beta}$ for a particular parametrization. Profits are maximized at $\bar{k} = \tilde{k}$. If $\bar{k} < \tilde{k}$, relaxing the constraint raises π_2 on the margin; if $\bar{k} > \tilde{k}$, it lowers profits.

This conclusion does not require knowing the source of wedges. Even when wedges generate profits, a marginal reduction in wedges can raise profits. In this example:

$$\begin{aligned} &\mu_{1l} = 1, \quad \mu_{2l} = 1, \quad \mu_{1k} = 1, \\ &\mu_{2k} > 1 \quad \left[= \left(\frac{\beta\bar{k}}{(1-\beta)(1-\bar{k})}\right)^{-\frac{1}{1-(1-\alpha)(1-\theta)}}\right]. \end{aligned}$$

We now derive how equilibrium objects respond to a marginal change in wedges,

$$\begin{aligned} d\ln\pi_2 &= d\ln\lambda_2 + \frac{\lambda_2}{\pi_2}\Omega_{2k}\, d\ln\mu_{2k}, \\ d\ln\lambda_2 &= d\ln b_2 = (1-\theta)b_1\left(d\ln p_2 - d\ln p_1\right). \end{aligned}$$

Using $d\ln p_1 = \alpha d\ln\lambda_l + (1-\alpha)d\ln\lambda_k$ and $d\ln p_2 = (1-\alpha)d\ln\mu_{2k} + \alpha d\ln\lambda_l + (1-\alpha)d\ln\lambda_k$, it follows that

$$d\ln\lambda_2 = (1-\theta)b_1(1-\alpha)\, d\ln\mu_{2k}.$$

Substituting back,

$$\frac{d\ln\pi_2}{d\ln\mu_{2k}} = (1-\theta)b_1(1-\alpha) + \frac{\lambda_2}{\pi_2}\Omega_{2k} = \lambda_2(1-\alpha)\left[(1-\theta)\frac{1-\lambda_2}{\lambda_2} + \frac{1}{\pi_2\mu_{2k}}\right].$$

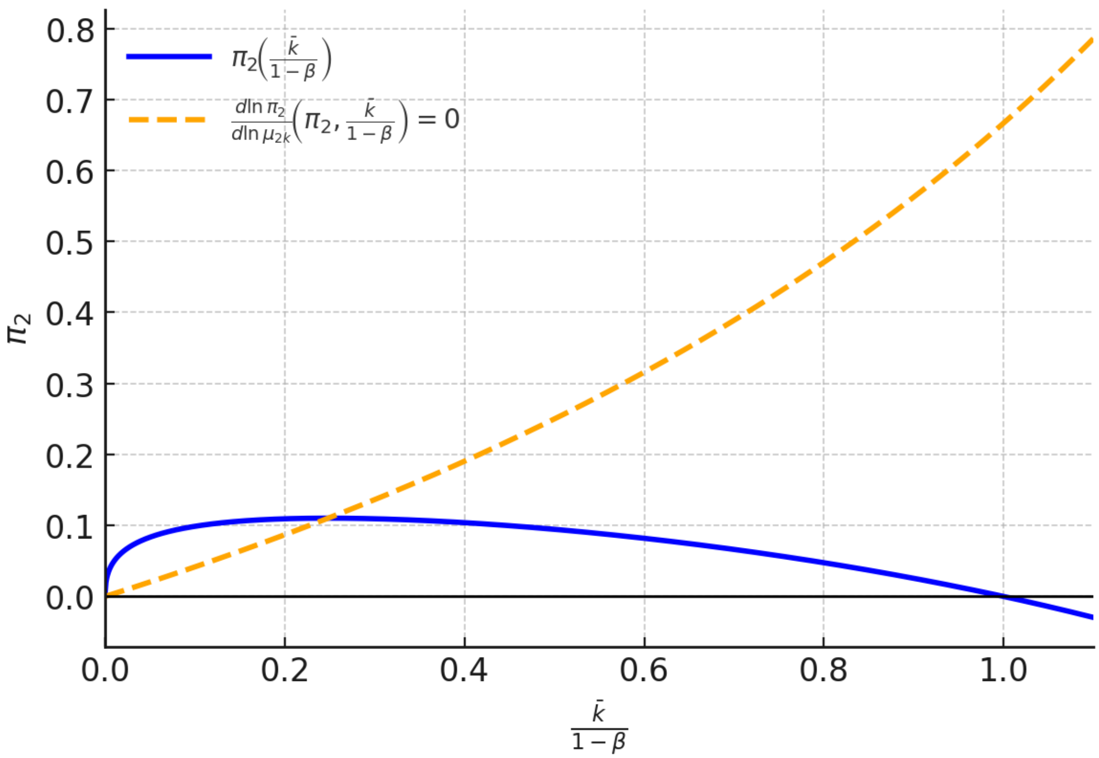

Figure 15: Firm profits in an economy with financial constraints and distortionary wedges. Parameter values: $\theta = 2$, $\beta = 0.6$, $\alpha = 0.3$.

As shown in Figure 15 (dashed line), this derivative equals zero precisely at $\bar{k} = \tilde{k}$, confirming that if the baseline economy is highly constrained, reducing wedges increases firm 2's profits on the margin—even when wedges are their ultimate source.

B Data Appendix: Raw Datasets

In this appendix we introduce the many data sources we use and describe the key variables in these data sources. Except where otherwise noted, we utilize data from the year 2022, the latest year for which all the administrative data sources are accessible.

The raw datasets we used can be divided into three types: administrative, survey and other auxiliary datasets. We describe each in turn.

B.1 Administrative datasets

We use administrative datasets from Chile's Servicio de Impuestos Internos (SII), the equivalent of their Internal Revenue Service (henceforth, IRS) and the Central Bank of Chile. These datasets cover the entire formal private sector in Chile. All individuals and all formal firms in Chile are assigned a unique tax ID that is consistently recorded across the datasets. We use this tax ID to merge all datasets.

- Domestic Firm-to-Firm Electronic Transaction Records: This dataset records all domestic transactions made between two Chilean firms and is gathered through the Factura Electronica system of the SII (SII 2022c). It is at the buyer-origin-seller-destination-product-transaction level, and for each firm-location-product pair (firms may have several production sites), the dataset reports the value of the transaction from origin to destination, month and year of transaction, the price and the non discounted price. Data is available from 2014 onwards (with only the largest firms reporting data prior to 2017). There is no reporting threshold at either the firm level or the transaction level. To match the (firm-specific and non-harmonized) product descriptions in each transaction to product classifications we use the machine learning algorithm developed in Acevedo et al. 2025. Three product classifications are available, all produced by the Central Bank of Chile using this machine learning algorithm: COICOP (Classification of Individual Consumption by Purpose), CPC (Central Product Classification) and CUP (*Clasificación Uniforme de Productos*). CPC is the most disaggregated with 2,563 products. CUP is a Chilean version of CPC, which is used for national accounts. Finally, COICOP has 1,187 products. Although it is less disaggregated, it allows for matching to the consumption survey described below.

- Customs Firm-to-Firm Transaction Records: This dataset reports international trade transactions at the domestic firm-country-foreign firm-product-transaction level. For each firm, it reports the total value of imports and exports as well as quantities, year of transaction, HS product classification, origin/destination country and foreign partner. Products are identified by 6-digit Harmonized Standard (HS) system codes, which describes 5,613 product categories.

- Firm-to-Individual Electronic Transaction Records: This dataset covers transactions between firms and individuals from 2022 to the present, and it is gathered through the Boletas Electronicas system of the SII (SII 2022d). Whenever a consumer makes a final good purchase at a formal firm, the firm electronically submits the itemized receipt to the tax authorities with the electronic system requesting the tax ID of the consumer. The tax id is mandatory for online orders as well as certain big ticket items. For other items, while the tax ID is not mandatory it is commonly requested by sellers and provided by consumers. Many retailers link the tax ID to their loyalty programs and the tax ID also facilitates returns, and the tax ID is not considered confidential by consumers. Due to the size of the dataset, the Central Bank of Chile currently only has access to transaction records for the 2,124 largest retailers, predominantly large retail chains. For these firms, we know whether a transaction has a tax ID attached. As seen in Figure 1 of the main text, 98% of transactions at these larger retailers report a tax ID. This high coverage ratio is consistent with anecdotal evidence that chain stores aggressively enroll customers into their loyalty programs (with the major reward programs covering multiple brands), and instruct cashiers to collect tax IDs on every transaction. For each individual, product and firm, we observe the total value of transactions, number of transactions and the average price. There are no reporting thresholds. Both CPC and COICOP product classifications are available, and as above, they are matched to transactions using the algorithm developed in Acevedo et al. 2025. Coverage is incomplete for two two types of transaction. First, we do not observe purchases from informal retailers. Second, we do not observe transaction level purchases linked to tax IDs for purchases from smaller retail stores whose data have not been transmitted to the Central Back of Chile (some of these smaller firms may collect tax IDs from consumers although we many likely do not based on our visits to small stores in Chile). In lieu of transaction level data from these smaller formal stores, the IRS provides the Central Bank with aggregate sales at the firm level. Figure 1 provides the breakdown of formal expenditure covered by tax ID-identified data. We address both of these data issues by using consumption surveys as described below.

- Employer-Employee Records: This data is gathered by the Chilean tax authority, Servicio de Impuestos Internos (SII), through affidavits 1887 and 1879 (SII 2022a). This dataset captures details about each job held by a worker within a specific year and it is at the individual-firm-year level. It has information on the amount earned by the employee in that year and dummies showing in which months the individual worked. Affidavit 1887 reports formal contracts, while affidavit 1879 reports invoices of labor services.

- Firm Ownership Records: We use ownership linkages of firms from tax records, which are collected through IRS tax affidavits 4415 and 4416 collected by the SII (SII 2022b). This dataset includes, for each firm, the complete list of owners, with both firms and individuals being able to own a firm. It is at the owner-firm-year level. It contains a dummy showing if the owner is an individual or a firm, the initial investment, the ownership share, the share of the firm's capital, and a dummy indicating if the invested amount has changed. Shares are displayed as a value between zero and one hundred. If the total shares of a given firm do not sum up to one, including the shares owned by pension funds discussed next, we assign the missing shares to the firm itself, ensuring that all firm's ownership shares sum up to one.

- Individual Pension Holdings: Chile has a mandatory private pension system which requires formal workers to hold pensions with one of seven private pension fund administrators, known as AFPs. Each AFP must offer five different funds that vary in their level of risk exposure. The portfolios of each of these 35 funds are publicly available at the Superintendencia de Pensiones website (SP 2022), detailing every Chilean firm in which the funds invest. The pension holdings dataset reports at the individual-pension type-year level the accumulated pension holdings with the individual identified by their tax id.

- Individual Income Declaration Forms and Government Transfers: We make use of government-to-household data by merging the universe of money transfer records with another dataset on income tax payments, enabling the creation of direct net transfers. This data is collected by the SII (SII 2022f). The income tax dataset is at the individual-year level. For each individual, it reports his or her annual income declaration: labor, capital and total income. However, not all income is reported since there is a threshold on the minimal earned income that requires declaration. Below this threshold, people do not appear in this form. The government transfer form is at the individual-year-type of transfer level. It reports the amount of all the different monetary transfers the government provides to individuals and describes the purpose of the transfer.

- Civil Registry Form: It reports for each individual, the date of birth, the gender, the tax ID of the mother and the tax ID of the father, and it is gathered by Servicio de Registro Civil e Identificación (SRCI 2022). This allows us to build demographics of individuals such as age and gender as well as to link together household members into household units.
- Firms Form 22 and 29: We measure additional firms' characteristics with tax forms 22 and 29, both gathered by the SII (SII 2022e). Form 22 is at the firm-year level. This dataset reports the value of declared capital, assets, and liabilities a firm has in a particular year. Form 29 is at the firm-month level. This dataset contains firms' sector (a 5 digit number with 674 sectors), total sales, total exports, total materials, total imports and investment.
- Unemployment Insurance Dataset: Contains information on the education level of all formal workers in the economy hired by private firms.

B.2 Survey data

While 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 middle-income countries such as Chile. Thus, we use three large-scale government surveys that capture informal transactions to complement administrative data (as well as to shed light on the products purchased at smaller stores for which we do not have Firm-to-Individual Electronic Transaction Records).

- Consumption Survey: We use the *Encuesta de Presupuestos Familiares* or EPF for its acronym in Spanish (INE 2022a). This is a consumer survey run by the Chilean statistical agency (*Instituto Nacional de Estadistica* or INE) and is a key input for measuring inflation and poverty. Each observation is at the individual-transaction level. For each individual, the survey reports the full list of consumption expenses (categorized into 1,187 products using the COICOP product classification) during the survey recall period, including the associated prices, quantities, product description, and the precise store it was purchased from. Store brands are further classified into store types. Consumption of housing is also calculated, either directly via rent payments or through owner imputed rent. The survey is conducted every 5 years.

We use the waves VIII from 2017 and IX from 2022. In 2022, there are 15,134 households surveyed which includes individual-level survey covering 44,680 individuals 15 and over. In 2017, 48,308 individuals were surveyed in 15,239 households. The survey also has modules recording, among other things, incomes from both formal and informal employment, transfers, pension fund participation, and demographic information all at the individual level.

- Employment Survey: We use a Chilean labor force survey that reports labor market activity for a sample of representative workers in the economy, the *Encuesta Suplementaria de Ingresos* or ESI for its acronym in Spanish (INE 2022b). This is a survey run once a year and is conducted by the INE. This survey complements the monthly employment survey INE runs, providing richer information about the sources of income of each individual, both formal and informal, as well as demographics.
- Micro-enterprise Survey: We use the survey *Encuesta de Microemprendimiento* or EME for its acronym in Spanish (INE 2022d). The EME surveys small firms and is designed by the Chilean government to measure the informal sector. We use the waves from 2017, 2019 and 2022, which in total survey over 20,000 small business owners. Each business owner reports detailed characteristics of their business including whether they are formal, their labor expenses, characteristics of the workers employed, materials expenses, and capital expenses. This allows us to measure wedges of these firms, their profits, and their labor relationships.
- Socioeconomic Survey: We use the survey *Encuesta de Caracterización Socioeconómica Nacional* or CASEN for its acronym in Spanish (University of Chile 2022). CASEN surveys people residing in private homes located in urban and rural areas of the country's 16 regions. CASEN provides detailed information on household demographics, income sources, labor market activities, education, health coverage, and a comprehensive set of government transfers received by individuals and households. We use CASEN to construct measures of individual-level transfer receipts, complemented by fiscal transfer data on the municipality level.

B.3 Other data sources

We use several auxiliary datasets that provide aggregates rather than firm- or individual-specific data.

- Price Deflators: We include in the analysis several price deflators provided by the Central Bank of Chile (Chile 2022). First, we use the Capital Gross Price Index for each year since 2013 to deflate the firms' capital expenditure and stock. Second, we use the CPI time series which goes back to 1980 and reports the price index, including sub indexes for major product categories.

- Sectors Description: Dataset produced by the Central Bank of Chile containing sector classification of firms (INE 2012). Several levels are available: sector (674 categories), subclase (485 categories), division (88 categories), seccion (21 categories) and 1-digit sector (9 categories)[22].

- List of Municipalities: Dataset provided by the Central Bank of Chile reporting the GPS coordinates, population, density, region (16 categories) and province (55 categories) for the 345 municipalities of Chile (INE 2022c). These are called "municipalidades" in Chile. They perfectly overlap with communas, the smallest administrative division, except for the case of two communas—Cabo de Hornos and Antártica—that fall under a single municipality government.

- WIOD Socio Economic Account: University of Groningen provides nominal and real sectoral value added, gross output, intermediate and final expenditure from the year 2000 to 2014 for 43 countries (Groningen 2014b). We use 2014 data from this dataset to measure cost shares from the US and to populate the rest of the world.

- US National IO Table: University of Groningen releases national Input Output Tables (NIOT) for 43 countries (Groningen 2014a). These data describe input and output linkages and transaction values from 2000 until 2014. The values are denoted in million of dollars. We use 2014 data from this dataset to measure cost shares from the US and to populate the rest of the world.

- World KLEMS Dataset: Measures output, inputs and productivity at a detailed industry level across countries, including Chile (KLEMS n.d.). To assist in our capital measurement, we use information on capital disaggregation into equipment versus structures across 1-digit sectors in Chile and also measures of capital depreciation across sectors in Chile.

- US input-output table from the Bureau of Economic Analysis (BEA): This is an input-output table released by the BEA with information on 402 industries in the US BEA 2017. We use 2017 data to compute US cost shares with higher granularity when compared to the NIOTs.

[22] 1-digit sectors were constructed from the 21 ISIC sectors. Some ISIC sectors were aggregated into a single sector due to small sample

- BEA/BLS Integrated Industry-level Production Account (KLEMS): This contains data on capital expenses for 63 sectors in the US BEA/BLS 2017. We use 2017 data with the IO tables of the previous bullet point to compute US cost shares with higher granularity when compared to the NIOTs.

- Chile's National Municipal Information System (SINIM): This is an official platform that reports standardized annual fiscal and operational information for all Chilean municipalities. We use 2022 data to construct per-capita public education and health transfers and create corresponding public service firms in any municipality where no such firm exists.

C Data Appendix: Cleaning, Combining and Transforming the Data

In this section, we document the cleaning processes we carry out on the raw datasets above, as well as how we combine and transform various datasets to obtain the exposure matrices that play a central role in the incidence formula derived in the theoretical sections of the paper.

C.1 Preliminary steps

C.1.1 Censoring outliers

To reduce the prevalence of outliers, we censor the labor payments of the top 0.01% of employer-employee payments. Firm-to-firm and firm-to-individual transactions are censored both at the upper and lower 0.01% of the distribution of flows.

C.1.2 Identify capital transactions

To separate consumption from investment, we need to distinguish between durable and non-durable transactions. In the CUP product descriptions file, 290 products are flagged as durable. For COICOP products, 34 of the 303 products are classified as durable in the product classification file.

C.1.3 The rest of the world

Since we allow for international trade, we need to keep track of transactions, both imports and exports, between Chileans and the rest of the world. We create a representative individual living in the rest of the world and a representative firm operating in the rest of the world. We have two sources of data for these agents. The first one is international transactions with Chilean firms. Whenever we observe a Chilean firm exporting in the customs data, we interpret this as sales from the Chilean firm to the rest of the world individual. Similarly, any Chilean firm import, conditional on the product being non-durable from the customs data, is interpreted as a material purchase of the Chilean firm from the rest of the world firm. This international output is assumed to be produced with a technology that directly transforms labor into output, hence the international firm does not generate any profit from this sale. The second source of international data is the WIOD. From this dataset, we obtain the rest of the world consumption exclusive of Chilean exports, as well as labor expenses, capital stock and sales to all countries other than Chile.

C.1.4 Factor markets

We define a labor market as a region of Chile (recall there are 16 regions). We also define a residual labor market in Chile, to which we assign labor usage by firms when we do not know their region. Finally, we define one more labor market, corresponding to labor used by the rest of the world firm.

We define two capital markets. The first one is for every firm that is located in Chile, regardless of the region. We define a second capital market used by firms in the rest of the world.

C.1.5 Multiproduct firms

In our incidence formulas, all firms are assumed to be single product firms. We turn multi-product firms into single product firms as follows.

For all firms, we first compile a list of products that each firm sells to both other firms and households. For each product that a firm sells, we create a new entity—referred to as a "subfirm", which represents the combination of the original firm and each product they sell. We then allocate the firm's inputs to these subfirms. As wedges may differ for each output-input pair ij, allocating inputs (including capital and labor) proportionately to revenues would be consistent with a specific set of wedges $\mu_i \tau_{ij}$ uninformed by the data. Instead we make the assumption that $\mu_i \tau_{ij}$ is equal for all subfirms i. That is, the product of output and input wedges is identical for all the different products a firm makes. This assumption implies that firms will allocate inputs based on the output elasticities $\eta_{i,j}$ for each pair ij, as they would if there were no distortions. Thus, we allocate inputs across subfirms proportionately to revenues multiplied by output elasticties, $\kappa_{ij} = \frac{y_i \eta_{i,j}}{\sum_{i \in \text{firm}} r_i \eta_{i,j}}$ where κ_{ij} is the share of input j purchased by the firm that is allocated to subfirm i and y_i are the revenues of each subfirm.

We amend this process slightly for wholesalers and retailers, whose primary business involves selling untransformed inputs. For these firms, we first allocate any materials bought from other firms that match the product codes of one of the newly created subfirms exclusively to that subfirm. In other words, we treat this as the stock of goods they sell untransformed to buyers. Inputs that do not match any of the product codes of goods sold by the firm are assumed to be shared among the subfirms proportionately to revenues multiplied by output elasticities in the same manner as above. For example, we would split a large department store into multiple retailers with any men's shoes purchased by the firm being allocated to their men shoe retailer subfirm and any unmatched inputs being spread across their subfirms.

C.2 Creating a firms' dataset

We first create a dataset that lists all firm-product pairs operating in 2022 that satisfy the following conditions:

- Employ at least two workers.
- Report some capital.
- Purchase material from other firms.
- Report some sales.
- Report their operating sector.

One key challenge in our data construction is to always identify both sides of any given any economic transaction. This subsection describes how we build all the relevant firm-specific information from bilateral administrative datasets.

C.2.1 Geography, sector, geography and product information

- Geography: A firms' main location comes from form 29 with firms assigned to one of 345 municipalities, 55 provinces and 16 regions. Municipalities are categorized by the SII as urban or rural, and assigned macroregions (either North, South or Santiago).
- Sectors: Using form 29 sector information, we assign to each firm a sector and a division, taking 674 and 88 possible values, respectively.
- Product: We obtain the product from the operating sector or sectors of the firm. As described above we convert multiproduct firms into single-product "subfirms".

C.2.2 Firms labor, capital, material, output

- Labor: Information on labor payments comes from the employer-employee records. We collapse the amount earned by each individual at the firm-year level and count the number of employees by firm in each year. Thus, we construct a dataset of total number of employees and total annual earnings paid by each firm.

- Capital: Information on firms' capital comes from forms 22, 29 and the capital gross price index dataset. Firms' capital stock is built using the perpetual inventory method. Using a depreciation rates from KLEMS for Chile, we assume capital has the law of motion: $K_{i,t} = (1 - \delta)K_{i,t-1} + I_{i,t}$, where $K_{i,t}$ is firm i's capital in year t, $I_{i,t}$ is i's investment in year t and δ is the depreciation rate. We measure the initial condition for capital, $K_{i,0}$, using the first report of fixed assets from form 22 that a firm reports over the 2005–2022 period. Investment is measured from form 29. We deflate capital and investment using the gross capital price index. Since the measure of investment is better at capturing capital expenses in machines and equipment, rather than structures and buildings, we adjust firms' capital by the ratio of structures to equipment by sector reported in KLEMS. After measuring the capital stock, we measure the user cost of capital as the risk-free nominal interest rate of 10% plus the sectoral depreciation rate measured from KLEMS.

- Materials: Material purchases are obtained by combining non-durable purchases from Chilean firms through the firm-to-firm electronic transactions records and from the rest of the world through the customs data.

- Output sales: There are four sources of output sales: annual sales from form 29, annual exports from customs data, sales to other firms from the firm-to-firm electronic transaction records, and sales to individuals from firm-to-individuals electronic records. We are thus able to observe output sold at a disaggregated level. Total sales of each firm are taken from form 29. For that reason, output which is not sold to firms, individuals or the foreign market, i.e. the difference between the form 29 total sales and total exports as well as intermediate sales to individuals and other firms, is allocated as output sold to a residual individual. We obtain for each firm the total sales value and a breakdown of these sales by buyer: other domestic firms, consumers, the rest of the world and the residual individual. In order to avoid the input-output matrix being singular, we do not consider output sales of firms for which we only observe transactions to themselves. This could happen if a firm has several sites and the only sales of one plant are to other plants in the same firm. Foreign firms' output equals final sales to the foreign individual, input sales to firms in Chile, and input sales to the foreign firm.

C.2.3 Residual firm

For measuring wedges, we restrict attention to firms for which we can measure their labor, capital, material, output sales, and 5 digit sector. We further restrict attention to firms with at least two employees to guard against tax IDs which are being used for consumption purposes. Omitting firms that do not meet these two sets of criteria breaks the circularity of the economy because there are sales for which we only have one side of the transaction, for example when a firm in our sample sells to one who does not possess any capital and so is omitted. To deal with this possibility, we create an additional firm, that we label the "residual firm", to which we attribute all sales and purchases from firms that we exclude from our analysis—ensuring no economic transactions are lost.

C.3 Creating an individuals' dataset

Next we construct a dataset of Chilean individuals present in the administrative data in 2022. As with the firms above, we first create a dataset of individuals which we have measures for the following key variables:

- Income (including informal income, capital ownership and profits).
- Consumption (including from informal stores).
- Age of at least 15.
- Gender.

We describe in this subsection how we obtain and measure formal income and consumption measured from tax records. In Section C.4 we describe how we attach informal income and consumption to individuals in the administrative data by statistically matching variables in the administrative data with those from household survey and employment surveys.

A prerequisite for measuring formal income is observing the ownership matrix of firms in the economy, which is needed to distribute capital payments and profits from firms to individuals.

C.3.1 Computing the ownership matrix

To allocate individuals their capital and profit income, we use the ownership records. For each firm, we know the tax id of every shareholder and the share of profits their ownership stake entitles them to. Some of these owners are other firms not individuals. Thus, we first create two matrices.

- Direct ownership: Rows are individuals, while columns are the firms they own.
- Indirect ownership: Rows are firms that own other firms, columns are the firms they own.

The full ownership matrix is the product of the direct ownership matrix and the Leontief inverse of the indirect ownership matrix. Therefore,trace all ownership shares back to individuals, including direct ownership or indirect ownership through another firm. There are three types of firm for which further adjustments are required:

- Residual firm: Its ownership shares are computed using individuals' ownership shares of the firms that make up this residual firm. We measure each firm's "contribution" to the aggregate residual firm by computing their profits relative to aggregate profits. Because these contribution weights sum up to one, we can multiply initial ownership shares by the weights to recover the residual firm's ownership shares.
- Rest of the world firm: We define a rest of the world individual that owns the entire rest of the world firm.
- Firms not appearing in ownership records: We allocate the ownership of these firms to the residual individual. Thus, all firms are owned entirely by individuals.

C.3.2 Individual's consumption, labor income, capital income, profit income

- Consumption: We define two types of consumption. The first is "direct" consumption. This is consumption of final goods that we capture in the firm-to-individuals electronic records and via the household surveys described below. The second type of consumption we call "indirect" consumption. This corresponds to the capital expenses from the firm-to-firm electronic transaction records. Since we do not allow for savings and investment, we treat capital expenses as consumption by individual owners of firms who purchase the capital and rent it to the firms they own. Thus, we rebate firms' expense on durable goods as consumption by individual owners using the full ownership matrix described above. Thus, total consumption of an individual is the sum of his or her final goods purchases from firms, and the capital that the firms they own purchase from other firms multiplied by his or her ownership share of the firm. As above, we make adjustments for the two special types of individual:

  – Rest of the world individual: Foreign individual consumption is equal to total domestic exports plus the rest of the world consumption.
  – Residual individual: When allocating each firm's output, we assume that any difference between a firm's annual output (as measured by their total sales reported in their tax filing) and their total sales to individuals, firms and foreigners was purchased by a ″residual individual″. This additional individual consumes the amount of output that ensures all firm output is accounted for.

- Labor income: From the employer-employee records, we compute annual labor earnings for each individual. We also define the rest of the world labor income to be the labor expenses from the rest of the world reported in the WIOD.
- Capital income: We calculate capital income accruing to individuals as their share (based on the full ownership matrix) of the user cost of the firm's capital calculated above. The user cost is the normal rate of return plus depreciation. The normal rate of return is assumed to be the nominal interest rate of the Central Bank of Chile at the beginning of 2022 of 4% plus the equity risk premium of 6%. The depreciation rate is obtained at the sectoral level from the KLEMS data for Chile.
- Profit income: We compute firms profits as sales minus costs using output, labor, capital, and materials measured in the various ways described above. Data on labor expenditures comes from employer-employee records, on materials expenditures it comes from firm-to-firm electronic transaction records, and on sales and capital expenditures it comes from Forms 22 and 29. We distribute these firm profits to individuals using the full ownership matrix.

C.3.3 Additional individuals

As with firms, we create additional individuals to ensure that the economy is circular and that standard aggregate identities hold. Here, we recap the additional individuals we create while measuring income and consumption:

- Rest of the world individual: A "foreign" individual who consumes all of domestic exports. They have labor income, which is all the revenue the foreign firm makes (i.e. total domestic imports). They also consume the rest of the world production that is not sold to Chilean firms.

- Residual individual: The residual individual consumes all firm output that remains after taking into account sales to other firms (domestic and foreign) and to individuals (domestic and foreign, with and without tax IDs). It has characteristics that match the averages of individuals for which we do not observe income, consumption, age or gender (for example, when calculating incidence by gender, we allocate some fraction of the residual individual to the female bin based on the share of dropped individuals that are female).

C.3.4 Other characteristics of individuals

- Age and gender: We obtain age and gender from the civil registry dataset. We drop individuals below age 15. Each individual is either classified as female or male and in an age category: 15-25, 25-35, 35-45, 45-55, 55-65 or 65+.
- Location: We record the location of each individual from the most common location from which they purchase final goods as reported in the firm-to-individual electronic transaction records.

C.3.5 Grouping individuals into households

For the subset of incidence calculations that we report at the level of households we use civil registry data to group individuals into households.

C.4 Statistical matching between administrative data and survey data

Despite the richness of the administrative tax data, there are two limitations. First, informality plays a role in a country like Chile, where in 2022 around 27% of workers are employed informally (National Statistics Agency or INE for its acronym in Spanish).[23] Second, the firm-to-individual transactions are not exhaustive because we cannot observe these transactions for the smaller formal firms that are not included in the extract of the tax authorities firm-to-individual database transferred to the Central Bank of Chile. Thus, we use several surveys of both individuals and firms to improve the measurement of income and consumption. Since these surveys do not report the tax IDs used in the administrative data, we use statistical matching techniques to associate survey responses to individuals in the administrative data that are similar along multiple dimensions.

[23] The informality rate in Chile can be found at https://stat.ine.cl/Index.aspx?DataSetCode=INF_TOI

C.4.1 Matching strategy

We use the consumption (EPF), employment (ESI), micro-enterprise (EME) and socioeconomic (CASEN) surveys to incorporate consumption at small and informal stores, labor income from working in informal firms, profit and capital income from the ownership of informal firms, and transfers from the government.

For each of the surveys, we first construct a distinct ID for each individual (or firm). Then, we group these individuals into several broad demographic and geographic bins ("bin variables") before applying statistical matching techniques within each bin. To carry out the match, we utilize multiple "matching variables" that are common in both the administrative sources and the surveys (with the specific variables used for each match detailed in the subsequent sections). An optimal transport algorithm minimizes the distance between the values of these variables for the survey respondents and individuals/firms in the administrative data with each tax ID in the administrative data assigned one survey ID (survey IDs are typically allocated to many tax IDs). Optimal transport algorithms have previously been used for building distributional national accounts (see Blanchet, Saez and Zucman, 2023).

The key benefit of statistical matching in our context is that, once the variables missing from the administrative data—the "target variables"—are brought across from the survey data, the resulting dataset preserves the joint distribution of the target variables in the survey (after correctly applying the survey weights). In addition, since each tax id is assigned an actual survey response, this matching technique avoids the extrapolation (and possibility of extreme outliers) that arises using polynomial-based prediction models.

After implementing the statistical matching approach, we check how well each match performs by comparing the original distribution of the target variable with the matched one. With the match assigned, we carry across the outcomes missing from the administrative data but present in the surveys to the administrative dataset described above.

To more accurately measure the public education and health transfers received by each eligible individual, we complement the CASEN survey data with administrative expenditure records from Chile's municipal information system (SINIM). For each municipality, we take the total executed spending on public education (health) and allocate it on a per-capita basis to all individuals in CASEN who reported attending public education (health) in that municipality.

In total, we build five bridges described in the subsections below: EME firms to administrative data, EME employees to administrative data, EME employers to ESI, ESI to administrative data, EPF to administrative data, and finally CASEN to administrative data. The variables used for creating bins, for matching and the target variables of each statistical matching procedure are summarized in Table 2.

C.4.2 Informal firms

In order to capture the informal side of the Chilean economy, we create informal firms. Unlike formal firms, informal ones do not pay taxes. Just like our formal firms, we assign them labor, capital, materials and output values and measure the linkages and transactions with other agents in the economy that are consistent with those firm-level values. Thus, the main challenge for us is to connect these new firms with the rest of the economy while ensuring economic identities hold and that the economy remains circular.

We start with the microenterprise survey, the EME, to obtain region, labor, capital, materials , output and 1-digit sector for informal firms directly from the EME survey. We duplicate firms in this sample using the attached survey weights that are designed to be representative of small informal firms in the Chilean economy. To provide finer detail on the sector and product than is present in the EME, we statistically match the informal firms to similar small formal firms in the same region, sector and firm size bin, minimizing the distance between their labor expenses and employment with optimal transport algorithms. We then expand the informal firms using survey weights in EME.

Next, we create linkages and transactions with other agents in the economy that are consistent with the labor, materials, capital and output values of our informal firms. For labor payments, we statistically match the workers in the EME to individuals in the administrative data and allocate informal firm's labor payments to these individuals. For the firm-to-firm transactions that underlie material expenses, given that we do not have information about suppliers of informal firms, we once again use the match between EME firms and small firms in the administrative data to determine how many suppliers informal firms are likely to have within each section-region-firm-size bin (based on the suppliers for similar small formal firms). We then pull those suppliers at random (without replacement) from the set of suppliers to formal firms in the administrative data, and allocate expenditure shares of them randomly while preserving each informal firm's expenditure share in each section-region-firm-size bin. Then, we append these firms to the formal firms' dataset and create unique firm IDs for these newly added firms. The ownership dataset that measures the ownership of firms is also updated accordingly, using the ownership information from the EME survey.

For firm-to-individual transactions that underlie values of final consumption output, we assume that informal firms are downstream, and all their output is sold to final consumers. As with suppliers, we use the EME firm to admin match to link informal firms with similar formal firms that only sell to final consumers and sample final consumers conditional on keeping the total number of consumers the same. We do this allocation considering the product and region allocation described above. That is, we draw the same number of consumers within product and regions as formal firms, and allocate revenues following the same shares of revenues formal firms have across consumers.

C.4.3 Public education and health firms

We include the role of public education and health services by creating firms that provide those services. Some municipalities already have tax IDs that represent those firms. For the municipalities that do not have those tax IDs, we create one public education (health) firm in the education (health) sector producing those services. To populate revenues and costs of those firms, we use information from SINIM, which provides total revenues, total labor expenses, and intermediate input expenses. Since it does not report the specific workers, specifics suppliers and capital expenditures, we statistically match these newly created firms with formal firms in the administrative data to determine the firms' capital expenditure and how many suppliers public education (health) firms are likely to have within each section-region-firm-size bin. Similar to informal firms, we then pull those suppliers at random (without replacement) from the set of suppliers to formal firms in the administrative data, and allocate expenditure shares of them randomly while preserving each education (health) firm's expenditure share in each section-region-firm-size bin. We do the same with links with workers. Then, we append these firms to the formal firms' dataset and create unique firm IDs for these newly added firms.

For municipalities that already have a tax ID reporting those services, we update their operational characteristics by allocating values in SINIM evenly to these firms.

C.4.4 Individual's income

Having expanded our set of firms to include informal firms, we adjust individuals' income in two ways. First, we redistribute factor payment from informal firms to their owners and workers. For informal labor income, we use information on the employer and employee from the statistical matching described in the previous subsection. For capital and profit income, we define ownership of the informal firms. To do so, we first statistically match EME employers (i.e. small firm owners) with their ESI counterparts, and then statistically match the ESI individuals with the administrative data. This is preferable to matching the EME employers directly to the administrative data as the ESI contains additional matching variables not present in the EME.

Second, we obtain information on a more complete set of government transfers than we have access to via SINIM and CASEN. We statistically match the CASEN with the administrative data. This allows us to identify the amount of individuals receiving government transfers as the CASEN has a detailed section of its survey on sources of transfers.

The resulting size of the matched transfers are smaller than those recorded in government expenditures from administrative datasets. To make identities hold, we proportionally increase transfers to individuals such that the government's budget is balanced. That is, we increase transfers to make the government's expenditures equal to tax revenues, which is the sum of taxes on inputs, taxes on factor income, and corporate profit taxes.

C.4.5 Individual's consumption

Our final objective is to build a complete firm-to-individual consumption matrix, drawing on both administrative and survey sources. The consumption data from the firm-to-individual electronic transaction records in the administrative data does not include informal consumption or consumption at small formal stores while EPF covers all consumption, formal or informal, including services and housing. In a similar manner to our treatment of individual income, we use statistical matching to bring the complete consumption data from the EPF survey into our administrative datasets. Individuals are first classified into demographic bins based on their macroregion, age group and income group. Age groups are drawn using 25, 45 and 65 as cutoffs. We divide individuals into five income groups based on quantiles computed within each gender-by-age-group cell. Individuals with zero income or zero consumption are grouped separately. Individuals within each macroregion $\times$ age bin $\times$ income bin in the survey are matched to the closest individual in the administrative that minimizes the distance between matching variables using optimal transport algorithms that preserve the distribution of the target outcomes (consumption in this case).

Matching variables include gender, total income, labor income share, profit income share and 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 rich consumption surveys which contain the store name for each purchase). In the EPF, we classify store types using two complementary approaches: first a keyword-based classification that uses store names (matching brand names for the major retailers in Chile and category names such as shoe shop, hardware store, hairdresser for smaller stores) and then a vector-based classification that compares the distribution of a store's product sales to the typical product mix observed within each sector for the stores not matched using keywords. Then, stores are matched to five store types: supermarkets, department stores, traditional stores, specialized stores and other.

Since EPF does not include detailed measures of public education and health transfers, we use CASEN and SINIM as additional sources. From CASEN, we identify which individuals consume public health and education services. We assume individuals consume those services locally in the municipality they live in. Given this, we take the total revenue from the firms that provide those services, and we allocate the revenue per individual consuming those services as transfers of education and health. We then, allocate those transfers as consumption by identifying the relevant products corresponding to these public services (e.g., elementary school service and secondary school service for public education) and allocate the total transfer value across these products using each individual's consumption shares as weights. For individuals with no recorded consumption in these products, we allocate the transfer value equally across the relevant products.

After the statistical matching is complete and allocating consumption from the additional public services transfer income, we have a full record of individual consumption by product and store type. Next, we build the firm-to-individual consumption matrix, i.e. we identify which firms each individual purchases from. For each region-product-firm type cell, we first compute residuals between matched consumption and observed administrative consumption. Then, we allocate this residual to firms in the corresponding region-product-firm type cell that have no-tax ID transactions. We impose that total expenditure at each formal firm equals their final consumption sales in the administrative records (recall that we know every formal store's final goods sales). Additionally, for computational feasibility in subsequent steps, we cap the number of firms each consumer buys from and the number of consumers each firm sells to. These caps are informed by observed patterns in the EPF survey and firm-to-individual administrative data. We use iterative proportional fitting to impose both restrictions. For region-product-firm type cells that only appear in the administrative data, we apportion them evenly to a random sample of consumers in the same region. For region-product-firm type cells that only appear in the consumption survey, they are classified as informal consumption and are allocated to informal firms created in Section C.4.2.

Given that we lack detailed product information for informal firms, similarly to how we treated firm-to-firm transactions, we match informal firms to small formal firms (those in the bottom 40 percent of the firm size distribution), as we do above, based on shared characteristics including sector and region. The matching variables are total sales, material expenses, labor expenses, employment, and capital expenses. We then apply the same procedures used for formal firms to allocate consumption across these informal firms and construct the complete firm-to-individual consumption matrix.

D Measurement and Computation Appendix

The previous two data appendices introduced our raw datasets and described how we transform and combine those datasets to obtain a fully circular dataset of Chilean individuals and firms, as well as a representative foreign country and individual. In this appendix

Table 2: Summary of Statistical Matching Approaches

Bridges	Target Variables	Common Variables
EME employer to ESI	Informal firms' profit	**Bin variables:** 1-digit sector, firm size, informality status **Matching variables:** macroregion, gender, education, occupation, total income
EME employee to Admin	Informal firms' labor income	**Bin variables:** gender, no income dummy **Matching variables:** age, total income
EME firm to Admin	Informal firms' number of suppliers, the distribution of their expenditures across suppliers; product and sector classification	**Bin variables:** 1-digit sector, region, firm size **Matching variables:** sales, material expenses, labor expenses, capital expenses, total employment
ESI to Admin	Informal income	**Bin variables:** macroregion, age bin, gender, household size, income, no income dummy **Matching variables:** number of household members within an age group, age
CASEN to Admin	Transfers from government	**Bin variables:** region, age, gender, household size, income, no income dummy **Matching variables:** number of household members within an age group, age, formal income
EPF to Admin	Detailed consumption	**Bin variables:** macroregion, age, income, no income dummy, no consumption dummy **Matching variables:** gender, consumption in supermarkets, income, labor and profit income share

Note: Bin variables are used to create strata prior to matching; matching variables are used for within-stratum matching. Only individuals/firms in the same bin will be matched.

we describe how those datasets are fed into the various exposure matrices that lie at the heart of our incidence formula, explain how we obtain or estimate the remaining objects (price elasticities and wedges), and how we navigate the computational challenges that arise when trying to solve a linear system featuring millions of firms and consumers.

Recall we need to solve a system of equations in an iterative manner. These equations depends on three types of object: (1) the full matrices of individual-level exposure to (static) distortions; (2) the initial size of each $\mu_i \tau_{ij}$ distortion; and (3) the elasticities determining how the economy responds to their reduction. We discuss these three sets of objects in turn.

D.1 Populating the exposure matrices

For convenience, we reproduce the key equations from the linear system.

$$d \ln p_i = \sum_{j \in \mathcal{N}} \widetilde{\Psi}_{ij} d \ln \mu_j + \sum_{k \in \mathcal{N}, \mathcal{F}} \widetilde{\Psi}_{ij} \widetilde{\Omega}_{jk} d \ln \tau_{jk} + \sum_{f \in \mathcal{F}} \widetilde{\Psi}_{if} d \ln w_f. \tag{30}$$

$$d\chi_c = \sum_{f \in \mathcal{F}} \Phi_{cf} L_f dw_f + \sum_{i \in \mathcal{N}} \Phi_{ci} d\pi_i + dT_c, \tag{31}$$

$$d\pi_i = \left(\frac{\pi_i + T_i}{\lambda_i} \right) d\lambda_i + \lambda_i \sum_{j \in, \mathcal{N}, \mathcal{F}} \Omega_{ij} (d \ln \mu_i + d \ln \tau_{ij} - d \ln \widetilde{\Omega}_{ij}) - dT_i, \tag{32}$$

$$\begin{aligned} d\lambda_i = &- \sum_{l \in \mathcal{C}, \mathcal{N}; m \in \mathcal{N}} \lambda_l \Omega_{lm} (d \ln \mu_l + d \ln \tau_{lm}) \Psi_{mi} \\ &+ \sum_{k \in \mathcal{C}, \mathcal{N}} \mu_k^{-1} \lambda_k \text{Cov}_{\widetilde{\Omega}^{(k)}} (d \ln \widetilde{\Omega}^{(k)}, diag(\tau^{(k)})^{-1} \Psi_{(i)}) + \sum_{c \in \mathcal{C}} d\chi_c \sum_{k \in \mathcal{N}} b_{ck} \Psi_{ki}, \end{aligned} \tag{33}$$

$$\begin{aligned} dT_i = &\sum_{j \in \mathcal{N}, \mathcal{F}} \lambda_i \Omega_{ij} dt_{ij} + (t_{ij} - 1) \Omega_{ij} \left(d\lambda_i - \lambda_i (d \ln \mu_i + d \ln \tau_{ij} - d \ln \tilde{\Omega}_{ij}) \right) \\ &+ \pi_i dt_i^p + (t_i^p - 1) d\pi_i \end{aligned} \tag{34}$$

$$dT_c = \Phi_{cg} \sum_{i \in \mathcal{N}} \left(dT_i + dv_i \sum_{c'} \chi_{c'} b_{c'i} + (v_i - 1) \sum_{c'} b_{c'i} \left(d\chi_{c'} + \chi_c (d \ln \tilde{\Omega}_{c'i} - d \ln v_i) \right) \right) \tag{35}$$

$$d\ln\widetilde{\Omega}_{ij} = (1-\theta_i)\left(d\ln p_j^I - \sum_{k\in\mathcal{F},\mathcal{N}} \widetilde{\Omega}_{ik} d\ln p_k^I\right). \tag{36}$$

Given elasticities θ_i and wedges $\mu_i\tau_{ij}$, the system of equations in (30)-(36) constitutes a linear system that can be solved for any given values of what we term the exposure matrices—the input share $\widetilde{\Omega}$ and Leontief-inverse Ψ matrices, and values of the consumption shares b_{ci}. The solution to this system is then used in each step (updating the allocation each time) of a simple iterative algorithm that provides an exact solution for arbitrary reductions in wedges.

We now describe the precise data objects that we use to construct the various exposure matrices.

Ψ_{ij} and $\widetilde{\Psi}_{ij}$: These are the standard and distortion-inclusive Leontief inverse matrices, respectively. They show how a shock to costs in sector j propagates through the supply chain, taking into account both its direct and indirect effects. $\Psi = (I-\Omega)^{-1}$ and $\widetilde{\Psi} = (I-\widetilde{\Omega})^{-1}$ are constructed using the input-cost share matrices without and with distortions, respectively. That is, $\Omega_{ij} = \frac{p_j^I x_{ij}}{p_i y_i}$ and $\widetilde{\Omega}_{ij} = \tau_{ij}\mu_i\Omega_{ij}$. Materials shares are measured using firm-to-firm electronic transaction records, labor shares are measured using employer-employee records, and capital shares are measured using Forms 22 and 29 and the capital gross price index dataset.

Ω_{ij} and $\widetilde{\Omega}_{ij}$: These are the input-cost share without and with distortions, respectively. As mentioned above, they are measured using firm-to-firm electronic transaction records, employer-employee data, Forms 22 and 29 and the capital gross price index dataset.

π_i: Firm-level after-tax profits are defined by $\pi_i = p_i y_i - \sum_{j\in\mathcal{F},\mathcal{N}} p_j^I x_{ij} - T_i$. As above, cost of materials are measured using firm-to-firm electronic transaction records, cost of labor using employer-employee records, cost of capital using Forms 22 and 29 and the capital gross price index dataset. The measurement of taxes T_i is explained below.

T_i: Total net taxes rebated by firm i to the government, $T_i = \sum_{j\in\mathcal{N},\mathcal{F}}(t_{ij}-1)p_j^I x_{ij} + (t_i^p - 1)\pi_i$. Corporate profit tax t_i^p is described below. Taxes on materials and factors t_{ij} are charged according to $t_{ji} = 1 + (z-1)\mathbf{1}_{j\in N_f,\, i=\text{labor}} + (\textit{vatRate}-1)\mathbf{1}_{j\in N_i,\, i\in N_f}$, where z is labor income tax and *vatRate* is the VAT tax v_i. That is, if the firm is hiring labor, it pays the labor income tax z, if the firm is buying materials from an informal firm the tax is 0, and if it is buying from a formal it pays the VAT tax.

λ_i: The Domar weights are the share of firm i's sales in total economy-wide consumption. Sales are measured as the maximum between sales data documented in Form 29 and the sales measured as the sum of sales to other firms from the firm-to-firm electronic transaction records, to consumers from the firm-to-individuals transaction data, and exports from the customs data.GDP is normalized to 1 so that the Domar weights are measured simply as $\lambda_i = p_i y_i$.

b_{ck}: This is the share of household c's expenditure spent on product/firm k, $\frac{p_i x_{ci}}{\chi_c}$. There are two sources of consumption, direct consumption on final goods and indirect consumption, the latter being capital expenses by firms owned by individuals, as explained in section C.3.2. Direct consumption is measured using the firm-to-individuals electronic records and the household surveys (direct consumption), and indirect consumption using the firm-to-firm transaction and ownership data.

t_i^p: Corporate profit tax t_i^p is at the firm-level, and is computed by dividing taxes paid by the taxable income from Form 22. Small and medium enterprises (SME) that satisfy some requirements, like a revenue ceiling, can enter a tax regime called Regimen Pro Pyme in which the profit tax rate is lower. Thus we separate firms in two groups, SMEs that operate under this regime and other firms that don't, which tend to be larger. The average tax rate for SMEs is 10%, and for large firms is 27%.

Φ_{cg}: The share of government net transfers allocated to household c. A description of how these are measured is contained in subsection C.4.4 on statistical matching to measure individual's income.

χ_c: This is total after-tax nominal income of household c, $\chi_c = \sum_{f \in \mathcal{F}} \Phi_{cf} w_f L_f + \sum_{i \in \mathcal{N}} \Phi_{ci} \pi_i + T_c$. This includes labor income (from employer–employee records), capital income (from ownership records and pension holdings), profits from ownership in firms, and net government transfers (see statistical matching for a description of this object).

v_i: This is the VAT tax rate, which is 19% in Chile.

θ_i: These are the elasticities of substitution across inputs in firm i's technology, or across products in household i's preferences. On the production side, we assume Cobb–Douglas technologies, hence the elasticity of substitution is 1. On the consumer side, we obtain the lower-nest estimates (within-sector elasticities of substitution) from Gervais and Jensen 2019, and for the upper-nest (across-sector elasticities of substitution) from Redding and Weinstein 2024.

D.2 Estimating wedges

We now describe the various data inputs required to estimate wedges $\mu_i \tau_{ij}$ via the methodology described in Section 4.3.

D.2.1 Sector-specific cost shares for the US

Our first set of inputs are sector-specific cost shares for the United States that we take as our undistorted baseline. We estimate US cost shares in two ways, at two levels of aggregation. For the more aggregated level we compute these cost shares by combining the WIOD Socio Economic Account database with the National Input-Output Tables (NIOT) US National IO table. Both datasets have information on input use of 56 sectors. From the WIOD Socio Economic Account, we obtain the total labor, capital and material expenses in 2014. From the NIOT US National IO table we obtain energy expenditures. With this information we compute the cost share of each input.

For the more disaggregated level we compute cost shares by combining data from the US Input-Output table from the BEA and the BEA/BLS Integrated Industry-level Production Account (KLEMS). The BEA Input-Output tables contain information on labor and material expenses for 402 sectors, and the KLEMS datasets contain information on capital expenses for 63 industries, and we use data for 2017. However, the Chilean classification system has 674 sectors, thus we concord these three sectoral classifications systems to compute the cost shares.

D.2.2 Assigning each Chilean firm the input costs shares of its American counterpart

We create a concordance to map the 674 disaggregated sectors that are used to categorize firms into the 56 sectors available in the detailed US IO table/WIOD merge. We apply the relevant US sector-level cost share to each Chilean firm.

D.2.3 Input Wedges

We compute three types of input wedges for each firm: labor (τ_i^l), capital (τ_i^k) and materials (τ_i^m)—with inputs from each of the 674 sectors having their own firm-specific input wedge. Note that the approach documented in Section 4.3 cannot separate input wedges from output wedges. Hence, our "input wedges" are actually the product of the firm-factor-specific input wedge and the firm-specific output wedge. The combined wedge is equal to the product of the concorded US sector's input cost share (α_{sf}^{US}) and value of output produced ($p_i y_i$) divided by value of input consumed ($f_i p_{if}$):

$$\tau_i^f \mu_i = \alpha_{sf}^{US} \frac{p_i y_i}{f_i p_{if}}.$$

D.3 Elasticities of substitution

Preferences have a nested CES structure with two tiers. We denote the elasticity of substitution across products within sector s by θ_s, and the elasticity of substitution across sectors θ_U. For the elasticities θ_s we use estimates from Gervais and Jensen 2019, which provides estimates for roughly 70 sectors at approximately the 3-digit NAICS level using US data. The estimates range from 1.6 to 23, with a mean of 7.14 and a standard deviation of 3.20. For the elasticity θ_U we use the estimate from Redding and Weinstein 2024, which draws on Chilean trade data to estimate elasticities of substitution. The value for this estimate is 1.36.

On the production side, we assume firms have Cobb-Douglas production functions that use capital, labor and materials (with each industry providing different material inputs). Thus, the elasticity of substitution is 1 across all inputs.

D.4 Navigating computational challenges

Solving for the incidence of distortions requires computing operations involving massive matrices. For instance, we need to compute Leontief inverses of input-output matrices, $\Psi = (I - \Omega)^{-1}$. The input-output matrix Ω is defined at the firm-product level, and there are 3.5 million firm-products. Depending on which type of object we need to compute, we employ several techniques to reduce the computational burden of the operations.

In general, we exploit three recurrent features of the operations involving large matrices in our incidence formula. First, the matrices are sparse. In the case of Ω, less than 0.01% of cells have non-zero values. Given this, we leverage computational advancements in operations using sparse matrices, which are much more efficient than computing the operation with the entire matrix. Second, a common object used in computing incidence formulas are Leontief inverse matrices, which we approximate with a power series, e.g., $\Psi = I + \Omega + \Omega^2 + \Omega^3 + \ldots$. Finally, we leverage linear algebra by computing smaller operations iteratively. For instance, take the case of computing Leontief inverses. Consider a column vector b of dimension $N \times 1$, with N=number of firm-products. One can use b to reduce the burden of multiplying Ω times Ω, e.g., $\Psi b = b + \Omega b + \Omega(\Omega b) + \Omega(\Omega(\Omega b)) + \ldots$. This avoids computing operations between two matrices of dimension $n \times n$ and instead allows for computing operations between $n \times n$ and $n \times 1$ matrices, which is much faster. Taken together, we can compute the Leontief inverse of the 3.5 million column and row input-output matrix in less than 10 seconds.

This last feature can be applied to all operations involving large matrices. One object for which this is also useful are the covariance matrices in our incidence formula. One can reorganize the covariance so that it has a similar structure as the case of the Leontief inverse matrix. Consider the case where the covariance below is computed for one factor only. Then one has the following:

$$\begin{aligned}
\sum_{k\in\mathcal{N},\mathcal{C}} \lambda_k Cov_{\tilde{\Omega}^{(k)}}\left(\Psi_{(f)},\Psi_{(i)}\right) &= \sum_{k\in\mathcal{N},\mathcal{C}} \lambda_k \left(\sum_{l\in\mathcal{N},\mathcal{F}} \tilde{\Omega}_{kl}\Psi_{lf}\Psi_{li} - \left(\sum_{l\in\mathcal{N},\mathcal{F}} \tilde{\Omega}_{kl}\Psi_{lf}\right)\left(\sum_{l\in\mathcal{N},\mathcal{F}} \tilde{\Omega}_{kl}\Psi_{li}\right)\right) \\
&= \sum_{k\in\mathcal{N},\mathcal{C}} \lambda_k \left(\sum_{l\in\mathcal{N},\mathcal{F}} \tilde{\Omega}_{kl}\bar{\Psi}_{li} - \left(\sum_{l\in\mathcal{N},\mathcal{F}} \bar{\tilde{\Omega}}_{kl}\right)\left(\sum_{l\in\mathcal{N},\mathcal{F}} \tilde{\Omega}_{kl}\Psi_{li}\right)\right) \\
&= \sum_{k\in\mathcal{N},\mathcal{C}}\sum_{l\in\mathcal{N},\mathcal{F}} \lambda_k\tilde{\Omega}_{kl}\bar{\Psi}_{li} - \sum_{k}\lambda_k\bar{\tilde{\Omega}}_k\left(\sum_{l\in\mathcal{N},\mathcal{F}} \tilde{\Omega}_{kl}\Psi_{li}\right) \\
&= \sum_{l\in\mathcal{N},\mathcal{F}}\left(\sum_{k\in\mathcal{N},\mathcal{C}} \lambda_k\tilde{\Omega}_{kl}\right)\bar{\Psi}_{li} - \sum_{k}\sum_{l\in\mathcal{N},\mathcal{F}} \lambda_k\bar{\tilde{\Omega}}_k\tilde{\Omega}_{kl}\Psi_{li} \\
&= \sum_{l\in\mathcal{N},\mathcal{F}}\left(\sum_{k\in\mathcal{N},\mathcal{C}} \lambda_k\tilde{\Omega}_{kl}\right)\bar{\Psi}_{li} - \sum_{l\in\mathcal{N},\mathcal{F}}\left(\sum_{k} \lambda_k\bar{\tilde{\Omega}}_k\tilde{\Omega}_{kl}\right)\Psi_{li}
\end{aligned}$$

where $\bar{\Psi}_{li} = \Psi_{li}\Psi_{lf}$, $\bar{\tilde{\Omega}}_{kl} = \tilde{\Omega}_{kl}\Psi_{lf}$, $\bar{\tilde{\Omega}}_k = \sum_l \bar{\tilde{\Omega}}_{kl}$. This strategy is presented for a given factor and it delivers a covariance column vector that varies at the i level. To compute for all factors, one needs to loop over $f \in \mathcal{F}$. That is, compute the above covariance for each f and then concatenate the resulting column-vector. Thus, the result will be a matrix of dimensions $N \times F$.

D.5 Iteratively solving for incidence

Let $y \in R^m$ denote the vector of endogenous variables, $x \in R^k$ the vector of wedges, and $z \in R^p$ the vector of fixed parameters. The mapping of interest is $y = y(x; z)$. In particular, we seek to compute the change in y as wedges x moves from baseline x to counterfactual x'. We use the following relationship from the Fundamental Theorem of Calculus:

$$y(x', z) - y(x, z) = \sum_{i=0}^{n-1} y(x_{i+1}) - y(x_i) \approx \sum_{i=0}^{n-1} \frac{dy}{dx}_{y=y_i, x=x_i}(x_{i+1} - x_i)$$

where $x_0 = x$ and $x_n = x'$.

The endogenous vector y consists of firm and factor income shares $\lambda_{i\in\mathcal{N},\mathcal{F}}$, household income shares $\chi_{c\in\mathcal{C}}$, expenditure shares $b_{c\in\mathcal{C},i\in\mathcal{N}}$[24], input–output coefficients $\Omega_{i\in\mathcal{N},j\in\mathcal{N},\mathcal{F}}$, and firm profits $\pi_{i\in\mathcal{N}}$. Wedges are defined as $x = (\ln\mu_i\tau_{ij})_{i\in\mathcal{N},\, j\in\mathcal{N}\cup\mathcal{F}}$. The fixed parameter

[24] We actually work with $B_{ci} = b_c B_{ci}^{direct} + (1 - b_c)B_{ci}^{indirect}$. With Cobb–Douglas preferences, $(1 - b_c)B_{ci}^{indirect} = \frac{1}{\chi_c}\sum_j F_{cj}K_{ji}^{consump}$. Because this involves a Leontief inverse, we do not compute the object explicitly outside the loop, but keep components separate. χ_c and $K_{ji}^{consump}$ are fixed at baseline whenever they enter in this construction.

vector z includes consumption-side parameters $\theta_{i\in\mathcal{N},c\in\mathcal{C}}$; production-side coefficients $\tilde{\Omega}_{ij}$; factor and ownership shares $\Phi_{c\in\mathcal{C},i\in\mathcal{N}\cup\mathcal{F}}$; and transfers T_c.

Define the local update operator

$$T(y, x, \Delta x; z) := \frac{dy}{dx}(y, x; z)\, \Delta x,$$

which gives the approximate change Δy associated with perturbation Δx around state (y, x). Fix a partition $(x_i)_{i=0}^{n}$ with $\sum_{i=0}^{n-1}(x_{i+1} - x_i) = x_n - x_0$. The algorithm proceeds as follows:

1. Initialize with baseline data (y_0, x_0).
2. For each $i = 0, \ldots, n-1$:
   (a) Compute $\Delta y(i) = T(y_i, x_i, x_{i+1} - x_i; z)$.
   (b) Update $y_{i+1} = y_i + \Delta y(i)$.
3. Return the total change: $y_n - y_0$.