Payments Under the Table: Estimation and Implications for Optimal Tax Policy

Ana Paula Franco*
University of Michigan

[Click here for the latest version]

November 26, 2025

Abstract

This paper studies how firm–worker agreements to underreport wages affect tax policy design. In most tax systems, firms act as third-party reporters to limit evasion, but this mechanism weakens when employers pay part of workers' wages off the books—what I refer to here as payments under the table (PUT). I develop a model of optimal income taxation in which PUT arises endogenously from firm–worker interactions, and a corporate tax is also in place. The model shows that the corporate tax modifies the standard optimal income tax formula through two channels: (i) a fiscal externality, as PUT cannot be deducted from firms' taxable profits and thus raise corporate tax revenue; and (ii) a behavioral response, as higher corporate taxes reduce firms' incentives to pay under the table and narrow the scope for workers to evade, thereby lowering the elasticity of taxable income (ETI). Together, these effects imply that the optimal income tax rate increases with the corporate tax rate. Empirically, I use matched administrative and survey data from Peru—linked using optimal transport—and find that PUT is widespread, especially among higher-income workers. Exploiting a staggered corporate tax reform, I show that the ETI declines when the corporate tax rate increases, consistent with the model's predictions.

*E-mail: apfranco@umich.edu. I am extremely grateful to James Hines, Joel Slemrod, and Damian Vergara for their invaluable and supportive guidance throughout this project. For comments that have improved this paper, I thank Hoyt Bleakley, Jon Bakija, Marianne Bitler, Ashley Craig, Luciana de la Flor, Javier Feinmann, Luciana Galeano, Daniel Jacqua, Dylan Moore, Hugo Ñopo, Daniel Reck, Jordan Rosenthal-Kay, Ana Reynoso, Dean Yang, and Eric Verhoogen. I am also grateful to many participants at the Public Finance Seminar at the University of Michigan, Labor Lunch at the University of Michigan, Development Seminar at the University of Michigan, NBER Summer Institute 2025, RIDGE-LACEA Workshop in Public Economics 2025, IIPF Annual Conference 2025, AEA CSQIEP Mentoring Conference 2025, Michigan Tax Invitational 2025, LACEA-LAMES Annual Meeting 2023.

1. Introduction

Governments use income taxes to redistribute resources from richer to poorer individuals, with the goal of promoting greater equity (Mirrlees, 1971; Diamond and Saez, 2011). But taxes change behavior. People may work less when taxes increase (real responses) or try to reduce their reported income without changing their actual work effort (avoidance or evasion responses) (Saez, 2001; Kleven et al., 2011). To limit evasion, governments often rely on firms to act as third-party reporters. By requiring employers to report workers' earnings, tax authorities can use firms' information and administrative capacity to improve compliance (Kleven et al., 2016). However, the effectiveness of third-party reporting depends on firms' truthful reporting of wages. When employers and workers agree to underreport earnings, the informational role of firms as tax collectors is undermined (Slemrod, 2008). This problem is particularly relevant in countries with limited enforcement capacity and large informal sectors.

A growing body of evidence from low- and middle-income countries shows that many employees receive a higher wage than the one officially reported by their *employer*. The official wage is the amount declared by firms to tax and social security authorities, while the difference—paid off the books, often in cash—is referred to as an envelope wage or payment under the table (PUT). These workers are registered on a firm's payroll, which is why part of their wage is reported. They are not workers hired entirely outside the formal system, who are typically referred to as 'informal workers' in the labor economics literature. Instead, PUT represent a hybrid form of employment: the job itself is formally registered, but compensation is only partially declared. In this sense, PUT can be understood as a partial informality.

The existing evidence on PUT has primarily focused on measurement, as evasion is inherently difficult to observe directly. The literature has relied either on surveys that ask workers directly whether they receive payments under the table from their employers in addition to their reported wages (Williams and Padmore, 2013; Feinmann et al., 2025b), or on inferring PUT indirectly by exploiting policy or enforcement changes that affect reported but not true wages (Bergolo and Cruces, 2014; Kumler et al., 2020; Gavoille and Zasova, 2023). Despite this growing empirical evidence, the implications of PUT for tax policy remain largely unexplored.

This paper makes two contributions. First, it introduces an optimal transport approach to improve the measurement of payments under the table (PUT) using linked administrative and survey data. Second, it incorporates PUT into a model of optimal income taxation in an economy with a corporate tax, to study how tax policy should adapt when evasion takes this form.

In the standard framework, the optimal income tax balances the gains from collecting more taxes—more revenue to transfer to poorer individuals (more redistribution)—against the efficiency costs of taxation, captured by how reported income responds when tax rates increase, either because people work less or evade more. This responsiveness is summarized by the elasticity of taxable income (ETI). I show that this familiar trade-off continues to hold even when PUT are widespread. The novelty lies in the role of the corporate tax. Because PUT cannot be deducted from firms' taxable profits, a higher corporate tax generates a positive fiscal externality that raises revenue and increases the redistribution benefits of taxation. At the same time, a higher corporate tax reduces firms' incentives to pay under the table, which limits workers' ability to evade and therefore lowers their response to tax changes — that is, it reduces the ETI. Both mechanisms strengthen the case for a higher optimal income tax compared to settings where evasion is purely a worker decision and firms play no role. I test these predictions using

Peruvian data and a reform that simultaneously changed both income and corporate tax rates, and I find empirical responses consistent with the model.

Getting this main result involves three steps, which form the building blocks of the paper. In the first part, I bring optimal transport (OT)—a novel econometric approach to linking datasets—to the evasion estimation literature. The goal is to provide governments with a practical tool to assess whether PUT are a significant concern without relying on costly or potentially biased worker surveys. OT constructs a mapping between datasets based on common covariates, minimizing the distance between their joint distributions and effectively reweighting one dataset to match the covariate distribution of the other. I apply this method to data from Peru, linking administrative employer–employee records (reported by firms to the tax authority) with household surveys that contain self-reported wages. I assume that the administrative data capture the true wage net of PUT, while survey data reflect the true wage plus measurement error.

Each individual in the administrative dataset is matched to a self-reported wage, and PUT is defined as the difference between the two. Because OT preserves the joint distribution of worker characteristics, it allows me to recover how PUT underreporting varies systematically along observable dimensions such as income, occupation, and firm size. This is crucial, as noncompliance need not be homogeneous across the income distribution. Empirical and theoretical work shows that the extent and nature of tax evasion vary with income (Kopczuk, 2001; Slemrod and Yitzhaki, 2002; Alstadsæter et al., 2019). Higher-income individuals often have greater opportunities to engage in complex avoidance or evasion, while lower-income workers have lower probabilities of being audited. PUT may exhibit similar heterogeneity, with potential implications for whether PUT underreporting is regressive or progressive—an issue central to the distributional incidence of taxation and the design of enforcement policy.

I show that PUT are prevalent and more salient among individuals with higher education and income in Peru. On average, workers receive 30% of their wages under the table. These findings are consistent with survey-based evidence on PUT behavior across Latin America. I also estimate the prevalence of PUT (but not the distributional effects) using OLS, propensity score matching, and inverse probability weighting, obtaining similar results across methods. However, these approaches cannot capture distributional patterns accurately, as they match only on an index or means.[1] I perform several robustness checks for classical and nonclassical measurement error and find the results to be stable. Given that high-income individuals are more likely to underreport wages or not respond to surveys at all, underreporting of PUT is more plausible than overreporting.

In the second part of the paper, I develop a model of the optimal personal income tax in an economy where PUT arise endogenously in labor market equilibrium and where a corporate tax is also in place. Workers' personal income is subject to a linear income tax, and firms' profits are taxed under a linear corporate tax.[2] PUT matter because they are not subject to personal taxation (a benefit for workers)

[1] As an illustration, consider a highly educated worker who reports earning the minimum wage (at, say, the 10th percentile) but whose true wage lies at the 90th percentile. In this case, the administrative wage distribution would appear shifted to the left. If PUT were measured as the difference in wages at a given percentile, this worker's reported wage at the 10th percentile would be compared to others at the 90th percentile, overstating the estimated PUT. The true PUT would instead be the gap between the self-reported 10th percentile wage and the actual 90th percentile wage. This mismatch arises because percentile-based comparisons are not one-to-one. Propensity score matching and nearest-neighbor methods are particularly problematic because they may not fully balance composition differences. They allow estimation of self-reported wages at the individual level but do not preserve the joint distribution of covariates.

[2] I assume a single personal income tax that encompasses both the payroll and income tax.

but are also non-deductible from firms' taxable profits (a cost for firms).

I start by modeling a labor market in which workers choose whether to supply labor under a fully formal arrangement (no PUT) or to accept an arrangement that includes a share of PUT. Workers are identical in productivity but differ in an exogenous informality cost, which can be interpreted as the valuation they place on payroll-linked benefits. Workers self-select into formal or PUT arrangements according to this cost. Firms are indifferent between hiring formal and PUT workers as long as the total cost of hiring is equal across arrangements. The equilibrium wage in each arrangement adjusts to clear the market. Because PUT cannot be deducted from taxable profits, wages in the PUT sector are lower to offset the higher effective cost to firms. Workers in the PUT arrangement therefore receive part of their compensation under the table, while fully formal workers report their entire wage. Given these labor market outcomes, the government chooses the income tax rate (and lump-sum transfer) that maximizes a social welfare function, subject to a balanced-budget constraint and a fixed corporate tax.

Using this model, I derive the conditions that characterize the optimal income tax rate. The solution to the planner's maximization yields an optimal income tax formula aligned with the well-known Mirrlees–Saez framework, balancing the equity–efficiency trade-offs of a change in the income tax rate. The elasticity of taxable income (ETI) remains the key sufficient statistic. The novelty is that the formula now includes two channels through which the corporate tax directly affects the optimal income tax. First, PUT introduces a fiscal externality in the numerator. When wages are paid under the table, personal income tax revenue falls, but the same under-the-table payments increase firms' taxable profits and therefore corporate tax revenue. The government thus recovers part of the lost revenue automatically through the corporate tax base, which lowers the effective revenue cost of underreporting and raises the optimal income tax relative to a setting without this offset. Second, the ETI in the denominator becomes a negative function of the corporate tax. In equilibrium, firms are indifferent between hiring workers formally or under the table, as wages adjust to offset the non-deductibility of PUT payments. A higher corporate tax increases the cost of PUT arrangements, compresses the wage gap between formal and PUT contracts, and reduces the scope for PUT underreporting. As a result, workers' reported income becomes less elastic. Similar to Slemrod and Kopczuk (2002), the ETI is endogenously shaped by worker behavior and firm incentives through the corporate tax. The magnitude of these two channels depends on the ratio of PUT to reported income in the economy.

Together, these two channels imply that the optimal income tax rate increases with the corporate tax rate. When the corporate tax is zero, the model collapses to the standard evasion case, where underreporting decisions fall entirely on the worker side, the ETI is higher, and the optimal income tax is lower. As the corporate tax rises, PUT become less attractive for firms, reducing the ETI and strengthening the fiscal externality, both of which raise the optimal income tax. When the corporate tax rate exceeds the personal income tax rate, incentives for PUT disappear, and the model converges to full compliance. These results depends on the non-deductibility of PUT. If firms were able to deduct these payments, the cost of PUT would vanish, returning the model to a case closer to standard evasion, where all incentives lie on the worker side. In practice, anecdotal evidence suggests that firms deduct only part of these payments, placing real economies somewhere between the two extremes of full and zero deductibility. In this intermediate case, the corporate tax continues to play a central role by shaping both the fiscal externality and the ETI—and, therefore, the optimal income tax rate.

In the final part of the paper, I bring the model to the data using two recent reforms in Peru's labor tax regimes and then use numerical simulations calibrated to Peruvian moments to explore the policy implications. I focus on the key theoretical prediction that the elasticity of taxable income (ETI) is a decreasing function of the corporate tax.

The first reform, introduced in 2021, increased the payroll tax for workers employed under the Agricultural Regime (AgR) but left the corporate tax unchanged. The policy applied only to field workers hired under the AgR, who serve as the treated group. In contrast, administrative employees within the same agricultural firms could not be hired under the AgR and had to be employed under the General Regime (GR). These workers, exposed to the same firm- and sector-specific shocks but unaffected by the tax change, form the control group. The second reform, implemented in 2023, simultaneously increased both the payroll and corporate tax rates for firms under the AgR, while the GR remained unaffected. This differential exposure allows me to examine how changes in corporate taxation affect the responsiveness of reported income to labor taxes.

Using a continuous difference-in-differences design, I estimate the elasticity of reported income with respect to the net-of-tax rate for workers in the AgR. I find large ETIs of around 2 during the first reform, consistent with weaker enforcement and greater scope for underreporting in developing economies. However, the response is attenuated in 2023, when the corporate tax also increased, consistent with the model's prediction that higher corporate taxation reduces the attractiveness of PUT and lowers the sensitivity of reported income to labor taxation. Because the magnitude of tax changes varied across firms of different sizes, I exploit this staggered structure to construct alternative control groups and verify the robustness of the results. The validity of the ETI estimates relies on the standard parallel-trends assumption. Event-study analyses confirm that pre-trends were parallel and show a discrete decline in reported income following the 2021 payroll tax hike but little response after the 2023 reform.

Numerical simulations calibrated to Peruvian moments for PUT and the ETI further illustrate the policy implications. At the effective corporate tax rate that best matches the Peruvian data, the model delivers an optimal top income tax rate of about 25%, well below the roughly 40% observed in the actual schedule. Taken at face value, this gap suggests that current policy is closer to a 'no-evasion' benchmark in which the government effectively ignores PUT when setting the income tax rate, treating third-party reported wages as true income and therefore perceiving a lower ETI than it would if the PUT margin were internalized. By contrast, if the government behaved as if PUT followed a classical evasion model driven solely by workers' decisions—so the corporate tax plays no deterrent role—the ETI would be higher and the optimal income tax rate in the simulations would be closer to 15%.

Related literature. This article contributes to several strands of literature, the primary one being the analysis of tax evasion and optimal taxation by explicitly incorporating payments under the table (PUT) into the discussion.

The existing literature on PUT or 'envelope wages' has primarily focused on documenting and measuring its prevalence. Two main approaches have been used. The first approach relies on survey data, where workers are asked directly whether they receive part of their pay off the books. In Europe, (Barth and Ognedal, 2018) present survey evidence that around 5 percent of employees in the European Union receive part of their wages off the books, though this share varies widely—from 15 percent in Romania and 10 percent in Bulgaria to 5–7 percent in Spain (Di Nola et al., 2019). Williams and Padmore (2013) use Eurobarometer data across 28 European countries and similarly find that about 3 percent of workers report receiving underreported salaries, with much higher prevalence among low-skilled and Eastern European workers. In Latin America, Feinmann et al. (2025b) find that 25 percent of formal employees in Brazil admit to receiving PUT for about 20 percent of their wages, with higher-wage workers evading a larger share. In Feinmann et al. (2025a), we conduct a multi-country survey in Argentina, Peru, Mexico, Colombia, and Chile, finding an average of 15 percent of formal employees receiving PUTs for roughly 23 percent of their earnings, with substantial heterogeneity—from more than one-quarter of employees in Argentina to less than 4 percent in Chile. Across contexts, this literature consistently documents that PUT are widespread and increase with workers' earnings.

The second approach infers underreporting indirectly, typically by exploiting reforms that affect reported—but not actual—compensation. In Uruguay, Bergolo and Cruces (2014) show that after a reform reinforcing the link between reported wages and future social security benefits, reported earnings rose while total labor costs remained unchanged—consistent with a reduction in PUT. Kumler et al. (2020) find similar evidence in Mexico following a 1997 pension reform, with the largest increases in reported wages occurring among younger workers and small firms. (Bíró et al., 2022) detect comparable responses in Hungary, finding that stronger enforcement reduces bunching at contribution thresholds. (Bjørneby et al., 2021) identify underreporting in Norway by comparing employer-reported wages to third-party financial information, while (Gavoille and Zasova, 2023) document that increased monitoring in Latvia led to higher reported earnings and smaller discrepancies relative to actual pay. Together, these studies provide compelling evidence that the extent of wage underreporting responds systematically to enforcement and contribution incentives.

An exception to these two approaches is López-Luzuriaga et al. (2023), who detect envelope wages in Ecuador by combining electronic invoicing records with payroll tax returns. Their method uses billing data as a quasi–third-party verification device to infer unreported wage components. This is a compelling innovation, but it depends on full coverage of e-invoicing systems—an infrastructure still unavailable in many low-capacity developing countries. By contrast, my method is explicitly designed to work with administrative and household survey data that governments already collect, making it more scalable and immediately applicable in such settings.

This literature has largely focused on documenting the existence and magnitude of PUT, without examining how their presence should change the design of tax policy. This paper advances the literature in two main ways. First, it develops a new method that preserves the full distribution of earnings—crucial for distributional analysis—to measure PUT using administrative data combined with survey data where available. This avoids the high costs and potential biases of direct survey methods while ensuring cross-country comparability. Second, it goes beyond measurement to study the implications of PUT for tax policy. I develop a theoretical model in which firms and workers jointly determine how much of the wage is reported and derive optimal tax formulas that incorporate this behavioral margin. This is the first paper to link the measurement of PUT to its implications for optimal tax design, bridging two strands of literature—empirical work on wage underreporting and normative analyses of tax policy—that have evolved largely in isolation.

This paper also contributes to the optimal income taxation literature (Mirrlees, 1971; Saez, 2001; Diamond and Saez, 2011), which studies how governments can redistribute income while minimizing efficiency costs. While this literature does allow for tax evasion, it typically models evasion as an individual worker decision and treats third-party reporting by firms as a silver bullet that effectively closes the evasion margin. In other words, once wages are reported by the employer, compliance is assumed to be complete. I depart from this assumption by allowing firms and workers to jointly determine how much of the wage is reported, thereby endogenizing the possibility that third-party enforcement may fail. In doing so, the paper connects the optimal tax tradition to the foundational evasion literature—beginning with Allingham and Sandmo (1972) and Yitzhaki (1974), who model individual non-compliance under audit risk—and to more recent work on third-party reporting and enforcement capacity (Kleven et al., 2011; Kleven et al., 2016). But whereas these models assume that firm-reported wages are truthful by construction, this paper allows for underreporting even in the presence of employer reporting and shows that the existence of PUT fundamentally alters optimal income and corporate tax design.

Beyond the tax evasion literature, this paper also contributes to the broader work on labor informality, which studies earnings that escape taxation but typically frames informality as a binary firm or sectoral choice. A large empirical literature documents the prevalence, determinants, and consequences of informality (Williams and Padmore, 2013; La Porta and Shleifer, 2014; Ulyssea, 2020), emphasizing its implications for productivity, welfare, and tax capacity. However, informality in this literature is generally modeled as binary decision—either a job is formal or informal—rather than as a continuous choice over how much income is reported. This paper instead focuses on the intensive margin of informality, studying partial underreporting within formal employment (i.e., PUT), and analyzes how this continuous evasion margin affects optimal tax design.

Finally, this paper contributes to the growing empirical use of optimal transport (OT) in economics. In labor economics, Kim and Lee (2024) use OT to recover the joint distribution of worker skills and job characteristics from observed matches, and Arellano and Bonhomme (2023) apply an OT-based matching framework to recover the latent distribution of income shocks from panel earnings data. In industrial organization, Li et al. (2019) use inverse OT to back out the implicit matching cost function underlying observed college admissions and labor market allocations. In trade, Gaskin et al. (2024) use OT with neural networks to learn counterfactual trade cost structures from observed global trade flows, while Leite and De Bacco (2024) use an OT-inspired algorithm to reconstruct realistic transportation networks from minimal origin–destination data. These examples illustrate a broader shift toward using OT to infer hidden economic primitives from observed distributions. In that spirit, this paper is the first to apply OT to recover the latent firm–worker underreporting distribution, by transporting survey and administrative wage data into a revealed structure of payments under the table.

The rest of the paper is organized as follows. Section 2 describes the institutional setting, focusing on the Peruvian tax system. Section 3 estimates payments under the table (PUT) using linked administrative and survey data through an optimal transport approach. Section 4 develops the model of optimal income taxation with PUT. Section 5 brings the model to the data. Section 6 concludes.

2. Institutional Setting

Payments under the table (PUT) are wage underreporting arrangements that require coordination between firms and workers. This is what differentiates them from classical forms of evasion, where workers misreport income unilaterally. Because PUT depend on joint incentives, they interact with both sides of the tax system: the worker side (personal income and payroll taxes) and the firm side (corporate taxation and deductibility rules). Understanding this institutional structure is essential to analyzing PUT.

Most countries, including Peru, have three main tax instruments relevant for labor earnings: (i) a personal income tax on workers, (ii) payroll taxes that finance social security benefits (borne by workers and/or firms), and (iii) a corporate income tax on firms' taxable profits. PUT interact with all three. For workers, PUT are not declared to tax or social security authorities, and therefore escape both personal income and payroll taxation. This is the benefit of PUT. For firms, however, these payments cannot be deducted as labor expenses, and thus increase the corporate tax base. This is the cost of PUT.

Before describing the tax system in detail, I first explain why Peru is particularly relevant for analyzing PUT. Peru has one of the most informal labor markets in Latin America: around 70 percent of workers are employed entirely outside the formal system. Importantly, informality in Peru is not merely the result of exclusion from formal jobs; it is often a deliberate choice. A large body of evidence shows that workers and firms opt out of formality when the perceived benefits of contributing are low, especially given weak enforcement and the limited value many individuals attach to future social security benefits (Yamada, 1996; Loayza, 1996; Perry et al., 2007). Even during periods of strong growth, informality has remained persistent and largely voluntary (Bosch and Maloney, 2010; Levy, 2018). These strong preferences for informality suggest that workers and firms may also have high willingness to engage in PUT arrangements.

Enforcement capacity is limited and biased toward detecting who is registered, not whether reported wages are truthful. As of 2023, Peru had fewer than 800 labor inspectors overseeing nearly 20 million workers (U.S. DOL, 2024; SUNAFIL, 2023). In practice, the tax authority (SUNAT) prioritizes uncovering non-registration or omitted (fully informal) workers (SUNAT, 2023). This institutional environment creates conducive environment for formal-but-underreported employment, the margin this paper studies.

Moreover, household surveys in Peru are accurate in capturing both wages and formality status. Given that informality is a central policy concern, these surveys are the primary instrument to measures labor market conditions (as administrative data do not capture informal workers). The national statistics office (INEI) has invested heavily over decades to ensure truthful reporting: the survey is anonymized, widely trusted, and known by workers to be used for statistical purposes (not enforcement). Administrative payroll declarations, on the other hand, are legal compliance documents. Tax liability is calculated using what is reported there. The survey is therefore the more reliable source for measuring true earnings and formality status (subject, of course, to standard survey nonresponse), while administrative records reveal what firms actually report. It is precisely this wedge between reported and true wages that I will use to estimate PUT in Section 3.

I now describe the tax and contribution system relevant for firm–worker decisions over wage reporting. The objective is to highlight the institutional features of the Peruvian tax system that incentivize PUT.

Personal Income Taxation and Payroll Taxes Table 1 summarizes income and payroll taxes. On the worker side, labor income is subject to a progressive personal income tax (PIT), withheld monthly by the employer. The tax schedule applies only above a standard deduction of 7 tax units[3], and marginal rates rise with income from 8%, 14%, 17%, 20% to 30%. Employees must also contribute 13% of earnings to the pension system, either the public pay-as-you-go pension system or the private savings-funded pension system. On the firm side, employers are required to contribute 9% of reported payroll to EsSalud,

[3]There is a unique tax unit per year. Between 2020 and 2023, it was 4,300 ($1,200), 4,400 ($1,250), 4,600 ($1,300), and 4,950 ($1,400) Peruvian soles, respectively.

the national health insurance system, and approximately 9.7% in mandatory severance (CTS), which is an individual savings account for each worker (that serves as unemployment insurance). Under special regimes, some of the payroll taxes are lower (see Table 2)

Corporate Taxation Peru has multiple firm tax regimes with different corporate income tax (CIT) rates and payroll burdens, generating differentiated costs of reporting wages by sales and sector. Table 2 summarizes them. First, there is a general regime, firms face a 29.5% CIT, and all reported wages and payroll contributions are fully deductible. Second, there is an agricultural regime. Firms under this regime face lower payroll contributions of 4–8% and a 15-20% CIT. Finally, there is a Small Firms regime, which applies to firms with sales below 1,700 tax units. These firms face 10% CIT for sales below 15 tax units and 29.5% for sales above 15 tax units.

Table 1: Personal Income Tax and Payroll Taxes Levied on Workers, MTR

	Personal Income Tax (PIT) (1)	Pensions (ONP/AFP) (2)
Taxable unit:	Income	Labor income
Income brackets:		
(in multiples of min. wage, $\bar{w}$)		
$[\bar{w}, 2.5\bar{w})$	0%	13%
$[2.5\bar{w}, 4\bar{w})$	8%	13%
$[4\bar{w}, 9\bar{w})$	14%	13%
$[9\bar{w}, 14.5\bar{w})$	17%	13%
$[14.5\bar{w}, 18\bar{w})$	20%	13%
$18\bar{w}+$	30%	13%
Levied on:	Workers	Workers

Notes: Table shows marginal tax rates (MTR). $\bar{w}$ is the monthly minimum wage (930 soles in 2021). PIT and the worker pension contribution (13%) are independent of the employer's tax regime (General, Agricultural, MYPE). Health insurance (EsSalud) is levied on firms and appears in Table 2. Unemployment insurance (CTS) is levied on the firm side as in Table 2. For the purpose of this paper, I assume that the only source of taxable income is labor income.

Table 2: Corporate Income Tax and Payroll Taxes Levied on Firms in 2020-2023, MTR

(1)	Health Insurance (EsSalud) (2)	Severance Pay (CTS) (3)	Corporate Income Tax (CIT) (4)
Taxable unit:	Labor income	Labor income	Profits
General Regime:			
All firms	9%	9.7%	29.5%
Small Firms Regime:			
Sales $\leq$1700 Tax Units	9%	9.7%	10% for sales $\leq$15 TU 29.5% for sales >15 TU
Agricultural Regime:			
$<$ 100 workers	4% (2020)	9.7%	15% (2020)
	6% (2021-23)	9.7%	15% (2021-23)
$\geq$ 100 workers	4% (2020)	9.7%	15% (2020)
	7% (2021-22) 8% (2023)	9.7%	15% (2021-22) 20% (2023)
Levied on:	Firms	Firms	Firms

Notes: The table shows marginal tax rates (MTR). If a rate is not accompanied by a year, it remained unchanged between 2020 and 2023. Payroll contributions refer to EsSalud (public health insurance). CTS (Compensación por Tiempo de Servicios) is a mandatory severance pay system equivalent to roughly one month's wage per year of service. Because deposits are calculated based on wages plus bonuses, the effective rate is approximately 9.7%. Pension contributions (13%) are uniform across regimes (see Table 1). The 2020 agricultural reform raised both payroll and corporate income taxes for firms with 100 or more workers. The MYPE regime is a special regime based on reported sales. Firms with annual sales below 1,700 tax units qualify as small enterprises and are eligible for this regime. The value of the tax unit varies by year—between 2020 and 2023, it was 4,300 ($1,200), 4,400 ($1,250), 4,600 ($1,300), and 4,950 ($1,400) Peruvian soles, respectively. The threshold to enter the MYPE regime is approximately $2 million.

PUT Incentives PUT are attractive when the taxes avoided on underreported wages exceed the workers' valued benefits foregone from being formal and the firms' foregone corporate tax deductions.

Firm Side. Under Peruvian tax law, expenses without proper invoices or paystubs cannot be deducted from taxable profits. This includes wages paid under the table. Thus, the foregone corporate tax deduction is the marginal corporate tax rate (e.g., in the general regime, for each $1 paid under the table, the foregone deduction is $0.295). In practice, however, firms willing to misreport wages are often also willing to misreport deductible expenses (e.g., via fictitious invoices). This behavior is widely documented in developing-country tax systems (Best et al., 2015; Carrillo et al., 2017) and, if used, effectively reduces the corporate tax penalty of PUT. Thus, the true non-deductibility cost is likely fractional rather than full.

Worker Side. The personal income tax is a pure tax: reporting a higher wage and paying more tax does not grant access to additional benefits. So, the tax savings from receiving $1 under the table equal the marginal personal income tax rate. Payroll contributions, by contrast, include a contributory component (such as pensions and severance) that accrues with reported wages, and a pure-tax component (like health care) that does not. Workers may not fully value the contributory component. As a result, the additional tax savings from receiving $1 under the table include the marginal rate of the pure-tax component of payroll contributions. The cost, on the other hand, is the value the worker places on the forgone payroll benefits.

In the model, I capture heterogeneity in the valuation of payroll benefits with parameter $\theta \in [0, 1]$.[4]. Let τ be the statutory tax rate (PIT + payroll tax) on reported wages, t the corporate tax rate, and $d \in [0, 1]$ the fraction of PUT effectively deducted via disguised expenses. Then, per dollar shifted to PUT, the joint surplus is as follows:

$$\Delta S = \underbrace{\tau}_{\text{tax saving}} - \Big(\underbrace{(1-d)t}_{\text{foregone corporate deduction}} + \underbrace{\theta}_{\text{value of forgone benefits}} \Big). \tag{1}$$

PUT is chosen for that dollar whenever $\Delta S > 0$. All else equal, higher θ (more valuation of payroll benefits) reduces the incentive to underreport, while higher d (more deductibility) increases it[5].

I formalize these PUT incentives in the optimal income tax model in Section 4. But, before turning to theory, I first show that PUT can be empirically recovered in the data. Section 3 develops a distribution-matching approach based on optimal transport to estimate the extent and structure of wage underreporting.

3. Estimating PUT

In this section, I estimate the prevalence of PUT using a structural distribution-matching approach based on Optimal Transport (OT). I link administrative employer–employee records (filed by firms with the tax authority) to household surveys that report workers' self-declared wages. OT is used to match each individual in the administrative data to a self-reported wage, and PUT is measured as the difference between reported and matched wages.

Before introducing the method formally, I begin with standard comparisons (raw differences, propensity score matching, inverse probability weighting) to show that the core patterns are not driven by sample composition or survey noise, and to motivate the need for OT.

3.1. Data

I use two main datasets. The first is a matched employer–employee administrative dataset collected jointly by the Peruvian Ministry of Labor and the tax authority (SUNAT). It records all formal private-sector employment contracts reported through the electronic payroll system. I refer to this as the administrative dataset. It includes monthly information on wages, hours worked, tenure, social security benefits, firm characteristics (location, tax regime, firm size), and worker characteristics (sex, age, education). The data cover workers only, it does not include sales or revenue data. The system has been collected electronically since 2010 and has been publicly available since 2020, covering roughly five million workers per year. I will use 2022 for the matching procedure and 2020-2023 for the estimation of elasticities.

[4]Hiding $1 saves τ in taxes but forfeits a marginal valued benefit equal to the slope of the forgone-benefit schedule at the current hiding level. For the per-dollar entry margin used here (starting from full formality), I normalize that slope to θ, so the equation below replaces the marginal term by θ. If one evaluates at other hiding levels, replace θ by the appropriate marginal value at that level (which will also be a function of θ).

[5]This static surplus comparison is a reduced-form version of the more general environment in Section 4, where the cost of underreporting enters through a convex function of hidden earnings.

The second dataset is the National Household Survey of Peru (ENAHO), a nationally representative survey fielded annually since 1990 across all 25 regions of the country. It includes a detailed Work and Income Module (Module 500) that collects employment and earnings information for individuals aged 15 and above. The module contains information on labor force status, hours worked, wages, social security coverage, and related characteristics. Approximately 40,000 individuals respond to the labor module each year.

I harmonize the survey and administrative samples to ensure both contain the same population of formal private-sector workers. I drop fully informal workers from the survey—those without pension provision from their employer, as they are not present in the administrative payroll. I also exclude public-sector workers, as well as part-time workers (fewer than 20 hours per week), since the payroll dataset only includes full-time private-sector employees. Following standard practice in the survey–administrative literature (Bound et al., 2001; Hokayem et al., 2015), I trim the top 5 percent of administrative wages to align the support where both datasets are informative, since very high earners are underrepresented or top-coded in surveys. The resulting 2022 harmonized sample consists of 5.6 million administrative wage earners and 7,300 survey respondents.

Figures 1 show the 2022 wage distributions for workers by education group. The orange line corresponds to the administrative payroll distribution, the black line to the survey distribution, and the dashed line marks the minimum wage. For workers with less than secondary education, both distributions are similarly concentrated around the minimum wage. For those with secondary education, the payroll distribution remains tightly clustered around the minimum wage, while the survey distribution is slightly shifted to the right. Among workers with higher education, the payroll distribution is still compressed around the minimum wage, but the survey distribution is centered at a clearly higher level.

The patterns in Panels A and B of Figure 1 are consistent with expectations: low-skilled workers in Peru often earn close to the minimum wage. The deviation appears in Panel C. In Peru, only about 30 percent of individuals complete higher education, and the labor market returns to education are substantial. The household survey shows that workers with a bachelor's degree earn roughly 50 percent more than those with only secondary education, and workers with a master's degree or higher earn close to 190 percent more than those with a bachelor's degree. However, in the administrative payroll data, the wage distribution for higher-educated workers remains concentrated around the minimum wage, with no apparent education premium. A plausible explanation is that part of their compensation is being paid under the table.

Consistent with this, Mincer regressions estimated separately in both datasets show sizable and statistically significant returns to schooling. In the survey data, schooling explains 22% of the variation in wages, while in the administrative dataset it explains only 8% (see Appendix MA1).

Figure 1: Wage distribution in administrative and survey by education

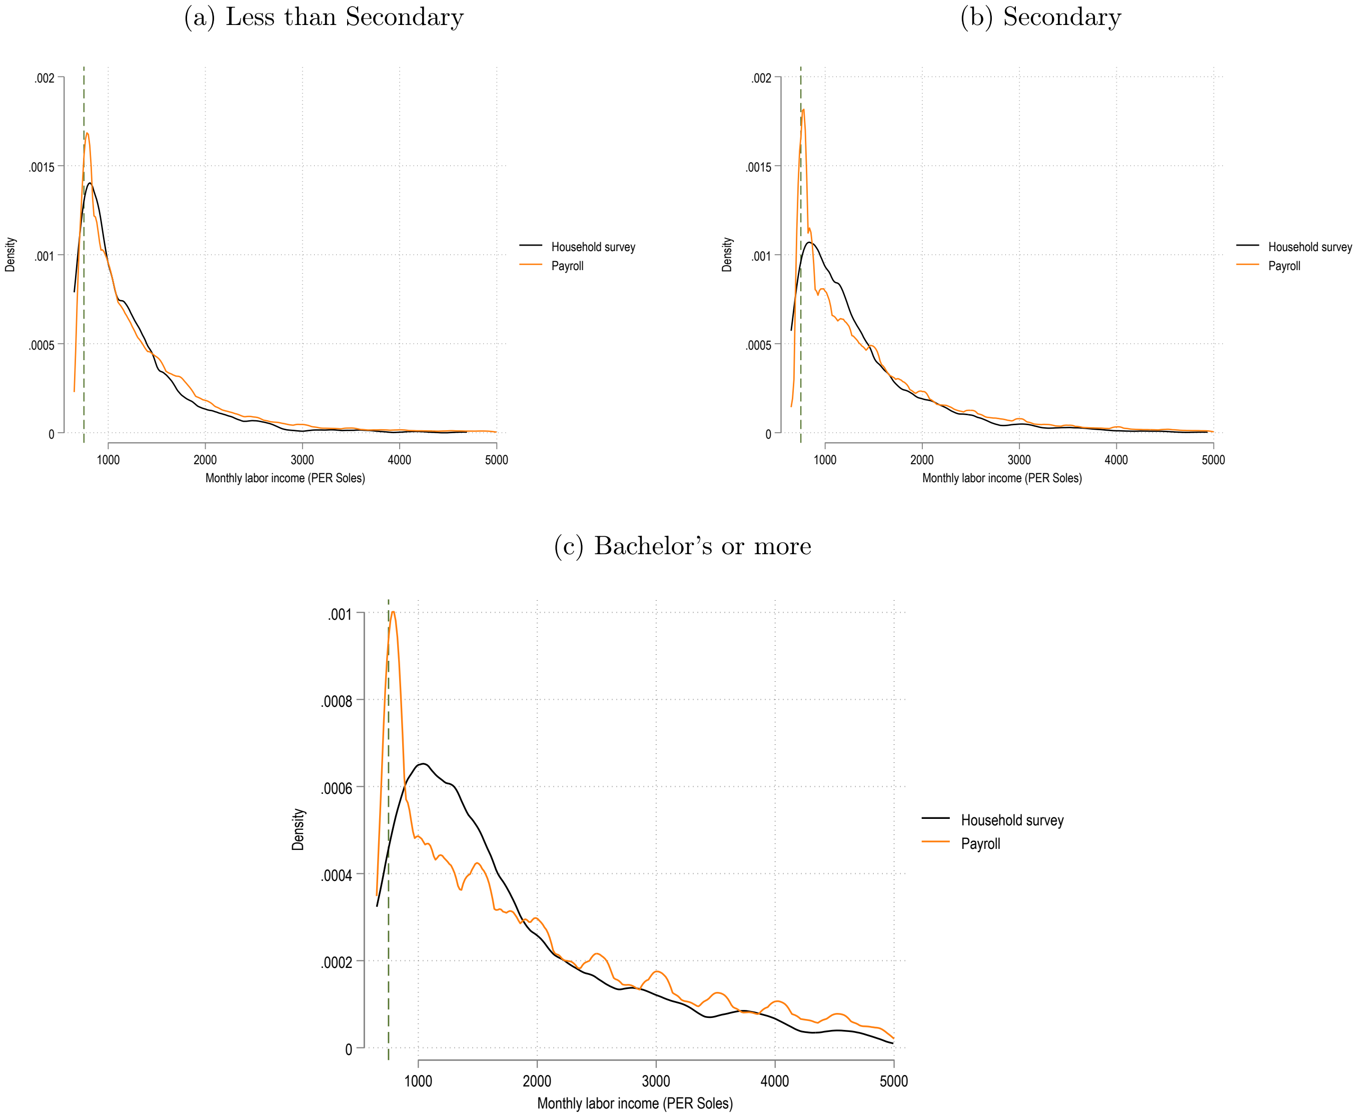

3.2. Evidence of PUT using Standard Methods

In Feinmann et al. (2025a), we document the presence of PUT in Peru and show that workers are willing to self-report engaging in such arrangements. In addition, the survey wage distributions shown above align more closely with what one would expect in a labor market with substantial education premia. These two facts motivate the core assumption maintained throughout the paper: survey wages reflect true earnings plus measurement error, whereas administrative wages reflect true earnings net of PUT. Under this assumption, any positive survey–administrative gap can be interpreted as evidence of underreporting.

Formally, let $\bar{w}$ denote the third-party (administrative) wage, and w the true wage. I define payments under the table (PUT) as $\text{PUT} = w - \bar{w}$, so that $w = \bar{w} + \text{PUT}$. The survey contains self-reported wages

$\tilde{w}$, which I assume follow: $\tilde{w} = w + e$, where e is measurement error[6].

I begin by comparing unconditional mean and median wages across the two datasets. In the harmonized data, mean survey wages exceed administrative wages by about 18 percent, and the median gap is over 30 percent. These are large differences, but they may partly reflect differences in observable worker or firm composition. To account for these compositional differences, I adjust for observables using two standard procedures. First, I use propensity score matching (PSM), where each administrative worker is matched to observably similar survey workers (based on education, age, gender, firm size, industry, and region). This improves local comparability but does not guarantee alignment of overall covariate means. Second, I use inverse probability weighting (IPW), where the administrative sample is reweighted to match the covariate means of the survey — consistent with the assumption that survey wages reflect true wages up to classical, mean-zero measurement error.

Table 3: Survey–Admin Wage Differences
(Peruvian Soles, S/.)

	Panel A: Means				**Panel B: Medians**			
	Admin	Survey	Diff		Admin	Survey	Diff	
Method:	(1)	(2)	(2–1)	(%)	(3)	(4)	(4–3)	(%)
Raw	1,893	2,230	337.0***	17.8***	1,500	1,967	467***	31.2***
PSM	1,974	2,230	256.1***	13.0***	1,569	1,967	398***	25.4***
IPW	2,148	2,230	82.5***	3.8***	1,671	1,967	296***	17.7***

Notes: Outcome is monthly wage. Positive coefficients indicate higher wages in the survey relative to administrative data. "Diff" is Survey minus Admin, and "Diff (%)" is 100 × (Survey−Admin)/Admin. Significance levels: $^{***}p < 0.01$. Means and medians are computed using the respective weights for each adjustment (unweighted, propensity-score matching, and inverse-probability weighting). $1 is equivalent to 3.4 Peruvian Soles, so the differences are between $25-$100. Minimum monthly wage in Peru is $320.

As shown in Table 3, under PSM, the wage gap narrows — the mean gap falls from 18 to 13 percent, and the median gap from 31 to 25 percent — but remains sizable. After IPW, mean survey wages remain 4 percent higher than administrative wages, and the median gap remains 18 percent. This suggests that administrative wages are still systematically lower even after accounting for observable composition. Figure A2 confirms that covariate balance improves progressively: raw standardized mean differences are sizable, shrink under PSM, and fall below the conventional 0.1 threshold under IPW. After IPW, the covariate *means* are statistically indistinguishable across the two datasets.

One could assume that underreporting is uniform across the wage distribution, in which case the mean and median comparisons above would be sufficient. However, there are strong reasons—both theoretical (e.g., enforcement nonlinearity, bargaining with firms) and empirical (e.g., Feinmann et al. (2025a))—to expect that PUT may have heterogeneous incidence across the distribution. I therefore examine where the underreporting occurs within the wage distribution.

[6] If the measurement error is classical (mean zero, and independent of w and covariates), the survey wage is an unbiased proxy for the true wage, and any positive survey–administrative gap can be interpreted as PUT. I will address about measurement error in the next subsection. Here the goal is to fix the ideas.

Figure 2: Standardized Mean Differences in Covariates

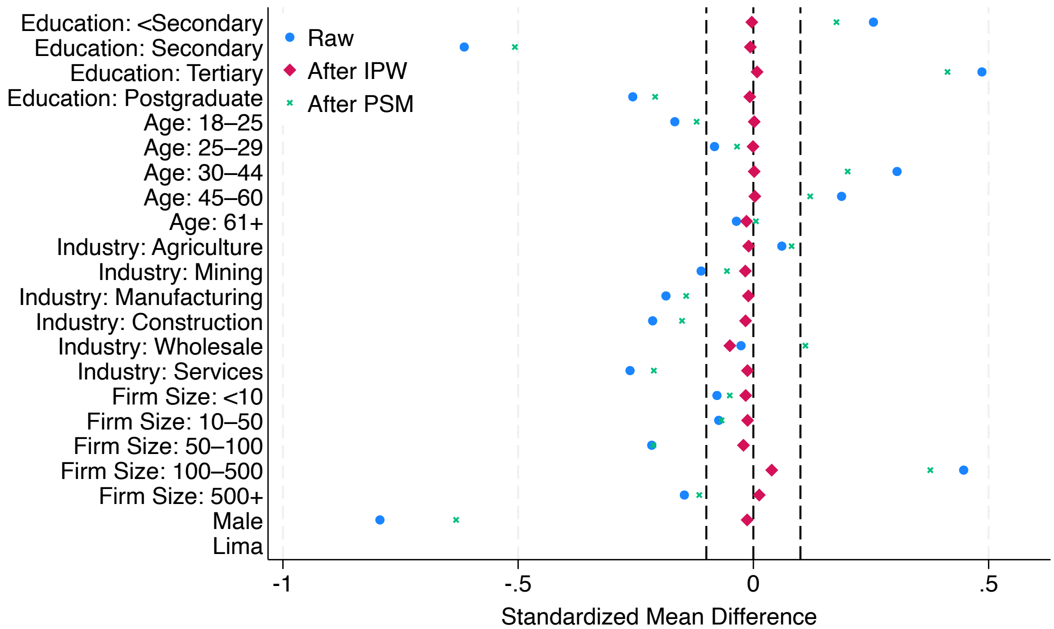

Specifically, after IPW reweighting—so that the two samples are comparable in terms of observable characteristics—I compare survey and administrative wages at different points of the distribution using two complementary diagnostics. Panel A of Figure 3 reports quantile gaps, computed as $\frac{Q_q^{\text{survey}} - Q_q^{\text{admin}}}{Q_q^{\text{admin}}} \times 100, \quad q \in \{10, 25, 50, 75, 90\}$. For each dataset, I obtain the wage corresponding to quantile q and then compute the relative difference using the expression above. A positive gap indicates that, at that quantile—for example, the 25th percentile among tertiary-educated workers—the survey wage is higher than the administrative wage. Panel A shows that for workers with less than secondary or secondary education, these gaps are close to zero and roughly constant across the distribution. In contrast, for tertiary-educated workers, the gaps are largest at the lower quantiles — meaning that many high-skill workers appear too far down in the administrative data relative to where they appear in the survey. In other words, the administrative distribution is shifted to the left, with the 10th, 25th, and 50th percentiles consistently lower in the administrative data than in the survey.

Panel B of Figure 3 reports CDF gaps, computed at the percentile cutoff q, c_q, for the pooled dataset as $F^{\text{admin}}(c_q) - F^{\text{survey}}(c_q)$. A positive value indicates that a larger share of workers in the administrative data fall below cutoff c_q than in the survey — in other words, the administrative distribution lies to the left of the survey distribution (i.e., it reports lower wages).[7] For workers with less than secondary or secondary education, the excess mass is small and relatively flat across the distribution. But for tertiary-educated workers, the excess mass rises starting around the median and peaks near the 75th percentile, implying that high-skill workers are disproportionately missing from the upper part of their own wage distribution in the administrative data (as compared to the survey distribution), consistent

[7]By construction, the CDF of each dataset must converge to 1 at the 100th percentile, so the CDF gap necessarily returns to zero in the limit. Also note that these are global (pooled) percentiles — not percentiles within each education group — so a value like the'75th percentile' corresponds to a common wage cutoff applied to both datasets.

with systematic downward misreporting.

These patterns are consistent with higher-earning workers having stronger incentives to underreport, making them disproportionately appear as low-wage workers in the administrative data. This is why only looking at specific moments (such as means or medians) is not enough, which is precisely what standard approaches like OLS, PSM, or IPW do. These methods align moments. However, the PUT issue is not a moments problem, it is a mapping problem: the administrative data systematically places mass in the wrong part of the distribution.

To see the limitation clearly, suppose a master's-educated worker truly belongs at the 90th percentile of the tertiary wage distribution — but is reported near the minimum wage in the administrative data. IPW ensures that the average characteristics of tertiary workers are the same across datasets, but it does not ensure that this particular worker is placed at the correct point in the distribution. So if I compare the 25th percentile in the administrative data to the 25th percentile in the survey, that misreported worker is effectively being matched to genuinely low-wage tertiary workers in the survey, not to their true position higher up the distribution. This would potentially underestimate PUT.

Optimal Transport is designed exactly for this problem. Instead of comparing percentiles or moments, OT directly learns where each worker should be placed in the distribution, that is, it recovers the wage that the misreported worker would have had in the absence of PUT. In the next subsection, I use OT to construct this one-to-one mapping from reported administrative wages to their imputed true wages, while preserving the distribution of observable characteristics.

Figure 3: Gap by education

(a) Wage Gap (mean) (b) CDF Gap

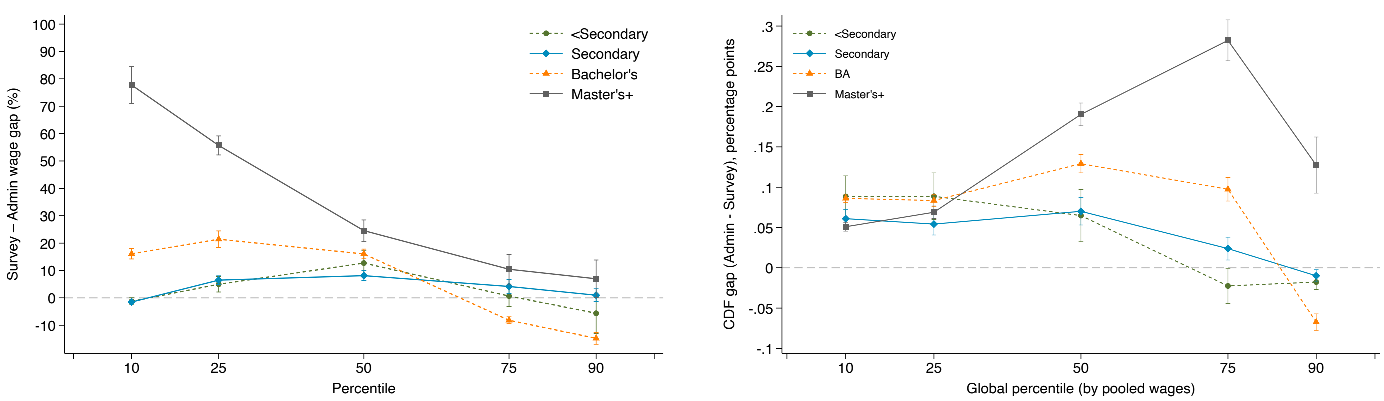

3.3. Recovering PUT via Optimal Transport

To recover the true wage w for individuals in the administrative dataset, I estimate a distribution-level matching function between administrative and survey wages, conditional on observable characteristics X. The goal is not merely to align moments—as in IPW or PSM—but to recover the wage each administrative worker would have if they were placed in the correct location of the survey wage distribution, given their covariates. In a sense, OT can be viewed as extending IPW from merely reweighting marginals to fully aligning the entire conditional wage distribution.

OT asks 'given this worker's observable characteristics, where would an observationally identical worker appear in the survey distribution, and what wage would they report there?' OT then transports that wage back to the administrative worker's existing rank position, preserving the survey distribution (Gunsilius, 2025; Galichon, 2016). Crucially, OT does not select a single nearest match. If multiple survey workers share the same covariates but earn different wages, OT splits the weight across them, allocating fractions in the way that minimizes the cost of transporting the entire distribution. The resulting imputed wage is thus a transport-weighted average, reflecting the full conditional distribution rather than a single moment or neighbor (Galichon, 2016).

Formally, the optimal transport problem finds the most efficient way to map the survey wage distribution onto the administrative one, subject to preserving the empirical distribution of observable characteristics X in both datasets. I define a distance between administrative observation i and survey observation j in the joint space of observable characteristics, using the L^1 norm:[8]

$$d_{ij} = \sum_k |X_{i,k} - X_{j,k} \tag{2}$$

The matching covariates X include firm size (5 categories), industry (6 categories), gender, schooling (4 education levels), region (Lima vs. non-Lima), age group (5 categories), and pension regime (public vs. private). These are the same covariates used for PSM and IPWP.

The optimal transport problem is:

$$\min_{\Gamma} \sum_{i=1}^{n} \sum_{j=1}^{m} \Gamma_{ij}\, d_{ij} \quad \text{subject to} \quad \Gamma\mathbf{1} = u, \quad \Gamma'\mathbf{1} = v, \quad \Gamma_{ij} \geq 0, \tag{3}$$

where $\Gamma \in \mathbb{R}^{n\times m}$ is the transport matrix, with Γ_{ij} the fraction of survey observation j assigned to administrative observation i. The vectors u and v are the empirical weights of the administrative and survey datasets, respectively (uniform for admin, sampling weights for survey).

Under mild regularity conditions (strictly positive marginals and a continuous cost function), this problem admits a unique solution Γ^* (Villani, 2009). Importantly, OT preserves the marginal distribution of observable characteristics X, ensuring that workers are only compared within the same covariate profile, but it allows the wage distribution to be fully reallocated (Gunsilius and Xu, 2023).

After solving the OT problem, I impute the true wage for each administrative worker as the transport-weighted average of survey wages:

$$\widehat{w}_i = \sum_{j=1}^{m} \Gamma_{ij}\, \tilde{w}_j \tag{4}$$

This preserves the full information contained in the optimal transport plan and yields a distribution-consistent estimate of the wage that each worker would report in the absence of PUT.

Assumptions. The identification of true wages via OT requires two assumptions. First, survey wages reflect true wages up to classical measurement error: $\tilde{w} = w + e$, where e has zero mean and is independent of w and X. Second, the conditional independence assumption (CIA) holds: conditional on X, the distribution of self-reported wages in the survey is the same as the distribution of true wages for

[8] Categorical variables are transformed into dummies so their absolute differences contribute to the distance.

observationally equivalent workers in the administrative data.

I assume Assumption 1 holds for now and focus on evaluating the credibility of the CIA, the empirically stronger requirement, as household surveys are known to systematically underrepresent high-income individuals. [9] Earlier, for the raw survey–administrative comparisons, I excluded the very top of the administrative wage distribution, where survey coverage visibly deteriorates. I now formalize and strengthen that step using a data-driven procedure to identify the precise upper bound of the administrative wage distribution where credible common support with the survey remains. Rather than imposing an arbitrary cutoff, I estimate this threshold directly from the data.

Given the size of the administrative dataset (about 4 million observations), I draw 1% random subsamples (approximately 20,000 observations each) for computational tractability. For each subsample, I iteratively restrict the data to observations below wage percentiles $p \in \{70, \ldots, 100\}$, in 1 p.p. increments. For each percentile threshold, I solve the OT matching problem and compute the maximum absolute standardized mean difference (SMD) across covariates. The optimal threshold is defined as the percentile that minimizes the maximum absolute SMD, that is, the point at which the survey and administrative datasets are most comparable conditional on observables. To assess robustness, I repeat this procedure across 100 independent subsamples of the administrative data.

Figure 4 reports the mean and standard error of the maximum absolute SMD at each candidate threshold. The results consistently point to an optimal threshold around the 96th percentile of the administrative wage distribution. I therefore restrict the administrative data to wages below the 96th percentile when performing the final OT mapping and constructing PUT estimates. This endogenous overlap restriction is fully consistent with best practice in applied econometrics, where trimming observations outside common support is routinely recommended to strengthen the plausibility of CIA.[10]

With CIA now empirically supported on the restricted sample, the formal properties of the estimator follow directly. Under Assumptions 1 and 2, the OT-imputed wage is a consistent estimator of the true wage:

$$\widehat{w}_i \xrightarrow{p} w_i \tag{5}$$

Moreover, the corresponding PUT estimate,

$$\widehat{\mathrm{PUT}}_i = \widehat{w}_i - \bar{w}_i, \tag{6}$$

inherits consistency and—because OT minimizes transport cost over the full joint distribution—it is asymptotically more efficient than nearest-neighbor, PSM, or regression-based imputation methods. Formal proofs are provided in Appendix MA1.

[9] I will later show that my results are robust to non-classical measurement error.

[10] See Imbens and Rubin (2015), Dehejia and Wahba (2002), Heckman et al. (1997), among others. These papers use data-driven trimming rules—e.g., discarding treated units whose propensity scores lie outside the support of the controls—to ensure credible overlap and improve confidence in CIA.

Figure 4: SMD for different wage percentiles

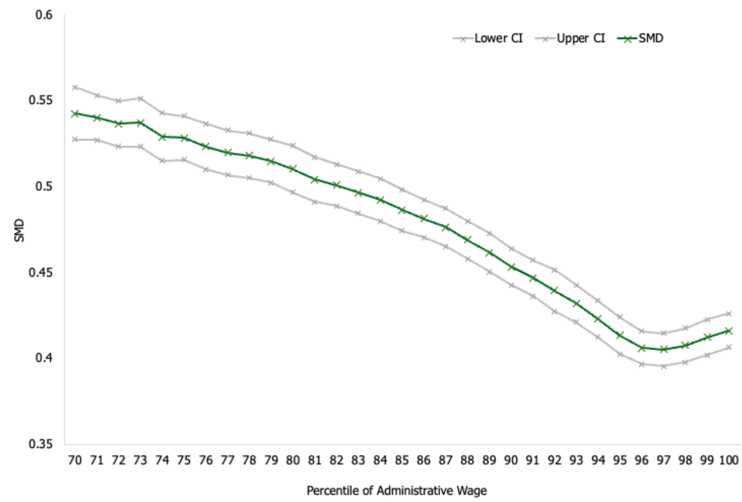

A methodological validation using synthetic data—where the true PUT is known by construction—is presented in Section 9.2 to illustrate the behavior of the OT estimator under controlled underreporting scenarios.

3.4. Results

I now estimate the extent of underreported wages in the Peruvian labor market. Recall that individual-level payments under the table are defined as

$$\widehat{\mathrm{PUT}}_i = \widehat{w}_i - \bar{w}_i, \tag{7}$$

where $\widehat{w}_i$ is the OT-imputed true wage and $\bar{w}_i$ is the third-party reported wage from the administrative records. I begin by documenting the overall prevalence and average magnitude of PUT among formal private-sector workers and then examine how it varies across the wage distribution and observable characteristics.

Figure 5 presents estimates of PUT prevalence which is the ratio of PUT to reported income. To compute the standard errors of the PUT estimator $\hat{\mathrm{PUT}}$, I use a bootstrap procedure, resampling the household survey dataset 100 times. In each iteration, I reapply the OT method and calculate the transported (attributed) wage for each individual in the payroll dataset.

On average, PUT accounts for 33% of reported wages. Panel A of Figure 5 presents PUT estimates by education level. PUT prevalence is positive and statistically significant across all education groups. The results show a strong positive relationship between education and the share of wages paid under the table: PUT ranges from 24% for workers with lower education to 113% for those with postgraduate degrees. However, it's important to note that workers with a Master's degree or higher make up less than 1% of the sample. Panel B shows that results are broadly similar across gender. Women with secondary education underreport more, while among those with a bachelor's degree, men underreport more.

Figure 6 shows the estimates of PUT prevalence by firm size and regime type (general, agricultural, and the special micro- and small-enterprise regimes). In general, PUT prevalence is lower in firms with more workers, consistent with (Kleven et al., 2016), as collusion is easier among fewer parties. When disaggregated by regime, PUT prevalence is lowest in the general and agricultural regimes, the two regimes where the gap between the aggregate income tax and the corporate tax is smallest. In contrast, in the special regimes—where the corporate tax is significantly lower and the gap with the income tax is larger—PUT prevalence is higher, even among larger firms.

Figure 5: PUT Prevalence

(a) by education

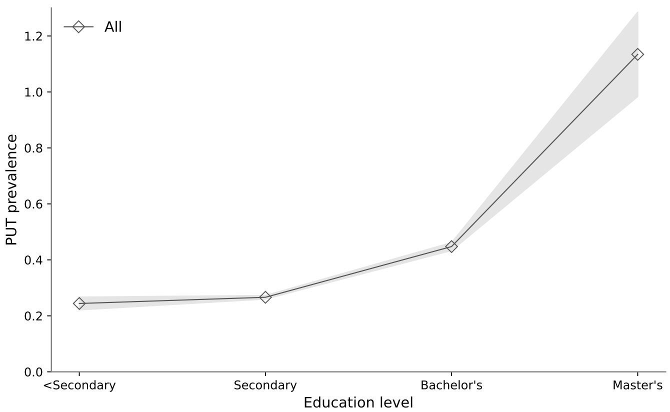

(b) by education & gender

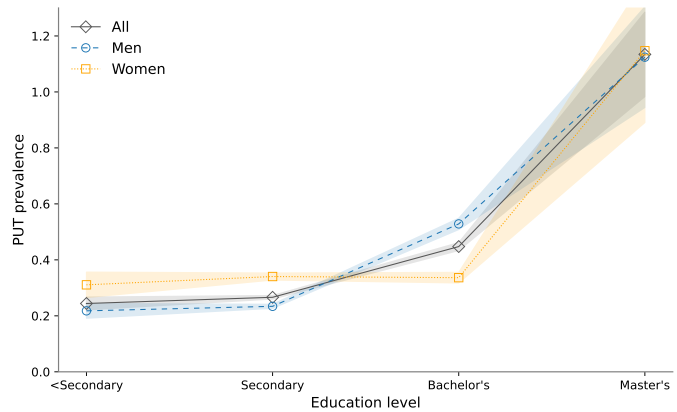

Figure 6: PUT Prevalence

(a) by firm size & regime

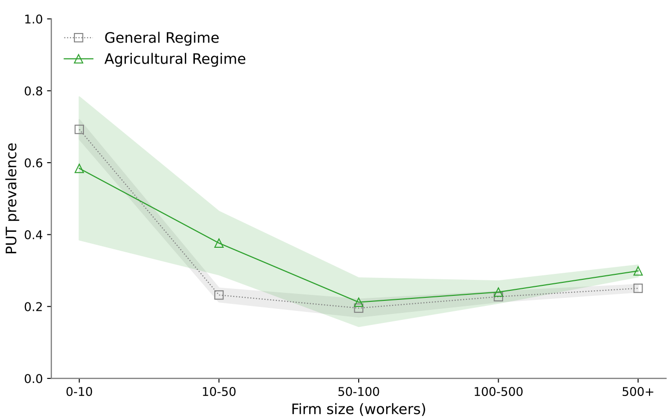

(b) by firm size & regime

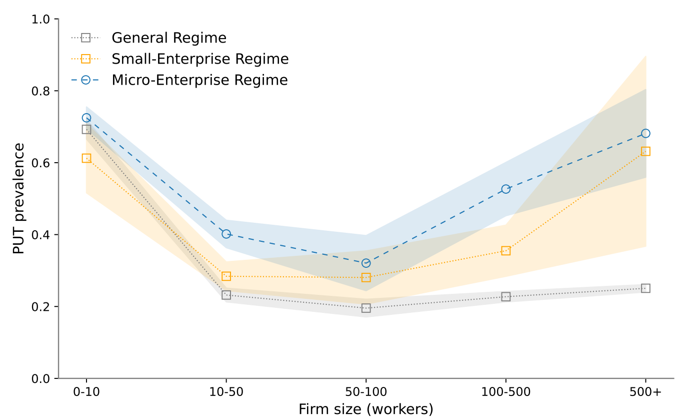

Results are robust to stress tests for measurement error. For classical error, I add mean-zero noise to survey incomes with standard deviation $\sqrt{\lambda}\,\hat{\sigma}_e$, where $\hat{\sigma}_e$ is the residual SD from an OLS of survey income on the covariates used in OT. Neither the average PUT nor the schooling gradient loses significance for $\lambda \leq 10$. In other words, survey noise would need to be at least $\sqrt{10}\,\hat{\sigma}_e \approx 3.16\,\hat{\sigma}_e$ to overturn the results (see Figure A1). For non-classical error, I impose systematic misreporting in the survey top p by scaling reported income to $(1 \pm \theta)y$, then recompute PUT with the OT plan held fixed. To erase the average PUT via overreporting, the top 30% would need to overreport by about 50% ($p = 0.30,\ \theta \approx 0.50$), or the top 50% by about 35% ($p = 0.50,\ \theta \approx 0.35$) (see Figure A2). Underreporting works in the opposite direction: because the plan depends only on X, scaling survey wages down lowers transported income and therefore lowers measured PUT, so our estimates are conservative lower bounds if survey underreporting is present.

Overall, the evidence shows that PUT is quantitatively important, systematically related to both worker and firm characteristics, and concentrated near key institutional thresholds such as the minimum wage. Given the prevalence, this behavior should be considered by governments when designing tax policy. These stylized facts provide the empirical grounding for the model developed in the following section.

4. Optimal Tax Theory with PUT

4.1. Model

I develop a static general-equilibrium model in which PUT arise endogenously in equilibrium, and use it to show how its presence alters the standard optimal labor income tax formula.

The population is normalized to one. Individuals differ in productivity ψ, which for simplicity is assumed to have only mass in $\{0, 1\}$ with exogenous shares s_0 and s_1. Individuals with $\psi = 0$ do not supply labor and consume only a universal lump-sum transfer T. Individuals with $\psi = 1$ are workers. They may work either formally or in an arrangement that involves underreporting in the form of PUT. Their choice determines labor supply, and how much of their wage is concealed from the government. Firms are competitive, choose labor demand, and take wages as given. The government sets a linear labor income tax on reported earnings, a linear corporate tax on firm profits, and finances the transfer T. The income tax includes both the personal income tax and the payroll taxes. There are no labor market frictions. In equilibrium, wages and allocations are pinned down by market clearing in both the formal and PUT sectors.

Workers Workers are equally productive but differ in how much they value the benefits of formality. I index this heterogeneity by θ, distributed with cdf Q and pdf q. A higher θ means the worker places greater value on accruing those benefits and therefore faces a higher perceived loss from underreporting.

Workers decide whether to work in a formal or a PUT arrangement.[11] For each arrangement, they first solve for optimal hours, and, in PUT, the hidden wage. Then, they pick the arrangement with the higher indirect utility. If a worker works in a formal arrangement, she gets utility $u^F = (1 - \tau)wl + T - v(l)$, where w is the wage, l are the hours of work supplied, τ is the linear personal income tax, T is the universal lump-sum transfer, $(1 - \tau)wl$ is the after-tax income (consumption). v is a strictly increasing function in labor supply. Workers choose l^F that maximizes u^F. The optimization of a worker the formal arrangement for a given wage w^F yields the following first order condition[12]:

$$\max_{l}\{(1 - \tau)w^F l + T - v(l)\} \tag{8}$$

[11] I refer to these as 'arrangements' rather than 'submarkets' because they coexist within the same labor market and differ mainly in compliance conditions; wage-setting will nonetheless differ across them.

[12] Throughout the paper, l denotes individual hours of labor supplied, and L denotes *aggregate* hours supplied or demanded, integrating across all workers in a given arrangement. The aggregate labor supply is thus a combination of the extensive margin (how many workers select each arrangement, determined by the threshold (θ^*) and the intensive margin (hours per worker, l).

$$l: \quad (1-\tau)w^F = v'(l) \tag{9}$$

and the corresponding indirect utility function is $V^F(\tau) = (1-\tau)w^F l^F - v(l^F)$, where l^F solve the FOC.

If a worker chooses a PUT arrangement, she gets utility $u^P = wl - \tau(w - \text{PUT})l + T - v(l) - g(\theta, \text{PUT}\cdot l)$. $\text{PUT} \in [0,\, w^P]$ is the amount of earnings hidden or paid under the table wage. g is a strictly increasing and convex cost function in the hidden wage. The term $g(\theta, \text{PUT}\cdot l)$ is interpreted as the perceived loss of contributory benefits when earnings are not reported. In the PUT arrangement, the personal income tax is only levied on the reported wage which is equivalent to $w - \text{PUT}$. Given w^P, the worker chooses (l, PUT) subject to $0 \le \text{PUT} \le w^P$. The optimization of a worker in the PUT arrangement for a given wage w^P yields the following first order conditions:

$$\max_{l,h}\{w^P l - \tau\big(w^P - \text{PUT}\big)l + T - v(l) - g(\theta,\, \text{PUT}\cdot l)\} \tag{10}$$

$$l: \quad (1-\tau)w^P + \tau\text{PUT} = v'(l) + g'_l(\theta, \text{PUT}\, l) \tag{11}$$

$$\text{PUT}: \quad \tau = g'_{\text{PUT}}(\theta, \text{PUT}\cdot l) \tag{12}$$

and the corresponding indirect utility function is $V^P(\tau,\theta) = (1-\tau)w^P l^P + \tau\,\text{PUT}^* l^P - v(l^P) - g\big(\theta,\, \text{PUT}^*\, l^P\big)$, where (PUT^*, l^P) solve the FOCs.

A worker chooses the PUT arrangement if $V^P(\tau,\theta) > V^F(\tau)$, which defines a threshold θ^*. Workers choose PUT only if $\theta < \theta^*$. The aggregate hours supplied in the PUT arrangement are $L^{S,P} = s_1 \int_0^{\theta^*} l^P(\tau,\theta)\, q(\theta)\, d\theta$, which sums hours over all workers who select PUT. Here s_1 is the share of individuals who are workers. The aggregate hours supplied in the formal arrangement are $L^{S,F} = s_1\big(1 - Q(\theta^*)\big)\, l^F$. In the formal arrangement, hours are the same across workers because productivity and wages are identical. In the PUT arrangement, however, individual hours l$^P canvarywith\theta$, since workers take into account their own cost of informality (i.e., the value they place on lost payroll benefits). At a given wage, those with higher θ effectively receive lower compensation and may be willing to supply fewer hours.

Non-workers Individuals who do not work receive a lump-sum transfer T financed by the labor income tax and the corporate tax, so $u^N = T$. Let s_0 denote the share of non-workers in the population; the aggregate transfer to non-workers is $s_0 T$.

Firm There is a representative firm with production function $\phi(\cdot)$ that uses labor and a fixed (short-run) capital stock to produce output. The firm is a price-taker in the product market (output price normalized to 1) and takes wages (w^F, w^P) as given, determined in labor-market equilibrium. The function ϕ is strictly concave, so marginal product of labor falls with scale and the firm earn positive profits. Let total labor be $L = L^F + L^P$, where L^F and L^P are the hours employed in the formal and PUT arrangements, respectively. Pre-tax profits are

$$\pi^{\text{before-tax}} \;=\; \phi(L) - w^F L^F - w^P L^P \tag{13}$$

Payroll expenses are deductible only when reported. I assume PUT wages are non-deductible.[13] $\bar{\text{PUT}}$

[13]If a fraction $d \in [0,1]$ of hidden payroll is effectively deducted via disguised expenses, $\bar{\text{PUT}}$ is replaced by $(1-d)\bar{\text{PUT}}$ in the

denote the average earnings paid under the table to workers in PUT jobs (taken as given by the firm). Reported profits are

$$\widehat{\pi} \;=\; \phi(L) - \big(w^P - \bar{\text{PUT}}\big)L^P - w^F L^F \tag{14}$$

so with a linear corporate tax rate t, after-tax profits are

$$\Pi^{\text{after-tax}} \;=\; \pi - t\,\widehat{\pi} \;=\; (1-t)\left[\phi(L) - w^F L^F - w^P L^P\right] \;-\; t\,\bar{\text{PUT}}\,L^P \tag{15}$$

The firm chooses L^F and L^P:

$$\max_{L^F, L^P} \;(1-t)\big[\phi(L) - w^F L^F - w^P L^P\big] - t\,\bar{\text{PUT}}\,L^P \tag{16}$$

with first-order conditions

$$L^F:\;(1-t)\big[\phi'(L) - w^F\big] = 0 \;\;\Rightarrow\;\; \phi'(L) = w^F \tag{17}$$

$$L^P:(1-t)\big[\phi'(L) - w^P\big] - t\,\bar{\text{PUT}} = 0 \;\;\Rightarrow\;\; \phi'(L) = w^P + \frac{t}{1-t}\,\bar{\text{PUT}} \tag{18}$$

Combining,

$$w^F \;=\; w^P + \frac{t}{1-t}\,\bar{\text{PUT}} \tag{19}$$

I call equation 19, the PUT-formal wage wedge. The firm takes $\bar{\text{PUT}}$ as given when choosing labor demand: it depends on workers' hiding decisions and is determined in equilibrium. As workers are equally productive, the marginal product of labor $\phi'(L)$ is the same irrespective of the arrangement. Because only reported payroll is deductible, hiring a PUT hour imposes an extra non-deductible cost of $\frac{t}{1-t}\bar{\text{PUT}}$ on the firm. For PUT to exist in equilibrium, the PUT wage must be lower than the formal wage. Otherwise, the firm would have no incentive to hire workers under the table [14]. Given $\bar{\text{PUT}}$, the conditions above pin down the firm's labor demands in each arrangement, $L^{D,F}\big(w^F, t\big)$ and $L^{D,P}\big(w^P, t, \bar{\text{PUT}}\big)$. After-tax profits at the optimum are $V^{\pi}(t) \;=\; (1-t)\big[\phi(L^{D,F} + L^{D,P}) - w^F L^{D,F} - w^P L^{D,P}\big] - t\,\bar{\text{PUT}}\,L^{D,P}$.

Equilibrium The equilibrium $(w^F, w^P, L^F, L^P, \theta^*)$ satisfies workers' and firm's optimality conditions and simultaneously clears the formal and PUT labor arrangements. The equilibrium conditions are given by $L^{D,P}(w^P, t, \bar{\text{PUT}}) \;=\; L^{S,P}(w^P, \tau, \theta^*) \;=\; L^P$ and $L^{D,F}(w^F, t) \;=\; L^{S,F}(w^F, \tau, \theta^*) \;=\; L^F$, where $L^{S,P}$ and $L^{S,F}$ represent total labor hours supplied by workers, and $L^{D,P}$ and $L^{D,F}$ represent total labor hours demanded by the firm in each arrangement. The market clearing conditions determine the equilibrium wages (w^P, w^F) and the equilibrium allocations (L^P, L^F).

Aggregate PUT are given by the integral of PUT wages across workers hired in the PUT arrangement:

$$\text{PUT}^{\text{agg}} \;=\; s_1 \int_0^{\theta^*} \text{PUT}^*(\tau, \theta; w^P) l^P(\tau, \theta; w^P)\, dQ(\theta). \tag{20}$$

Thus, equilibrium also determines the threshold θ^* that divides workers across formal and PUT arrangements. Changes in the personal income tax τ shift labor supply and the threshold θ^*, while changes in

FOC.

[14]This does not imply that the worker's take-home pay is lower in the PUT arrangement. Since a portion of the wage is not subject to taxation, their effective compensation could still exceed the net-of-tax formal wage.

the corporate tax t shift firm labor demand directly and, via $\bar{\text{PUT}}$, widen the PUT-formal wage wedge.

Planner The social planner chooses τ (and the transfer T) to maximize a generalized utilitarian social welfare function (SWF). The SWF is given by:

$$\begin{aligned} SWF = \; & s_0 \omega_L G(T) \\ & + s_1 \omega_L \int_0^{\theta^*} G\Big((1-\tau) w^P l^P + \tau h l^P - v(l^P) - g(\theta, \text{PUT}\, l^P)\Big) dQ(\theta) \\ & + s_1 \omega_L \int_{\theta^*}^{\infty} G\Big((1-\tau) w^F l^F - v(l^F)\Big) dQ(\theta) \\ & + \omega_\pi G\Big((1-t)[\phi(L^P + L^F) - w^P L^P - w^F L^F] + t\, \text{PUT}\, L^P\Big) \end{aligned} \tag{21}$$

where G is an increasing and concave function, which which, together with the exogenous Pareto weights on individuals Ω_L and the firm Ω_π summarizes the social preferences for redistribution. L^F and L^P are market-clearing labor demands/supplies, and $\text{PUT}^*(\theta)$ denotes optimal PUT per hour for a worker of type θ (so total PUT dollars are $\text{PUT}^*(\theta)\, l^P(\theta)$). The planner internalizes that $(w^F, w^P, L^F, L^P, \theta^*)$ are equilibrium objects, determined by workers' and firm's FOCs and market clearing. With reported earnings $w^F L^F + (w^P - \bar{\text{PUT}}) L^P$ and reported profits $\widehat{\pi} = \phi(L^F + L^P) - (w^P - \bar{\text{PUT}}) L^P - w^F L^F$, the balanced-budget condition is

$$T \;=\; \tau\Big[w^F L^F + \big(w^P - \bar{\text{PUT}}\big) L^P\Big] \;+\; t\,\widehat{\pi} \tag{22}$$

Let λ be the multiplier on the government budget constraint. The average social welfare weight of non-workers, workers in PUT arrangement, workers in formal arrangement, and firms, are defined respectively as:

$$\begin{aligned} g_N = \frac{\omega_L\, G'(T)}{s_0 \lambda}, \qquad & g_L^P \;=\; \frac{\omega_L \displaystyle\int_0^{\theta^*} G'\big(V^P(\theta)\big)\, dQ(\theta)}{s_1\, Q(\theta^*)\, \lambda} \\ g_L^F \;=\; \frac{e \omega_L \int_{\theta^*}^{\infty} G'\big(V^F\big)\, dQ(\theta)}{s_1\, (1 - Q(\theta^*))\, \lambda}, \qquad & g_\pi = \frac{\omega_\pi\, G'(V^\pi)}{\lambda} \end{aligned} \tag{23}$$

The welfare weights represent the social value of the marginal utility of consumption normalized by the social cost of raising public funds, thus measuring the social value of redistributing one dollar across a group of individuals. At the optimum, the planner is indifferent between giving one more dollar to an individual in group i or having g_i more dollars of public funds.

4.2. Optimal Income Tax

Before deriving the optimal income tax, I define the elasticities (all with respect to the log net-of-tax rate):

$$e_1^i \;=\; \frac{\partial \log L^i}{\partial \log(1-\tau)}, \quad i \in \{P, F\}; \qquad e_2 \;=\; \frac{\partial \log \overline{\text{PUT}}}{\partial \log(1-\tau)} \tag{24}$$

e_1 is the elasticity of labor supply and e_2 is the elasticity of average PUT with respect to the net-of-tax rate (intensive margin of PUT evasion). These are general-equilibrium elasticities that include all quantity adjustments in their respective arrangements.

Two remarks about elasticities. First, e_1^i aggregates all labor-supply margins: a change in $1-\tau$ shifts both the number of workers in arrangement i (selection) and their hours (intensive response). I do not separate these components. Second, the sign of e_1^F is positive. The sign of e_1^P is a priori ambiguous. When $1-\tau$ increases, workers supply more hours in both arrangements due to the Frisch elasticity of labor supply, which pushes e_1^P in the positive direction. However, a higher net-of-tax wage also makes hiding less attractive, so fewer workers choose the PUT arrangement and L^P declines, pulling e_1^P down. I assume that the Frisch elasticity is smaller, so the net effect is negative, implying $e_1^P < 0$.

The wage elasticities with respect to the net-of-tax rate are defined as follows[15]:

$$\eta^P = \frac{\partial \log w^P}{\partial \log(1-\tau)}, \qquad \eta^F = \frac{\partial \log w^F}{\partial \log(1-\tau)} \tag{25}$$

Because a higher τ pushes workers toward PUT, formal wages rise and PUT wages fall. Since elasticities are with respect to $\log(1-\tau)$, under the same assumptions as for labor-supply elasticities, this implies $\eta^F < 0$ and $\eta^P > 0$.

Aggregate income in arrangement i is $Z^i = w^i L^i$. Aggregate reported income is

$$\hat{Z} = \underbrace{w^F L^F}_{\hat{Z}^F} + \underbrace{\left(w^P L^P - \mathrm{PUT}^{\mathrm{agg}}\right)}_{\hat{Z}^P}, \tag{26}$$

and define the shares (of the reported base)

$$\alpha^F = \frac{Z^F}{\hat{Z}}, \qquad \alpha^P = \frac{Z^P}{\hat{Z}}, \qquad \gamma = \frac{\mathrm{PUT}^{\mathrm{agg}}}{\hat{Z}}, \quad \text{so} \quad \alpha^F + \alpha^P - \gamma = 1 \tag{27}$$

I refer to γ as PUT prevalence. Using the model, I derive the conditions for an optimal personal income tax in this economy. The Appendix shows the full derivation. From the first-order conditions and using the envelope theorem for the workers and the firm maximization, I get the following expression for the optimal labor income tax:

$$\tau^* = \frac{1 - \bar{g} - \Omega - t\,\gamma\Big((1-g_\pi)e_2\Big)}{1 - \tilde{g} - \Omega + \underbrace{\alpha^P(e_1^P + \eta^P) + \alpha^F(e_1^F + \eta^F) - \gamma(e_1^P + e_2)}_{\text{ETI}}} \tag{28}$$

$\bar{g} = g^P \alpha^P + g^F \alpha^F$, where g_L^P and g_L^F are the average social marginal welfare weights on PUT and formal workers, respectively. g_π the weight on firm profits. The term $t\,\gamma\,(1-g_\pi)\,e_2^P$ captures the corporate–evasion interaction: changes in τ change PUT, and because PUT payroll is non-deductible, this shifts the CIT base; $(1-g_\pi)$ converts that into social value.

Equation (28) is the classical efficiency–equity tax formula with evasion and a fiscal externality, as in Piketty et al. (2014). It balances redistributive gains (numerator) against efficiency costs (denominator), with the novelty that both are shaped by general-equilibrium responses that are functions of the corporate tax rate. In the numerator, a higher welfare weight on workers ($\bar{g}$) lowers the optimal tax rate, as in the standard case. The new element is the positive fiscal externality: PUT wages increase taxable profits

[15] These are also general-equilibrium elasticities: η^i captures all equilibrium adjustments in wages in arrangement i when the net-of-tax rate changes.

and therefore corporate tax revenue. This makes the corporate tax rate t enter directly, with the effect scaled by PUT prevalence (γ), the elasticity of evasion ($e_2 < 0$), and the welfare weight on firms (g_π). The higher the welfare weight on firms, the smaller the social value of this externality.

Ω captures the welfare effect of wage adjustments induced by τ:

$$\Omega \;=\; \alpha^F\Big(\tfrac{1-t}{1-\tau}g_\pi - g_L^F\Big)\eta^F \;+\; \alpha^P\Big(\tfrac{1-t}{1-\tau}g_\pi - g_L^P\Big)\eta^P \tag{29}$$

When w^i moves by dw^i in arrangement $i \in \{F, P\}$, workers gain $(1-\tau)L^i\,dw^i$ while the firm loses $(1-t)L^i\,dw^i$. If the planner values workers more than the firm after tax (i.e. $g_L^i > \frac{1-t}{1-\tau}g_\pi$), the formal term is positive while the PUT term is negative. Hence $\Omega > 0$ when the formal-side response dominates (decreasing τ^*), and $\Omega < 0$ when the PUT-side response dominates (increasing τ^*).

Define the ETI as the log elasticity of aggregate reported income $\hat{Z}$ with respect to the net-of-tax rate $1-\tau$. A small increase in $1-\tau$ changes $\hat{Z}$ through: (i) labor supply and wages in each arrangement, and (ii) the amount of income hidden under the table. Aggregating these channels yields:

$$\text{ETI} \;=\; \alpha^P(e_1^P + \eta^P) + \alpha^F(e_1^F + \eta^F) - \gamma(e_1^P + e_2) \tag{30}$$

where the first two terms capture formal and PUT income responses and the last term captures the effect of moving income under the table

The ETI combines labor-supply elasticities e_1^i, wage elasticities η^i, and the PUT evasion elasticity e_2. These are general equilibrium elasticities and are weighted by income shares and PUT prevalence. Equation (28) shows that the ETI captures all the efficiency costs associated with raising the income tax rate. In this sense, the ETI is a sufficient statistic of this model. Once the ETI is known, the welfare effect of changing $(1-\tau)$ can be computed without specifying the full microstructure of preferences or informality costs. The difference is that the ETI depends on structural parameters related to PUT behavior between workers and firms, as well as wage effects that vary with the corporate tax rate.

Special Case. If wages do not respond ($\eta^P = \eta^F = 0$) and the corporate tax is zero ($t = 0$), (28) reduces to

$$\tau^* \;=\; \frac{1-\bar{g}}{1-\bar{g} \;+\; \underbrace{\alpha^P e_1^P + \alpha^F e_1^F - \gamma(e_1^P + e_2^P)}_{\text{ETI}}}, \tag{31}$$

which is identical to the well-known standard optimal linear tax rate in Mirrleesian frameworks (Piketty and Saez, 2013). The Appendix 9.3 shows that deriving (28) with perturbation methods—considering a marginal increase in τ—yields the same result.

ETI as a function of the corporate tax rate t The corporate tax affects both the equilibrium allocation between formal and PUT arrangements and the relevant elasticities. To show the mechanisms, I differentiate the ETI with respect to t and denote derivatives as $x' = \partial x/\partial t$:

$$\begin{aligned}\frac{d\,\text{ETI}}{dt} = &\underbrace{\alpha^{P\prime}(e_1^P + \eta^P) \;+\; \alpha^{F\prime}(e_1^F + \eta^F) \;-\; \gamma'(e_1^P + e_2)}_{\text{Composition}} \\ &+ \underbrace{\alpha^P(e_1^{P\prime} + \eta^{P\prime}) \;+\; \alpha^F(e_1^{F\prime} + \eta^{F\prime}) \;-\; \gamma(e_1^{P\prime} + e_2')}_{\text{Behavioral}} < 0\end{aligned} \tag{32}$$

Appendix 9.4 shows the proof of $\frac{d\,\text{ETI}}{dt} < 0$. Here I develop the intuition. A higher corporate tax lowers the ETI through two forces. First, it shifts the composition of the labor force: fewer workers choose the PUT arrangement. A higher t raises the effective cost of PUT labor (since hidden wages are not deductible), lowering w^P relative to w^F and shifting employment toward the formal sector. Thus, for a given informality cost, the PUT arrangement is less attractive to workers, making the share of workers in PUT fall ($\alpha^F \downarrow,\ \alpha^P \downarrow,\ \gamma \downarrow$). As a result, when the income tax changes, there are fewer workers in the PUT arrangement who can adjust their level of evasion, making the ETI less responsive.

Second, a higher corporate tax alters how workers respond to changes in the income tax, which directly affects the ETI. It raises the effective cost of hiring PUT workers and of paying a dollar under the table. As a result, firms are less willing to offer these contracts, and workers become less able to shift into the PUT arrangement when the income tax rate changes. Since evasion requires agreement from both sides of the market, the behavioral elasticities related to concealment fall ($e_1^P \downarrow$, $e_2 \downarrow$). This makes reported income less responsive to income tax changes, and therefore lowers the ETI. Together, a higher t both shrinks the size of the PUT sector and reduces its elasticity, making reported income less responsive to changes in the income tax (ETI $\downarrow$).

Mapping t- to τ-elasticities. A key feature of the model is that all relevant behavioral margins—labor supply in each arrangement, hidden wages, and equilibrium wages—are jointly determined through a single surplus condition. Define the employer–employee surplus wedge as

$$S \;=\; \tau - \frac{t}{1-t} \tag{33}$$

where the first term captures the personal income tax burden and the second term reflects the effective non-deductibility of hidden pay.

Totally differentiating the equilibrium system shows that first-order responses of all endogenous variables are proportional to changes in dS. Appendix 9.5 show the algebra. A change in the corporate tax t and a change in the personal tax τ affect the economy only through their joint impact on S:

$$dS = d\tau - \frac{1}{(1-t)^2}dt \tag{34}$$

$$(1-\tau)\,d\log(1-\tau) \;=\; \frac{1}{1-t}\,d\log(1-t) \tag{35}$$

In other words, a one-percent increase in firms' deductibility $(1-t)$ generates the same change in the surplus as a $(1-\tau)(1-t)$-percent increase in workers' net-of-tax rate $(1-\tau)$. Thus t-elasticities can be directly mapped into τ-elasticities as

$$e^{1-\tau} \;=\; (1-\tau)(1-t)\,e^{1-t} \tag{36}$$

This mapping shows that t-elasticities are also sufficient statistics for the efficiency costs of taxation, once scaled by $1/[(1-\tau)(1-t)]$. This gives the planner an alternative way of estimating the ETI even in the absence of income tax reforms: observed responses to corporate-tax reforms contain the same information about efficiency costs as income tax reforms, after applying this conversion.

5. ETI Estimates

In this section, I take the theoretical model to the data. As shown in the previous section, the ETI remains a sufficient statistic in an economy with PUT, but it becomes a negative function of the corporate tax rate. This section estimates the ETI and empirically tests its relationship with the corporate tax rate.

5.1. Identification

To estimate the ETI, I use a staggered difference-in-differences (DiD) design and an event-study framework. The identifying variation comes from the 2021 reform to Peru's Agricultural Regime.

Since 2001, *eligible* workers in agricultural firms could be hired under a special agricultural regime in which they were subject to a 4% healthcare (EsSalud) contribution, compared to 9% under the general regime. Eligible workers include those engaged in farming, livestock, and agro-industrial activities carried out in the field. However, administrative and technical support staff who do not perform their duties primarily in the firm's fields are excluded and must be hired under the general regime. The corporate income tax rate for all agricultural firms was also reduced to 15%, compared to 29.5% under the general regime. This Agricultural Regime was originally set to expire in December 2021.[16]

In 2019, health authorities repeatedly noted that the cost per worker hired under the agricultural regime far exceeded their contributions and called for an increase.[17] In December 2020, Congress passed a reform that increased payroll tax rates (EsSalud contributions) for workers in agricultural firms hired under the Agricultural Regime, with rates differentiated by firm size. Firms with fewer than 100 workers saw their healthcare contribution rate rise to 6% starting in 2021, while firms with 100 or more workers faced a staggered increase: 7% in 2021–2022 and 8% from 2023 onward. The corporate tax rate remained unchanged during 2021–2022 but was raised to 20% for large firms starting in 2023. Table 4 summarizes these changes.

Table 4: Marginal Tax Rates

	Agricultural Regime (AgR)				General Regime (GR)	
	Healthcare		Corporate Tax		Healthcare	Corporate Tax
Year	< 100 workers (1)	> 100 workers (2)	< 100 workers (3)	> 100 workers (4)	All (5)	All (6)
2020	4%	4%	15%	15%	9%	29.5%
2021	6%	7%	15%	15%	9%	29.5%
2022	6%	7%	15%	15%	9%	29.5%
2023	6%	8%	15%	20%	9%	29.5%

As in the theoretical model, I treat the effective tax on labor income as the sum of the personal income tax and the payroll tax $\tau = \tau^{\text{PIT}} + \tau^{\text{payroll}}$. Healthcare (EsSalud) coverage and service entitlements are not

[16] These benefits were initially set to expire on December 31, 2010, but their validity was later extended until December 31, 2021.

[17] EsSalud reported that as of December 2019, for every S/100 contributed by agricultural firms, workers consumed S/262 in services—meaning healthcare costs were 2.6 times their contributions. These workers were subject to a 4% contribution rate, rather than the standard 9% for other sectors (RPP Noticias, 2020).

a function of the employer contribution rate at the margin; workers receive the same benefits irrespective of whether $\tau^{\text{payroll}} = 4\%, 6\%, 7\%, 8\%$. Hence the contribution behaves like a tax on labour rather than a priced benefit. Throughout, I therefore interpret changes in the payroll tax as changes in τ.

To quantify how the mechanical tax changes (driven by changes in the tax laws) affect reported income, I estimate a continuous difference-in-differences model that exploits two sources of intensity variation within the agricultural sector: (i) whether a worker is hired under the special agricultural regime or the general regime (I call this the 'regime' variable), and (ii) whether the firm has fewer or more than 100 employees (I call this the 'firm-size' variable). Restricting to agriculture rules out cross-sector productivity differences.

Following Kleven and Schultz (2014), I first estimate the following specification on 2020–2022:

$$\log \hat{z}_{ijt} = \alpha + \varepsilon \, \log(1 - \tau_{jt}) + X_{ijt}\gamma + \delta_j + \delta_t + \nu_{ijt}. \tag{37}$$

Here, $\hat{z}_{ijt}$ is reported labor income of worker i in firm-size$\times$regime cell j and year t, and τ_{jt} is the statutory income-tax rate for cell j and year t. In this setting the statutory rate τ_{jt} is flat within a cell$\times$year, so $\log(1 - \tau_{jt})$ does not depend on individual income $\hat{z}_{ijt}$. Unlike in nonlinear schedules, there is no virtual-income term and no need for simulated IV; the regressor is mechanically determined by law. Identification comes from continuous DiD variation in the dose $\log(1 - \tau_{jt})$ across regime$\times$size cells over time.

I include group fixed effects δ_j (firm-size$\times$regime) and year fixed effects δ_t, plus regional fixed effects. $\varepsilon = \partial \ln \hat{z} / \partial \ln(1 - \tau)$ is the ETI: a 1% increase in $1 - \tau$ is associated with an ε% increase in reported income. The 2020–2022 estimates capture the elasticity of taxable labor income when there are no changes in the corporate tax. I then re-estimate the same specification on 2022–2023, but restricting the sample only to firms in the AgR. ε still measures the ETI but in a setting where the corporate tax also increase. In 2023, large firms in the agricultural regime faced a 1 p.p. increase in the income tax and a concurrent increase in the corporate tax from 0.15 to 0.20. This setting lets me estimate the ETI off a 1 p.p. change in the net-of-income-tax rate while corporate tax also moved. Small firms in the agricultural regime saw no change in either tax. Under the assumption that the ETI is stable over this short window, any difference from the 2020–2022 estimate is informative about how the concurrent corporate-tax change interacts with the income-tax response. Under the 'surplus-wedge' logic of the model, the concurrent corporate-tax increase should attenuate the measured response to labor income-tax changes.

5.2. Results

Table 5 shows the results. Column (1) reports the ETI for 2020–2022. Firms in the Agricultural Regime (AgR) are treated with the change in the net-of-tax rate $1 - \tau$, while firms in the General Regime (GR) operating in agriculture serve as controls because their income-tax schedule does not change over this window. The ETI is identified from within-unit changes for AR between 2021–2022 relative to the pre period in 2020, net of the change for GR over the same calendar months. Estimates imply that a 1% *decrease* in the net-of-tax rate is associated with about a 2.3% *decrease* in reported labor income (implied ETI ≈ 2.3). This is the pooled effect across small and large firms. This result is robust to the inclusion of different controls and fixed effects. Columns (2) and (3) split the sample by size; the elasticity is larger for large firms.

Table 5: ETI Estimates

	Dep. Var.: Log Reported Labor Income			
	2020–2022			2022–2023
Treated firms:	All	Large	Small	Large
	(1)	(2)	(3)	(4)
$\ln(\tau_{jt})$	2.274***	2.614***	1.620***	0.748*
	(0.410)	(0.684)	(0.254)	(0.337)
Treatment	AR	Large AR	Small AR	Large AR
Control Firms	GR	GR	GR	Small AR
Time FE	Yes	Yes	Yes	Yes
Yes				
(month-year)				
Group FE	Yes	Yes	Yes	Yes
Yes				
(firm size × regime)				
Region FE	Yes	Yes	Yes	Yes
Yes				
Worker controls†	Yes	Yes	Yes	Yes
Yes				
Observations	10,003,142	8,774,570	1,228,572	4,051,641
R^2	0.223	0.216	0.230	0.116

Notes: AR is Agricultural Regime and GR is General Regime. Dependent variable is $\ln(\hat{z}_{ijt})$, reported labor income. Column (1) excludes 2023, capturing the response to income-tax increase with *constant* corporate-tax deductibility. Column (2) uses only 2022–23, when the corporate tax increases for high-sales firms, dampening the income-tax response. Standard errors (parentheses) are clustered at the firm–size × regime level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. † Controls: gender, education, age, occupation dummies, pension-regime dummies, and firm-size dummies.

Column (4) focuses on 2022–2023 where large firms are the only treated ones. Large firms are treated in 2023, while small firms face no change and act as controls. With unit, month, and department fixed effects, the coefficient on $\ln(1-\tau)$ is identified by the within-unit change for large AR firms in 2023 relative to 2022, net of the change for small AR firms in the same months. In this case the estimates are smaller and barely statistically significant. This suggests that the 1 p.p. income-tax increase, which coincided with an increase in the corporate tax rate, did not measurably change reported labor income over this short window. Under the assumption that the ETI is stable across this period, the lack of a detectable response to the 2023 income-tax change is consistent with the corporate tax offsetting part of the income-tax effect.

ETIs of around 2 may seem large compared to ETIs in developed economies, where estimates usually fall between 0.1 and 0.4 for labor income (Chetty, 2012; Saez, Slemrod, and Giertz, 2012). These imply that labor supply responses to changes to the net-of-tax rate are small. But it is well established that the ETI is not a structural parameter pinned down only by preferences or technologies. It also reflects the design of the tax system and the scope for avoidance and evasion (Kopczuck and Slemrod, 2002).

In developing countries, where informal sectors are large and enforcement is weaker, higher ETIs are expected. Although there are fewer studies in these contexts, existing work has documented ETIs in the range of 2–5.[18] These findings line up closely with my estimates.

Besides the regression, I use event studies to do two things: (i) test for parallel trends in the pre-period and (ii) see whether there is a discrete jump right when the policy takes effect. The coefficients plot ATT at each event month relative to the month just before treatment (December the previous year). These figures are complementary. These do not replace the ETI regressions, since the elasticity is identified from the slope with respect to $\ln(1-\tau)$. The Appendix shows the details of the event studies.

Figure 7 plots dynamic DiD estimates around each reform. Panel A compares AgR firms to GR firms around the 2021 change. Pre-treatment coefficients are centered around zero, which is consistent with parallel trends. After the reform, the path turns negative and stays there: the ATT settles around -0.3 to -0.4 log points within the first year and remains persistently below zero. This lines up with the ETI estimates for 2020-2022 shown in Column (1) of Table 5. Panel B restricts to the AgR firms and compares large AgR firms to small AgR firms around the 2023 change. Pre-period coefficients are close to zero, and post-treatment estimates are small and imprecise, centered near zero. That pattern matches the ETI estimates for AgR firms in 2022-2023 in Column (4) of Table 5, where the point estimate is modest and not precise.

Figure 7: ETI Estimates: Event Studies

(a) First Treatment

(b) Second Treatment

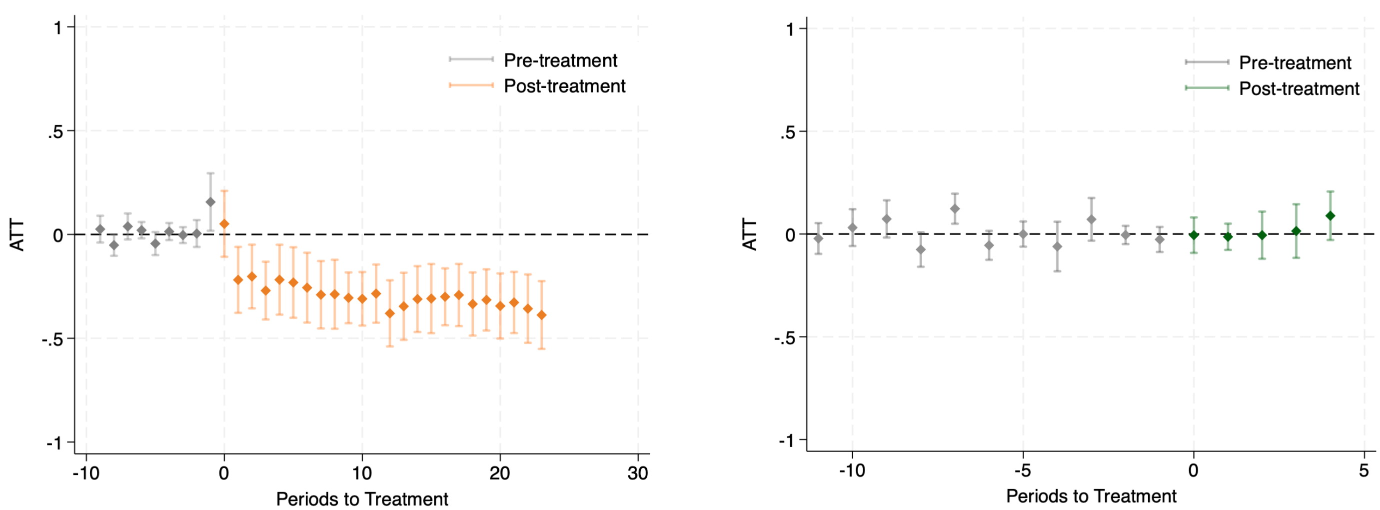

I use the same windowing as in the regressions of Table 5, and the bins are defined relative to the month before treatment. GR firms (control units) never activate in the event dummies; they only help pin down the time effects in Panel A.

Taken together, the first-stage estimates and the flat response after the 2023 reform are consistent with the model's prediction that the elasticity of taxable income is a decreasing function of the corporate tax rate: ETIs are high when the corporate tax is low and attenuate sharply when the corporate tax increases.

[18] Bachas and Soto (2021) estimate elasticities of taxable profits of 3–5 in Costa Rica. Waseem (2013) finds ETIs of about 2–3 in Pakistan using a large PIT reform. Yan and Gao (2023) estimate an ETI of around 2 in China.

6. Numerical Simulations

A key result in the model is that the corporate tax t shifts the optimal labor income tax. To illustrate the comparative statics at the optimum, I build a simple numerical example. The exercise is meant to show the qualitative mechanics rather than estimate a full policy rule or provide a structural calibration for Peru.

Workers draw a type θ from a truncated exponential distribution on $[\theta_{\min}, \theta_{\max}]$ with parameter λ. Hours have quadratic disutility, so for any wage w the optimal hours are

$$\ell(w) \;=\; \frac{(1-\tau)w}{\kappa}.$$

In the PUT arrangement, a worker can hide a per–hour amount $\text{PUT} \in [0, w^P]$ and pays a convex cost

$$\frac{1}{\eta}\,\theta\,(\text{PUT}\cdot\ell)^{\eta},$$

with $\eta = 2$ in the baseline. The firm has Cobb–Douglas technology $Y = A\,L^{\beta}$ and labor demand $\beta A L^{\beta-1} = w^F$.

Reported income, PUT prevalence, and the elasticity of taxable income (ETI) at a given (τ, t) are defined as

$$\widehat{Z} \;=\; w^F L^F \;+\; \big(w^P L^P \;-\; \text{PUT}^{\text{agg}}\big), \qquad \gamma \;=\; \frac{\text{PUT}^{\text{agg}}}{\widehat{Z}}, \qquad \text{ETI}(\tau,t) \;=\; \frac{d\log\widehat{Z}}{d\log(1-\tau)}.$$

The unknown object is the distribution of the idiosyncratic cost of forgone payroll benefits, captured by θ. I discipline this distribution by choosing parameters so that the model reproduces three qualitative patterns that emerge from the data and the existing evidence: (i) PUT prevalence is higher when corporate taxes are lower (see Section 3), (ii) ETIs in low–enforcement environments lie in the 1–2 range, in line with my own estimates and with the literature for developing countries, and (iii) ETIs decrease with the corporate tax rate. In practice, feeding the statutory corporate tax rate directly into the model makes it difficult to match simultaneously the observed levels of PUT and ETI without assuming implausibly low valuations of tax deductions on the firm side. I therefore interpret t in the simulations as an *effective* corporate tax rate on PUT—a reduced–form wedge that can be lower than the statutory rate because some envelope payments may be disguised as other deductible expenses or not enforced at the statutory margin.

To keep the exercise simple, I work with three values for this effective corporate tax rate, $t \in \{0.10, 0.15, 0.20\}$. These are loosely inspired by the small–firm, agricultural, and general regimes (see Section 3). For each t, I compute PUT prevalence $\gamma(t)$ and the ETI $\text{ETI}(t)$ at a fixed personal income tax rate. As a convenient target profile that encodes the patterns above, I take

$$\big(\gamma_{0.10}, \gamma_{0.15}, \gamma_{0.20}\big) \approx (0.40,\, 0.30,\, 0.20), \quad \text{and} \quad \big(\text{ETI}_{0.10}, \text{ETI}_{0.15}, \text{ETI}_{0.20}\big) \approx (2.0,\, 1.5,\, 1.0).$$

These values come from the PUT estimates and ETIs that I view as most informative for Peru, but here they should be interpreted as stylized calibration targets rather than exact empirical moments: the goal is to produce realistic, decreasing ETIs and higher PUT prevalence at lower corporate rates, not to fully discipline the model with the point estimates in Section 5.

For the calibration, I fix the personal income tax at $\tau = 0.40$, which is close to the average marginal tax rate in the data, and compute $\gamma(t)$ and $\text{ETI}(t)$ at this reference rate for each t. To characterize social preferences I set $\omega_L = 1$, $\omega_\pi = 0.25$, and specify the social welfare function as $G(x) = \frac{x^{1-\rho}-1}{1-\rho}$ with $\rho = 2$. I fix $\kappa = 10$ and $A = 4$, and sweep across $\beta \in [0.85, 0.95]$.[19]

The free parameters are $p = (\lambda, \theta_{\min})$, with $\theta_{\max} = 100$ held fixed. For each p and each corporate tax rate $t \in \{0.10, 0.15, 0.20\}$, I solve for equilibrium wages (w^P, w^F), arrangement allocations (L^P, L^F), and hours worked, and then record $\gamma(t)$ and $\text{ETI}(t)$. I explore values $\lambda \in \{0.10, 0.15, 0.25\}$ and $\theta_{\min} \in [0.10, 0.50]$, and rank parameter vectors using a joint loss function that assigns twice the weight to matching $\gamma(t)$ relative to $\text{ETI}(t)$. The best–fit region includes $\lambda \approx 0.25$ and $\theta_{\min} \in [0.10, 0.15]$. At these values, the model reproduces the main qualitative facts: high PUT prevalence when $t = 0.10$, a sizable drop at $t = 0.15$, and low prevalence with a smaller ETI when $t = 0.20$.

Using the calibrated parameters that match the target profiles for PUT prevalence and ETIs, I then treat τ as the policy instrument and simulate the optimal income tax at different corporate tax rates. Figure 8 shows the results. Panel A plots the optimal income tax rate across values of the corporate tax in an economy calibrated to Peruvian data, where PUT is prevalent given the current valuation of payroll benefits. When the corporate tax rate is zero, the optimal income tax rate is about 12%. In this case, the ETI is high because firms face no deterrent to underreporting wages, so worker responses are strong and the optimal income tax is relatively low. As the corporate tax rate increases from zero to $t = 0.20$, the optimal income tax rises monotonically. This is expected: a higher corporate tax generates a positive fiscal externality and lowers the ETI, both of which push the optimal income tax upward.

However, once the corporate tax approaches the upper end of the range where PUT is no longer profitable for firms, the optimal income tax starts to fall back toward the level observed when the corporate tax is zero. This is the point at which PUT disappears (see Panel B). At that level, firms no longer have incentives to pay under the table because the cost is too high, even for individuals who place little value on the benefits of formality. At this point, the fiscal trade-off becomes one of collecting revenue through the corporate tax or the income tax. If the corporate tax is high, the income tax no longer needs to be.

Panels C and D of Figure 8 illustrate two contrasting "naive government" cases. In Panel C, the orange line represents the optimal income tax when the government assumes that PUT behaves like classical evasion: workers decide whether to underreport by comparing their own costs of forgone payroll benefits to the income tax rate, and the firm plays no role in this decision. In this scenario, the optimal income tax is lower than in the full PUT model because workers are more responsive to tax changes (since evasion depends solely on them), so the ETI is higher. Panel D shows the opposite extreme, where the government assumes that third-party reporting fully eliminates evasion and ignores the PUT margin when computing the optimal tax. In the government's perceived model, all reported income reflects true income and workers cannot adjust reported income through PUT, so the only behavioral response in the ETI comes from labor supply. As a result, the perceived ETI is smaller and the optimal income tax is higher.[20]

For the general regime, at the effective corporate tax rate that best matches the Peruvian data,

[19] I also explore values $\kappa \in [8, 11]$, $A \in [3.5, 7]$, and $\beta \in [0.70, 0.95]$.

[20] The gap widens further because the government takes observed (underreported) wages as true incomes. Average wages therefore appear lower in the economy than they actually are, increasing the perceived need for redistribution and pushing the optimal tax even higher.

$t \approx 0.20$, the model delivers an optimal income tax of about 25%. This is well below the average marginal income tax rate observed in the data of around 40% (with the lowest bracket at 30%). Taken at face value, this gap suggests that the actual tax schedule is closer to the 'no–evasion' benchmark in Panel D than to the full PUT model. In other words, current policy looks as if the government effectively ignored PUT when choosing the income tax rate: it behaves more like a planner who believes third–party reporting shuts down evasion and therefore perceives a low ETI, rather than like a planner who internalizes the PUT margin and chooses a lower optimal rate.

Figure 8: Simulations of Optimal Income Tax Rates

(a) Optimal Income Tax Rate

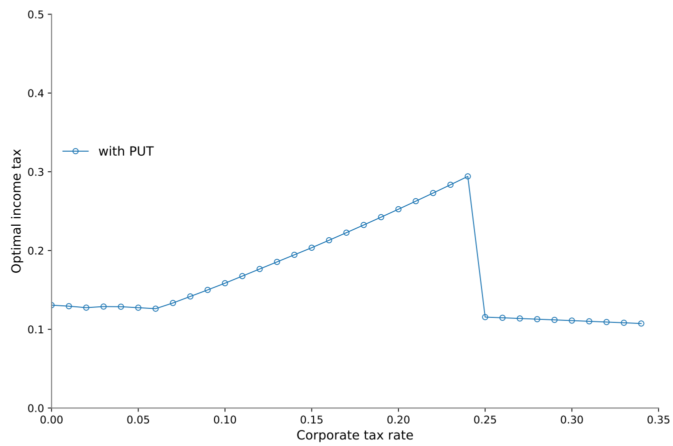

(b) PUT Prevalence

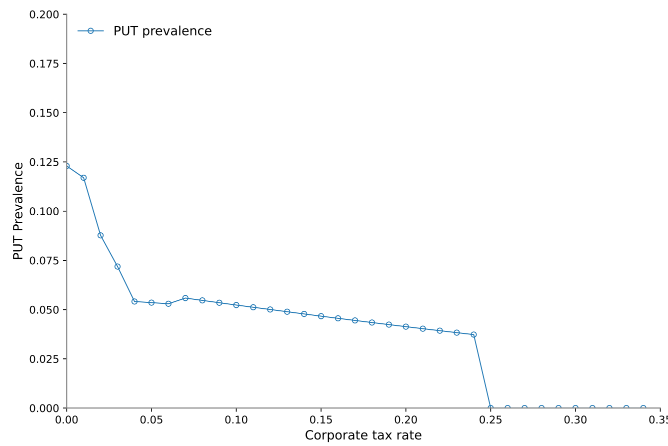

(c) Optimal Income Tax Rate with One-Sided Evasion

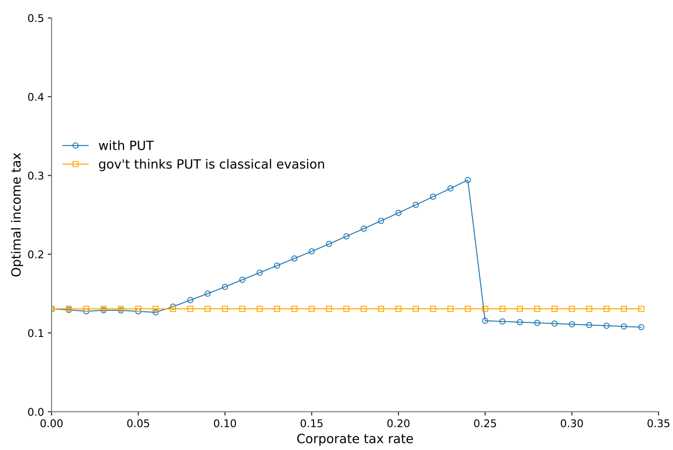

(d) Optimal Income Tax Rate with No Evasion in the Government's Model

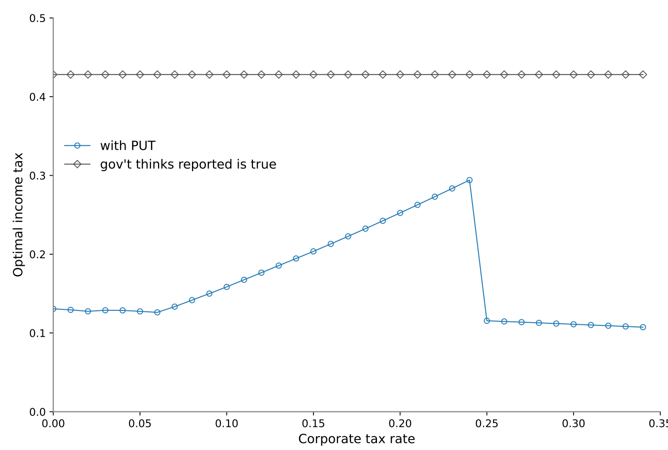

7. Concluding Remarks

This paper studies payments under the table (PUT), a form of wage underreporting in which firms and workers collude to report only part of the wage to the tax and social security authorities and pay the rest off the books. This type of wage underreporting has been documented in the literature, mostly using surveys, and is especially prevalent in Latin America and Eastern Europe. Yet its implications for tax policy design remain largely unexplored.

The paper makes two contributions to the study of PUT, one methodological and one in optimal tax design. On the methodological side, I introduce optimal transport, a novel econometric technique to statistically match two wage distributions using common covariates, into the tax underreporting literature. I combine datasets that governments already have: matched employer–employee administrative records with firm-reported wages and household surveys with worker-reported wages. A government that wants to assess the extent of PUT could directly implement this approach. I assume the survey wage equals the true wage plus measurement error. Optimal transport finds, for each worker in the administrative data, where a worker with the same characteristics would be ranked in the survey distribution and what wage they would earn, and then maps that wage back into the administrative data. PUT is defined as the difference between this 'transported' wage and the reported wage. I find that, on average, 30 percent of wages in Peru are paid under the table, and that the share of PUT increases with income.

On the tax design side, I develop a general equilibrium model of optimal income taxation in an economy where there is also a corporate tax. The key result is that because PUT cannot be deducted from firms' taxable profits, the response of workers to changes in the income tax (the ETI) depends negatively on the corporate tax rate. The ETI is lower when the corporate tax is higher, which allows governments to set a higher optimal income tax rate than they would in a setting with only worker-side evasion. To test this mechanism in the data, I leverage a tax reform in Peru that affected the contracts of *some* agricultural workers and unfolded in two stages: first, only the income tax increased; second, both the income and corporate taxes increased. The reform applied only to agricultural workers in some regions (treated workers), while others remained under a regime in which tax rates did not change (control workers). Using the administrative dataset, I estimate the ETI for all agricultural workers and show that it is higher in the first stage, when the corporate tax did not change.

Numerical simulations calibrated to Peruvian moments for PUT and the ETI further illustrate the policy implications. At the effective corporate tax rate that best matches the Peruvian data, the model delivers an optimal top income tax rate of about 25 percent, well below the roughly 40 percent observed in the actual schedule. This gap suggests that current policy is closer to a 'no–evasion' benchmark in which the government effectively ignores under–the–table payments when setting the income tax rate, rather than to a planner who recognizes the PUT margin and chooses a lower optimal rate.

References

Allingham, Michael G. and Agnar Sandmo (1972). "Income Tax Evasion: A Theoretical Analysis". In: *Journal of Public Economics* 1, pp. 323–338 (cit. on p. 7).

Arellano, Manuel and Stéphane Bonhomme (2023). "Recovering Latent Variables by Matching". In: *Journal of the American Statistical Association* 118.541, pp. 693–706 (cit. on p. 7).

Barth, Erling and Tone Ognedal (2018). "Tax Evasion in Firms". In: *LABOUR* 32.1, pp. 23–44 (cit. on p. 5).

Bergolo, Marcelo and Guillermo Cruces (2014). "Work and tax evasion incentive effects of social insurance programs: Evidence from an employment-based benefit extension". In: *Journal of Public Economics* 117, pp. 211–228 (cit. on pp. 2, 6).

Best, Michael, Anne Brockmeyer, Henrik Jacobsen Kleven, Johannes Spinnewijn, and Mazhar Waseem (2015). "Production versus Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan". In: *Journal of Political Economy* 123.6, pp. 1311–1355 (cit. on p. 10).

Bíró, Anikó, Dániel Prinz, and László Sándor (2022). "The minimum wage, informal pay, and tax enforcement". In: *Journal of Public Economics* 215, p. 104728 (cit. on p. 6).

Bjørneby, Kristine et al. (2021). "Wage underreporting and collusive tax evasion in small firms". Working Paper (cit. on p. 6).

Bosch, Mariano and William Maloney (2010). "Comparative Analysis of Labor Market Dynamics Using Markov Processes: An Application to Informality". In: *Labour Economics* 17.4, pp. 621–631 (cit. on p. 8).

Bound, John, Charles Brown, and Nancy A. Mathiowetz (2001). "Measurement Error in Survey Data". In: *Handbook of Econometrics, Volume 5*. Ed. by James J. Heckman and Edward Leamer. Amsterdam: Elsevier, pp. 3705–3843 (cit. on p. 12).

Carrillo, Paul, Dina Pomeranz, and Monica Singhal (2017). "Dodging the Taxman: Firm Misreporting and Limits to Tax Enforcement". In: *American Economic Journal: Applied Economics* 9.2, pp. 144–164 (cit. on p. 10).

Di Nola, Alessandro, Georgi Kocharkov, and Aleksandar Vasilev (2019). "Envelope wages, hidden production and labor productivity". In: *The B.E. Journal of Macroeconomics.* (cit. on p. 5).

Diamond, Peter and Emmanuel Saez (2011). "The Case for a Progressive Tax: From Basic Research to Policy Recommendations". In: *Journal of Economic Perspectives* 25.4, pp. 165–190 (cit. on pp. 2, 6).

Feinmann, Javier, Ana Paula Franco, Pablo Garriga, Nathalie Gonzalez-Prieto, Roberto Hsu Rocha, and Maximiliano Lauletta (2025a). "Payments Under the Table in Latin America". In: (cit. on pp. 6, 13, 14).

Feinmann, Javier, Roberto Hsu Rocha, and Maximiliano Lauletta (2025b). "Payments Under the Table: Employer-Employee Collusion in Brazil". In: (cit. on pp. 2, 6).

Galichon, Alfred (2016). *Optimal Transport Methods in Economics*. Princeton University Press (cit. on p. 17).

Gaskin, Thomas, Marie-Therese Wolfram, Andrew Duncan, and Guven Demirel (2024). *Modelling Global Trade with Optimal Transport*. Tech. rep. (cit. on p. 7).

Gavoille, Nicolas and Anna Zasova (2023). "Minimum wage spike and income underreporting: A back-of-the-envelope-wage analysis". In: *Journal of Comparative Economics* 51.1, pp. 372–402 (cit. on pp. 2, 6).

Gunsilius, Florian (2025). "A Primer on Optimal Transport for Causal Inference with Observational Data". In: *arXiv preprint arXiv:2503.07811.* (cit. on p. 17).

Gunsilius, Florian and Yuliang Xu (2023). "Matching for causal effects via multimarginal unbalanced optimal transport". Working Paper (cit. on p. 17).

Hokayem, Charles, Christopher R. Bollinger, and James P. Ziliak (2015). "The Role of CPS Nonresponse in the Measurement of Poverty". In: *Journal of the American Statistical Association* 110.511, pp. 935–945 (cit. on p. 12).

Kim, Dongwoo and Young Jun Lee (2024). *Semi-nonparametric Models of Multidimensional Matching: An Optimal Transport Approach*. Tech. rep. (cit. on p. 7).

Kleven, Henrik, Martin B. Knudsen, Claus T. Kreiner, Soeren Pedersen, and Emmanuel Saez (2011). "Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark". In: *Econometrica* 79.3, pp. 651–692 (cit. on pp. 2, 7).

Kleven, Henrik, Claus Kreiner, and Emmanuel Saez (2016). "Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries". In: *Economica* 83, pp. 219–236 (cit. on pp. 2, 7, 20).

Kleven, Henrik and Esben Anton Schultz (Nov. 2014). "Estimating Taxable Income Responses Using Danish Tax Reforms". In: *American Economic Journal: Economic Policy* 6.4, pp. 271–301 (cit. on p. 29).

Kumler, Todd, Eric Verhoogen, and Judith Frías (2020). "Enlisting Employees in Improving Payroll Tax Compliance: Evidence from Mexico". In: *Review of Economics and Statistics* 102.5, pp. 881–896 (cit. on pp. 2, 6).

La Porta, Rafael and Andrei Shleifer (2014). "Informality and Development". In: *Journal of Economic Perspectives* 28, pp. 109–126 (cit. on p. 7).

Leite, Daniela and Caterina De Bacco (2024). "Similarity and economy of scale in urban transportation networks and optimal transport-based infrastructures". In: *Nature Communications* 15, p. 7981 (cit. on p. 7).

Levy, Santiago (2018). *Under-Rewarded Efforts: The Elusive Quest for Prosperity in Mexico*. Inter-American Development Bank (cit. on p. 8).

Li, Ruilin, Xiaojing Ye, Haomin Zhou, and Hongyuan Zha (2019). "Learning to Match via Inverse Optimal Transport". In: *Journal of Machine Learning Research* 20, pp. 1–37 (cit. on p. 7).

Loayza, Norman V. (1996). "The economics of the informal sector: a simple model and some empirical evidence from Latin America". In: *Carnegie-Rochester Conference Series on Public Policy* 45, pp. 129–162 (cit. on p. 8).

López-Luzuriaga, Andrea, Monica Calijuri, Carola Pessino, Simeon Schächtele, Ubaldo Gonzalez, and Carla Chamorro (2023). "Detecting envelope wages with e-billing information". In: *International Tax and Public Finance* 31, pp. 26–65 (cit. on p. 6).

Mirrlees, James A. (1971). "An Exploration in the Theory of Optimum Income Taxation". In: *Review of Economic Studies* 38, pp. 175–208 (cit. on pp. 2, 6).

Perry, Guillermo, William Maloney, Omar Arias, Pablo Fajnzylber, Andrew Mason, and Jaime Saavedra-Chanduvi (2007). *Informality: Exit and Exclusion*. World Bank Latin American and Caribbean Studies (cit. on p. 8).

Saez, Emmanuel (2001). "Using Elasticities to Derive Optimal Income Tax Rules". In: *Review of Economic Studies* 68, pp. 205–229 (cit. on pp. 2, 6).

Slemrod, Joel (2008). "Does It Matter Who Writes the Check to the Government? The Economics of Tax Remittance". In: *National Tax Journal* 61.2, pp. 251–275 (cit. on p. 2).

SUNAFIL (2023). *Reporte anual de inspecciones laborales y procedimientos sancionadores laborales*. URL: https://lpderecho.pe/reporte-anual-inspecciones-laborales-procedimientos-sancionadores-laborales/ (cit. on p. 8).

SUNAT (2023). *Memoria Institucional 2023*. Superintendencia Nacional de Aduanas y de Administración Tributaria (SUNAT). URL: https://www.sunat.gob.pe/cuentassunat/planestrategico/memoria/memoria2023.pdf (cit. on p. 8).

U.S. DOL, Department of Labor (2024). *Child Labor in Peru: Findings from the U.S. Department of Labor*. URL: https://www.dol.gov/sites/dolgov/files/ILAB/child_labor_reports/tda2024/Peru.pdf (cit. on p. 8).

Ulyssea, Gabriel (2020). "Informality: Causes and Consequences for Development". In: *Annual Review of Economics* 12, pp. 239–263 (cit. on p. 7).

Villani, Cédric (2009). *Optimal Transport: Old and New*. Springer (cit. on p. 17).

Williams, Colin and Jo Padmore (2013). "Envelope wages in the European Union". In: *International Labour Review* 152.3–4, pp. 411–430 (cit. on pp. 2, 5, 7).

Yamada, Gustavo (1996). "Urban Informal Employment and Self-Employment in Developing Countries: Theory and Evidence". In: *Economic Development and Cultural Change* 44.2, pp. 289–314 (cit. on p. 8).

Yitzhaki, Shlomo (1974). "A Note on Income Tax Evasion: A Theoretical Analysis". In: *Journal of Public Economics* 3, pp. 201–202 (cit. on p. 7).

8. Appendix: Figures and Tables

Table MA1: Mincer Equation in Administrative Dataset and Household Survey

	Survey	Administrative
Education:		
Secondary	214.16***	167.63***
	(31.56)	(3.14)
Bachelor	930.81***	806.99***
	(35.04)	(3.20)
Masters+	2029.49***	1358.81***
	(79.76)	(6.10)
Constant	1484.75***	1507.73***
	(25.38)	(3.08)
Observations	9,907	4,468,287
R-squared	0.218	0.082

Notes: Base category is less than secondary. Robust SEs in parentheses.
*** $p<0.01$.

Table MA2: Covariate Balance Before and After IPW

	Raw			After IPW		
Covariate	Survey mean	Admin mean	Std. diff	Admin mean (w)	Diff (w)	Std. diff (w)
Schooling						
< Secondary	0.085	0.027	0.255	0.0856	−0.0003	−0.0011
Secondary	0.339	0.633	−0.614	0.3421	−0.0028	−0.0058
Tertiary	0.575	0.340	0.486	0.5723	0.0031	0.0062
Age groups						
18–25	0.068	0.147	−0.256	0.0697	−0.0017	−0.0068
25–29	0.125	0.186	−0.167	0.1247	0.0005	0.0016
30–44	0.384	0.424	−0.082	0.3846	−0.0010	−0.0021
45–60	0.338	0.203	0.305	0.3367	0.0009	0.0019
61+	0.086	0.040	0.188	0.0843	0.0013	0.0046
Industry groups						
Agriculture	0.071	0.081	−0.036	0.0746	−0.0036	−0.0140
Mining	0.034	0.024	0.060	0.0360	−0.0018	−0.0098
Manufacture	0.096	0.132	−0.111	0.1012	−0.0049	−0.0166
Construction	0.040	0.085	−0.186	0.0421	−0.0020	−0.0102
Wholesale	0.091	0.162	−0.214	0.0957	−0.0047	−0.0160
Services	0.502	0.516	−0.026	0.526	−0.0233	−0.0466
Firm size						
<10	0.108	0.202	−0.262	0.112	−0.0037	−0.0117
10–50	0.138	0.165	−0.077	0.143	−0.0052	−0.0149
50–100	0.059	0.078	−0.074	0.062	−0.0013	−0.0111
100–500	0.115	0.192	−0.216	0.121	−0.0061	−0.0190
500+	0.580	0.363	0.447	0.563	0.0176	0.0356
Other						
Male	0.580	0.651	−0.147	0.574	0.0057	0.0116
Lima	0.200	0.558	−0.794	0.205	−0.0050	−0.0125

Notes: Differences are Survey minus Admin. "Std. diff" is the standardized mean difference. Rule-of-thumb threshold: $|\text{SMD}| < 0.10$. After IPW, all covariates satisfy the threshold except Other services (0.115).

Table MA3: Survey–Admin Median Wage Gap by Education (Raw, IPW, PSM)

	<Secondary	Secondary	BA	Master's+
Medians (currency units)				
Survey median	1,358.5	1,525.3	2,320.7	3,108.9
Admin median (Raw)	1,200.0	1,378.0	1,852.5	2,500.0
Admin median (IPW)	1,205.2	1,410.8	2,000.0	2,496.0
Admin median (PSM)	1,383.8	1,400.2	2,000.0	3,102.5
Gaps: Survey − Admin (levels)				
Raw	158.5	147.3	468.3	608.9***
IPW	153.3	114.5	320.7	612.9***
PSM	−25.3	125.1	320.7	6.4
Gaps: Survey − Admin (% of Admin)				
Raw (%)	13.2***	10.7***	25.3***	24.4***
IPW (%)	12.7***	8.1***	16.0***	24.5***
PSM (%)	−1.8	8.9***	16.0***	0.2

Notes: Percent gaps are computed as (Survey − Admin)/Admin × 100. Stars (***) indicate statistically significant differences ($p < 0.01$), based on magnitude and consistency with mean-based results.

Table MA4: Survey–Admin Mean Wage Gap by Education Level

	<Secondary	**Secondary**	**BA**	**Master's+**
Mean wages (monthly, local currency)				
Survey (mean)	1,522.4	1,783.5	2,458.0	3,430.5
Admin (raw)	1,507.7	1,675.4	2,314.7	2,866.5
Admin (IPW)	1,518.1	1,744.6	2,455.5	2,818.5
Admin (PSM)	1,553.2	1,687.5	2,402.7	3,221.0
Differences (Survey – Admin)				
Raw difference	14.7 (1.5%)	108.1*** (6.4%)	143.3*** (6.2%)	563.9*** (19.7%)
After IPW	4.3 (0.3%)	38.9*** (2.2%)	2.6 (0.1%)	612.0*** (21.7%)
After PSM	–30.9 (–2.0%)	96.0*** (5.7%)	55.4*** (2.3%)	209.5*** (6.5%)

Notes: Outcome is monthly wage. Gaps are computed as the difference between survey and administrative means within each education level. Percentages in parentheses express the gap as a share of the corresponding administrative mean. Stars indicate significance based on two-sample t-tests: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table MA5: Survey–Admin Mean Wage Gap by Firm Size

	1–5 workers	**6–10**	**11–50**	**51–250**	**250+**
Mean wages (monthly, local currency)					
Survey (mean)	1,589.0	1,841.7	2,040.9	1,997.7	2,454.2
Admin (raw)	1,371.1	1,818.0	1,999.4	2,118.8	2,076.2
Admin (IPW)	1,349.5	1,850.8	2,119.6	2,297.6	2,388.3
Admin (PSM)	1,389.5	1,813.8	2,092.6	2,169.8	2,189.0
Differences (Survey – Admin)					
Raw difference	218.0*** (15.9%)	23.7 (1.3%)	41.5 (2.1%)	–121.1*** (–5.7%)	378.1*** (18.2%)
After IPW	239.4*** (17.7%)	–9.1 (–0.5%)	–78.8*** (–3.7%)	–299.9*** (–13.1%)	65.9 (2.8%)
After PSM	199.5*** (14.4%)	27.9 (1.5%)	–51.8*** (–2.5%)	–172.2*** (–7.9%)	265.2*** (12.1%)

Notes: Outcome is monthly wage. Gaps are computed as the difference between survey and administrative means within each firm-size category. Percentages in parentheses express the gap as a share of the corresponding administrative mean. Stars indicate significance based on approximate two-sample t-tests using reported standard deviations: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table MA6: Survey–Admin Median Wage Gap by Firm Size

	1–5 workers	**6–10**	**11–50**	**51–250**	**250+**
Median wages (monthly, local currency)					
Survey (median)	1,311.2	1,502.9	1,748.1	1,676.6	2,303.8
Admin (raw)	1,025.0	1,486.1	1,666.1	1,756.7	1,636.6
Admin (IPW)	1,025.0	1,500.0	1,793.0	1,906.7	1,904.1
Admin (PSM)	1,025.0	1,499.3	1,750.0	1,811.5	1,785.2
Differences (Survey – Admin)					
Raw difference	286.2*** (27.9%)	16.8 (1.1%)	82.0*** (4.9%)	–80.1*** (–4.6%)	667.1*** (40.8%)
After IPW	286.2*** (27.9%)	2.9 (0.2%)	–45.0*** (–2.5%)	–230.1*** (–12.1%)	399.7*** (21.0%)
After PSM	286.2*** (27.9%)	3.6 (0.2%)	–2.0 (–0.1%)	–134.9*** (–7.4%)	518.6*** (29.1%)

Notes: Outcome is monthly wage. Gaps are computed as the difference between survey and administrative medians within each firm-size category. Percentages in parentheses express the gap as a share of the corresponding administrative median. Stars indicate significance based on approximate two-sample t-tests using reported standard deviations: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure A1: Robustness to classical measurement error

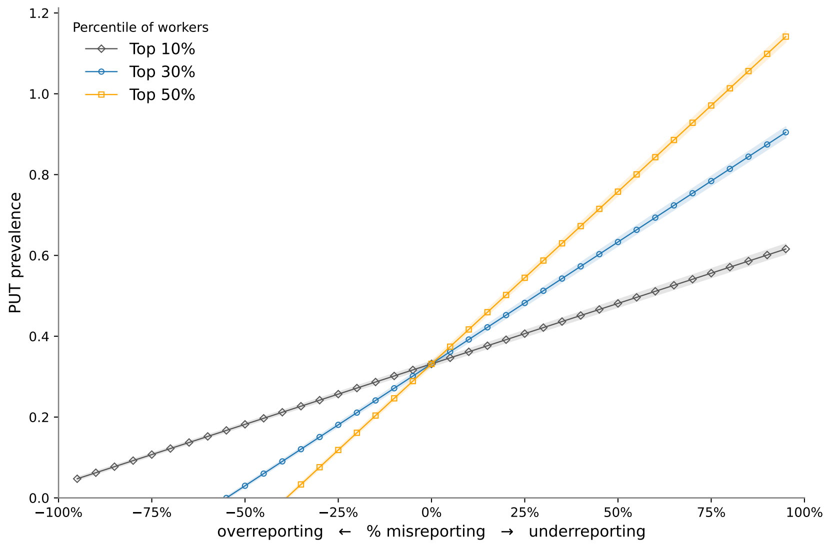

Figure A2: Robustness to non-classical measurement error

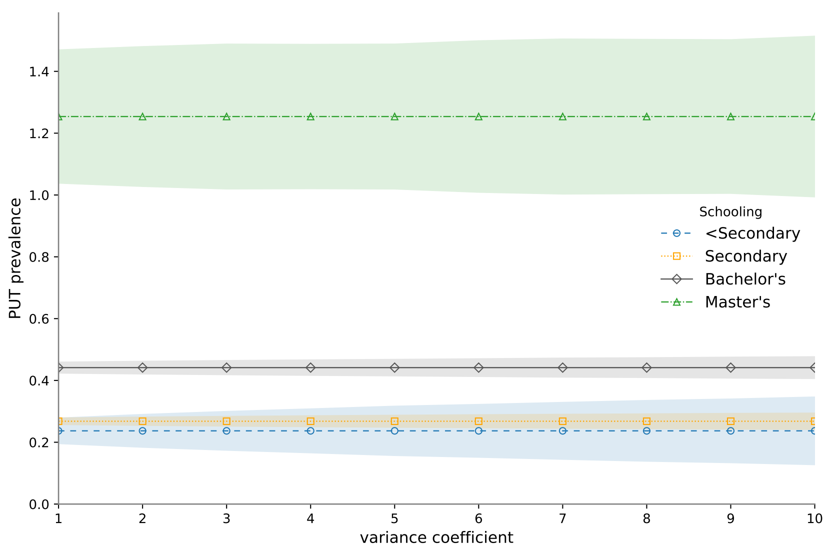

9. Appendix: Methodological

9.1. OT: Asymptotic Properties of the PUT Estimator

PUT Estimator The PUT estimator uses the imputed wage $\widehat{w}_i = \sum_{j=1}^{m} G_{ij}\tilde{w}_j$ and the third-party reported wage $\bar{w}_i$:

$$\begin{aligned}\widehat{\text{PUT}}_i &= \widehat{w}_i - \bar{w}_i \\ &= \sum_{j=1}^{m} G_{ij}\tilde{w}_j - \bar{w}_i.\end{aligned} \tag{38}$$

Since $\tilde{w}_j = w_j + e_j$, and e_j is classical measurement error (zero mean, independent of w_j and X, and with finite variance), $\widehat{\text{PUT}}_i$ represents a consistent estimator of the true PUT, under the assumptions discussed above.

Consistency Under the CIA and classical measurement error assumptions, the PUT estimator is consistent:

$$\widehat{\text{PUT}}_i \xrightarrow{p} \text{PUT}_i \quad \text{as } n, m \to \infty.$$

The consistency argument follows standard econometric reasoning: Gunsilius and Xi (2023) show that the OT solution G converges to a limit G^*, and since $\tilde{w}_j = w_j + e_j$ with zero-mean error, large-sample averaging ensures $\widehat{w}_i \xrightarrow{p} w_i$. This mirrors the logic behind classical matching and imputation methods, where large samples and unbiased noise ensure consistent estimates. Thus, the OT-based estimator is well-grounded in standard econometric principles of convergence and identification.

The consistency of the PUT estimator relies on the assumption that measurement error is classical. If this assumption is violated—such as in cases where self-reported wages systematically understate or overstate true wages, or where measurement error is correlated with covariates—the transported measurement error term $\sum_{j=1}^{m} G_{ij}e_j$ may introduce bias into the estimator. To address this concern, I will perform robustness checks in Subsection 4.3, exploring scenarios with systematic biases and bounded-error structures. These checks assess the sensitivity of the PUT estimates to deviations from the classical measurement error assumption, ensuring that the results remain valid under plausible alternative error structures.

Efficiency OT is designed to minimize the global cost of matching distributions, and under mild conditions, it is asymptotically efficient relative to less structured matching procedures such as nearest-neighbor matching, propensity score matching, or regression-based imputation methods. These more traditional methods may not preserve the joint distributions as effectively, potentially leading to higher asymptotic variances under the same conditions (Gunsilius and Xi, 2023). For the PUT estimator, efficiency implies lower asymptotic variance when the CIA and classical measurement error assumptions hold. Specifically:

$$\text{Var}(\widehat{\text{PUT}}_i) = \text{Var}(w_i) + \sigma_e^2 \sum_{j=1}^{m} G_{ij}^2,$$

where $\text{Var}(w_i)$ represents the inherent variability in the underlying population from which the individual i is drawn, and σ_e^2 is the variance of the measurement error. The term $\sum_{j=1}^{m} G_{ij}^2$ measures how concentrated the weight distribution is. A more balanced G-distribution reduces the variance contribution of σ_e^2, reflecting standard econometric intuition: spreading risk (or noise) across many observations attenuates variance, just as more dispersed matching reduces the influence of a single noisy match.

9.2. OT: Methodological Validation with Synthetic Data

To validate the OT estimator in a controlled environment with known underreporting behavior, I construct synthetic administrative and survey datasets where the true level of PUT is observable.

Synthetic Data Generation The synthetic data comprises two components: an *administrative dataset* and multiple *survey datasets*, each designed to reflect different sampling strategies.

The administrative dataset consists of $n_{\text{admin}} = 2000$ individuals and includes the following variables:

- True Income (w): True incomes are drawn from a truncated normal distribution with a mean of 2000, a standard deviation of 500, and a minimum wage of 1000. This distribution represents the actual income distribution in the economy.
- Underreported Income ($\bar{w}$): A subset of individuals (40%) underreport their incomes to the minimum wage, simulating real-world wage underreporting behavior.
- Covariates (X): Covariates include gender (binary), schooling (primary, secondary, higher), and age group (young, middle-aged, senior), generated to approximate realistic demographic distributions.

The administrative dataset provides both the true incomes (w) and underreported incomes ($\bar{w}$), enabling the calculation of the true PUT ($\textsc{put}_i^{\text{True}} = w_i - \bar{w}_i$) for validation purposes.

Survey datasets are sampled from the *true income distribution* (w) of the administrative dataset. Each survey simulates a different sampling strategy by selecting individuals below specific percentiles of the true income distribution. The surveys are defined as follows:

- Survey 1: Individuals with incomes in p_0-p_{100} (full population).
- Survey 2: Individuals with incomes in p_0-p_{80}.
- Survey 3: Individuals with incomes in p_0-p_{85}.
- Survey 4: Individuals with incomes in p_0-p_{90}.
- Survey 5: Individuals with incomes in p_0-p_{95}.

Each survey includes 200 randomly sampled individuals from the respective percentile range, ensuring representativeness across covariates. These survey datasets serve as inputs for both the OT and OLS methods, enabling direct comparisons under varying sampling assumptions.

Conditional Independence and Threshold Selection The Conditional Independence Assumption (CIA) requires that, conditional on the covariates (X), the distribution of survey incomes ($\tilde{w}$) reflects the distribution of true incomes (w) for individuals in the administrative dataset. To ensure that CIA holds, I iteratively test multiple thresholds for restricting the administrative dataset based on the underreported income distribution ($\bar{w}$).

For each threshold p, I define a subset of the administrative dataset ($\mathcal{A}_p$) as:

$$T_p = \text{Quantile}_p(\bar{w}), \quad \mathcal{A}_p = \{i : \bar{w}_i \leq T_p\}.$$

For each subset ($\mathcal{A}_p$), I calculate the standardized mean differences (SMDs) for the covariates (X) to measure the similarity between the administrative subset and the survey dataset. The SMD for a covariate k is calculated as:

$$\mathrm{SMD}_k = \frac{\mu_{\mathcal{A}_p,k} - \mu_{\mathcal{S},k}}{\sqrt{\frac{\sigma^2_{\mathcal{A}_p,k}+\sigma^2_{\mathcal{S},k}}{2}}},$$

where μ and σ^2 are the mean and variance of covariate k in the administrative subset ($\mathcal{A}_p$) and the survey dataset ($\mathcal{S}$). The optimal threshold p^* minimizes the maximum absolute SMD across all covariates:

$$p^* = \arg\min_p \max_k |\mathrm{SMD}_k|.$$

The optimal threshold identifies the administrative subset ($\mathcal{A}_{p^*}$) most comparable to the survey dataset.

Statistical Matching with Optimal Transport Using the administrative subset defined by the optimal threshold ($\mathcal{A}_{p^*}$), I perform statistical matching with the survey dataset using OT. The OT method aligns the distributions of the two datasets while preserving their marginal distributions. For each administrative record, OT imputes a self-reported wage (w^*) from the survey dataset, and the PUT is calculated as:

$$\text{PUT}_i^{\mathrm{OT}} = w_i^* - \bar{w}_i.$$

OLS Estimates To compare OT with OLS, I fit an OLS regression model for each survey dataset. The model estimates the relationship between self-reported wages ($\tilde{w}$) and the covariates (X):

$$\tilde{w}_i = \beta_0 + \sum_k \beta_k X_{i,k} + \epsilon_i.$$

Using the estimated coefficients ($\hat{\beta}_k$), I predict self-reported wages (w^{OLS}) for all individuals in the administrative dataset:

$$w_i^{\mathrm{OLS}} = \hat{\beta}_0 + \sum_k \hat{\beta}_k X_{i,k}.$$

The OLS-based PUT estimate is then calculated as:

$$\text{PUT}_i^{\mathrm{OLS}} = w_i^{\mathrm{OLS}} - \bar{w}_i.$$

This process is repeated for each survey dataset, yielding multiple sets of OLS-based PUT estimates.

Results and Comparison The following table summarizes the key statistics for OT and OLS across all survey datasets:

	Optimal Transport			OLS			True PUT
	Mean PUT	SE PUT	N	Mean OLS	SE OLS	N	
Survey1	400.33	17.37	1966	340.33	14.06	2000	414.52
Survey2	506.04	15.48	1606	319.12	14.07	2000	507.44
Survey3	406.55	16.63	1926	352.11	14.11	2000	423.13
Survey4	413.38	16.54	1920	338.32	14.04	2000	424.46
Survey5	394.99	17.47	1966	362.04	14.10	2000	414.52

Table MA7: Summary Statistics of PUT Estimates from OT and OLS

The following figure compares the mean PUT estimates from OT, OLS, and the true PUT across all survey datasets. Error bars indicate the standard errors for OT and OLS estimates:

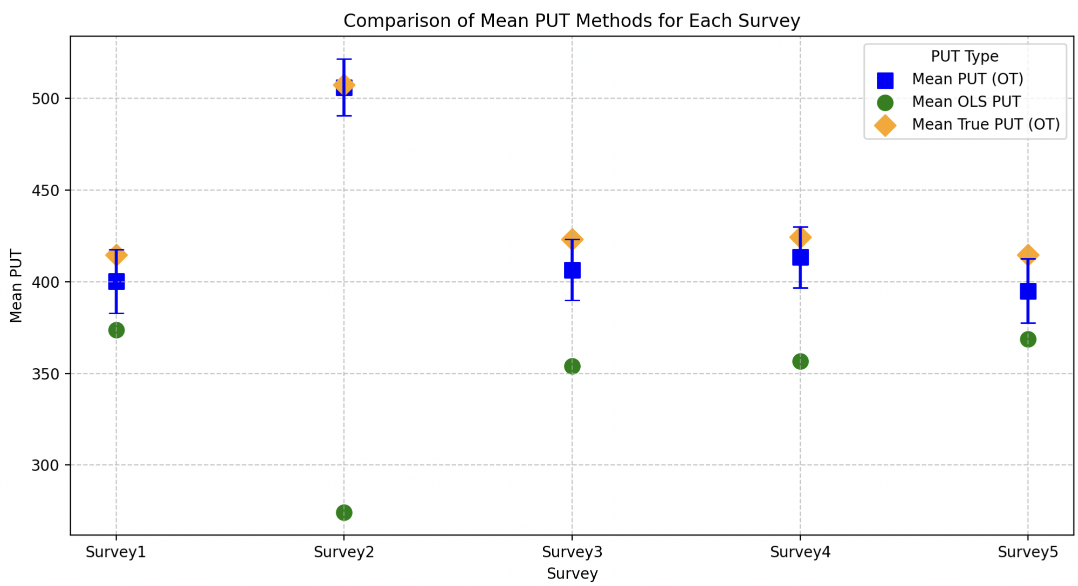

Figure A3: Comparison of Mean PUT Methods for Each Survey

Taken together, the table and figure confirm that the OT estimator reliably reconstructs true PUT across survey sampling scenarios where OLS systematically fails. This exercise demonstrates that OT recovers the correct wage distribution even when the survey is truncated at the top—precisely the setting relevant for real data. The synthetic validation thus provides direct support for the identification strategy pursued in the empirical analysis.

9.3. Optimal Income Tax: Perturbation Method

An alternative derivation of equation (28) in aggregate terms is to consider a marginal increase in the income tax rate ($d\tau > 0$) and analyze its effects on welfare. At the optimum, the total effect must be zero. First, there is a mechanical effect (ME): the government collects more revenue from workers in proportion to their reported wages and redistributes it to all workers. This increases welfare when redistribution preferences are non-zero:

$$\text{ME} = (1-\bar{g})\hat{Z}d\tau \tag{39}$$

where $\bar{g}$ is the average welfare weight on workers (across arrangements), 1 is the normalized welfare weight of all individual (non-workers and workers). That is, the government values \$1 to an individual the same as $\bar{g}$ dollars to a worker. $\hat{Z}$ is total reported labor income.

Second, there is a fiscal externality (FE): the increase in the income tax rate generates behavioral responses that affect government revenue. First, it reduces reported labor income in two ways: workers reduce their labor supply (captured by the aggregate elasticity e_1), and they shift more of their earnings into PUT (captured by the aggregate elasticity e_2). Both channels reduce the income tax base, lowering revenue. At the same time, however, greater underreporting increases the firm's taxable profits. This is PUT are not deductible from corporate taxes. So, when a worker receives more of their wage informally, the firm reports lower labor costs and thus higher profits, on which it must pay the corporate tax rate t. From the planner's perspective, only a fraction of this additional profit tax revenue counts toward welfare, depending on how much the planner values firm profits (through the welfare weight g_π).

The net effect on revenue is:

$$\text{FE} = -\tau\frac{d\hat{Z}}{d\tau} + t(1-g_\pi)\frac{d\hat{\Pi}}{d\tau} = \left[-\tau\left(e_1 - \underbrace{\frac{\text{PUT}}{\hat{Z}}}_{\gamma}e_2\right) + t(1-g_\pi)\underbrace{\frac{\text{PUT}}{\hat{Z}}}_{\gamma}e_2\right]\hat{Z} \tag{40}$$

where $\hat{Z}$ is total reported labor income and $\hat{\Pi} = F(L) - \hat{Z}$ are reported profits, and PUT is total underreported wage income. $\gamma = \frac{\text{PUT}}{\hat{Z}}$ is PUT prevalence, defined as before. $e_1 = \frac{\partial \log L}{\partial \log(1-\tau)} = e_1^P(\alpha^P - \gamma) + e_1^F\alpha^F > 0$ is the labor supply elasticity of all workers in both arrangements weighted by their reported income as share of total income in the economy. It aggregates intensive and extensive margin responses in labor supply. $e_2 = \frac{\partial \log \text{PUT}}{\partial \log(1-\tau)} < 0$ is the intensive margin PUT elasticity only for workers in the PUT arrangement. Note that $e_1 - \gamma e_2$ is the ETI defined as before.

At the optimum, the marginal welfare gain from redistribution equals the marginal revenue loss (or gain) from behavioral responses:

$$\text{ME} + \text{FE} = 0 \tag{41}$$

Solving this condition yields the optimal income tax rate:

$$\tau^* = \frac{1-\bar{g}-t(1-g_\pi)e_2\gamma}{1-\bar{g}+e_1-e_2\gamma} \tag{42}$$

This expression is exactly the same as derived before and shows the trade-off introduced by PUT: PUT reduces income tax revenue (denominator) but raises corporate tax revenue (numerator).

9.4. Optimal Income Tax: Comparative statics of ETI with respect to the corporate tax

The elasticity of taxable income is

$$\text{ETI} \;=\; (\alpha^P-\gamma)e_1^P \;+\; \alpha^F e_1^F \;-\; \gamma e_2 \;+\; \alpha^P\eta^P \;+\; \alpha^F\eta^F, \tag{43}$$

where $\alpha^i = Z^i/\hat{Z}$ are income shares, $\gamma = \text{PUT}/\hat{Z}$ is PUT prevalence, e_1^i are labor-supply elasticities, e_2 is the PUT (hiding) elasticity, and η^i are wage elasticities (all with respect to $\log(1-\tau)$).

Differentiating with respect to the corporate tax t and writing $x' \equiv \partial x/\partial t$,

$$\begin{aligned}\frac{d\,\text{ETI}}{dt} &= \underbrace{\alpha^{P\prime}(e_1^P+\eta^P) \;+\; \alpha^{F\prime}(e_1^F+\eta^F) \;-\; \gamma'(e_1^P+e_2)}_{\text{Composition}} \\ &\quad+\; \underbrace{\alpha^P(e_1^{P\prime}+\eta^{P\prime}) \;+\; \alpha^F(e_1^{F\prime}+\eta^{F\prime}) \;-\; \gamma(e_1^{P\prime}+e_2')}_{\text{Behavioral}}\end{aligned} \tag{44}$$

Recall the assumptions so far: (i) PUT is non-deductible; (ii) convex PUT costs.; (iii) concave production, and (iv) small Frisch elasticities (hours response is locally small), so adjustments in aggregate labor supply come primarily from worker reallocation across arrangements)

There are two mechanisms through which t affects the ETI. First, there is a composition mechanism. A higher t raises the non-deductibility wedge, making PUT relatively more expensive. Firms substitute toward formal contracts, so the share of workers in the PUT arrangement decreases and the share in the formal arrangement increases. This also reduces aggregate PUT dollars in the economy (i.e., $\alpha^{P\prime} < 0,\ \alpha^{F\prime} > 0,\ \gamma' < 0$). Using $\alpha^{F\prime} = -\alpha^{P\prime} + \gamma'$ (since $\alpha^F + (\alpha^P - \gamma) = 1$), the composition term can be written as

$$\alpha^{P\prime}\big[(e_1^P+\eta^P)-(e_1^F+\eta^F)\big] \;+\; \gamma'\big[(e_1^F+\eta^F)-(e_1^P+e_2)\big]. \tag{45}$$

Under small Frisch elasticities and smooth wage pass-through (from concave production and the firm wedge), the first bracket is second-order. These effects are small in magnitude and are also proportional—one sector expands while the other contracts—so they partially offset, making the first bracket even smaller. The second bracket is positive because it contains the PUT elasticity $e_2 < 0$ and $\eta^F \le 0$; with $\gamma' < 0$ this makes (i) negative. In short: the reallocation away from PUT reduces ETI. [21]

The second mechanism is the behavioral one. This is the more intuitive channel, as this captures how the elasticities themselves move with t:

$$\alpha^P(e_1^{P\prime}+\eta^{P\prime}) \;+\; \alpha^F(e_1^{F\prime}+\eta^{F\prime}) \;-\; \underbrace{\gamma(e_1^{P\prime}+e_2')}_{\text{PUT attenuation}}\;. \tag{46}$$

With convex PUT costs, a higher t strengthens the non-deductibility wedge and makes the PUT margin less responsive to $1-\tau$: the PUT labor-supply elasticity moves toward zero ($e_1^{P\prime} > 0$) and the hiding elasticity moves toward zero ($e_2' > 0$). Hence the term $-\gamma(e_1^{P\prime}+e_2')$ is strictly negative. This

[21] Formally, with $L^P = sL$, $L^F = (1-s)L$ and $a \equiv 1-\tau$, small intensive-hours responses imply $e_1^P \approx \frac{a}{s}\frac{ds}{da}$ and $e_1^F \approx -\frac{a}{1-s}\frac{ds}{da}$, so $e_1^P + e_1^F$ is modest. Smooth, wedge-linked pass-through keeps η^P and η^F of comparable order with opposite signs. Hence $(e_1^P+\eta^P)-(e_1^F+\eta^F) = a\,\frac{d}{da}\log\big(\frac{w^P s}{w^F(1-s)}\big) \approx 0$. In the second bracket, $(e_1^F+\eta^F)-(e_1^P+e_2) = (e_1^F-e_1^P)+(\eta^F-e_2)$ is typically > 0 because $e_1^F - e_1^P > 0$ (reallocation) and $-e_2 > 0$ dominates the smooth $\eta^F \le 0$ under the baseline.

is the PUT attenuation term. The remaining pieces are second-order. Under concave production and the firm wedge $w^F = w^P + \frac{t}{1-t}\,\overline{\text{PUT}}$, wage pass-through adjusts smoothly: as the PUT margin shrinks, pass-through into w^P is dampened ($\eta^{P\prime} \le 0$). Using

$$\eta^F \;=\; \frac{w^P}{w^F}\,\eta^P \;+\; \frac{\frac{t}{1-t}\,\overline{\text{PUT}}}{w^F}\,e_2, \tag{47}$$

and noting that both η^P and e_2 shrink in magnitude as t rises, we obtain that $\eta^{F\prime}$ is small/ambiguous. These wage terms enter as $\alpha^P\eta^{P\prime} + \alpha^F\eta^{F\prime}$ and do not overturn the negative PUT-attenuation term. Moreover, the $e_1^{P\prime}$ pieces combine to $(\alpha^P - \gamma)e_1^{P\prime}$ (since $-\gamma e_1^{P\prime}$ sits opposite $+\alpha^P e_1^{P\prime}$), which is positive but typically small. In short: higher t makes the key PUT margins less elastic, so the behavioral component lowers ETI.

9.5. Optimal Income Tax: Surplus Derivation

This appendix shows that every equilibrium object is a function of a single surplus wedge in (τ, t)-space. At any interior equilibrium the five key conditions are

$$\begin{aligned}
\text{Firm FOC (indifference):}\quad & w^F = w^P + \frac{t}{1-t}\,\bar{\text{PUT}}, && \text{(A)}\\
\text{MPL–wage (formal):}\quad & \phi'(L) \;=\; w^F, \qquad L = L^P + L^F, && \text{(B)}\\
\text{Worker FOC for PUT:}\quad & \tau = \theta\bar{\text{PUT}}\, g'\big(\theta\bar{\text{PUT}}\, l\big), && \text{(C)}\\
\text{PUT labor market-clearing:}\quad & L^{D,P}(w^P, t, \bar{\text{PUT}}) = L^{S,P}(w^P, \tau, \bar{\text{PUT}}), && \text{(D}_P\text{)}\\
\text{Formal labor market-clearing:}\quad & L^{D,F}(w^F, t, \bar{\text{PUT}}) = L^{S,F}(w^F, \tau, \bar{\text{PUT}}). && \text{(D}_F\text{)}
\end{aligned}$$

I linearise these conditions around a given interior equilibrium

$$(\tau, t, \bar{\text{PUT}}, w^P, w^F, L^P, L^F),$$

and study small policy changes $(d\tau, dt)$. The argument proceeds in three steps: (i) linearise the five equilibrium conditions, (ii) derive the first-order responses of wages to $(d\tau, dt)$, and (iii) show that all equilibrium responses load on a single wedge $S(\tau, t)$.

9.5.1. Linearization of the five equilibrium conditions

Firm FOC (A). Total differentiation of (A) gives

$$\begin{aligned}
dw^F - dw^P &= \frac{\partial}{\partial t}\left(\frac{t}{1-t}\right)\bar{\text{PUT}}\,dt + \frac{t}{1-t}\,d\bar{\text{PUT}} && (48)\\
&= \frac{\bar{\text{PUT}}}{(1-t)^2}\,dt + \frac{t}{1-t}\,d\bar{\text{PUT}}. && \text{(A-diff)}
\end{aligned}$$

MPL–wage (B). Differentiating $\phi'(L) = w^F$ and using $L = L^P + L^F$,

$$\phi''(L)\big(dL^P + dL^F\big) = dw^F. \tag{B-diff}$$

Worker FOC for PUT (C). Differentiating (C) with respect to τ and $\mathrm{P\bar{U}T}$ gives

$$d\tau - \underbrace{\Big[g'(\,\cdot\,) + \theta \mathrm{P\bar{U}T}\, g''(\,\cdot\,)\Big]}_{G>0} d\mathrm{P\bar{U}T} - \underbrace{\theta \mathrm{P\bar{U}T}^2 g''(\cdot)\, dl}_{\text{Frisch elasticity term}} = 0. \tag{49}$$

We assume the intensive Frisch elasticity of hours is small, so in first-order terms we hold hours fixed, $dl \simeq 0$. This implies that most of the change in labor supply comes from switching between arrangements, not from adjusting hours. Then:

$$d\tau - G\, d\mathrm{P\bar{U}T} = 0, \qquad G \equiv \frac{\partial \tau}{\partial \mathrm{P\bar{U}T}} > 0, \tag{50}$$

so

$$d\mathrm{P\bar{U}T} = \frac{d\tau}{G}. \tag{C-diff}$$

The explicit formula for G is not needed; it is only important that it is positive and evaluated at the initial equilibrium.

PUT labor market clearing (D_P). Define

$$H^P(w^P, t, \tau, \mathrm{P\bar{U}T}) \equiv L^{D,P}(w^P, t, \mathrm{P\bar{U}T}) - L^{S,P}(w^P, \tau, \mathrm{P\bar{U}T}) = 0 \tag{51}$$

at equilibrium. Total differentiation gives

$$H^P_w\, dw^P + H^P_t\, dt + H^P_\tau\, d\tau + H^P_{\mathrm{P\bar{U}T}}\, d\mathrm{P\bar{U}T} = 0. \tag{52}$$

Define the net slopes (demand minus supply) at the equilibrium as

$$L^P_w = H^P_w, \quad L^P_t = H^P_t, \quad L^P_\tau = H^P_\tau, \quad L^P_{\mathrm{P\bar{U}T}} = H^P_{\mathrm{P\bar{U}T}}. \tag{53}$$

Then the linearized PUT market-clearing condition is

$$L^P_w\, dw^P + L^P_t\, dt + L^P_\tau\, d\tau + L^P_{\mathrm{P\bar{U}T}}\, d\mathrm{P\bar{U}T} = 0. \tag{D$_P$-diff}$$

Formal labor market clearing (D_F). Similarly, define

$$H^F(w^F, t, \tau, \mathrm{P\bar{U}T}) = L^{D,F}(w^F, t, \mathrm{P\bar{U}T}) - L^{S,F}(w^F, \tau, \mathrm{P\bar{U}T}) = 0. \tag{54}$$

Total differentiation gives

$$L^F_w\, dw^F + L^F_t\, dt + L^F_\tau\, d\tau + L^F_{\mathrm{P\bar{U}T}}\, d\mathrm{P\bar{U}T} = 0, \tag{D$_F$-diff}$$

By definition, L^P_w and L^F_w are net (demand minus supply) slopes with respect to wages, and $L^\bullet_t, L^\bullet_\tau, L^\bullet_{\mathrm{P\bar{U}T}}$ capture the demand/supply shifts induced by t, τ, and $\mathrm{P\bar{U}T}$.

So far there are five linearized equilibrium conditions:

$$\text{(A-diff)}, \quad \text{(B-diff)}, \quad \text{(C-diff)}, \quad (\mathrm{D_P}\text{-diff}), \quad (\mathrm{D_F}\text{-diff}),$$

which jointly pin down the responses $\big(dw^P, dw^F, d\mathrm{P\bar{U}T}, dL^P, dL^F\big)$ to $(d\tau, dt)$.

9.5.2. Wages as functions of the policy shocks

The worker FOC gives

$$d\bar{\text{PUT}} = \frac{1}{G}\,d\tau. \tag{55}$$

Substitute this into the firm FOC (A-diff):

$$dw^F - dw^P = \frac{\bar{\text{PUT}}}{(1-t)^2}\,dt + \frac{t}{1-t}\,\frac{1}{G}\,d\tau. \tag{56}$$

Now use the market-clearing conditions to express how wages respond to the policy shocks. Start from ($\mathrm{D_P}$-diff):

$$L^P_w\,dw^P + L^P_t\,dt + L^P_\tau\,d\tau + L^P_{\bar{\text{PUT}}}\,d\bar{\text{PUT}} = 0.$$

Solve this for dw^P:

$$dw^P = -\frac{L^P_t}{L^P_w}\,dt - \frac{L^P_\tau}{L^P_w}\,d\tau - \frac{L^P_{\bar{\text{PUT}}}}{L^P_w}\,d\bar{\text{PUT}}. \tag{57}$$

Define the reduced-form slope coefficients

$$\beta_1 \equiv -\frac{L^P_\tau}{L^P_w}, \qquad \beta_2 \equiv -\frac{L^P_t}{L^P_w}, \qquad \beta_3 \equiv -\frac{L^P_{\bar{\text{PUT}}}}{L^P_w},$$

so that

$$dw^P = \beta_1\,d\tau + \beta_2\,dt + \beta_3\,d\bar{\text{PUT}}. \tag{58}$$

Similarly, from ($\mathrm{D_F}$-diff):

$$L^F_w\,dw^F + L^F_t\,dt + L^F_\tau\,d\tau + L^F_{\bar{\text{PUT}}}\,d\bar{\text{PUT}} = 0,$$

one obtains

$$dw^F = \gamma_1\,d\tau + \gamma_2\,dt + \gamma_3\,d\bar{\text{PUT}}, \tag{59}$$

where

$$\gamma_1 \equiv -\frac{L^F_\tau}{L^F_w}, \qquad \gamma_2 \equiv -\frac{L^F_t}{L^F_w}, \qquad \gamma_3 \equiv -\frac{L^F_{\bar{\text{PUT}}}}{L^F_w}.$$

Next, eliminate $d\bar{\text{PUT}}$ using equation (55). Plugging (55) into (58)–(59) gives:

$$dw^P = \Big(\beta_1 + \frac{\beta_3}{G}\Big)d\tau + \beta_2\,dt, \qquad \equiv A_\tau\,d\tau + A_t\,dt, \tag{60}$$

$$dw^F = \Big(\gamma_1 + \frac{\gamma_3}{G}\Big)d\tau + \gamma_2\,dt, \qquad \equiv B_\tau\,d\tau + B_t\,dt. \tag{61}$$

Thus the four coefficients $(A_\tau, A_t, B_\tau, B_t)$ are determined functions of the underlying slopes $L^P_\cdot, L^F_\cdot$ and the PUT FOC parameter G.

Now impose the firm arbitrage condition from equation (56). Substituting (60)–(61) into (56) gives:

$$\underbrace{(B_\tau - A_\tau)}_{\text{coeff. on } d\tau}\,d\tau + \underbrace{(B_t - A_t)}_{\text{coeff. on } dt}\,dt = \frac{\bar{\text{PUT}}}{(1-t)^2}\,dt + \frac{t}{(1-t)G}\,d\tau.$$

Equating coefficients on $d\tau$ and dt,

$$B_\tau - A_\tau = \frac{t}{(1-t)G}, \qquad B_t - A_t = \frac{\bar{\mathrm{PUT}}}{(1-t)^2}. \tag{62}$$

At this point, A_τ, A_t, B_τ, B_t are expressed in terms of the slopes $(L_\cdot^P, L_\cdot^F)$, the PUT FOC parameter G, and the equilibrium $\bar{\mathrm{PUT}}$ and t.

9.5.3. A single surplus wedge in (τ, t)

The MPL condition (B-diff) links wages and employment:

$$\phi''(L)\big(dL^P + dL^F\big) = dw^F$$

Since L^P and L^F are themselves equilibrium objects, they must also have linear first-order responses to $(d\tau, dt)$. By the market-clearing conditions $(\mathrm{D_P})$–$(\mathrm{D_F})$, their levels are determined by the same demand and supply primitives that generated the slopes $L_\cdot^P$ and $L_\cdot^F$. Hence, for small perturbations $(d\tau, dt)$ around the initial equilibrium there exist reduced-form responses

$$dL^P = C_\tau^P\, d\tau + C_t^P\, dt, \qquad dL^F = C_\tau^F\, d\tau + C_t^F\, dt, \tag{63}$$

for some coefficients $C_\tau^P, C_t^P, C_\tau^F, C_t^F$ that are functions of the same underlying slopes $L_\cdot^P, L_\cdot^F$ and of the PUT FOC parameter G. Substituting these expressions into (B-diff) and using $dw^F = B_\tau\, d\tau + B_t\, dt$ gives:

$$\phi''(L)\big[(C_\tau^P + C_\tau^F)\, d\tau + (C_t^P + C_t^F)\, dt\big] = B_\tau\, d\tau + B_t\, dt \tag{64}$$

Then, matching coefficients gives

$$\phi''(L)(C_\tau^P + C_\tau^F) = B_\tau, \qquad \phi''(L)(C_t^P + C_t^F) = B_t \tag{65}$$

and from equation (65) it follows that:

$$\frac{B_t}{B_\tau} = \frac{C_t^P + C_t^F}{C_\tau^P + C_\tau^F} \tag{66}$$

The coefficients $C_\tau^P, C_t^P, C_\tau^F, C_t^F$ are obtained by applying the implicit-function theorem to the same five-equation system (A)–(D_F) that determines the wage coefficients $A_\cdot$ and $B_\cdot$. Hence they are functions of the same underlying primitives: the net slopes $L_\cdot^P, L_\cdot^F$, the PUT FOC parameter G, and the tax parameters $(\tau, t, \bar{\mathrm{PUT}})$.

Substituting the reduced-form expressions for $C_\tau^P, C_t^P, C_\tau^F, C_t^F$ (together with (62)) into the ratio above generates that all terms involving labor-demand and labor-supply slopes $L_\cdot^P, L_\cdot^F$ cancel. The remaining dependence on the primitives is

$$\frac{B_t}{B_\tau} = -\frac{1}{(1-t)^2}, \tag{67}$$

so the formal wage response can be written as

$$dw^F = B_\tau \Big[d\tau - \frac{dt}{(1-t)^2}\Big]. \tag{68}$$

An analogous computation shows that every other endogenous variable $X \in \{w^P, w^F, \text{P\bar{U}T}, L^P, L^F\}$ has the same policy direction, i.e.

$$dX = \kappa_X \left[d\tau - \frac{dt}{(1-t)^2} \right], \tag{69}$$

for some scalar κ_X evaluated at the equilibrium.

Since the vector $(1, -1/(1-t)^2)$ is a gradient, there exists a scalar function $S(\tau, t)$ such that

$$\partial_\tau S(\tau, t) = 1, \qquad \partial_t S(\tau, t) = -\frac{1}{(1-t)^2} \tag{70}$$

Integrating with respect to τ and t gives

$$S(\tau, t) = \tau - \frac{t}{1-t} + C_0, \tag{71}$$

and dropping the irrelevant constant C_0,

$$S(\tau, t) = \tau - \frac{t}{1-t}, \qquad dS = d\tau - \frac{dt}{(1-t)^2}. \tag{72}$$

Thus, all equilibrium objects are functions of this surplus wedge $S(\tau, t)$ rather than of τ and t separately:

$$dX = \kappa_X \, dS \tag{73}$$

9.6. ETI Estimation: Event Studies

I use event studies to complement my ETI estimates. These allow me to do two things: (i) test for parallel trends in the pre-period and (ii) see whether there is a discrete jump right when the policy takes effect.

I label the 2020–2021 change as Event A and the 2022–2023 change as Event B. Event A raises payroll (healthcare) tax rates for Agricultural Regime (AgR) firms; General Regime (GR) firms are unaffected. Event B applies only to large AgR firms in 2023; small AgR firms are unaffected. I estimate dynamic effects with monthly event time, using the month before the start of the event as the reference month.

To estimate the effects of Event A, I run the following regression for the period 2020-2022:

$$y_{it} = \alpha_i + \lambda_t + \sum_{s \neq -1} \beta_s \, \mathbf{1}\{t - g_i^A = s\} \, A_i + \varepsilon_{it}, \qquad t \in [2020, 2022], \tag{74}$$

where $A_i = \mathbf{1}\{\text{AR}\}$ takes the value of one for firms affected by Event A (the treatment, the Agricultural Regime firms). Controls are GR firms ($A_i = 0$). g_i^A is the first treated month (January 2021 for all $A_i = 1$). α_i are regime×size fixed effects (and I also include department FE in some specifications), and λ_t are month fixed effects. The coefficients β_s are dynamic DiD effects at event time s relative to $s = -1$:

$$\beta_s = \big[y_{it}(s) - y_{it}(-1) \bigm| A_i{=}1 \big] - \big[y_{it}(s) - y_{it}(-1) \bigm| A_i{=}0 \big].$$

To estimate the effects of Event B, I run the following regression only among firms in the Agricultural

Regime:

$$y_{it} = \alpha_i + \lambda_t + \sum_{s \neq -1} \theta_s \, \mathbf{1}\{t - g_i^B = s\} \, B_i + u_{it}, \qquad t \in [2022, 2023], \tag{75}$$

where the sample is restricted to AR firms. $B_i = \mathbf{1}\{\text{large AR}\}$ takes the value of one for firms affected by Event B (large AR firms), small AR firms serve as controls ($B_i = 0$). g_i^B is the first month after Event B starts (January 2023 for $B_i = 1$). Here θ_s are dynamic effects of moving from the first to the second schedule, relative to $s = -1$, netting out contemporaneous changes for small AR firms.