

ENLISTING EMPLOYEES IN IMPROVING PAYROLL TAX COMPLIANCE: EVIDENCE FROM MEXICO

Todd Kumler, Eric Verhoogen, and Judith Frías*

Abstract—Comparing two sources of wage data in Mexico—firms’ reports to the social security agency and individuals’ responses to a household survey—we document extensive underreporting of wages by formal firms, with compliance better in larger firms. We also present evidence that the 1997 Mexican pension reform, which tied pension benefits more closely to reported wages for younger workers, led to a relative decline in underreporting for younger age groups. The results suggest that giving employees incentives and information to improve the accuracy of employer reports can be an effective way to improve payroll tax compliance.

I. Introduction

A growing body of research suggests that lack of fiscal capacity—in particular, difficulty in raising taxes to fund public goods—is an important constraint on economic performance in developing countries (Burgess & Stern, 1993; Besley & Persson, 2013). Developing countries generally have low ratios of tax revenues to GDP and large informal sectors. Mexico is no exception: it has the lowest tax revenue share of GDP in the OECD, between 15% and 20% during the period we study, and the informal sector has been estimated to make up 40% or more of total output (OECD, 2011; IMF, 2010; Schneider & Enste, 2000). Given weak enforcement institutions and widespread evasion, the task of improving fiscal capacity in developing countries is a difficult one, and there is acute interest among researchers and policymakers in potential remedies.

Evidence from developed countries suggests that having firms report employees’ wages—often referred to as *third-party reporting*—greatly reduces noncompliance. Careful studies of the “tax gap” by the U.S. Internal Revenue Service indicate that in 2001, about 57% of nonfarm proprietor income but only 1% of wages and salaries went unreported (Internal Revenue Service, 2006; Slemrod, 2007). Saez (2010) finds significant bunching around the first kink point of the earned income tax credit, suggesting misreporting, only among the self-employed. In a randomized audit

study in Denmark, Kleven et al. (2011) find little evasion when incomes are reported by employers or other third parties. The view that having firms report wages is effective in ensuring compliance is widespread among practitioners and government agencies (see Plumley, 2004, and OECD, 2006).

To what extent does the accuracy of firms’ wage reports carry over to developing countries, with their weaker enforcement regimes? In this paper, we draw on rich microdata from Mexico to estimate the extent of wage underreporting by formal firms and how it responded to an important change in the Mexican social security system. To measure underreporting, we compare two sources of detailed wage information: firms’ reports of individuals’ wages to the Mexican social security agency and individuals’ responses to a household labor force survey. We construct three measures of evasion for different demographic groups, based on the median and mean wage differences between the two data sets and the excess mass in the social security data to the left of a given cutoff in the household data.

In cross-sectional results, we document substantial underreporting of wages by formal firms. We also find that evasion is declining in firm size.¹ We believe that this paper is the first to show systematic differences in wage underreporting by firm size using direct information on firms’ wage reports. This finding is consistent with a simple partial-equilibrium model of endogenous compliance by heterogeneous firms, which we summarize briefly in section III and present in full in online appendix B. The cost of evasion is assumed to be increasing in both the unreported part of the wage per worker and in firm output, for reasons that may include the greater difficulty of maintaining collusion in larger firms (as in Kleven et al., 2016), or simply the greater visibility of larger firms to auditors. The finding that compliance is increasing in firm size is consistent with the idea that the burden of taxation in developing countries falls more heavily on larger firms and that this is part of the explanation for the disproportionately large number of small firms (Hsieh & Klenow, 2014; Hsieh & Olken, 2014).

We also provide evidence that evasion responded to an important change in the Mexican social security system in the way that our simple economic model would predict. We focus on a pension reform that introduced a system of personal retirement accounts, passed by the Mexican Congress on December 21, 1995, and implemented on July 1, 1997. As discussed in more detail below, prior to the reform, the social security benefits of most workers were insensitive to the

Received for publication April 27, 2015. Revision accepted for publication August 14, 2019. Editor: Asim I. Khwaja.

*Kumler: Cornerstone Research; Verhoogen: Columbia University; Frías: Consejo de la Judicatura Federal.

We thank (without implicating) Emma Aguila for sharing information on tax rates and benefit formulas; Matthew Clifford for research assistance; and David Card, David S. Kaplan, Wojciech Kopczuk, Sebastian Galiani, François Gerard, Roger Gordon, Santiago Levy, Suresh Naidu, Ben Olken, Kiki Pop-Eleches, Bernard Salanié, Monica Singhal, many seminar participants, and three anonymous referees for helpful comments. E.V. gratefully acknowledges the support of the National Science Foundation (SES-0721068). The views expressed in this paper are our own and do not necessarily reflect the views of the Instituto Mexicano del Seguro Social or Cornerstone Research. This paper was previously circulated under the title “Enlisting Workers in Monitoring Firms: Payroll Tax Compliance in Mexico” (Kumler, Verhoogen, & Frías, 2012).

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/rest_a_00907.

¹As discussed in the online data appendix, neither main data set draws a very clear distinction between firms and establishments. In Mexico, a large majority of firms are single establishment (INEGI, 2009). Here we treat the terms *firm* and *establishment* as interchangeable and mainly refer to firms.

wages reported by firms on their behalf as long as the firms reported at least the minimum allowable wage. The reform tied individual pensions more closely to firms' wage reports and made it easier for employees to observe those reports. Workers already in the traditional system prior to July 1, 1997, retained the right to choose, at the time of retirement, the pension that they would have received under the prereform regime. Because older workers had little time to accumulate sufficient balances in their personal accounts, their expected pension was higher under the old regime. Younger workers had a greater expectation of being better off under the new regime and hence had stronger incentives to ensure accurate reporting. We use this differential impact by age as the basis for a difference-in-differences estimation strategy. Consistent with our theoretical model, evasion declines relatively more for younger age groups. This result suggests that giving employees incentives and information to improve the accuracy of employer reports can be an effective way to improve payroll tax compliance.

A key limitation of the data we use is that the household survey does not contain firm identifiers, and we are not able to construct measures of evasion at the firm level. Instead, we construct measures of evasion at the level of cells defined by different combinations of metropolitan areas, sectors, firm-size categories, and age groups, depending on the specification. A second limitation is that it is difficult to separate the effects of the change of incentives and the change of information with the pension reform (discussed in more detail in section II). The model we develop presumes that evasion is collusive, and in that case, the information has little effect in addition to the change in incentives. But in models where workers are less than perfectly informed about firms' reports, the information itself could have an effect on wages and evasion in equilibrium. The "experiment" we consider combined the two elements, and the effects we estimate should be interpreted as the joint effect of both.²

This paper is related to a number of different literatures. Research in development economics on the noncompliance of firms with tax regulations has tended to focus on firms failing to register with tax authorities or hiring labor off the books, which we might term extensive margins of compliance (Gordon & Li, 2009; McKenzie & Sakho, 2010; de Mel, McKenzie, & Woodruff, 2013; Bruhn & McKenzie, 2014; Ulysea, 2018; Rocha, Ulysea, & Rachter, 2018; De Giorgi, Ploenzke, & Rahman, 2018). In this paper, by contrast, we focus on an intensive margin of compliance: the extent of compliance by formally registered firms reporting wages for formally registered workers.

This paper is also related to a large and growing literature seeking to measure the extent of evasion of taxes on incomes, reviewed recently by Slemrod (2019) and previously by Andreoni, Erard, and Feinstein (1998), Slemrod and Yitzhaki

(2002), and Saez, Slemrod, and Giertz (2012). Besides doing direct audits, a common approach in this literature has been to compare reported incomes with "traces" of true incomes. In a classic paper, Pissarides and Weber (1989) use food expenditures as an indicator of true incomes and compare to reported incomes separately for employees and self-employed people, where the latter are presumed to be more likely to evade.³ By contrast, we use responses to a household survey as an indicator of true incomes to compare to reported incomes in administrative data. This paper appears to be the first in the literature to do so.⁴ The focus on systematic differences in compliance by firm size is also a distinctive contribution. Other recent work on salary misreporting by firms includes Nyland, Smyth, and Zhu (2006), Bérigolo and Cruces (2014), and Mao, Zhang, and Zhao (2013).⁵

This paper appears to be the first to analyze how tying benefits more closely to reported wages can contribute to improved payroll tax compliance.⁶ This idea is related to recent work on incentivizing decentralized agents to improve tax enforcement. Kopczuk and Slemrod (2006), Keen and Lockwood (2010), and Pomeranz (2015) argue that value-added taxes (VATs) have attractive enforcement properties in part because they give each party in a supply-chain transaction greater incentive to ensure that the other reports accurately. A recent paper by Naritomi (2019) analyzes a Brazilian program to give consumers incentives to ask for receipts from retail establishments and finds positive effects on compliance.⁷

More broadly, this paper is in the spirit of a growing empirical literature in development economics examining how corruption and other forms of illegal behavior respond to economic incentives, recently surveyed by Olken and Pande (2012). It is part of a small but growing literature using administrative records from developing countries to document various aspects of taxpayer behavior (Pomeranz, 2015; Kleven & Waseem, 2013; Best et al., 2015). It is also related to an active

³Feldman and Slemrod (2007) pursue a similar approach using charitable contributions. Also related are papers by Braguinsky, Mityakov, and Liscovich (2014) and Tonin (2011), which use car registries and food consumption, respectively.

⁴A recent paper by Paulus (2015), which appeared after our paper had been circulated, compares individual tax reports to individual survey responses and does not focus on underreporting by firms. Papers using the general strategy of comparing information from more than one data source to infer illicit behavior (in other contexts) include Fisman and Wei (2004), Olken (2006), Gorodnichenko, Martinez-Vazquez, and Peter (2009), Marion and Muehlegger (2008), Hurst, Li, and Pugsley (2014), and Niehaus and Sukhtankar (2013).

⁵Bérigolo and Cruces (2014) consider individuals' responses to an extension of a child benefit in the Uruguayan social security system and find an effect on salary misreporting by firms as reported by employees, among other effects. Best (2014) considers heterogeneity across firms based on firms' and individuals' reports to the tax authority in Pakistan. But it does not have a source of wage information not subject to misreporting incentives, and to the extent that firms and workers collude, the difference between the firms' and workers' responses is likely to be an inaccurate measure of the extent of misreporting.

⁶Bailey and Turner (2001) suggest verbally that tying pension benefits to contributions would have the effect of reducing evasion.

⁷In other related work, Khan, Khwaja, and Olken (2016) find that randomized incentives to tax inspectors in Pakistan increase tax revenues.

²Our argument should not be interpreted as advocating a system of personal accounts per se; one could imagine a change in pension benefits under the traditional pay-as-you-go system that would have had similar effects.

recent literature on the role of firms in tax systems (Kopczuk & Slemrod, 2006; Gordon & Li, 2009; Dharmapala, Slemrod, & Wilson, 2011).

II. Institutions: The Mexican Social Security System

Because our empirical strategy relies crucially on incentives in the Mexican social insurance system, this section describes the system and the pension reform in some detail. Other salient dimensions of the Mexican tax system are discussed briefly in appendix A.1. Because of data constraints, discussed in more detail below, we focus on the years 1988 to 2003. Because the incentives and empirical patterns for women are complicated by rapidly changing labor force participation, we focus primarily on men in the empirical analysis.⁸

A. General Features

The Instituto Mexicano del Seguro Social (IMSS), the Mexican social security agency, is the primary source of social insurance for private sector workers in Mexico.⁹ Beginning with its creation in 1944, IMSS operated as a pay-as-you-go (PAYGO) scheme financed by payroll taxes. By the late 1980s, rising health care costs and an increase in the number of pensioners relative to the working-age population led to projected shortfalls in the IMSS financial accounts. On December 21, 1995, in part because of concerns about the financial viability of the system, the Congress enacted a comprehensive pension reform, to take effect on July 1, 1997.¹⁰ This reform replaced the PAYGO system with a system of personal retirement accounts (PRA). More extensive discussions of the reform are provided in Grandolini and Cerda (1998), Sales-Sarrapy, Solis-Soberon, and Villagomez-Amezcuca (1996), and Aguila (2011).

IMSS requires contributions from both employers and employees based on reported wages; these are supplemented by government contributions. The contribution schedules reflect a complicated set of formulas determining contributions to the various components of the IMSS system, principally health care, pensions, and child care. Full details are pre-

sented in appendix A.2. For our purposes, the key point is that the changes in contributions were the same for all age groups, and their effects will be differenced out in our difference-in-differences procedure.

Any worker in the system is entitled to free health care at IMSS hospitals and clinics, for himself or herself and immediate family, as well as child care benefits and a number of other nonpension benefits, independent of the reported wage, as detailed in appendix A.3. Although it is difficult to estimate workers' valuations of these benefits, the health care and child care benefits did not change with the 1997 pension reform, and the valuations will arguably also be differenced out in our difference-in-differences procedure.¹¹

Neither before nor after the reform was there a reward to employees for revealing evasion by their employers beyond ensuring accurate reporting of their own wages. (See appendix A.4.) Although we argue that evasion has been widespread, at least one aspect of IMSS reporting requirements does appear to have been strictly enforced. By law, firms in Mexico are required to pay the relevant minimum wage (of which there are three in Mexico, with the highest in Mexico City and the lowest in rural areas) and a holiday bonus called an *aguinaldo*, worth two weeks of salary—approximately 4.5% of annual earnings. In order to avoid fines, establishments are required to report wages of at least the corresponding minimum wage plus 4.5% throughout the year.¹²

B. Prereform Pension System

Under the prereform PAYGO regime, workers became vested in the system after ten years of contributions and were then entitled to receive at least the minimum pension. Pensions were calculated on the basis of the final average wage, defined as the average nominal wage in the five years preceding retirement. Figure 1A illustrates the expected daily pension as a function of the final average wage for workers with ten, twenty, and thirty years of contributions in selected years. The schedules combine a minimum pension guarantee with a benefit proportional to an individual's wage. At first glance, the pension values illustrated in panel A appear to be sensitive to the reported final average wage, but in the years leading up to the reform, inflation had severely eroded the real value of wages and pensions, such that a large majority of workers had final average wages in the region in which the minimum was binding. Inflation exceeded 50% in every year in the volatile 1982–1988 period and exceeded 100% in 1987 and 1988; it was above 25% in a number of subsequent years (1990–1991 and 1995–1996). (See appendix table A4.) In response to public pressure, the Mexican Congress

⁸In addition to the changing labor force participation, the patterns for women are complicated by the facts that many women receive IMSS benefits through their spouses and that because of relatively low labor force participation by older women, sample sizes in the ENEU household survey (described below) are often inadequate when looking separately by metropolitan area and firm size. Results for women are presented in the appendix to the working paper version of the paper (Kumler, Verhoogen, & Frías, 2015). Briefly, the cross-sectional patterns are robust for women, but the difference-in-differences results are not, in part for the reasons just discussed.

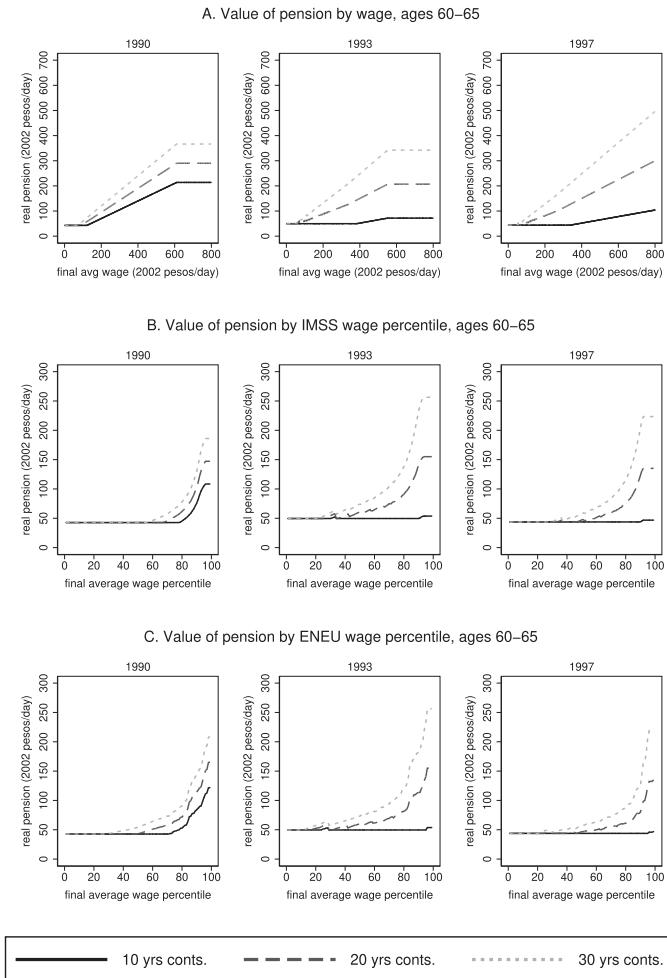
⁹Public sector workers and workers for PEMEX, the state-owned oil company, are covered by separate systems. In 2003, the government created an alternative system, Seguro Popular, which provides basic health coverage for all individuals and is not tied to formal employment. In this paper, we focus on the IMSS system and sectors with minimal government employment.

¹⁰This change followed an unsuccessful partial reform in 1992, described in appendix A.5.

¹¹IMSS also provides an individual savings account for housing expenditures, which in some cases can be used to contribute to an individual pension. See appendix A.3 for details.

¹²Prior to 1991, there are a scattered few reports of wages below this level; beginning in 1991, IMSS stepped up enforcement of this rule, and such wages have no longer been observed.

FIGURE 1.—PENSION VALUES BY YEARS OF CONTRIBUTIONS, SELECTED YEARS



Final average wage (2002 pesos/day) is the average nominal daily wage over five years prior to retirement, deflated to constant 2002 pesos. The figure indicates pension values for individuals with ten, twenty, and thirty years of contributions to IMSS. In panel B, we calculate the nominal wage at each quantile of the IMSS wage distribution for 60- to 65-year-old men in each year and take the average for that quantile over the preceding five years. Panel C is constructed similarly using wage distributions from the ENEU baseline samples. See section IV for details of samples and sections IIB and IIC for details on pension benefits. Average 2002 exchange rate: 9.66 pesos/dollar.

in 1989 increased the minimum pension to 70% of the minimum wage in Mexico City and subsequently indexed it to the minimum wage going forward, without raising the value of pensions greater than the minimum.¹³ The Congress subsequently raised the value of the minimum pension relative to the minimum wage, until it reached 100% of the minimum wage in Mexico City in 1995.

As a consequence of the erosion of the real value of pensions above the minimum and the legislative interventions to raise the minimum, the fraction of workers who expected to receive the minimum pension remained high throughout the prereform period. Figure 1B plots the real value of the pension for male workers with ten, twenty, or thirty years

¹³That is, if a worker's final average wage was twice the minimum wage in 1991, the pension payment in 1992 was calculated on the basis of twice the minimum wage. The real minimum wage declined steadily over the period (see appendix table A4), so the slowing of the erosion of pensions as a result of this change was modest.

of contributions against the final average wage percentile of men 60 to 65 years old in the IMSS data, for selected years.¹⁴ In 1990, approximately 80% of male retirees with ten years of contributions received the minimum pension. The corresponding numbers for male workers with twenty to thirty years of contributions were 70% and 60%, respectively. In 1997, just prior to the implementation of the pension reform, nearly all workers with ten years of contributions, roughly 50% of those with twenty years, and 40% of those with thirty years could expect to receive the minimum pension.¹⁵ Unfortunately, the data to which we have access do not contain total years of contributions by each individual worker, and hence we are not able to calculate the precise number of workers receiving the minimum pension. But analysts with access to this information report that approximately 80% of retirees were receiving the minimum pension prior to the reform (Grandolini & Cerda, 1998).¹⁶

Strictly speaking, pension values were insensitive to final wages only for inframarginal workers whose true final wage corresponded to the minimum pension. If wages were underreported to IMSS, as we argue below, then the graphs in figure 1B likely overstate the fraction of workers whose pensions were insensitive to underreporting. To address this, in figure 1C, we plot similar graphs using final average wage percentiles calculated from the ENEU household data (described in section IV), which should not be subject to underreporting. We see that somewhat smaller fractions of workers with ten, twenty, and thirty years of contributions would have received the minimum pension. But the key point is that the graph for 1997 resembles quite closely the corresponding graph in panel B: essentially all workers with ten years of contributions would have received the minimum pension, as well as more than 40% of workers with twenty years and more than 20% of workers with thirty years.

C. Postreform (Personal Retirement Accounts) System

Under the postreform personal retirement account (PRA) system, employees, employers, and the government are required to make contributions to workers' personal retirement accounts in each period. Each worker is required to choose an investment institution, known as an Administrador de

¹⁴To calculate the final average wage percentile, we calculate the nominal wage at each percentile of the IMSS wage distribution for men 60 to 65 years old in each of preceding five years, then take the average for each percentile.

¹⁵In addition, there was a penalty for retirement before age 65 of 5% per year (i.e., a worker who retired at age 60 would have his or her pension reduced by 25%), but this penalty was not allowed to reduce the pension below the minimum. This reduced the disincentive to retire early for workers with pensions near the minimum (Aguila, 2011).

¹⁶In addition, because pensions were calculated only on the basis of the last five years of employment, any worker who was certain that he or she would work for more than five years in covered employment could also be certain that the current reported wage would not affect the pension benefit. In unreported results, we have investigated whether we see an increase in reported wages five years before retirement, as one might expect if workers were being sophisticated in adjusting strategically to the five-year rule, but we do not find a significant change.

TABLE 1.—PENSION WEALTH SIMULATION, BY AGE IN 1997, MALE WORKER WITH 35 YEARS OF EXPECTED CONTRIBUTIONS

Age in 1997	Years of Expected PRA Contributions	Plan	Real Daily Wage					
			43	100	200	300	500	1079
25	35	PRA	398.6	<i>815.0</i>	<i>1,626.2</i>	<i>2,437.3</i>	<i>4,059.7</i>	<i>8,751.9</i>
		PAYGO	398.6	398.6	603.8	890.2	1,483.6	3,200.1
30	30	PRA	398.6	<i>523.4</i>	<i>1,044.3</i>	<i>1,565.3</i>	<i>2,607.1</i>	<i>5,620.5</i>
		PAYGO	398.6	398.6	603.8	890.2	1,483.6	3,200.1
35	25	PRA	398.6	<i>398.6</i>	<i>659.1</i>	<i>987.8</i>	<i>1,645.3</i>	<i>3,546.9</i>
		PAYGO	398.6	398.6	603.8	890.2	1,483.6	3,200.1
40	20	PRA	398.6	<i>398.6</i>	<i>403.9</i>	<i>605.4</i>	<i>1,008.4</i>	<i>2,173.9</i>
		PAYGO	398.6	398.6	603.8	890.2	1,483.6	3,200.1
45	15	PRA	398.6	<i>398.6</i>	<i>398.6</i>	<i>398.6</i>	<i>586.6</i>	<i>1,264.7</i>
		PAYGO	398.6	398.6	603.8	890.2	1,483.6	3,200.1
50	10	PRA	398.6	<i>398.6</i>	<i>398.6</i>	<i>398.6</i>	<i>398.6</i>	<i>662.6</i>
		PAYGO	398.6	398.6	603.8	890.2	1,483.6	3,200.1
55	5	PRA	398.6	<i>398.6</i>	<i>398.6</i>	<i>398.6</i>	<i>398.6</i>	<i>398.6</i>
		PAYGO	398.6	398.6	603.8	890.2	1,483.6	3,200.1

Values are real present discounted value of the future stream of pension benefits in thousands of 2002 pesos, for a male worker who began contributing at age 25 and expects to continue until age 60. Numbers in italics indicate that a personal retirement account (PRA) has a higher expected payoff than the prereform pension (PAYGO). Average 2002 exchange rate: 9.66 pesos/dollar. Forty-three pesos is real daily minimum wage (in Mexico City) in 1997, and 1,079 pesos is the topcode we impose (corresponding to the lowest real value of IMSS topcode over study period). See sections IIB and IIC and appendix A.6 for further details.

Fondos de Ahorro para el Retiro (AFORE; retirement savings fund administrator), to manage his or her account.¹⁷ The reform also specified a minimum pension equal to the minimum wage on July 1, 1997, with further increases in the minimum pension indexed to the Consumer Price Index. Eligibility for the minimum pension was raised from 10 years of contributions to 25 years of contributions.

The establishment of the new pension regime created two categories of workers: “transition” workers who first registered with IMSS before July 1, 1997, and new workers who first registered after July 1, 1997. At retirement, transition workers are given a choice between receiving pension benefits under the PAYGO scheme or the PRA scheme. The PAYGO pension is calculated as if workers’ postreform contributions were under the old regime. If a transition worker opts for the PAYGO pension, IMSS appropriates the balance of his or her personal retirement account. The only option for new workers is the PRA.

To illustrate the impact of the reform on pension wealth, we conduct a simulation of pension wealth under the two regimes, based on a similar simulation by Aguila (2011). In carrying out the simulation, we choose a relatively optimistic annual return on the personal accounts: 8.59%, the average return from 1998 to 2002, as in the more optimistic of the two scenarios considered by Aguila (2011). We also assume that participants expected the real value of the minimum wage to decline, as it had done for more than a decade (see appendix table A4). Assumptions of lower interest rates and less rapid declines in the real minimum wage would be less favorable to the PRAs. Details of the simulation are in appendix A.6.

One way to see the differences in incentives by age in the system is to compare pension wealth for workers of different ages in 1997. Table 1 displays the real present value of pension wealth by wage level for male workers of different ages

in 1997, all of whom began working at age 25 and expect to continue working until age 60, assuming real wages are constant over their lifetimes. Numbers in italics indicate where the PRA pension is more valuable than the PAYGO pension.

This simulation should be treated with caution, for two reasons. First, in carrying it out, we have made a number of assumptions about the future paths of wages, investment returns, and the minimum wage. The calculations in the table are sensitive to these assumptions. Second, what matters for the mechanism we are trying to highlight are employees’ expectations of the relative future value of the different pensions and how they respond to reported wages. We do not know how individuals form these expectations or how well informed they are. For both reasons, we believe that the values of the PRA and PAYGO pensions in table 1 should be interpreted as quite noisy indicators of individuals’ true expectations.

With those caveats, the basic message of the simulation seems clear: the PRA pension is expected to be relatively more valuable to younger workers, who expect to contribute to the personal account for more years. Within an age group, the PRA pension is relatively more attractive for higher-wage workers.¹⁸ We believe that participants understood this basic message at the time of the reform, even though we do not believe that we can draw sharp predictions about the precise

(2012) investigate the role of behavioral issues in employees’ choices of AFORES.

¹⁸Another way to see the effect of the reform is to consider the values of the pensions for different numbers of years of expected contributions, for a worker who entered the system on June 30, 1997, as presented in appendix table A5. Note that workers with fewer than 10 years of contributions are better off under the new regime, since they receive no pension under the old regime but a small pension under the new regime. But conditional on a worker having at least 10 years of contributions, we again see that the attractiveness of the PRA pension is increasing in the number of years of contributions and the wage. The median wage for male workers is just above 100 pesos/day, and for a worker at this level, the PRA becomes more attractive only if he expects to contribute for more than 25 years.

¹⁷The AFORE management fees are in many cases substantial, and it is not clear that workers choose AFORES optimally. Duarte and Hastings

age above which the traditional PAYGO pension dominates the personal account (PRA) pension.

Another important aspect of the pension reform is that the law requires AFOREs to send an account statement to each holder of a personal retirement account every four months. A redacted example of such an account statement appears as appendix figure A3. The account statement reports previous balances (*saldo anterior*), new contributions (*aportaciones*), withdrawals (*retiros*), interest earned (*rendimientos*), AFORE commissions charged (*comisiones*), and final balances (*saldo final*) for the pension account, as well as for the voluntary savings account (see appendix A.2) and the housing savings account (see appendix A.3). The bottom section reports three-year returns and commissions for each AFORE, as well as the average five-year net return (at left). It appears that these account statements made it significantly easier for workers to figure out how much employers were contributing on their behalf. This mechanism would not be expected to reduce evasion if employers and employees were colluding in underreporting wages, but it may have reduced evasion in cases in which workers were unaware that their employers were underreporting their wages.

III. Conceptual Framework

To organize our empirical analysis, we have developed a simple partial-equilibrium model of the compliance decisions of heterogeneous firms, in which employees and firms collude in underreporting (as in Yaniv, 1992) and firms are monopolistically competitive and differ in productivity (as in Melitz, 2003). The model shares with a number of existing models the feature that less able entrepreneurs, whose firms are smaller, comply less than more able entrepreneurs (Rauch, 1991; Dabla-Norris, Gradstein, & Inchauste, 2008; De Paula & Scheinkman, 2011; Dharmapala et al., 2011; Galiani & Weinschelbaum, 2012) but differs in that we consider partial compliance: wage underreporting by formally registered firms, as opposed to a binary decision about whether to register.¹⁹ To save space, we have put the full model in appendix B; here we briefly summarize the main ideas.

Let w_r be the pretax wage reported by a firm to the government, w_u the unreported wage (paid under the table), and τ the tax rate (the sum of firm and worker contributions). Then the net take-home wage workers receive is $w_{net} = w_u + (1 - \tau)w_r$. Rearranging,

$$w_u = w_{net} - (1 - \tau)w_r. \quad (1)$$

¹⁹Three other recent papers discuss heterogeneity of firms' tax-compliance decisions. Kleven, Kreiner, and Saez (2016) consider a particular mechanism that generates greater compliance among larger firms—the increasing difficulty of maintaining collusion as the number of employees increases—but do not focus on differential responses to tax or benefit changes. Besley and Persson (2013) note that if compliance costs depend on firm size, then firm heterogeneity will matter for compliance, without taking a position on the source of the firm heterogeneity or on the implication for responses to tax changes. Dharmapala et al. (2011) consider the optimal taxation of firms in a setting with firm heterogeneity and the implications for firm size distributions but do not focus on wage underreporting.

In the empirics, w_r will correspond to the pretax wage reported by the firm in the administrative records of the social security agency, $(1 - \tau)w_r$ to the posttax reported wage, and w_{net} to the take-home pay reported by workers in the ENEU household survey. The difference between the (average) ENEU take-home wage and the (average) IMSS posttax wage, which we call the wage gap, corresponding to w_u in equation (1), will be our primary measure of evasion.

We assume that the cost of evasion is given by $xc(w_u)$, where $c'(w_u) > 0$, $c''(w_u) > 0$ and x is the output of the firm. One justification for this assumption is simply that auditors are more likely to audit larger firms because their operations are more visible, as suggested by Besley and Persson (2013). Another is the argument of Kleven et al. (2016) that collusion in underreporting is more difficult to sustain in larger firms. Whatever the underlying mechanism, the assumptions on the cost-of-evasion function give us our first key theoretical implication: in equilibrium, more productive firms, which are larger, evade less.

In our static setting, we model the per-period value of the future pension benefit as bw_r , where we call b the “benefit rate.”²⁰ We assume that $b < \tau$, which arguably corresponds to the Mexican institutional setting. This assumption means that there is a rent to not reporting wages at the margin (some of which may be shared with employees); firms weigh their share of this rent against the costs of evasion. In developed countries, firms' incentives to underreport wages to evade payroll taxes are often offset by incentives to overreport wages in order to reduce corporate taxes. But corporate tax evasion and avoidance is common in Mexico (see appendix A.1), and the social security agency and tax authority had no data-sharing agreement over almost all of our study period, so it appears that firms' countervailing incentives to overreport (or not underreport) wages were weak.

Given how we model pension benefits, the total effective wage, which we denote by w_e , is, using equation (1):

$$w_e = w_{net} + bw_r = w_u + (1 - (\tau - b))w_r. \quad (2)$$

We assume that the labor market is competitive and that workers' labor supply responds to the effective wage, w_e .²¹ It can be shown that an increase in the benefit rate, b , will lead firms to rely more heavily on the reported wage, w_r , in the compensation package to achieve a given market-clearing effective wage. This is our second key theoretical implication: an increase in the pension benefit rate will lead to a decrease in the unreported wage, w_u , within each firm. The model considers homogeneous workers, but could be easily extended to consider other types of workers, who differ in the benefit

²⁰One could also introduce a parameter to capture the fixed benefit of participation in the social security system. This parameter would affect workers' utility levels and participation decisions, but since it would not affect workers' incentives at the margin of how much of their wage is reported, it would not affect firms' evasion behavior, the primary object of interest.

²¹We assume that workers observe both w_{net} and w_r , and hence w_u and w_e . In this sense, workers collude in underreporting: they are aware of it and do not report it.

rate they face. We would then expect the unreported wage, w_u , to decline more for workers who face a greater increase in the benefit rate, b .²²

An important issue in this context is the incidence of the change in the pension benefit rate on wages. Theoretically, it is possible to show that for a finite labor-supply elasticity, the effective wage, w_e , is increasing in the benefit rate, b . If b rises, the government ends up paying a larger share of the effective wage, and some of this increased contribution redounds to workers. But in general it is not possible to sign the effects of the reform on the observable wage measures, the firm-specific reported wage, w_r , or the firm-specific take-home wage, w_{net} , for reasons discussed in the appendix. It is worth emphasizing, however, that in the model, the response of the unreported wage, w_u (i.e., evasion), to the policy change does not depend on the incidence of the policy change on w_e , w_r , or w_{net} . In this sense, the model suggests that it is reasonable to examine the effect of the policy change on evasion separately from the question of incidence, which is how we proceed in the empirical analysis.

IV. Data

The establishments' wage reports are drawn from IMSS administrative records. All private Mexican employers are in principle legally obligated to report wages for their employees and pay social security taxes on the basis of the reports. The IMSS data set contains the full set of wage reports for employees in registered, private sector establishments over the period 1985 to 2005.²³ The data set contains a limited set of variables: age, sex, daily wage, state, and year of the individual's first registration with IMSS, an employer-specific identifier, and industry and location of the employer. Wages are reported in spells (with a begin and end date for each wage level), and in theory we could construct a day-by-day wage history for each individual. To keep the data set manageable, we extract wages for a single day, June 30, in each year. Prior to 1997, records for temporary workers were not collected in digital form. To ensure comparability before and after 1997, we focus on workers identified in the IMSS data as permanent, defined as having a written contract of indefinite duration.

We select ages 16 to 65. To maintain consistency across years, we impose the lowest real value of the IMSS topcode for wage reporting (which occurred in 1991) in all years. We drop establishments with a single insured worker since these are likely to be self-employed workers. We include only the metropolitan areas in the ENEU samples (described below). We also focus on sectors for which we are confident that IMSS is the only available formal sector social insurance program: manufacturing, construction, and re-

tail/hotel/restaurants. Other broad sectors contain a substantial share of public employees, who are typically covered by a separate system.²⁴ We focus on men, for the reasons discussed in section II. When individuals have more than one job, we select the highest-wage job. We refer to the sample selected following these criteria as our IMSS baseline sample. We also construct a small, separate sample of male workers in the social security sector itself who are covered by the IMSS system, following the same selection criteria as for the IMSS baseline sample. Further details on sample selection and data processing are in appendix C.

The household data we use are from the Encuesta Nacional de Empleo Urbano (ENEU; National Urban Employment Survey), a household survey similar to the Current Population Survey (CPS) in the United States, collected by the Instituto Nacional de Estadísticas y Geografía (INEGI), the Mexican statistical agency. The original ENEU sample, beginning in 1987, focused on the sixteen largest Mexican metropolitan areas. Although the geographical coverage expanded over time, we focus on the original sixteen areas to maximize the number of prereform years. The ENEU samples households and surveys individuals within households about their employment. As in the IMSS data, we include male workers ages 16 to 65, focus on the second quarter of each year, exclude self-employed workers, impose the 1991 IMSS topcode in all years, and include only manufacturing, construction, and retail/hotels/restaurants. When individuals report having more than one job, we use the information only from their main job. When aggregating the ENEU information to the cell level, we use the sampling weights provided by INEGI.

A very useful feature of the ENEU for our purposes is that it asks respondents whether they receive IMSS coverage as an employment benefit. Beginning in the third quarter of 1994, the ENEU also asks respondents whether they had a written contract of indefinite duration, the legal definition of a permanent employee that IMSS used. Hourly wages are calculated as monthly wages divided by 4.3 times hours worked in the previous week and daily wages as 8 times hourly wages.²⁵ The ENEU wage measures are based on respondents' reports of take-home pay, after social security taxes have been paid. They also exclude bonuses paid less frequently than monthly, and hence exclude the yearly *aguinaldo* bonus. The differences between the IMSS and the ENEU wage measures are discussed further in appendix C.

Although the ENEU survey does not contain a firm identifier, it does ask respondents about the size of the firm at which they work. We use this information to generate a firm-size indicator taking on values 1 to 10, 11 to 50, 51 to 100,

²²An additional implication of the model is that a decrease in the tax rate, τ , has an analogous effect to an increase in the benefit rate on compliance; we return to this briefly in section VII.

²³The data have been used in several previous papers, including Castellanos, Garcia-Verdu, and Kaplan (2004) and Frias et al. (2018).

²⁴We focus on manufacturing, construction, and retail/hotel/restaurants in part so that we can be confident that respondents to the household survey are not mistaking coverage under the public sector system for IMSS coverage.

²⁵Respondents can choose the time frame over which to report their earnings: hourly, daily, weekly, biweekly, or monthly. The ENEU enumerators then convert them to monthly earnings.

TABLE 2.—COMPARISON OF IMSS BASELINE SAMPLE AND VARIOUS ENEU SAMPLES

	IMSS Baseline Sample (1)	Full ENEU Sample (2)	ENEU with IMSS (3)	ENEU without IMSS (4)	ENEU Permanent with IMSS (5)	ENEU Full-Time with IMSS (6)
A. 1990						
Real average daily posttax wage	121.02 (0.07)	163.88 (1.58)	172.98 (1.94)	143.88 (2.62)		166.73 (1.85)
Age	31.75 (0.01)	31.46 (0.15)	32.13 (0.17)	29.98 (0.29)		32.22 (0.17)
Fraction employed in establishments > 100 employees	0.52 (0.00)	0.43 (0.01)	0.55 (0.01)	0.18 (0.01)		0.55 (0.01)
<i>N</i> (raw observations)	1,691,417	16,169	11,592	4,577		10,978
<i>N</i> (population, using weights)	1,691,417	2,578,847	1,772,523	806,324		1,645,229
B. 2000						
Real average daily posttax wage	123.60 (0.07)	148.20 (1.31)	161.15 (1.60)	120.78 (2.16)	166.42 (1.80)	155.80 (1.59)
Age	32.70 (0.01)	32.22 (0.14)	32.82 (0.16)	30.94 (0.28)	33.22 (0.17)	32.88 (0.16)
Fraction employed in establishments > 100 employees	0.58 (0.00)	0.44 (0.01)	0.59 (0.01)	0.10 (0.01)	0.63 (0.01)	0.59 (0.01)
<i>N</i> (raw observations)	2,420,307	19,171	14,063	5,108	11,918	13,246
<i>N</i> (population, using weights)	2,420,307	3,509,828	2,384,267	1,125,561	2,042,988	2,225,318

All columns focus on wage-earning male workers ages 16 to 65 in manufacturing, construction, and retail/hotel/restaurant sectors in sixteen metropolitan areas from the original ENEU sample. Column 3 reports for employees in ENEU who report receiving IMSS benefit in current employment; column 4 for those who do not; column 5 reports for those with IMSS benefit and who have a written contract of indefinite duration; and column 6 for those with IMSS benefits who worked at least 35 hours in previous week. Standard errors of means in parentheses. Sampling weights from ENEU used in columns 2 to 6. For further details of data processing, see section IV and appendix C.

101 to 250, or 250 or more employees.²⁶ In addition, we drop workers with reported daily wages below 30 pesos (in 2002 constant pesos, approximately US\$3, about 50% of the lowest legal minimum wage). In principle, both the IMSS and the ENEU data are available over the 1987–2005 period. But there appear to be a number of data inconsistencies in the ENEU in 1987, the first year of the survey. In addition, the ENEU sampling scheme was redesigned in the third quarter of 2003. We therefore take as our study period 1988 to 2003, quarter 2.

Our goal in the preparation of the data sets is to construct samples in the IMSS and ENEU data that are as similar as possible. Table 2 presents summary statistics for the IMSS baseline sample and various ENEU samples for 1990 and 2000, for a set of variables that are common between the sources: daily (posttax) wage, age, and share in large establishments (with more than 100 employees).²⁷ Column 2 presents the full ENEU sample, containing all non-self-employed men satisfying the age and sector criteria. Comparing columns 3 and 4, we see that ENEU workers with IMSS coverage tend to be higher wage and more likely to work in large establishments than workers without IMSS coverage. Column 5 contains the sample that in principle should be the best match for the IMSS baseline sample: ENEU workers who report receiving IMSS coverage and having a written contract of indefinite duration (i.e., who are permanent employees). The average wage for this ENEU sample is greater than for

the IMSS baseline sample, consistent with our argument below that there is underreporting of wages in the IMSS data. Because the contract-type variable is available only beginning in 1994, however, we have prohibitively few years of prereform data for this sample. Instead, we will focus henceforth on the column 6 sample, ENEU workers who report receiving IMSS coverage and working full time (at least 35 hours in the previous week), which can be defined consistently over the entire period. We refer to the column 6 sample as our ENEU baseline sample.

We also construct a small, separate sample of ENEU male respondents in the social security sector itself who report being covered by the IMSS system, following the same selection criteria as for the ENEU baseline sample, to compare to IMSS records for the same sector.

It is important to note that there are a number of reasons why the IMSS and ENEU baseline samples may differ. Some temporary workers may work full time, and some permanent workers may work part time. It may also be that firms interpret “permanent” to mean something different from the legal definition (a written contract of indefinite duration) when reporting wages. In addition, patterns of nonresponse may differ between the IMSS and ENEU samples. It is well known, for instance, that richer households tend to be less likely to respond to income questions in household surveys (Groves & Couper, 1998). The weighted employment totals from the ENEU data in columns 5 and 6 of table 2 are below the IMSS totals in column 1; this may in part reflect such nonresponse.²⁸

²⁶The survey allows for eight responses, 1, 2 to 5, 6 to 10, 11 to 15, 16 to 50, 51 to 100, 101 to 250, and 250 or more. To ensure that sample sizes within cells are sufficiently large, we create the five categories listed above.

²⁷We focus on decennial population census years because (in unreported results) we have been able to validate the ENEU sample against the population censuses in those years.

²⁸Note, however, that nonresponse by richer households will tend to lead us to understate evasion, making it more difficult for us to pick up statistically significant differences in cross-section.

To further explore the employment discrepancy, appendix figure A4 plots employment totals over the 1988–2003 period for the same samples as in table 2. We see that over most of the period, the number of workers in the IMSS sample is slightly greater than the numbers in any of the ENEU samples. In addition to nonresponse in the ENEU, this difference likely reflects that fact that the IMSS sample is based on place of work, while the ENEU sample is based on place of residence; hence, people who commute into metropolitan areas are included in the IMSS data but not in the ENEU. Another possibility is that some respondents are unaware that they receive IMSS coverage from their employer, or believe that they are covered by the public sector social security agency (known by the acronym ISSSTE) when in fact they are covered by IMSS. For our purposes, however, the most important lesson of the figure is that there does not appear to have been a large change over time in the extent of the employment discrepancy between the IMSS and ENEU samples in response to the pension reform. Nor does it appear that there was a significant large inflow to (or outflow from) formal employment in response to the pension reform. Appendix C.3 presents further comparisons of the distributions of age and firm size in the IMSS and ENEU baseline samples and finds that the distributions are similar.

Overall, although the IMSS and ENEU samples are not identical, and caution is warranted in interpreting cross-sectional differences between them, the samples appear to be sufficiently similar that it is reasonable to use the wage discrepancy between them as a measure of evasion. We also note that any differences between the samples that are constant over time will be differenced out in our difference-in-differences analysis.

V. Cross-Sectional Patterns of Compliance


In this section, we consider cross-sectional differences in wage distributions between the IMSS and ENEU baseline samples in the prereform period, focusing on 1990. We begin with simple histograms to illustrate the main patterns, then explain the construction of our evasion measures and present simple cross-sectional regressions.

A. Simple Histograms

Figure 2 plots simple histograms of pretax daily wages from the IMSS data (gray bars) and daily take-home wages from the ENEU data (bars with black borders and no fill color), using bins 5 pesos wide. The three vertical lines at left indicate the three regional minimum wages in Mexico, with the right-most corresponding to the minimum wage in Mexico City.²⁹ For visual clarity, figure 3 plots similar his-


²⁹It is well known that the real value of the minimum wage eroded in the 1980s to the point where it was only binding in particularly low-wage rural areas (Bell, 1997). We see in the ENEU histogram that almost no workers earn at or below the minimum wage in the urban areas covered by the ENEU sample.

FIGURE 2.—WAGE HISTOGRAMS, 1990



Samples are IMSS and ENEU baseline samples of men, from the second quarter of 1990. IMSS wage is the real daily pretax reported wage from the IMSS administrative records. ENEU wage is the real daily take-home wage. Wages in 2002 pesos. Average 2002 exchange rate: 9.66 pesos/dollar. Vertical lines indicate minimum wages in the three minimum-wage zones in Mexico (A, B, C). Bins are 5 pesos wide. The right-most bin captures all individuals with reported wages at or above the minimum IMSS topcode over the study period (from 1991). ENEU histogram uses ENEU sampling weights. See section IV and appendix C for further details of data processing.

FIGURE 3.—WAGE HISTOGRAMS, 1990, LOW WAGE LEVELS



Histogram is similar to figure 2 but only includes male workers with wages less than 200 pesos/day (approximately \$20/day) in constant 2002 pesos. Bins are 2 pesos wide.

tograms using the same samples but using only observations below 200 pesos (approximately US\$20), with bins 2 pesos wide. The pattern is clear: the IMSS distribution lies largely to the left of the ENEU distribution,³⁰ and there is bunching in the IMSS sample slightly above the three minimum wages. These bunches correspond to 104.5% of the minimum wages in each zone, the minimum reports to IMSS that did not incur penalties. Note that the IMSS distribution would be even further to the left if we plotted the posttax IMSS wage (which

³⁰The exception to this generalization is at the far-right tail. In figure 2, we see that there is relatively more weight at the topcode in the IMSS sample; there is also slightly more weight at high-wage values just below the topcode. This appears to reflect nonresponse by high-income households in the ENEU, a common pattern in household surveys, as mentioned above.

is more directly comparable to the ENEU wage). The bunching and shift to the left of the IMSS distribution is precisely what one would have expected, given that for most workers, social security benefits were insensitive to reported wages as long as their firms made the minimum contributions on their behalf.

To check that the discrepancy between the IMSS and ENEU distributions is not an artifact of variable definitions or other idiosyncrasies of the data sets, we can compare the two distributions for a set of workers for whom we can be confident there is no underreporting: employees in the social security sector itself. There are several thousand employees in the IMSS data in this sector until 1994. (As noted above, most public sector workers in Mexico are covered by a different social security agency, ISSSTE.) In the ENEU data, we include workers who report that they work in the social security sector and are covered by the IMSS system. Appendix figure A5A plots the IMSS pretax wage distribution versus the ENEU take-home wage distribution, similar to figure 2, for male workers in these samples. We do not see stacking in the IMSS distribution at 104.5% of the minimum wages. Appendix figure A5B plots the IMSS posttax wage, which corresponds more closely to the take-home pay, against the ENEU wage distribution. Appendix figure A5C does the same for low-wage levels. Although we have relatively few observations for social security sector workers in the ENEU, approximately 180 men in 1990, it appears that the IMSS posttax and ENEU distributions coincide fairly closely.³¹ The number of employees in the social security sector in the IMSS data dropped sharply in 1995, and we do not have sufficient observations in the ENEU to construct measures of evasion separately by cells, so we are not able to carry out the full analysis in this sector. But the limited evidence in appendix figure A5 is nonetheless reassuring that the discrepancy between the IMSS pretax and ENEU distributions in figures 2 and 3 is not due solely to differences in variable definitions or other particularities of the data sets.

A key empirical implication of our model, as well as of the previous theoretical work by Kleven et al. (2016), is that there is less evasion in larger firms. Figure 4 presents figures similar to figure 3 (focused on daily wages below 200 pesos) separately for five firm sizes. Caution is warranted in interpreting these figures, since observed establishment size in the IMSS data may itself be affected by firms' compliance decisions. Subject to this caveat, it appears that there is less bunching on the minimum reportable wage at larger firm sizes, suggesting greater compliance. Even in establishments with 250 workers or more, however, there is evidence of bunching at the minimum allowable wage report, suggesting some underreporting even in quite large firms. The observed heterogeneity in compliance by firm size is consistent with the view of Hsieh and Klenow (2014) and Hsieh and

Olken (2014) that the payroll tax burden falls more heavily on larger firms, and this may be part of the explanation for the disproportionately large number of small firms in Mexico and other developing countries.

B. Measures of Evasion

To quantify the extent of noncompliance, we construct three measures of evasion. Recall from equation (1) that the unreported wage is the difference between the worker's net wage and the posttax wage reported by the firm: $w_u = w_{net} - (1 - \tau)w_r$. The ENEU survey asks individuals their take-home wage, w_{net} ; the IMSS administrative records contain the reported wage, w_r ; and we know the social security tax scheduled in each year (see appendix A.2). The ENEU data do not contain firm identifiers, but we can construct an estimate of w_u at the level of cells defined by metropolitan area, sector, firm size categories, and/or age groups.

At the cell level, our first measure of evasion is the log median ENEU take-home wage minus the log median IMSS posttax wage. Our second measure is defined analogously, using the mean instead of the median. We refer to these as the "wage gap (medians)" and the "wage gap (means)," respectively.

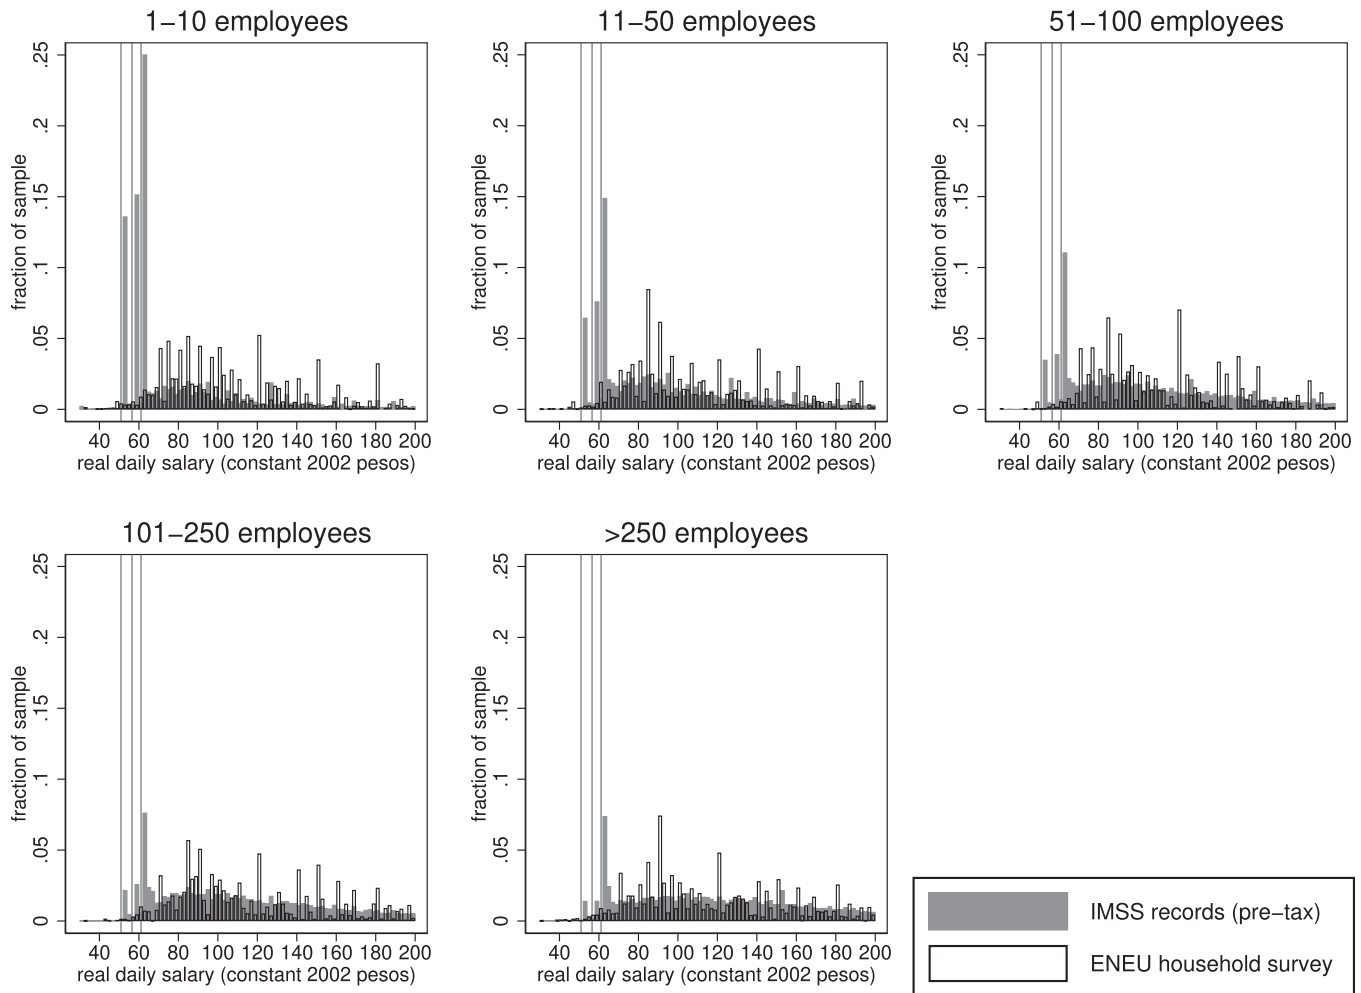
Our third measure of evasion is an estimate of the excess mass at the left tail of the IMSS distribution relative to the ENEU distribution. Appendix figure A6 illustrates the calculation. The dotted (blue) curve is a nonparametric estimate of the ENEU distribution, the one that underlies the hollow-rectangle histogram in figure 2. The solid (red) curve is a nonparametric estimate of the posttax IMSS distribution. We calculate the excess mass as the fraction of the IMSS sample minus the fraction of the ENEU sample to the left of the vertical line (the 15th percentile of the ENEU distribution).³² Intuitively, our excess mass measure reflects the share of the ENEU sample that would have to be moved from right to left across the vertical line in order to transform the dotted (blue) ENEU distribution into the solid (red) IMSS distribution. We emphasize that we see this excess mass measure as simply a complement to the wage gaps measures that is more sensitive to divergence in the left tail (as opposed central tendencies).

The level of aggregation is an important issue when constructing these evasion measures. Although sample size is not a severe constraint in the IMSS administrative records, the ENEU contains on the order of 10,000 to 14,000 raw observations on male full-time workers in each quarter in the country as a whole. When we divide these by age group, metropolitan area, firm size category, and sector, cell sizes in the ENEU can become prohibitively small. We cannot avoid doing some aggregation. As discussed above, we focus on

³¹The apparent lumpiness in the ENEU histogram is due to the facts that there are relatively few observations in the ENEU and that we use the reported ENEU sampling weights when constructing the histogram.

³²In choosing the critical value for the excess mass calculation, we face a trade-off. On the one hand, we want a value that is clearly to the right of the region of bunching in the IMSS data. On the other hand, we do not want a value so far to the right that it misses underreporting behavior. The results using other higher percentiles than the 15th are qualitatively similar, although slightly weaker in some years.

FIGURE 4.—WAGE HISTOGRAMS BY FIRM SIZE, 1990, LOW WAGE LEVELS



Histograms are similar to those in figure 3 but for different firm size categories. Vertical lines indicate minimum wages in the three minimum-wage zones in Mexico (A, B, C). Bins are 2 pesos wide. Average 1990 exchange rate: 9.66 pesos/dollar. See section IV and appendix C for further details of data processing.

five age categories and five firm size categories. We also aggregate four-digit industries into three broad sectors: manufacturing, construction, and retail/services. In addition, when constructing the evasion measures, we pool all four quarters within a given year in the ENEU data. In the next section, we present cross-sectional statistics at the metro area/sector/firm size category/age group level. In section VI, we move to the metro area/age group level, to increase precision of our evasion measures.

C. Cross-Sectional Regressions

Table 3 reports simple cross-sectional regressions of our three evasion measures on age group, firm size, and sector indicators in 1990. For each evasion measure, we report simple regressions on a set of age-group or firm size indicators without controls (columns 1–2, 4–5, and 7–8) and then a regression including sector indicators and metro area indicators

(columns 3, 6, and 9). The number of observations, 1,062, reflects the number of metro area/age group/firm size/sector cells for which we could construct the evasion measures in 1990. We weight cells in the regression by employment in the IMSS administrative records. Turning to the estimates, the level of evasion is highest for the youngest age group, ages 16 to 25 (the omitted category). The differences in coefficients among the over-25 age groups are generally not significantly different from one another. Evasion tends to decline in firm size, consistent with the pattern in the raw histograms in figure 4. There appears to be some nonmonotonicity for the intermediate size categories (51–100 and 101–250 employees), but evasion appears to be robustly lower in 11–50 employee firms than in 1–10 employee firms (the omitted category), and lower still in 250+ employee firms. The estimates are largely unaffected by controlling for age group, metro area, and sector, which suggests that the pattern we observed in the raw data in figure 4 is not due to differing age or metro area composition in different firm size categories. Finally, evasion

TABLE 3.—CROSS-SECTIONAL PATTERNS OF EVASION, 1990

	Wage Gap (medians)			Wage Gap (means)			Excess Mass (15th percentile)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Age 26–35	−0.131** (0.059)		−0.113*** (0.019)	−0.142*** (0.041)		−0.127*** (0.014)	−0.213*** (0.048)		−0.200*** (0.015)
Age 36–45	−0.164** (0.075)		−0.150*** (0.027)	−0.181*** (0.047)		−0.169*** (0.019)	−0.252*** (0.052)		−0.241*** (0.016)
Age 46–55	−0.166** (0.083)		−0.177*** (0.033)	−0.220*** (0.055)		−0.223*** (0.027)	−0.238*** (0.052)		−0.244*** (0.017)
Age 56–65	−0.176* (0.094)		−0.208*** (0.046)	−0.224*** (0.050)		−0.240*** (0.025)	−0.201*** (0.053)		−0.224*** (0.021)
11–50 employees		−0.307*** (0.053)	−0.315*** (0.032)		−0.121*** (0.042)	−0.138*** (0.025)		−0.135*** (0.030)	−0.146*** (0.016)
51–100 employees		−0.420*** (0.050)	−0.426*** (0.035)		−0.203*** (0.044)	−0.226*** (0.028)		−0.216*** (0.036)	−0.231*** (0.019)
101–250 employees		−0.440*** (0.053)	−0.447*** (0.038)		−0.248*** (0.042)	−0.280*** (0.027)		−0.258*** (0.039)	−0.277*** (0.020)
>250 employees		−0.563*** (0.055)	−0.582*** (0.034)		−0.294*** (0.046)	−0.337*** (0.025)		−0.348*** (0.044)	−0.385*** (0.019)
Construction			0.171*** (0.033)			0.095*** (0.035)			0.074*** (0.016)
Retail/services			−0.063** (0.025)			−0.104*** (0.016)			−0.044*** (0.012)
Constant	0.445*** (0.040)	0.741*** (0.041)	0.737*** (0.033)	0.427*** (0.024)	0.514*** (0.033)	0.582*** (0.026)	0.466*** (0.030)	0.542*** (0.018)	0.655*** (0.022)
Metro area effects	No	No	Yes	No	No	Yes	No	No	Yes
R-squared	0.06	0.37	0.69	0.13	0.20	0.65	0.27	0.33	0.82
N	1,062	1,062	1,062	1,062	1,062	1,062	1,062	1,062	1,062

Data are from IMSS and ENEU baseline samples of men, collapsed to metro area/age group/firm size category/sector level for 1990. Regressions are weighted by IMSS employment in each cell. The omitted category for age is 16 to 25, for firm size is 1 to 10 employees, and for sector is manufacturing. The wage gap (medians) is log median real daily take-home wage from the ENEU minus log median real daily posttax reported wage from IMSS. Wage gap (means) is analogous, using mean in place of median. Excess mass is calculated as described in section V and appendix figure A6. In calculating evasion measures, we pool ENEU data across quarters within year. ***1%, **5%, *10% level. See section IV and appendix C for further details of data processing.

follows a consistent pattern across broad sectors, with construction displaying the greatest extent of evasion, followed by manufacturing, followed by retail/services.

VI. Effect of Pension Reform on Compliance


We now consider how evasion varied over time by age group in response to the pension reform.³³ A simple graph illustrates the main finding. We first calculate the wage gap (medians) measures at the age group/metro area/year level, then regress them on a full set of metro area–year indicators, recover the residuals, and average them at the age group–year level. Figure 5 plots these deviated measures of evasion separately by age group over time. Given that we have netted out any changes in wage gaps common to all age groups within metro-years, the levels of the wage gap in this figure sum to zero in each year. Any changes in the wage gap for a given age group are relative to changes for the other age groups. There are two main points. First, as we saw in table 3, evasion was highest for the youngest age group. Second, and more important, the relative decrease in evasion (the rela-

tive increase in compliance) tended to be greater for younger age groups. The largest decline was for the 26–35 age group, and second largest for the 36–45 age group. The relative levels of evasion increased for the 46–55 and the 56–65 age groups, with a larger relative increase for the latter. For the oldest age group, it appears that the relative increase may have begun between 1995 and 1996, which corresponds to the passage of the pension reform by Congress; there was significant discussion of the pension reform in the popular press at the time, and it is possible that employees started paying attention to underreporting even before the reform took effect in July 1997. However, to be conservative, in our regression analysis below, we take July 1, 1997, as the date of the reform. Even using this date, it appears that there was a differential increase in evasion for the oldest two age groups relative to the younger ones. An important caveat about this figure is that the relative decline in evasion was not larger for the 16–25 age group than for the 26–35 or 36–45 age groups as our model would predict. Although our model does not consider differences of attentiveness to pension incentives, it seems plausible that 16–25-year-olds were less attentive to the pension system than the older age groups and that this in part explains the divergence of the empirical patterns from our theoretical expectation.³⁴

³³A possible alternative strategy would be to compare workers of different wage levels, since, as table 1 indicates, higher-wage workers had a higher expectation of using the personal retirement account and, hence, arguably a greater incentive to improve compliance. But our argument is that wage reports in the IMSS respond endogenously to reporting incentives. It is not clear how to construct comparable samples in the IMSS and ENEU data for workers of different wage levels. For this reason, we focus on a clearly predetermined variable, age.

³⁴Appendix figure A7 provides additional graphical evidence, plotting the full ENEU and posttax IMSS wage distributions for the different age groups in 1990, 1997, and 2003. This figure does not deviate from metro

FIGURE 5.—WAGE GAPS (MEDIANS) BY AGE GROUP, DEVIATED FROM METRO-YEAR MEANS



Wage gaps are log median net wage from the ENEU survey minus the log median posttax reported wage from the IMSS administrative records, using the ENEU and IMSS baseline samples of men (pooling ENEU data across quarters within year), calculated separately by age group-year-metro area. The gaps are then regressed on a full set of metro area-year dummies, weighting by IMSS employment in each cell, and the residuals are averaged at the age-group level. IMSS data are from June 30 of each year; the year tick marks should be interpreted as reflecting the middle of each year. The dashed vertical line indicates the date the pension reform was passed by Congress (December 21, 1995) and the solid vertical line the date the reform took effect (July 1, 1997). See section IV and appendix C for details of sample selection.

Table 4 reports difference-in-differences regressions that absorb the effects of additional covariates, including metro area/age group fixed effects. Motivated by the pension wealth simulation in table 1, we categorize the three younger age groups (16–25, 26–35, 36–45) into a “more treated” group and the two older groups (46–55, 56–65) into a “less treated” or “control” group. We interact an indicator for being in the less-treated group (age 45 or below) with year effects for each year over the 1988–2003 period, omitting the interaction with 1997, the implementation year. We control flexibly for metro area–year effects and age group–metro area effects. The number of observations, 1,280, reflects the number of metro area/age group/year cells ($16 \times 5 \times 16$). We weight each cell by employment in the IMSS administrative records. The key finding is that we see little evidence of a differential pretrend but robust evidence of a relative decrease in evasion for the younger age groups following the passage of the reform. We see robust negative coefficients on the 1(age

area–year means, and especially ENEU distributions (dotted blue lines) shift in part because of macroeconomic changes, including the peso crisis of late 1994 and 1995. (In addition, the ENEU distributions, estimated from modest sample sizes, reflect some heaping at round numbers and multiples of the minimum wage.) With these caveats, the broad pattern that the discrepancy between the distributions declines more for younger workers is again evident.

45 or below)-year interaction in the later years of the study period for the wage gap (medians) and excess mass measures. The estimates for the wage gap (means) measure are somewhat less robust, but the results for the means measure are still largely consistent with the patterns for the other two measures. Overall, we interpret the results as supportive of our second main theoretical implication: evasion for younger age groups declined relatively more than evasion for older groups. Additional results, reported in appendix D.1, suggest that these results are not driven by discrepancies in the reporting of employment in the two data sources.

The difference-in-differences results in this section should be treated with caution for a number of reasons. One is that it is possible that reporting behavior in the ENEU household survey has changed over time and differentially by age. It could be, for instance, that younger workers expecting to receive the personal account (PRA) pension might become more worried about government cross-checks of their survey responses and firms’ reports of their wages over time, and for that reason they reported wages to the ENEU more in line with the IMSS records.³⁵ Such a response would bias our results in the direction of the effect we find. Another caveat, already

³⁵Households’ survey responses are not subject to such cross-checks, but individuals may have believed that they are.

TABLE 4.—DIFFERENTIAL EFFECTS OF PENSION REFORM ON EVASION

	Wage gap (medians) (1)	Wage gap (means) (2)	Excess Mass (15th percentile) (3)
1(age <= 45) × 1988	0.015 (0.033)	0.034 (0.040)	0.011 (0.011)
1(age <= 45) × 1989	0.025 (0.027)	0.036 (0.025)	0.018 (0.016)
1(age <= 45) × 1990	0.033 (0.035)	0.018 (0.031)	0.016 (0.013)
1(age <= 45) × 1991	-0.011 (0.031)	0.027 (0.026)	0.001 (0.012)
1(age <= 45) × 1992	-0.011 (0.028)	-0.015 (0.026)	0.010 (0.012)
1(age <= 45) × 1993	0.027 (0.027)	0.033 (0.023)	0.003 (0.009)
1(age <= 45) × 1994	-0.005 (0.027)	-0.035 (0.026)	0.011 (0.009)
1(age <= 45) × 1995	-0.025 (0.031)	0.002 (0.022)	-0.006 (0.014)
1(age <= 45) × 1996	-0.020 (0.022)	-0.028 (0.030)	-0.007 (0.009)
1(age <= 45) × 1998	0.001 (0.034)	0.019 (0.039)	-0.023** (0.009)
1(age <= 45) × 1999	-0.014 (0.028)	-0.021 (0.026)	-0.023** (0.010)
1(age <= 45) × 2000	-0.062** (0.028)	-0.051** (0.022)	-0.027*** (0.010)
1(age <= 45) × 2001	-0.065** (0.025)	-0.030 (0.024)	-0.023** (0.011)
1(age <= 45) × 2002	-0.073*** (0.026)	-0.081*** (0.022)	-0.023** (0.010)
1(age <= 45) × 2003	-0.087*** (0.025)	-0.046 (0.028)	-0.025** (0.012)
Age group-metro area effects	Yes	Yes	Yes
Metro-year effects	Yes	Yes	Yes
R-squared	0.96	0.95	0.99
N	1,280	1,280	1,280

Data from IMSS and ENEU baseline samples of men, collapsed to metro area/age group/year level. Regressions weighted by IMSS employment in each cell. Dependent variables are as explained in notes to table 3. In calculating evasion measures, we pool ENEU data across quarters within year. Age group-metro area effects are defined for five age groups (16–25, 26–35, 36–45, 46–55, 56–65). ***1%, **5%, *10% level. See section IV and appendix C for further details of data processing.

noted, is that our argument relies on differences in individuals' expectations about which pension they would choose, but we do not directly observe individuals' expectations. Our division of age groups into more- and less-treated groups may not map precisely into the differences in expectations held by participants at the time.

Because of all these caveats, the difference-in-difference results should not be interpreted as definitive. More research is needed, perhaps in more controlled settings, to examine the role of employee incentives in payroll tax compliance by firms. But the results are broadly consistent with the predictions of our parsimonious model of compliance by heterogeneous firms and are suggestive that giving workers incentives and information to ensure accurate reporting by their firms is effective in improving compliance.

VII. Conclusion

Improving firms' compliance with tax regulations is a first-order policy issue in many developing countries. Much of

the debate has focused on how to induce firms to register with tax authorities—what we might call an *extensive* margin of noncompliance. In this paper, we have shown that underreporting of wages for formal workers among registered firms—noncompliance on an *intensive* margin—is also substantial. Having firms report employees' wages is no guarantee of accurate reporting in a low-enforcement context like Mexico. We have also presented evidence that firms' compliance responds to incentives and the availability of information facing workers in the social security system.

The results suggest that providing incentives to workers to ensure accurate reporting, as well as information about firms' reports, should be a consideration in the design of social insurance systems. Conceptually, our theoretical model suggests that an increase in such incentives and a commensurate reduction of payroll taxes should have equivalent effects on evasion, other things equal. But the effects on government revenues are decidedly nonequivalent. If the policy goal is to increase the fiscal capacity of the state, there are reasons to prefer tying benefits more closely to wage reports.

A number of interesting questions remain open. One is to what extent workers are aware of underreporting by their employers and, relatedly, to what extent the effects of the pension reform we observe are due to the change in incentives versus the change in information. Separating the two effects will require a research design in which incentives and information vary separately.

Another important open question is whether increased pressure on firms to report accurately (which increases compliance on the intensive margin) induces more firms to be informal (i.e., reduces compliance on the extensive margin). Because of the nature of the IMSS data, we are not able to distinguish a transition to informality from an establishment death, and we leave this question for a research setting in which firms can be tracked across such transitions. But clearly a full accounting of the costs and benefits of policies to increase intensive-margin compliance will have to take such a response into account.

REFERENCES

- Aguila, Emma, "Personal Retirement Accounts and Saving," *American Economic Journal: Economic Policy* 3:4 (2011), 1–24.
- Andreoni, James, Brian Erard, and Jonathan Feinstein, "Tax Compliance," *Journal of Economic Literature* 36 (1998), 818–860.
- Bailey, Clive, and John Turner, "Strategies to Reduce Contribution Evasion in Social Security Financing," *World Development* 29 (2001), 385–393.
- Bell, Linda, "The Impact of Minimum Wages in Mexico and Colombia," *Journal of Labor Economics* 15 (1997), s102–s123.
- Bérgolo, Marcelo, and Guillermo Cruces, "Work and Tax Evasion Incentive Effects of Social Insurance Programs: Evidence from an Employment-Based Benefit Extension," *Journal of Public Economics* 117 (2014), 211–228.
- Besley, Timothy, and Torsten Persson, "Taxation and Development," in Alan Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez, eds., *Handbook of Public Economics*, vol. 5 (Amsterdam: Elsevier, 2013).
- Best, Michael Carlos, "The Role of Firms in Workers' Earnings Responses to Taxes: Evidence from Pakistan," unpublished paper, Columbia University (2014).

- Best, Michael Carlos, Anne Brockmeyer, Henrik Jacobsen Kleven, Johannes Spinnewijn, and Mazhar Waseem, "Production vs. Revenue Efficiency with Limited Tax Capacity: Theory and Evidence from Pakistan," *Journal of Political Economy* 123 (2015), 1311–1355.
- Braguinsky, Serguey Mityakov, and Andrey Liscovich, "Direct Estimation of Hidden Earnings: Evidence from Administrative Data," *Journal of Law and Economics* 57 (2014), 281–319.
- Bruhn, Miriam, and David McKenzie, "Entry Regulation and the Formalization of Microenterprises in Developing Countries," *World Bank Research Observer* 29 (2014), 186–201.
- Burgess, Robin, and Nicholas Stern, "Taxation and Development," *Journal of Economic Literature* 31 (1993), 762–830.
- Castellanos, Sara, Rodrigo Garcia-Verdu, and David Kaplan, "Wage Rigidities in Mexico: Evidence from Social Security Records," *Journal of Development Economics* 75 (2004), 507–533.
- Dabla-Norris, Era, Mark Gradstein, and Gabriela Inchauste, "What Causes Firms to Hide Output? The Determinants of Informality," *Journal of Development Economics* 85:1–2 (2008), 1–27.
- De Giorgi, Giacomo, Matthew Ploenzke, and Aminur Rahman, "Small Firms' Formalisation: The Stick Treatment," *Journal of Development Studies* 54 (2018), 983–1001.
- de Mel, Suresh, David J. McKenzie, and Christopher Woodruff, "The Demand for, and Consequences of, Formalization among Informal Firms in Sri Lanka," *American Economic Journal: Applied Economics* 5 (2013), 122–150.
- de Paula, Aureo, and José Scheinkman, "The Informal Sector: An Equilibrium Model and Some Empirical Evidence from Brazil," *Review of Income and Wealth* 57 (2011), S8–S26.
- Dharmapala, Dhammika, Joel Slemrod, and John Douglas Wilson, "Tax Policy and the Missing Middle: Optimal Tax Remittance with Firm-Level Administrative Costs," *Journal of Public Economics* 95 (2011), 1036–1047.
- Duarte, Fabian, and Justine Hastings, "Fettered Consumers and Sophisticated Firms: Evidence from Mexico's Privatized Social Security Market," NBER working paper 18582 (2012).
- Feldman, Naomi E., and Joel Slemrod, "Estimating Tax Noncompliance with Evidence from Unaudited Tax Returns," *Economic Journal* 117 (2007), 327–352.
- Fisman, Raymond, and Shang-Jin Wei, "Tax Rates and Tax Evasion: Evidence from Missing Imports in China," *Journal of Political Economy* 112 (2004), 471–500.
- Frías, Judith A., David S. Kaplan, Eric Verhoogen, and David Alfaró-Serrano, "Exports and Wage Premia: Evidence from Mexican Employer-Employee Data," Columbia University unpublished paper (2018).
- Galiani, Sebastian, and Federico Weinschelbaum, "Modeling Informality Formally: Households and Firms," *Economic Inquiry* 50 (2012), 821–838.
- Gordon, Roger, and Wei Li, "Tax Structures in Developing Countries: Many Puzzles and a Possible Explanation," *Journal of Public Economics* 93 (2009), 855–866.
- Gordnichenko, Yuriy, Jorge Martinez-Vazquez, and Klara Sabirianova Peter, "Myth and Reality of Flat Tax Reform: Micro Estimates of Tax Evasion Response and Welfare Effects in Russia," *Journal of Political Economy* 117 (2009), 504–554.
- Grandolini, Gloria, and Luis Cerda, "The 1997 Pension Reform in Mexico: Genesis and Design Features," World Bank Policy Research working paper 1933 (1998).
- Groves, Robert M., and Mick P. Couper, *Nonresponse in Household Interview Surveys* (New York: Wiley, 1998).
- Hsieh, Chang-Tai, and Peter J. Klenow, "The Life Cycle of Plants in India and Mexico," *Quarterly Journal of Economics* 129 (2014), 1035–1084.
- Hsieh, Chang-Tai, and Benjamin A. Olken, "The Missing 'Missing Middle,'" *Journal of Economic Perspectives* 28 (2014), 89–108.
- Hurst, Erik, Geng Li, and Benjamin Pugsley, "Are Household Surveys Like Tax Forms? Evidence from Income Underreporting of the Self-Employed," this REVIEW 96:1 (2014), 19–33.
- International Monetary Fund, *Government Finance Statistics Yearbook* (2010).
- INEGI, "Censo Económico 2009," *Aggregate Results* (2009), <http://www.beta.inegi.org.mx/programas/ce/2009/default.html>.
- Internal Revenue Service, "Updated Estimates of the Tax Year 2001 Individual Income Tax Underreporting Gap. Overview" (Washington, DC, 2006).
- Keen, Michael, and Ben Lockwood, "The Value Added Tax: Its Causes and Consequences," *Journal of Development Economics* 92 (2010), 138–151.
- Khan, Adnan Qadir, Asim Khwaja, and Benjamin Olken, "Tax Farming Redux: Experimental Evidence on Incentive Pay for Tax Collectors," *Quarterly Journal of Economics* 131 (2016), 219–271.
- Kleven, Henrik, Martin Knudsen, Claus Thustrup Kreiner, Soren Pedersen, and Emmanuel Saez, "Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark," *Econometrica* 79 (2011), 651–692.
- Kleven, Henrik Jacobsen, Claus Thustrup Kreiner, and Emmanuel Saez, "Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries," *Economica* 83 (2016), 219–246.
- Kleven, Henrik J., and Mazhar Waseem, "Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan," *Quarterly Journal of Economics* 128 (2013), 669–723.
- Kopczuk, Wojciech, and Joel Slemrod, "Putting Firms into Optimal Tax Theory," *American Economic Review: Papers and Proceedings* 96 (2006), 131–134.
- Kumler, Todd, Eric Verhoogen, and Judith A. Frías, "Enlisting Workers in Monitoring Firms: Payroll Tax Compliance in Mexico," Columbia University Department of Economics discussion paper 1112-07 (June 2012).
- "Enlisting Workers in Improving Payroll-Tax Compliance: Evidence from Mexico," NBER working paper 19385 (2015).
- Mao, Jie, Lei Zhang, and Jing Zhao, "Social Security Taxation and Compliance: The Chinese Evidence," Tsinghua University unpublished paper (2013).
- Marion, Justin, and Erich Muehlegger, "Measuring Illegal Activity and the Effects of Regulatory Innovation: Tax Evasion and the Dyeing of Untaxed Diesel," *Journal of Political Economy* 116 (2008), 633–666.
- McKenzie, David, and Yaye Seynabou Sakho, "Does It Pay Firms to Register for Taxes? The Impact of Formality on Firm Profitability," *Journal of Development Economics* 91 (2010), 15–24.
- Melitz, Mark J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica* 71 (2003), 1695–1725.
- Naritomi, Joana, "Consumers as Tax Auditors," *American Economic Review* 109 (2019), 3031–3072.
- Niehaus, Paul, and Sandip Sukhtankar, "Corruption Dynamics: The Golden Goose Effect," *American Economic Journal: Economic Policy* 5 (2013), 230–269.
- Nyland, Chris, Russell Smyth, and Cherie Jinhua Zhu, "What Determines the Extent to Which Employers Will Comply with their Social Security Obligations? Evidence from Chinese Firm-Level Data," *Social Policy and Administration* 40 (2006), 196–214.
- OECD, *Tax Administration in OECD and Selection Non-OECD Countries: Comparative Information Series* (Paris: OECD, 2006).
- *Revenue Statistics, 1965–2010* (Paris: OECD, 2011).
- Olken, Benjamin A., "Corruption and the Costs of Redistribution: Micro Evidence from Indonesia," *Journal of Public Economics* 90 (2006), 853–870.
- Olken, Benjamin A., and Rohini Pande, "Corruption in Developing Countries," *Annual Review of Economics* 4 (2012), 479–509.
- Paulus, Alari, "Income Underreporting Based on Income Expenditure Gaps: Survey vs. Tax Records," Institute for Social and Economic Research working paper 2015-15 (2015).
- Pissarides, Christopher A., and Guglielmo Weber, "An Expenditure-Based Estimate of Britain's Black Economy," *Journal of Public Economics* 39:1 (1989), 17–32.
- Plumley, Alan, *Overview of the Federal Tax Gap* (Washington, DC: U.S. Department of the Treasury, Internal Revenue Service, NHQ Office of Research, 2004).
- Pomeranz, Dina, "No Taxation without Information: Deterrence and Self-Enforcement in the Value-Added Tax," *American Economic Review* 105 (2015), 2539–2569.
- Rauch, James E., "Modelling the Informal Sector Formally," *Journal of Development Economics* 35 (1991), 33–47.

- Rocha, Rudi, Gabriel Ulyssea, and Lasa Rachter, "Do Lower Taxes Reduce Informality? Evidence from Brazil," *Journal of Development Economics* 134 (2018), 2–49.
- Saez, Emmanuel, "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy* 2 (2010), 180–212.
- Saez, Emmanuel, Joel B. Slemrod, and Seth H. Giertz, "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review," *Journal of Economic Literature* 50 (2012), 3–50.
- Sales-Sarrapy, Carlos, Fernando Solis-Soberon, and Alejandro Villagomez-Amezcuca, "Pension System Reform: The Mexican Case," NBER working paper 5780 (1996).
- Schneider, Friedrich, and Dominik H. Enste, "Shadow Economies: Size, Causes, and Consequences," *Journal of Economic Literature* 38 (2000), 77–114.
- Slemrod, Joel, "Cheating Ourselves: The Economics of Tax Evasion," *Journal of Economic Perspectives* 21 (2007), 25–48.
- , "Tax Compliance and Enforcement," *Journal of Economic Literature* 57 (2019), 904–954.
- Slemrod, Joel, and S. Yitzhaki, "Tax Avoidance, Evasion and Administration," in Alan J. Auerbach and Martin Feldstein, eds., *Handbook of Public Economics*, vol. 3 (Amsterdam: Elsevier, 2002).
- Tonin, Mirco, "Minimum Wage and Tax Evasion: Theory and Evidence," *Journal of Public Economics* 95 (2011), 1635–1651.
- Ulyssea, Gabriel, "Firms, Informality, and Development: Theory and Evidence from Brazil," *American Economic Review* 108 (2018), 2015–2047.
- Yaniv, Gideon, "Collaborated Employee-Employer Tax Evasion," *Public Finance* 47 (1992), 312–321.