

# Redistribution with Performance Pay

Paweł Doligalski

Abdoulaye Ndiaye

Nicolas Werquin \*

April 22, 2022

## Abstract

Half of the jobs in the U.S. feature pay-for-performance. We derive novel incidence and optimum formulas for the overall rate of tax progressivity and the top tax rates on total earnings and bonuses, when such labor contracts arise from moral hazard frictions within firms. Optimal taxes account for the fiscal externalities and welfare consequences of two distinct forces: a direct crowding-out of private insurance and a countervailing crowding-in due to endogenous labor effort responses. These imply that the amount of pre-tax earnings risk to which the worker is exposed is roughly invariant to tax progressivity, whereas the (adverse) welfare consequences of the crowd-out outweigh those of the crowd-in. Quantitatively, the optimal tax policy with performance-pay contracts is close to that prescribed by standard models that treat pre-tax earnings risk as exogenous. Finally, we uncover an efficiency-based argument for taxing bonuses at strictly lower rates than base earnings.

---

\*Doligalski: University of Bristol. Ndiaye: New York University, Stern School of Business and CEPR. Werquin: Federal Reserve Bank of Chicago, Toulouse School of Economics, and CEPR. Opinions expressed in this article are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. Nicolas Werquin acknowledges support from ANR under grant ANR-17-EURE-0010 (Investissements d’Avenir program). We thank Árpád Ábrahám, Joydeep Bhattacharya, Katherine Carey, Ashley Craig, Alex Edmans, Emmanuel Farhi, Antoine Ferey, Xavier Gabaix, Daniel Garrett, Mike Golosov, Andreas Haufler, Louis Kaplow, Paul Kindsgrab, Helen Koshy, Etienne Lehmann, Iacopo Morchio, Sergio Feijoo Moreira, Yena Park, Laura Pilossoph, Morten Ravn, Kjetil Storesletten, Florian Scheuer, Karl Schulz, Ali Shourideh, Stefanie Stantcheva, Aleh Tsyvinski, H  l  ne Turon, Venky Venkateswaran, Philipp Wangner, and Larry White, for helpful comments and suggestions. Emmanuel Farhi passed away a few weeks after we completed the first version of this paper. While the feedback he gave us about this paper in particular was very helpful, our whole generation of public finance economists owes him a much larger debt for his fundamental contributions to the field, his kindness and generosity, and his unwavering support. He continues to be a source of inspiration. This paper is dedicated to his memory.

# Introduction

The dramatic increase in income inequality observed since the 1980s is in large part due to the explosion of performance-based forms of remuneration at the top of the income distribution, such as the rise in bankers' bonuses or CEOs' stock options (see, e.g., [Bell and Van Reenen 2013, 2014](#); [Lemieux, MacLeod, and Parent 2009](#); [Piketty and Saez 2003](#)). While performance-pay contracts are particularly prevalent for high earners, they are common throughout the income distribution and across occupations, from agricultural workers paid a piece rate to real estate brokers or retail workers who earn a commission on their sales ([Bandiera, Barankay, and Rasul 2005](#); [Levitt and Syverson 2008](#); [Shearer 2004](#)). Performance incentives can also be provided over time via promotions or salary raises.<sup>1</sup> [Lemieux et al. \(2009\)](#) estimate that in the U.S., almost half of all jobs, and three quarters within the top percentile of the income distribution, involve performance-based compensation. In the UK, [Bell and Van Reenen \(2014\)](#) find that bonuses are a feature of more than 80 percent of the contracts of the top 1 percent earners and can account for up to 35 percent of their compensation.

Given the importance of performance-based earnings and their contribution to growing income inequality, it is natural to ask how governments should tax them. What is the incidence of tax policy on performance-based contracts? Should the overall progressivity of the income tax schedule be modified to account for the existence of these jobs? How should income taxes on top earners be designed when their earnings are partly composed of bonuses? Should bonuses be treated separately from base earnings by the tax system, and if so, at what rates should they be taxed?

This paper provides answers to these questions. To do so, we set up a model of moral hazard in the labor market that gives rise to performance pay in equilibrium. In our model, income disparities arise from two distinct sources: namely, ex-ante ability differences and ex-post performance (or output) shocks. While the former cannot be insured by private firms, the impact of the latter on inequality is very much shaped in the labor market. In particular, if firms can observe workers' ability (say, their education level) and their realized output but not their labor effort, they design a contract that provides only partial insurance against output risk in order to give

---

<sup>1</sup>[Abrahám, Alvarez-Parra, and Forstner \(2016\)](#) study the implications of dynamic performance-pay contracts for wage inequality and wage dynamics.

incentives to workers to exert the optimal labor effort.<sup>2</sup> Importantly, the amount of earnings risk—the size of the bonus, or the reliance on stock options—is endogenous to tax policy. In our model, the government uses income taxes to redistribute between workers with different ex-ante abilities, taking into account their effect on insurance within firms.

We first take a positive standpoint and evaluate the incidence of nonlinear taxes on the design of performance pay. In our baseline model, based on Rogerson (1985), the output of each worker is stochastic and binary: either high or low. Workers can increase the probability of high output by exerting more (unobservable and costly) effort. Firms offer a bonus when realized output is high in order to motivate workers to exert the right amount of labor effort. Since workers are risk-averse, the optimal contract strikes a balance between such incentives for effort and insurance against output risk—providing better insurance is a less costly way for the firm to give workers their reservation utility. Workers’ earnings, including the bonus, are subject to an income tax. Both the size of the bonus and the frequency of receiving a bonus are endogenous to the tax system.

We first show in a parametric version of the model that an increase in the rate of tax progressivity causes two offsetting effects on the labor contract. First, it leads to a standard crowding-out of the private insurance provided by the firm, that is, a one-for-one spread of the pre-tax earnings distribution. In response to a change in social insurance, firms adjust earnings endogenously so that the workers’ incentives for effort and participation remain unchanged despite the reform.<sup>3</sup> Second, it also generates a crowding-in of the earnings distribution. This is due to the fact that, as in standard models of taxation—e.g., Mirrlees (1971)—higher marginal tax rates reduce optimal labor effort. In our setting, this is because higher tax rates on high earnings make high-powered incentives more costly to provide for the firm.<sup>4</sup> In turn, to elicit

---

<sup>2</sup>There is strong field-experimental evidence of moral hazard frictions in the workplace (see, e.g., the review by Lazear 2018). Performance pay can also be micro-founded by models of adverse selection, in which incentive pay is offered to attract workers with higher unobserved ability. Lazear (2000), Leaver, Ozier, Serneels, and Zeitlin (2021) and Brown and Andrabi (2020) find empirical support for both moral hazard and adverse selection, the former accounting for at least a half of the overall effect of performance pay on productivity.

<sup>3</sup>Theoretically, the crowd-out of private insurance has been shown to severely limit the ability of governments to provide social insurance—see Attanasio and Rios-Rull (2000), Golosov and Tsyvinski (2007) and Krueger and Perri (2011). Empirically, crowd-out has been observed in health insurance (Cutler and Gruber 1996a, 1996b; Schoeni 2002) and unemployment insurance (Cullen and Gruber 2000).

<sup>4</sup>This argument echoes the wage-cum-labor demand effect of Lehmann, Parmentier, and Van der

this lower effort level in the presence of moral hazard, the firm lowers the sensitivity of pre-tax earnings to performance—that is, it compresses the earnings distribution. This crowd-in mechanism counteracts the usual crowd-out.

Our first main insight is that, regardless of the size of the labor effort elasticity, these two effects are of the same order of magnitude. Taken separately these effects are both significant, but summing them implies that taxes barely affect the sensitivity of (pre-tax) pay to performance. This finding may help explain why empirical studies of the impact of income taxes on the structure of performance-pay contracts often fail to find significant crowding-out; see [Rose and Wolfram \(2002\)](#) and [Frydman and Molloy \(2011\)](#). This insight remains robust to alternative forms of performance pay: We derive the same result when optimal contracts are a linear function of performance such as piece rates or commissions ([Holmstrom and Milgrom 1987](#)), convex such as stock options ([Edmans and Gabaix 2011](#)), or when incentives are provided over time ([Edmans, Gabaix, Sadzik, and Sannikov 2012](#); [Sannikov 2008](#)). The general principle that makes our result robust to different forms of performance pay is that the local incentive constraints in all of these models of moral hazard share a common structure: namely, they imply that the sensitivity of earnings to performance is proportional to the marginal disutility of effort. As a result, in order to elicit a given increase in labor effort, the firm must raise the pass-through of output risk proportionately to the *inverse* of the (Frisch) labor effort elasticity. Therefore, the crowding-in of private insurance, equal to the product of the elasticity of labor effort and the change in performance-sensitivity required to elicit the desired effort adjustment—the inverse Frisch elasticity—is large and robust to the size of the labor effort elasticity.

The second part of our theoretical analysis takes a normative standpoint and characterizes optimal taxes in both our parametric and general environments. Even though taxes have little effect on the contract, because of offsetting crowd-in and crowd-out effects, the endogeneity of private insurance can potentially have large effects on welfare. Indeed, recall that the crowd-in channel operates via optimal labor effort decisions. By an envelope theorem, the resulting effect on the worker’s utility is at most second-order. By contrast, the crowding-out of private insurance has first-order welfare consequences. Consequently, although the amount of within-firm insurance appears insensitive to tax policy, the worker’s utility reacts *as if* private

---

[Linden \(2011\)](#), whereby tax progressivity tends to make high pre-tax wages less attractive (which, in their setting, leads to lower unemployment).

insurance was fully crowded out.

We derive three sets of results. First, we characterize the optimal rate of progressivity of the overall tax schedule. The endogeneity of performance-pay contracts affects optimal redistribution mainly via the welfare consequences of crowd-out. Greater progressivity triggers a crowd-out of within-firm insurance which offsets any gains from social insurance. This eliminates the role of the income tax as an insurance device. In addition, the crowding-out generates an additional adverse effect on welfare by making tax cuts less accurately targeted towards high-marginal-utility agents than in standard models. To see this, suppose that the tax burden decreases uniformly by a constant amount. The consumption gain, however, will not be uniform across income levels, as firms need to adjust gross earnings to preserve effort incentives. We show that the effective consumption gains are distributed regressively: High-performers of a given ability type—whose marginal utility of consumption is relatively low—receive a larger share of the tax cut than do low-performers. This dampens the welfare benefits of redistribution via progressive taxation. These two negative welfare effects of tax progressivity due to the crowd-out of private insurance are accompanied by new fiscal externalities: To the extent that the crowd-out and crowd-in do not offset each other perfectly, mean-preserving changes in earnings risk affect tax revenue when the income tax is non-linear. Accounting for all of these effects, we show analytically that the optimal tax progressivity in the presence of performance pay is strictly lower than in the model with exogenous earnings risk.

Second, we characterize the optimal tax rate on top income earners in the general model with arbitrarily non-linear taxes. We start with the case where tax rates apply to total (fixed plus variable) earnings. Our formula generalizes that of [Saez \(2001\)](#) by accounting for the fiscal and welfare effects of crowd-out and crowd-in, in addition to standard Pareto coefficients and labor effort elasticities. In particular, a higher top tax rate creates a fiscal spillover typically ignored by standard models: It affects the base earnings of agents who would have been in the top tax bracket had they earned a bonus. In contrast to [Chetty and Saez \(2010\)](#), whose approach is based purely on reduced-form sufficient statistics, we provide analytical expressions for all of the relevant parameters in our formulas—including the crowd-in and crowd-out of private insurance and the social marginal welfare weights.

Third, we ask whether there would be gains from taxing bonuses and base earnings separately. We first show that it is optimal to tax bonuses strictly less than base pay.

In particular, starting from a joint tax on total earnings, we construct a reform that (i) lowers the tax rate on bonuses and raises the tax rate on base pay, (ii) does not affect the expected utility of any agent, and (iii) generates more tax revenue—a Pareto improvement—by fostering labor effort. To understand this result, recall that the bonus is a constrained efficient instrument used by firms to overcome agency frictions in the labor market. Taxing bonuses erodes the firm’s ability to provide incentives to workers and is thus highly distortionary. Consequently, it is optimal to shift the tax burden from bonuses to base pay. Finally, we establish an analytical formula for the optimal top tax rate on bonuses, for a given tax schedule on base earnings. As in the case of the joint taxation of overall earnings, the optimal bonus tax rate depends on the fiscal externalities and welfare effects of the crowd-out and crowd-in.

To complement this theoretical analysis, we calibrate the model to U.S. data. Our model features both performance-pay and fixed-pay jobs and accurately matches the earnings distribution, including the incidence of performance-pay jobs by income quartiles. We first document that, in line with our theoretical findings, changes in tax progressivity have very little impact on the pay-performance sensitivity of labor contracts. Quantitatively, the crowd-in effect offsets more than 90 percent of the crowd-out in terms of the bonus-to-base-pay ratio, and more than 100 percent in terms of the variance of log earnings. Thus, while higher progressivity does lead to lower labor effort and lower mean earnings—with implied elasticities that match empirical evidence—it leaves earnings risk largely unaffected.

We then proceed to computing the optimal tax policy. Our robust finding is that ignoring the endogeneity of earnings risk due to performance pay when setting tax policy has a small welfare cost. In other words, applying standard tax formulas allows the policymaker to reap the vast majority of the welfare gains from the optimal tax reform. This result first holds when we optimize over the overall rate of tax progressivity, in which case the welfare loss from applying the standard—rather than optimal—tax formula is equivalent to a modest 0.29 percent fall in consumption. The result also holds when we focus on the top tax rates: Those implied by the benchmark formula of [Saez \(2001\)](#) fall within a 1 percentage point margin of the optimal ones.

Intuitively, since earnings risk is virtually insensitive to tax policy, the policymaker is not making a large mistake when calculating the fiscal effects of tax reforms using the standard formulas. In addition, the welfare effects of crowd-out, which are not accounted for by the standard formulas, are of limited magnitude. In the case of

reforming the overall tax progressivity, this is because only half of the jobs in the U.S. feature performance pay—if all workers had performance-pay contracts, the welfare cost of ignoring endogenous earnings risk would increase more than fourfold. The welfare cost of crowd-out is also small in the case of the top tax rate, since top earners are typically associated with small marginal welfare weights in the social objective.

Finally, we calculate the optimal top bonus tax rate for a given tax schedule on base pay. Consistent with our theoretical findings, the optimal bonus tax is lower than the tax on base pay, with a difference of 20 percentage points or more for a utilitarian or Rawlsian government.

**Literature Review.** Several papers study optimal taxation with labor markets constrained by agency frictions. [Golosov and Tsyvinski \(2007\)](#) analyze a model in which firms employ ex ante identical workers who are subject to private productivity shocks and can engage in hidden asset trades. They show that tax reforms generate a large crowd-out of within-firm insurance, which in turn reduces the benefits of social insurance. The government optimally refrains from providing any insurance and, instead, uses a capital tax to correct the externality generated by hidden trades. [Stantcheva \(2014\)](#) focuses on the adverse selection model of the labor market of [Miyazaki \(1977\)](#) and shows that the upward hours distortions (rat race) due to employer screening allow the government to redistribute more than in the standard Mirrlees model. [Scheuer \(2013\)](#) studies the role of profit taxes in correcting the choice between payroll employment and entrepreneurship, which is distorted by adverse selection in the credit market. In contrast to these papers, we study labor markets that are constrained by moral hazard frictions.

[Ferey, Haufler, and Perroni \(2022\)](#) also focus on moral hazard and show that globalization in the presence of performance pay can reconcile the joint evolution of rising earnings inequality and falling optimal tax progressivity. The main difference between our framework and theirs lies in the labor effort adjustment margin: While we allow for a continuous effort choice, [Ferey et al. \(2022\)](#) assume that effort is binary, which limits individual effort responses and the crowd-in effect within firms that is at the core of our analysis.<sup>5</sup> [Chetty and Saez \(2010\)](#) derive a reduced-form sufficient

---

<sup>5</sup>In their framework, tax reforms still have a distortionary impact on effort in the aggregate by affecting the measure of firms that choose to incentivize effort via performance-pay contracts—an extensive margin response. By contrast, in our framework, taxes distort individual labor effort and within-firm insurance in already existing performance-pay jobs—an intensive margin response.

statistics formula for the optimal linear tax in the presence of linear private insurance contracts that can be subject to agency frictions. By contrast, we provide an explicit and tractable structural microfoundation for the equilibrium labor contracts. This allows us to characterize analytically the effects of nonlinear government policy on private insurance contracts via crowd-out and crowd-in responses, and derive explicit theoretical formulas for the tax incidence and optimal taxes in terms of underlying structural parameters. [Kaplow \(1991\)](#) studies optimal linear government policy in a different (disaster relief) setting that has endogenous private insurance constrained by moral hazard. He finds that the optimal level of social insurance is zero. This result holds in our setting as well, but in our model workers are also ex-ante heterogeneous, thus giving the government a redistributive role.

We contribute to the growing literature on the optimal taxation and regulation of bonuses. [Besley and Ghatak \(2013\)](#) study the role of bonus taxes in correcting the excessive risk taking in the financial sector, while [Thanassoulis \(2012\)](#) and [Bénabou and Tirole \(2016\)](#) show that bonus caps can restrict inefficiencies stemming from the competition for talent. In contrast to these papers, we focus on the basic model of moral hazard in which performance pay is constrained efficient, and uncover a force towards a *lower* tax rate on bonuses than on base earnings: This policy alleviates the labor effort distortions implied by income redistribution. Two related papers, [Gietl and Haufler \(2018\)](#) and [Haufler and Nishimura \(2022\)](#), study bonus tax competition in settings where top managers are internationally mobile, a feature that we ignore in this paper.

An important strand of papers studies income taxation in the presence of endogenous consumption insurance, which can take form of private insurance markets ([Cremmer and Pestieau 1996](#); [Netzer and Scheuer 2007](#)), asset trades ([Ábrahám, Koehne, and Pavoni 2016](#); [Chang and Park 2017](#); [Park 2014](#)) or informal exchanges in family networks ([Attanasio and Rios-Rull 2000](#); [Heathcote, Storesletten, and Violante 2017](#); [Krueger and Perri 2011](#); [Raj 2019](#)). In contrast to these papers, it is *pre-tax earnings* risk—rather than *consumption* risk—that is endogenous to policy in our model. This distinction matters since changes in earnings risk have a direct impact on tax revenue when the income tax is non-linear (as a result of Jensen’s inequality), while changes in consumption risk with exogenous wages do not have a direct fiscal impact.<sup>6</sup>

---

<sup>6</sup>Naturally, there can be indirect effects from consumption insurance to tax revenue through precautionary labor supply, as in [Netzer and Scheuer \(2007\)](#).



Other papers endogenize earnings risk by focusing on human capital accumulation (Craig 2019; Findeisen and Sachs 2016; Kapička and Neira 2019; Makris and Pavan 2021; Stantcheva 2017), job search (Sleet and Yazici 2017), or wage randomization in response to excessive tax regressivity (Doligalski 2019).

Finally, our paper relates to the literature on redistributive taxation in environments with earnings uncertainty and moral hazard; see, e.g., Varian (1980) and, for models that also allow for ability differences, Eaton and Rosen (1980), Shourideh (2014), and Boadway and Sato (2015) (who also offer a more complete synthesis of the literature). In these papers, however, there is no layer of endogenous private insurance between workers and the government: Earnings risk is thus exogenous, and the government is the sole provider of insurance.

**Outline of the Paper.** The paper is organized as follows. In Section 1, we study a simple version of our model and characterize the optimal rate of tax progressivity. In Section 2, we study our general environment and characterize the optimal top tax rates on total earnings and on bonuses when they are taxed separately. The proofs of our results and various extensions are in the Appendix. In particular, we study alternative forms of performance pay in Appendix C.

## 1 Performance Pay and Tax Progressivity

In this section, we set up a simple version of our general model to derive some of the main insights most transparently. We study the incidence of tax progressivity on individual earnings and welfare in Section 1.2, characterize the optimal rate of progressivity in Section 1.3, and evaluate our findings quantitatively in Section 1.4. Finally, in Section 1.5, we show that our results apply to several alternative models of performance pay that are aimed at capturing different forms of variable compensation: bonuses, piece rates or commissions, stock options, and dynamic incentives.

### 1.1 Environment

There is a continuum of mass one of agents indexed by their exogenous ability  $\theta \in \mathbb{R}_+$  that is distributed according to the c.d.f.  $F(\theta)$  and density  $f(\theta)$ . Preferences over consumption  $c$  and labor effort  $\ell$  are represented by the utility function  $u(c) - h(\ell)$ ,

where  $u$  and  $h$  are twice continuously differentiable,  $u$  is strictly concave, and  $h$  is strictly convex.

A worker with ability  $\theta$  who provides effort  $\ell$  can produce two levels of output  $y$ :

$$y = \begin{cases} \theta & \text{with probability } \pi(\ell) \\ 0 & \text{with probability } 1 - \pi(\ell) \end{cases}$$

where  $\pi : \mathbb{R}_+ \rightarrow [0, 1]$  is continuously differentiable and concave. Without loss of generality, and unless otherwise stated, we normalize units of effort so that  $\pi(\ell) = \ell \in [0, 1]$ .<sup>7</sup>

Firms observe both the agent's ability  $\theta$  and realized output  $y$ , but not their effort  $\ell$ . As a consequence, the labor contract specifies earnings  $z$  as a function of observed performance. If output is low, the worker earns a base salary  $\underline{z}$ . If output is high, earnings take a larger value  $\bar{z} \geq \underline{z}$ . We define the bonus as the difference between the high-level pay and the base pay,  $b = \bar{z} - \underline{z}$ . We also introduce the bonus rate  $\beta$ , or the pass-through of output risk to log-earnings, defined such that  $\bar{z} = e^\beta \underline{z}$ .<sup>8</sup>

The government levies non-linear income taxes. A worker with earnings  $z$  consumes their after-tax earnings  $c = z - T(z) \equiv R(z)$ , where  $T : \mathbb{R} \rightarrow \mathbb{R}$  is the tax schedule and  $R$  is the retention function. Throughout the paper, we denote the utility over pre-tax earnings by  $v(z) \equiv u(R(z))$ . In this section, we restrict preferences and taxes to the following functional forms.

**Assumption 1** *The utility of consumption is logarithmic,  $u(c) = \log c$ . The tax schedule has a constant rate of progressivity (CRP):<sup>9</sup> There exist  $\tau \in \mathbb{R}$  and  $p \in (-\infty, 1)$  such that  $R(z) = \frac{1-\tau}{1-p} z^{1-p}$ .*

A firm that hires a worker with ability  $\theta$  takes the tax schedule and the worker's reservation value  $U(\theta)$  as given. It chooses the earnings contract  $\{\underline{z}(\theta), \bar{z}(\theta)\}$  to

---

<sup>7</sup>The disutility of effort must then be re-normalized as  $\tilde{h} \equiv h \circ \pi^{-1}$ . For clarity, we keep the notation  $h(\ell)$ , except in our calibration exercises that require positing functional forms for  $h(\cdot)$  and  $\pi(\cdot)$ .

<sup>8</sup>Intuitively, for a small bonus  $b$ ,  $\beta$  coincides with the ratio of bonus to base pay  $b/\underline{z}$ .

<sup>9</sup>The CRP tax code is a good approximation of the U.S. tax system, see for instance, [Heathcote et al. \(2017\)](#). The rate of progressivity  $p$  is equal to (minus) the elasticity of the retention rate  $R'(z)$  with respect to income  $z$ . Alternatively,  $1 - p$  is equal to the ratio of marginal retained income  $R'(z)$  to average retained income  $R(z)/z$ .

maximize its expected profit:

$$\Pi(\theta) = \max_{\underline{z}, \bar{z}} \ell\theta - \mathbb{E}[z], \quad (1)$$

where  $\mathbb{E}[z] = (1 - \ell)\underline{z} + \ell\bar{z}$ ,<sup>10</sup> subject to the following constraints.

First, the incentive constraint requires that the worker's effort level maximizes their expected utility over all possible effort choices, taking the labor contract as given:

$$\ell = \arg \max_{l \in [0,1]} (1 - l)v(\underline{z}) + lv(\bar{z}) - h(l). \quad (2)$$

Second, the participation constraint requires that the worker's expected utility is at least as large as their reservation value:

$$\mathbb{E}[v(z)] - h(\ell) \geq U(\theta), \quad (3)$$

where  $\mathbb{E}[v(z)] = (1 - \ell)v(\underline{z}) + \ell v(\bar{z})$ .

To pin down the workers' reservation value  $U(\theta)$ , we assume that there is free-entry of firms in the labor market for workers of type  $\theta$ . Thus, in equilibrium, expected profits are equal to zero:

$$\Pi(\theta) = 0. \quad (4)$$

**Moral Hazard and the Bonus Rate.** The incentive constraint (2) can be simplified by taking the first-order condition of the worker's maximization problem:

$$h'(\ell) = v(\bar{z}) - v(\underline{z}). \quad (5)$$

We can show that this first-order condition is necessary and sufficient to ensure global incentive compatibility as long as the equilibrium level of effort is interior.<sup>11</sup> Using Assumption 1, we can rewrite this expression as follows:

$$\beta = \frac{h'(\ell)}{1 - p}. \quad (6)$$

---

<sup>10</sup>Throughout the paper, unless explicitly stated the operator  $\mathbb{E}$  denotes the expectation of a random variable conditional on a given worker  $\theta$ , as opposed to a population average. We also let  $\mathbb{V}(\cdot | \theta)$  denote the variance conditional on productivity  $\theta$ .

<sup>11</sup>See Lemma 4, which proves this result in the general setting of Section 2.

This equation plays an important role in our analysis. Intuitively, inducing a worker to provide a given level of costly effort  $\ell$  requires a larger reward for performance (i.e., a higher bonus rate  $\beta$ ) if the marginal disutility  $h'(\ell)$  is higher. Since  $h$  is convex, this implies that the sensitivity of earnings to performance  $\beta$  is strictly increasing in labor effort  $\ell$ . This captures the key moral hazard insight that eliciting higher effort from the worker requires raising their exposure to output risk.

**Optimal Contract.** The following proposition characterizes the equilibrium labor contract between a firm and a worker with ability  $\theta$ .

**Proposition 1** *Suppose that Assumption 1 holds. Effort  $\ell$  is then independent of  $\theta$ , so that the bonus rate  $\beta$  is constant and given by (6). Equilibrium earnings are given by*

$$\underline{z}(\theta) = \frac{1}{1 + \ell(e^\beta - 1)}\ell\theta \quad \text{and} \quad \bar{z}(\theta) = \frac{e^\beta}{1 + \ell(e^\beta - 1)}\ell\theta. \quad (7)$$

*Optimal labor effort satisfies*

$$\theta = b + b\ell(1 - \ell)\frac{h''(\ell)}{1 - p}. \quad (8)$$

*Expected utility is given by*

$$U(\theta) = v(\ell\theta) - h(\ell) - (1 - p)\{u(\mathbb{E}z) - \mathbb{E}[u(z)]\}, \quad (9)$$

*with  $u(\mathbb{E}z) - \mathbb{E}[u(z)] = \log(1 + \ell(e^\beta - 1)) - \beta\ell$ .*

**Full-Insurance Benchmark.** To interpret Proposition 1, consider first the benchmark setting where the firm can perfectly monitor the worker's effort, so that the incentive constraint can be ignored. In this case, the firm provides full insurance against output risk: Workers with ability  $\theta$  who provide effort  $\ell$  earn their expected output regardless of their performance,  $\underline{z} = \bar{z} = \theta\ell$ . Their utility is then equal to  $\mathcal{U}(\theta) \equiv v(\theta\ell) - h(\ell)$ . Note that this setting is equivalent to the standard [Mirrlees \(1971\)](#) model, where the relationship  $z = \theta\ell$  between labor effort and income also holds.

**Equilibrium Earnings.** With moral hazard frictions, the previous relationship continues to hold on average: The free-entry condition (4) imposes that workers'

expected earnings  $\mathbb{E}[z]$  remain equal to their expected output  $\ell\theta$ . However, providing effort incentives implies that realized earnings must now be dispersed around their mean. The firm provides only partial insurance against output risk: we have  $0 < \underline{z} < \theta\ell < \bar{z} < \theta$ . At the heart of our paper lies the observation that the optimal degree of within-firm insurance—captured by the parameter  $\beta$  in equation (6)—is *endogenous* to the tax system. We analyze this dependence in the next section.

**Labor Effort.** To interpret the optimality condition for effort (8), suppose that the firm aims to elicit marginally higher effort from the worker. The expected output gain, on the left hand side, is  $\theta$ . Keeping the earnings structure  $\underline{z}, \bar{z}, b$  fixed, inducing higher effort costs  $b$  to the firm, since the worker receives the bonus more frequently.<sup>12</sup> This is the first term on the right hand side of (8). In addition, the firm needs to raise the earnings spread in order to incentivize the worker to actually exert this extra effort. By equation (6), the bonus rate must increase proportionally to the rise in the marginal disutility of effort,  $h''(\ell)$ . This creates an additional cost for the firm—the *marginal cost of incentives* (MCI)—given by the second term on the right hand side of (8): Intuitively, since workers are risk averse, exposing them to more earnings risk while keeping the participation constraint satisfied is costly because it requires increasing mean earnings.

**Expected Utility.** Equation (9) decomposes the worker’s expected utility into three components. The first is the utility they would attain under full insurance,  $\mathcal{U}(\theta) = v(\theta\ell) - h(\ell)$ . Second, the incompleteness of private insurance makes the risk-averse worker worse off: The utility loss associated with a given earnings lottery  $z$  is equal to the utility difference between expected earnings  $\mathbb{E}[z]$  and the certainty equivalent.<sup>13</sup> This difference is the term in curly brackets in expression (9). Third, this utility loss is weighted by  $(1 - p)$ : All else being equal, a higher level of tax progressivity reduces the variance of disposable income that the consumer faces, which dampens the welfare cost of earnings uncertainty. Thus, keeping earnings risk and the level of effort fixed, higher social insurance raises welfare.

---

<sup>12</sup>Note that, by equation (5), this is just sufficient to compensate the worker for the higher effort level and keep the participation constraint satisfied.

<sup>13</sup>Recall that the certainty equivalent  $z^{CE}$  is defined by  $u(z^{CE}) = \mathbb{E}[u(z)]$ .

## 1.2 Tax Incidence

In this section, we explore how the characteristics of the equilibrium contract: labor effort  $\ell$ , mean earnings  $\mathbb{E}z$ , earnings risk  $\beta$ , and expected utility  $U(\theta)$ , respond to a change in tax progressivity  $p$ .

**Effect on Average Earnings and Labor Effort.** It is immediate that the effect of taxes on mean earnings  $\mathbb{E}z = \ell\theta$  is given by:

$$\frac{\partial \log \mathbb{E}z}{\partial \log(1-p)} = \frac{\partial \log \ell}{\partial \log(1-p)} \equiv \varepsilon_{\ell,1-p},$$

where  $\varepsilon_{\ell,1-p}$  denotes the elasticity of labor effort (or of the frequency of receiving a bonus) with respect to (one minus) the rate of tax progressivity. In the full-insurance setting, we have  $\varepsilon_{\ell,1-p} = \frac{\varepsilon_{\ell}^F}{1+\varepsilon_{\ell}^F}$  where  $\varepsilon_{\ell}^F \equiv \frac{h'(\ell)}{\ell h''(\ell)}$  is the Frisch elasticity of labor supply.<sup>14</sup> The following lemma characterizes  $\varepsilon_{\ell,1-p}$  with endogenously incomplete insurance.

**Lemma 1** *Higher tax progressivity reduces optimal labor effort, i.e.,  $\varepsilon_{\ell,1-p} > 0$ . Suppose that the disutility of effort is isoelastic, so that the Frisch elasticity  $\varepsilon_{\ell}^F$  is constant. Then  $\varepsilon_{\ell,1-p} > \frac{\varepsilon_{\ell}^F}{1+\varepsilon_{\ell}^F}$ .*

Lemma 1 shows that, as in standard models of taxation, tax progressivity disincentivizes labor effort. This behavioral response plays a key role in our analysis. Intuitively, greater progressivity makes high-powered incentives more costly to provide. Indeed, when high levels of income are taxed away more heavily, eliciting marginally higher effort requires a larger increase in the dispersion of pre-tax earnings and, therefore, a larger cost for the firm. Formally, both the bonus  $b$  and the MCI term in equation (8) are increasing in  $p$ .<sup>15</sup> Furthermore, Lemma 1 shows that raising tax progressivity causes a stronger negative response of effort than in the benchmark setting without exposure to output risk. Intuitively, this is because the MCI term is absent when effort is observable.

---

<sup>14</sup>This elasticity is lower than the Frisch elasticity because of the income effect of taxes on labor supply.

<sup>15</sup>Note that it would not be possible to deliver incentives by lowering base earnings  $\underline{z}$  rather than raising the bonus, since this would violate the worker's participation constraint.

**Effect on Earnings Risk.** We infer from the local incentive constraint (6) that taxes affect earnings risk  $\beta$  through two channels that work in opposite directions. First,  $\beta$  is inversely proportional to  $1 - p$ , so that, *ceteris paribus*, a higher rate of progressivity leads to a steeper pre-tax earnings schedule. This direct effect of taxes is a standard *crowding-out* of private insurance by social insurance. Intuitively, by raising tax progressivity, the government compresses the disposable income distribution and, therefore, reduces the worker's exposure to output risk. The firm responds by spreading out pre-tax earnings (i.e., raising  $\beta$ ) in order to preserve the worker's incentives for effort. The elasticity of output risk to tax progressivity is given by

$$\varepsilon_{\beta,1-p} \equiv \frac{\partial \log \beta}{\partial \log(1-p)} = -1,$$

which means that, absent effort responses, the firm adjusts the contract so as to keep consumption insurance—the variance of log-consumption—fixed.

Second,  $\beta$  is proportional to  $h'(\ell)$ . As a result, tax progressivity affects earnings risk indirectly via the endogenous choice of labor effort. Recall that a higher rate of progressivity reduces the optimal labor effort. The strength of this effect is captured by the elasticity  $\varepsilon_{\ell,1-p} > 0$ . Now, the firm implements this lower level of effort by reducing the worker's exposure to risk  $\beta$ : that is, by compressing the pre-tax earnings distribution. This effect is a *crowding-in* of private insurance by social insurance. The elasticity of earnings risk to desired labor effort is given by

$$\varepsilon_{\beta,\ell} \equiv \frac{\partial \log \beta}{\partial \log \ell} = \frac{\ell h''(\ell)}{h'(\ell)} = \frac{1}{\varepsilon_{\ell}^F},$$

where  $\varepsilon_{\ell}^F > 0$  is the Frisch elasticity.

**Lemma 2** *The total impact of tax progressivity on earnings risk is given by*

$$\frac{d \log \beta}{d \log(1-p)} = \varepsilon_{\beta,1-p} + \varepsilon_{\beta,\ell} \cdot \varepsilon_{\ell,1-p} = -1 + \frac{\varepsilon_{\ell,1-p}}{\varepsilon_{\ell}^F}.$$

*It is negative (net crowding-out) if  $\varepsilon_{\ell,1-p} < \varepsilon_{\ell}^F$  and positive (net crowding-in) otherwise. By Lemma 1, a lower bound for the net effect is given by  $\frac{d \log \beta}{d \log(1-p)} > -\frac{\varepsilon_{\ell}^F}{1+\varepsilon_{\ell}^F}$ .*

**Crowd-Out and Crowd-In Approximately Offset Each Other.** We now argue that these two counteracting forces—crowd-out and crowd-in—are of the same order

of magnitude, so that tax progressivity has only a modest impact on the amount of within-firm insurance. The key insight is that the strength of the moral hazard friction is inversely proportional to the Frisch elasticity of labor supply,  $\varepsilon_{\beta,\ell} = 1/\varepsilon_\ell^F$ ; This is an immediate consequence of equation (6). If the elasticity of labor effort  $\varepsilon_{\ell,1-p}$  is roughly equal to the Frisch elasticity, then the crowd-in effect is equal to  $\frac{1}{\varepsilon_\ell^F} \varepsilon_{\ell,1-p} \approx 1$ : that is, about the same size (in absolute value) as the direct crowd-out  $\varepsilon_{\beta,1-p} = -1$ .

Importantly, this insight is robust to small values of the labor effort elasticity. To see this, suppose that effort hardly diminishes in response to an increase in tax progressivity, which happens when  $\varepsilon_\ell^F \rightarrow 0$ . Lemma 2 shows that the crowd-in effect is bounded from below by  $1/(1 + \varepsilon_\ell^F)$ , which converges to 1 as  $\varepsilon_\ell^F \rightarrow 0$ . Thus, in this case, the crowding-in of private insurance offsets (at least as much as) the entire crowd-out. Intuitively, it is precisely because of the inelastic behavior that the firm must dramatically reduce the worker’s exposure to output risk in order to implement the desired (tiny) reduction in labor effort: formally,  $\varepsilon_{\beta,\ell} = 1/\varepsilon_\ell^F \rightarrow \infty$ . As a result, the product of elasticities  $\varepsilon_{\beta,\ell} \varepsilon_{\ell,1-p}$  does not vanish in the limit: The crowd-in effect remains significant even when effort is almost inelastic.

This discussion is correct as long as the labor effort elasticity is indeed approximately equal to the Frisch elasticity. In practice, when the Frisch elasticity is bounded away from zero, this need not be exactly the case. Suppose that the disutility of (un-normalized) labor effort is isoelastic on  $\mathbb{R}_+$ ,  $h(\ell) = \ell^{1+1/e}$ , with an empirically realistic value of  $e = 0.5$ . The relevant Frisch elasticity  $\varepsilon_\ell^F$  in Lemma 2, however, is lower than  $e$ , since it is that of the function  $h \circ \pi^{-1}$ , where  $\pi$  is concave (see footnote 7). If  $\pi(\ell) = \sqrt{\ell}$ , for instance, we get  $\varepsilon_\ell^F = e/(2 + e) = 0.2$ . In this case, we find a lower bound for the crowd-in effect  $\varepsilon_{\beta,\ell} \varepsilon_{\ell,1-p} > \frac{1}{1+0.2} = 0.83$ . That is, the earnings risk adjustment due to labor effort responses offsets at least 83 percent of the crowd-out of private insurance caused by tax progressivity. Even when we raise the Frisch elasticity to  $e = 1$ , the high end of empirical estimates, the crowd-in effect offsets at least 75 percent of the crowd-out.

Finally, the above results are based on the assumption that  $u(c) = \log(c)$ , which implies a particular strength of the income effect on labor effort and a particular degree of risk aversion. Nevertheless, our prediction—that the crowd-in offsets most of the crowd-out—continues to hold when there is no income effect and workers are arbitrarily risk-averse. In Appendix C.1, we study the moral-hazard framework of



Holmstrom and Milgrom (1987) with preferences  $u(c, \ell) = -\frac{1}{\gamma}e^{-\gamma(c-h(\ell))}$ , which imply no income effect and an arbitrary absolute risk aversion controlled by the coefficient  $\gamma \geq 0$ . Assuming that the Frisch elasticity is constant, we show that the rate at which the crowd-in offsets the crowd-out is decreasing in risk aversion  $\gamma$ . When workers are risk-neutral, the crowd-in offsets the crowd-out exactly. Even in the limit where the risk aversion coefficient goes to infinity, the two effects remain of comparable magnitude. For instance, assuming a Frisch elasticity of 0.5, the crowd-in always offsets at least two-thirds of the crowd-out.

**Effect on Expected Utility.** In the full-insurance benchmark, the envelope theorem implies that expected utility  $\mathcal{U}(\theta) \equiv v(\theta\ell) - h(\ell) = \log \frac{1-\tau}{1-p} + (1-p) \log(\ell\theta) - h(\ell)$  changes in response to an increase in progressivity by  $\frac{d\mathcal{U}(\theta)}{dp} = \frac{1}{1-p} - \log(\mathbb{E}z)$ . In our setting, we obtain the following result.

**Lemma 3** *The impact of tax progressivity on expected utility is given by*

$$\frac{dU(\theta)}{dp} = \frac{d\mathcal{U}(\theta)}{dp} + [\log(\mathbb{E}z) - \mathbb{E}(\log z)] + \frac{b}{\beta\mathbb{E}z} \mathbb{V}(\log z | \theta) \varepsilon_{\beta, 1-p}, \quad (10)$$

where  $\mathbb{V}(\log z | \theta) = \beta^2 \ell(1-\ell)$  is the variance of log-earnings conditional on  $\theta$ . Note in particular that, to a first order, the worker's expected utility is unaffected by changes in labor effort (envelope theorem).

Higher progressivity  $p$  raises expected utility when earnings are uncertain by reducing the consumption spread that workers face. This is captured by the second term (in square brackets) on the right-hand side of (10). In addition, higher progressivity raises the dispersion of pre-tax earnings via the crowd-out elasticity  $\varepsilon_{\beta, 1-p} = -1$ . This causes a loss in expected utility proportional to the variance of log-earnings  $\mathbb{V}(\log z | \theta)$ .

Importantly, note that the crowding-in of private insurance,  $\varepsilon_{\beta, \ell} \varepsilon_{\ell, 1-p}$ , does not appear in formula (10). This is because this effect operates via optimal labor effort choices. The envelope theorem implies that its impact on welfare is only of second-order.<sup>16</sup> Therefore, while the crowd-out and crowd-in have offsetting effects on the

---

<sup>16</sup>More precisely: Consider an (equivalent) dual formulation of the firm's problem, which consists of maximizing the worker's expected utility subject to making non-negative profits. The envelope theorem applied to this problem implies that changes in labor effort do not have first-order effects on expected utility.

structure of pre-tax earnings, only the former affects the worker's utility. In other words, observing that the amount of pre-tax insurance embedded in the labor contract is insensitive to taxes does *not* imply that the endogeneity of private insurance is negligible for welfare.

### 1.3 Optimal Taxation

The government chooses the tax schedule, or the retention function  $R$ , to maximize a weighted utilitarian social welfare function:

$$\int_0^\infty \alpha(\theta) U(\theta) f(\theta) d\theta \quad (11)$$

subject to a budget constraint:

$$\int_0^\infty \mathbb{E}[z(\theta) - R(z(\theta))] f(\theta) d\theta \geq G, \quad (12)$$

where the Pareto weights  $\alpha(\theta) \geq 0$  satisfy  $\int_0^\infty \alpha(\theta) f(\theta) d\theta = 1$ , and where  $G \geq 0$  is an exogenous expenditure requirement. We denote by  $g$  the ratio of public spending to aggregate output.

**Optimal Tax Progressivity.** Recall that, in this section, we restrict the tax schedule to the CRP class. Theorem 1 characterizes the optimal rate of tax progressivity  $p$  under the additional assumptions that the distribution of ability types is lognormal, and that the social welfare objective is utilitarian.<sup>17</sup>

**Theorem 1** *Suppose that  $\log \theta \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$  and that  $\alpha(\theta) = 1$  for all  $\theta$ . The optimal rate of progressivity satisfies*

$$\frac{p}{(1-p)^2} = \frac{\sigma_\theta^2 + \kappa_1(1 + \varepsilon_{\beta, 1-p}) \mathbb{V}(\log z \mid \theta)}{\left[1 + \frac{g}{(1-g)^p} + \ell \kappa_3\right] \varepsilon_{\ell, 1-p} + \kappa_2 \varepsilon_{\beta, \ell} \varepsilon_{\ell, 1-p} \mathbb{V}(\log z \mid \theta)}, \quad (13)$$

with  $\kappa_1 = \frac{1}{\beta(1-p)} \frac{\bar{c}-c}{\mathbb{E}c}$ ,  $\kappa_2 = \frac{1-p}{\beta p} \left(\frac{b}{\mathbb{E}z} - \frac{\bar{c}-c}{\mathbb{E}c}\right)$ , and  $\kappa_3 = \frac{1-p}{p} \left(\frac{b}{\mathbb{E}z} - \frac{1}{1-p} \frac{\bar{c}-c}{\mathbb{E}c}\right)$ , where  $c, \bar{c}$  denote after-tax earnings in the low- and high-performance states. We have  $\kappa_1 > 0$ ,  $\kappa_2 > 0$ , and  $\kappa_3 = 0$  if  $p = 0$ .

---

<sup>17</sup>In the Appendix, we generalize formula (13) to social welfare weights given by  $\alpha(\theta) \propto \theta^{-a}$ , where  $a \geq 0$ .

To understand Theorem 1, it is helpful to decompose it by considering first the full-insurance benchmark, then the case of exogenously incomplete private insurance, and finally the general setting with endogenously incomplete private insurance.

### A. Full-Insurance Benchmark

Suppose first that firms provide complete insurance against output risk, or equivalently that there is no output risk as in Mirrlees (1971). That is,  $\underline{z} = \bar{z} = \ell\theta$ . Formula (13) then reduces to

$$\frac{p}{(1-p)^2} = \frac{\sigma_\theta^2}{\left[1 + \frac{g}{(1-g)p}\right] \varepsilon_{\ell,1-p}}.$$

The optimal rate of progressivity is increasing in inequality, measured by the variance of the log-ability distribution  $\sigma_\theta^2$ , and decreasing in the elasticity of labor effort  $\varepsilon_{\ell,1-p}$ , which captures the efficiency cost of distortionary taxation. Moreover, it is decreasing in the share of government expenditures in GDP  $g$ : This is because a marginal tax increase induces a larger deadweight loss if the tax burden is already large due to high spending needs.

### B. Exogenously Incomplete Insurance

Suppose next that the firm provides incomplete insurance against output shocks, so that  $\bar{z} > \underline{z}$ , but that earnings risk  $\beta > 0$  is exogenous. Thus, the elasticities of crowd-out and moral hazard  $\varepsilon_{\beta,1-p}$  and  $\varepsilon_{\beta,\ell}$  are both set equal to zero. Formula (13) then reads

$$\frac{p}{(1-p)^2} = \frac{\sigma_\theta^2 + \kappa_1 \mathbb{V}(\log z | \theta)}{\left[1 + \frac{g}{(1-g)p} + \ell\kappa_3\right] \varepsilon_{\ell,1-p}}. \quad (14)$$

There are two differences to the full-insurance benchmark. First, the dispersion of earnings in the population is now mechanically larger than that of ability types  $\theta$ . Tax progressivity thus plays two roles: redistribution across ex-ante ability differences (measured by the variance  $\sigma_\theta^2$ ) and social insurance against ex-post earnings risk (measured by the conditional variance of pre-tax earnings  $\mathbb{V}(\log z | \theta)$ ). Up to a second-order as  $\beta \rightarrow 0$ , the numerator of (14) equals the total variance of log-earnings in the population,  $\sigma_\theta^2 + \mathbb{V}(\log z | \theta)$ . That is, optimal progressivity is an increasing function of overall earnings inequality, regardless of whether it is driven by innate ability differences or idiosyncratic performance shocks.

Second, the labor effort distortions in the denominator of (14) are also modified. In contrast to the full-insurance benchmark, a change in taxes affects not only the income levels  $\underline{z}$ ,  $\bar{z}$  on the intensive margin, but also triggers a response on the frequency (or extensive) margin by altering the probability  $\ell$  with which the high income level  $\bar{z}$ , and hence the high tax payment, occurs. The former behavioral responses affect government revenue proportionally to the income-weighted marginal tax rates  $\mathbb{E}[T'(z)z]\varepsilon_{\ell,1-p}$ , while the latter affects revenue proportionally to the gap in total tax payments between the high- and low-performance states  $\ell(T(\bar{z}) - T(\underline{z}))\varepsilon_{\ell,1-p}$ . Accounting for these frequency responses gives rise to the additional term  $\ell\kappa_3$  in the optimum tax formula.<sup>18</sup>

### C. Endogenously Incomplete Insurance

We finally analyze the general case where earnings risk  $\beta$  is endogenous to taxes. The optimal progressivity formula (13) is modified by two novel terms. These account for fiscal externalities (in the denominator) and welfare effects (in the numerator) that we describe in turn.

**Fiscal Externalities from Crowd-Out and Crowd-In.** If the tax schedule is progressive, a spread of the pre-tax earnings distribution caused by the crowding-out analyzed in Section 1.2 generates a positive fiscal externality, i.e., a first-order gain in government revenue. Conversely, an earnings compression due to the crowding-in of private insurance induces a negative fiscal externality. These are consequences of Jensen's inequality: A progressive tax code generates more revenue for the government if earnings are more volatile, keeping their mean constant.

Formally, the crowd-out  $\varepsilon_{\beta,1-p} < 0$  and crowd-in  $\varepsilon_{\beta,\ell}\varepsilon_{\ell,1-p} > 0$  affect government revenue by  $-\frac{p}{(1-p)^2}\kappa_2(\varepsilon_{\beta,1-p} + \varepsilon_{\beta,\ell}\varepsilon_{\ell,1-p})\mathbb{V}(\log z \mid \theta)$ . If the tax schedule is linear ( $p = 0$ ), these fiscal externalities are equal to zero. If the tax schedule is progressive ( $p > 0$ ), the crowd-out (resp., crowd-in) raises (resp., lowers) government revenue and, hence, the optimal level of taxes. Note finally that the denominator of formula (13) features only the negative externality from crowd-*in*. This is because, as we

---

<sup>18</sup>When  $p = 0$ , we have  $\kappa_3 = 0$ . Recall that the fiscal impact of the intensive margin response depends on the marginal tax rate, while that of the frequency margin response depends on the average tax rate. Since the marginal and average tax rates are equal when  $p = 0$ , there is no need to correct the efficiency cost of taxes in this case.

argue next, the positive fiscal externality caused by the crowd-out is (more than) compensated for by its negative welfare impact.

**Welfare Effects of Crowd-Out and Envelope Theorem.** The crowd-out of private insurance affects the social welfare objective (11) by  $[\kappa_1 + \frac{p}{(1-p)^2} \kappa_2] \varepsilon_{\beta, 1-p} \mathbb{V}(\log z | \theta)$ . Thus, since  $\varepsilon_{\beta, 1-p} = -1$ , it exactly offsets both the insurance benefits of tax progressivity and the positive fiscal externality discussed in the previous paragraphs.

The fact that the crowd-out of private insurance offsets the welfare benefits of social insurance ( $\kappa_1 \mathbb{V}(\log z | \theta)$ ) is straightforward. Any attempt by the government to compress the distribution of disposable income leads the firm to raise the dispersion of pre-tax earnings one-for-one in order to preserve the worker’s effort incentives.<sup>19</sup> Thus, in contrast to the case of exogenous private insurance, the government should not provide any social insurance against performance shocks: The numerator of (13) reduces to the benefits of insuring exogenous ability disparities  $\theta$ , measured by the variance  $\sigma_\theta^2$ .

The additional negative welfare effect of the crowd-out is more subtle. It is due to the fact that tax cuts lead to endogenous changes in the pay structure that render them less precisely targeted than in the standard framework—we elaborate on this point in more detail in Section 2. To build intuition, suppose that the tax liabilities at the base pay  $T(\underline{z})$  and at the high-level pay  $T(\bar{z})$  are lowered by the same amount. In Section 2.2 we show that such a reform increases the worker’s reservation value  $U(\theta)$  by (say)  $\Delta U > 0$ . The incentive constraint (5) implies that, to maintain effort  $\ell$ , the ex-post utility of both low- and high-performers must increase by the same amount  $\Delta U$ .<sup>20</sup> Absent changes in gross earnings, however, the (uniform) tax cut would not by itself lead to a uniform rise in ex-post utility, since the marginal utility of consumption is strictly decreasing. Consequently, the firm must raise the bonus and, to ensure that profits remain non-negative, lower the base pay. This implies that high-performers capture a disproportionately large share of the tax cut. This regressive distribution of rents—away from individuals whose marginal utility of consumption is the highest—further reduces the welfare benefits of redistribution via progressive taxes.

---

<sup>19</sup>More precisely, recall that the firm does not actually keep labor effort unchanged. But by the envelope theorem, the welfare consequences of the corresponding crowding-in are second-order.

<sup>20</sup>This is a standard consequence of the separability of the utility function; see, e.g., Golosov, Kocherlakota, and Tsyvinski (2003). We derive this property formally in equations (16) and (17) below.

Finally, note that these welfare losses due to the crowding-out are not counteracted by corresponding gains from the crowding-*in* responses. As explained in the discussion following Lemma 3, this is a consequence of the envelope theorem: The crowding-in of private insurance in response to an increase in tax progressivity operates via adjustments in labor effort that is chosen optimally, leading to an (at most) second-order impact on welfare. Hence, the observation that earnings risk is almost invariant to taxes (Section 1.2) does *not* imply that optimal tax design can safely ignore the endogeneity of private insurance.

**Taking Stock: Endogenous Insurance Reduces Optimal Progressivity.** In sum, there are two channels through which endogenous earnings risk matters for tax progressivity: (i) the negative fiscal externality due to crowd-in; and (ii) the negative welfare effect of crowd-out, which exactly offsets the benefits of social insurance.<sup>21</sup> As a result, the optimal rate of progressivity (13) is *strictly lower* than in a setting with exogenous earnings risk.

## 1.4 Quantitative Analysis

To evaluate our results quantitatively, we extend our baseline model with the following elements. A share  $s_{pp}$  of workers have a *performance-pay job*, and the remaining share  $s_{fp} = 1 - s_{pp}$  have a *fixed-pay job*. Performance-pay jobs are subject to the moral hazard friction described above: output is stochastic, equal to the worker's ability  $\theta$  with probability  $\pi(\ell_{pp})$  and 0 otherwise, where  $\ell_{pp}$  is their effort level. The optimal contract specifies earnings as a function of the output realization according to the bonus rate  $\beta = h'(\ell_{pp})/(1 - p)$ . Fixed-pay jobs, by contrast, are not subject to agency frictions and guarantee a risk-free earnings level  $\theta\ell_{fp}$ , where  $\ell_{fp}$  is the worker's effort level. In equilibrium, all fixed-pay workers exert the same effort.

We treat the job type of a worker as exogenous. In the data, the share of performance-pay jobs increases with earnings (see Lemieux et al. 2009, Gittleman and Pierce 2013, Grigsby, Hurst, and Yildirmaz 2019). To account for this fact in the model, we allow for a positive correlation between job type and ability. Specifically, we assume that ability is drawn from the job-type-specific Pareto-lognormal

---

<sup>21</sup>In the general environment of Section 2, the positive fiscal externality and the additional negative welfare effect of crowd-out no longer exactly cancel out, and both appear explicitly in the optimal tax formulas.

distribution (Colombi 1990). Thus, conditional on the job type  $j \in \{fp, pp\}$ , log-ability is the sum of independently drawn Gaussian and exponential random variables:  $\log(\theta) = x_N + x_E$ , where  $x_N \sim N(\mu_{\theta,j}, \sigma_{\theta,j}^2)$  and  $x_E = Exp(\lambda_{\theta,j})$ .

**Calibration.** We calibrate to model to match empirical evidence on performance-pay jobs, earnings elasticities, and the overall earnings distribution in the U.S. The chosen parameter values are summarized in Table 1.

Lemieux et al. (2009) (LMP) use the Panel Study of Income Dynamics to show that the fraction of performance-pay jobs  $s_{pp}$  was 0.45 in 1998, the most recent year included in their analysis. We replicated their analysis and found that mean earnings were 58 percent higher in performance-pay jobs than in fixed-pay jobs in 1998. This value pins down  $\mu_{\theta,pp} - \mu_{\theta,fp}$ , the difference in mean log-abilities between the two types of jobs. We postulate that the probability of a high output realization is given by  $\pi(\ell) = \bar{\pi} \ell^\rho$ , with  $\rho \in (0, 1]$ . In the data, the average probability of receiving a bonus conditional on having a performance-pay job is 23%, which pins down  $\bar{\pi}$ .<sup>22</sup> The exponent  $\rho$  affects the magnitude of earnings risk due to performance pay. LMP report that the variance of log-earnings in performance-pay jobs is 42 percent higher than in fixed-pay jobs. This excess variance can be explained either by the additional earnings risk generated by stochastic bonuses (controlled by  $\rho$ ) or by a greater dispersion in the ex ante abilities of performance-pay workers ( $\sigma_{\theta,pp}^2 > \sigma_{\theta,fp}^2$ ). In our baseline calibration, we assume that the entire excess variance arises from the former channel; that is, we assume that log-abilities in the two types of jobs have the same dispersion ( $\sigma_{\theta,pp}^2 = \sigma_{\theta,fp}^2 \equiv \sigma_\theta^2$ ). This makes performance pay as powerful in affecting earnings dispersion as the data allows: in our baseline calibration, performance pay explains 30 percent of the cross-sectional variance of log-earnings of performance-pay workers. We show later that our main conclusions are robust to this assumption.

To calibrate the overall mean and variance of log-earnings in the economy, we turn to the Survey of Consumer Finances (SCF), which uses data from the IRS Statistics of Income program to accurately represent the distribution of high-income households. Based on the SCF, Heathcote and Tsujiyama (2021) report a mean household labor

---

<sup>22</sup>Based on Table II and footnote 15 in LMP, we calculate that (i) in the 1990s, 82 percent of performance-pay jobs used bonuses, while the rest used piece rates or commissions; and (ii) the average frequency of receiving performance pay, conditional on having a performance-pay job, was 37 percent. Given that piece rates and commissions are paid out with certainty, we calculate that the probability of receiving a bonus, conditional on having a job with bonuses, is 23 percent.

Table 1: Calibrated parameters

parameter	value	description	source or target
$s_{pp}$	0.45	share of performance-pay jobs	Lemieux et al. (2009)
$\mu_{\theta,fp}$	3.43	mean log-ability at fixed-pay jobs	mean earnings in the economy
$\mu_{\theta,pp}$	5.29	mean log-ability at perf.-pay jobs	diff. in mean earnings between job types
$\sigma_{\theta}$	0.29	normal variance of log-ability	variance of log-earnings in the economy
$\lambda_{\theta}$	2.2	tail parameter of log-ability	Heathcote and Tsujiyama (2021)
$\bar{\pi}$	45	level parameter of $\pi(\cdot)$ function	mean frequency of bonus payments
$\rho$	0.82	curvature of $\pi(\cdot)$ function	diff. in var. of log-earnings between job types
$\varepsilon_{\ell}^F$	0.5	Frisch elasticity of labor effort	Chetty et al. (2011), Keane (2011)
$p$	0.181	tax progressivity	Heathcote et al. (2017)
$G/Y$	0.188	share of government spending in GDP	Heathcote and Tsujiyama (2021)

Note: All the target moments are matched exactly.

income of \$77,325 and an overall variance of log-labor income of 0.618 in 2007. They also estimate that the tail parameter of the log wage distribution is equal to 2.2. We assume that the ability distributions in both types of jobs have a common tail parameter  $\lambda_{\theta,pp} = \lambda_{\theta,fp} = 2.2$  and choose  $\sigma_{\theta}^2 = 0.29$  to match the overall variance of log-earnings.

We model the disutility of labor effort as isoelastic:  $h(\ell) = \frac{\ell^{1+1/\varepsilon_{\ell}^F}}{1+1/\varepsilon_{\ell}^F}$ . A Frisch elasticity  $\varepsilon_{\ell}^F = 0.5$  implies a compensated elasticity at fixed-pay jobs of approximately 0.3. Both values are consistent with empirical evidence (Chetty, Guren, Manoli, and Weber 2011; Keane 2011). Regarding government policy, Heathcote et al. (2017) estimate a value of 0.181 for the U.S. rate of tax progressivity, and Heathcote and Tsujiyama (2021) report a ratio of government purchases to output of 18.8 percent.

The implied distribution of earnings and job types is depicted in Figure 1. Specifically, in panel (b), we compare the (untargeted) shares of performance-pay jobs by earnings quartiles in the data and in the model.<sup>23</sup> The calibrated model successfully matches the empirical prevalence of performance-pay jobs: Both in the data and in the model, the share of performance-pay jobs is approximately 40 percent in the bottom three quartiles and it rises to more than 60 percent in the top quartile.

### Incidence of a Large Tax Reform: From the Status Quo to the Optimum.

We extend the optimal progressivity formula in the Appendix to account for fixed-pay jobs and a Pareto tail of earnings (see the proof of Theorem 1). We find that the utilitarian optimum rate of progressivity is equal to 0.376. This is more than twice as

<sup>23</sup>The data shares are computed for the year 1998 from PSID using the methodology of LMP. We are grateful to Karl Schulz for computing these values.



Figure 1: Joint distribution of earnings and job types





Figure 2: Optimal progressivity and Self-Confirming Policy Equilibrium




high as the current rate of tax progressivity in the U.S., and the implied increase in social welfare is equivalent to a 3.7 percent increase in consumption (see panel (a) of [Figure 2](#)). In this paragraph, we analyze the impact of a large reform that implements the optimal rate of progressivity, while keeping the overall tax revenue unchanged.

The impact of the reform on earnings risk in performance-pay jobs is depicted in [Figure 3](#). Following a large increase in progressivity, the bonus rate  $\beta$  increases modestly from 1.08 to 1.11. As a result, the variance of log-earnings conditional on ability, equal to  $\beta^2\pi(\ell_{pp})(1 - \pi(\ell_{pp}))$ , actually slightly decreases, as the impact of a higher bonus rate  $\beta$  is dominated by the impact of a lower effort  $\ell_{pp}$ , distorted downwards by higher tax progressivity.

Underlying the weak response of earnings risk are two countervailing forces: the crowding-out and the crowding-in of private insurance. If firms attempted to motivate workers to maintain their original level of effort, better social insurance via tax progressivity would crowd-out private insurance one-for-one. For that to happen, the bonus rate  $\beta$  would need to increase from 1.08 to 1.42, raising the log-earnings risk

Figure 3: Optimal progressivity reform and earnings risk



Note: Tax progressivity rate is increased from 0.181 (status quo) to 0.376 (utilitarian optimum).

of performance-pay workers by 72 percent. However, in equilibrium firms choose to elicit a 9 percent lower effort level (which implies the same fall in mean earnings) and reduce the power of incentive-pay accordingly. This crowding-in effect counteracts the crowd-out and brings the bonus rate back to the vicinity of its original level. As a result, workers end up much better insured: The variance of log-consumption falls by 43 percent.

**Importance of Performance Pay for Optimal Tax Progressivity.** How important is it to account for endogenous performance pay when setting tax policy? To answer this question, we compare the optimal rate of progressivity with a rate chosen by the government that erroneously assumes that the entire earnings risk is exogenous. The latter is found by applying the formula for the optimal rate of progressivity from the model with exogenous partial insurance to our calibrated model economy—in which private insurance is actually endogenous. Since the rate of progressivity affects the perceived earnings dispersion, we iterate on the tax formula from the exogenous-insurance model until convergence to a fixed point. Following [C. Rothschild and Scheuer \(2016\)](#), we call the resulting allocation a Self-Confirming Policy Equilibrium (SCPE). The results are depicted in panel (a) of [Figure 2](#).

A policymaker who ignores the endogeneity of earnings risk chooses a rate of progressivity (0.43) that is higher than the optimum (0.38). While we allow such a planner to estimate the elasticity of earnings precisely, the rate of progressivity is too high, since the negative effects of crowd-out on welfare and crowd-in on tax revenue

are not internalized. Nevertheless, the welfare cost of the policy mistake is relatively small, equal to approximately 0.3 percent of consumption—this is significantly smaller than the gains from raising taxes from the status quo to the optimum. Recall that in the calibration we assumed that it is performance pay, rather than a greater dispersion of abilities, that explains the higher variance of earnings among performance-pay jobs. Allowing for more dispersed abilities at performance-pay jobs would reduce the magnitude of endogenous earnings risk due to stochastic bonuses. Thus, it would further reduce the already small welfare cost of ignoring the endogeneity of earnings risk. Hence, our baseline results provide an upper bound for the impact of endogenous insurance on social welfare.

We also simulate a counterfactual economy, assuming that all the jobs feature performance pay; see panel (b) of [Figure 2](#). In this economy, the difference between the rate of progressivity at the optimum and in the SCPE doubles (from 0.055 to 0.11), while the welfare loss from the policy mistake more than quadruples (from 0.3 percent to 1.4 percent). Thus, the finding that the cost of ignoring performance pay for the design of tax policy is modest is not a theoretical necessity. It is rather a quantitative result that is due to the fact that only half of the jobs in the U.S. feature performance pay. If the share of performance-pay jobs continues to increase in the future, accounting for the welfare effects of endogenous earnings risk may become important for tax policy.

## 1.5 Commissions, Stock Options, Dynamic Incentives

Our baseline model, in which earnings are binary, is well suited to analyzing contracts that consist of a base salary and a bonus. We argued above that these represent more than 80 percent of the performance-based earnings contracts. Nevertheless, it is important to evaluate whether our main insights carry over to other types of compensation. In [Appendix C](#), we set up several alternative frameworks. The first, based on [Holmstrom and Milgrom \(1987\)](#), provides conditions under which linear contracts are optimal: It is thus a natural setting to study piece rates or commissions. Stock options can be represented by contracts that are convex in performance: To analyze the impact and optimality of tax policy with such contracts in a tractable way, we build on [Edmans and Gabaix \(2011\)](#). Finally, high-powered incentives can be provided over time via, e.g., promotions or salary raises. To study such effects, we use

the frameworks of [Sannikov \(2008\)](#) or [Edmans et al. \(2012\)](#).

Despite their differences, all of these models share a common structure that is critical for our analysis: When the utility is logarithmic and the tax schedule is CRP, the slope of the earnings contract, which measures the sensitivity of earnings to output shocks, is given by  $h'(\ell)/(1-p)$ .<sup>24</sup> This is a direct consequence of the local incentive constraint, which is common to all of these models and states that the sensitivity of utility to output shocks must be equal to the marginal disutility of effort  $h'(\ell)$ . In turn, this general principle implies that our discussion of the incidence of tax progressivity continues to hold: A tax change creates crowd-out and crowd-in effects on the pre-tax earnings distribution captured by the elasticities  $\varepsilon_{\beta,1-p} = -1$  and  $\varepsilon_{\beta,\ell} = 1/\varepsilon_{\ell}^F$ . To the extent that the impact of the tax change on effort  $\varepsilon_{\ell,1-p}$  is close to the Frisch elasticity of labor supply, the contract remains insensitive to tax policy.<sup>25</sup> By contrast, taxes have a large impact on welfare since the crowd-in affects utility only to a second-order. As in [Theorem 1](#), these effects are reflected in the optimal tax formulas in these respective environments (see, e.g., [Theorems 4 and 5](#) in [Appendix C](#)).

## 2 General Analysis of Non-Linear Taxation

The model we studied in [Section 1](#) relies on several strong assumptions: The utility of consumption is logarithmic, the tax schedule is a function of total earnings and has a constant rate of progressivity, and the distribution of earnings is lognormal. We now relax these restrictions. In particular, we allow the government to tax base earnings and bonuses either jointly or separately in an arbitrarily nonlinear way. We characterize the incidence of tax reforms in [Section 2.2](#) and derive analytical formulas for the optimal tax rate on top earners—levied either on total earnings or on bonuses separately from the base pay—in [Sections 2.3 and 2.4](#). We evaluate these results quantitatively in [Section 2.5](#). Additional results are derived in [Appendix B](#).

<sup>24</sup>In the case of [Holmstrom and Milgrom \(1987\)](#), the theoretical restrictions on agents' preferences allow us to study only CARA utility function with affine taxes. In this case, the relevant notion of pay-performance sensitivity becomes  $\theta^{-1}h'(\ell)/(1-\tau)$ , where  $\tau$  is the tax rate. Clearly, our discussion continues to hold once we focus on reforms of the tax rate  $\tau$ .

<sup>25</sup>This discussion suggests that these results hold more generally in any moral-hazard model with a continuous choice of effort in which the first-order approach (or the one-shot deviation principle in the continuous-time setting of [Sannikov 2008](#)) is valid.

## 2.1 Environment

The environment is the same as in Section 1, except that the utility of consumption  $u$  is now a twice continuously differentiable and concave function, and the retention schedule  $R : \mathbb{R}_+^2 \rightarrow \mathbb{R}$  is an a.e. twice continuously differentiable function of the agent's base pay and realized bonus. Thus, a worker with high output realization consumes  $c = R(\underline{z}, b)$  and gets utility  $v(\underline{z}, b) \equiv u(R(\underline{z}, b))$ , while a worker with low output realization consumes  $c = R(\underline{z}, 0)$  and gets utility  $v(\underline{z}, 0) \equiv u(R(\underline{z}, 0))$ . This specification allows us to consider both the joint taxation of overall earnings, and the separate taxation of fixed and variable pay. We impose the following assumption, which ensures that the tax schedule is not “too regressive.”<sup>26</sup>

**Assumption 2** *The utility of pre-tax earnings  $v$  is concave on  $\mathbb{R}_+^2$ .*

The firm maximizes its expected profit (1) subject to, *mutatis mutandis*, the incentive constraint (2) and the participation constraint (3). The equilibrium reservation value is pinned down by the free-entry condition (4). The following lemma shows formally that we can replace the global incentive constraint with the corresponding first-order condition for effort (5).<sup>27</sup>

**Lemma 4** *Suppose that the equilibrium effort level is interior. The firm's problem is equivalent to maximizing (1) subject to (3) and the local incentive constraint:*

$$h'(\ell) = v(\underline{z}, b) - v(\underline{z}, 0). \quad (15)$$

As in the previous section, equation (15) along with the convexity of  $h$  implies that eliciting a higher effort level from workers requires raising their exposure to output risk. The following characterization of the optimal contract generalizes Proposition 1 and has an analogous interpretation.

---

<sup>26</sup>See Appendix for details. In particular, this condition always holds under Assumption 1 regardless of the value of  $p$ . It is a natural restriction: Doligalski (2019) shows that when this condition is violated, firms have incentives to offer stochastic earnings even in the absence of moral hazard frictions. Furthermore, a tax schedule that encourages such earnings randomization is Pareto inefficient.

<sup>27</sup>Rogerson (1985) credits an unpublished paper by Holmstrom (1984) for the first proof of validity of the first-order approach in such a setting.

**Proposition 2** *The base pay  $\underline{z}(\theta)$  and high-level pay  $\bar{z}(\theta)$  satisfy*

$$v(\underline{z}, 0) - h(\ell) = U(\theta) - \ell h'(\ell) \quad (16)$$

$$v(\underline{z}, b) - h(\ell) = U(\theta) + (1 - \ell)h'(\ell). \quad (17)$$

*The effort level  $\ell(\theta)$  exerted by the worker satisfies*

$$\theta = b + \left[ \frac{1}{v_2(\underline{z}, b)} - \frac{1 - \frac{R_1(\underline{z}, b)}{R_2(\underline{z}, b)}\ell}{1 - \ell} \frac{1}{v_1(\underline{z}, 0)} \right] \ell(1 - \ell)h''(\ell). \quad (18)$$

*Expected utility  $U(\theta)$  satisfies  $\underline{z} + \ell b = \theta \ell$ .*

## 2.2 Tax Incidence

In this section, we evaluate the impact of tax reforms on workers' labor contracts and welfare. For conciseness, we focus on the taxation of overall earnings and relegate the analysis of the separate taxation of bonuses to the Appendix (Lemma 9). Consider a given baseline retention schedule  $R : \mathbb{R}_+ \rightarrow \mathbb{R}$  and another function (“tax reform”)  $\hat{R} : \mathbb{R}_+ \rightarrow \mathbb{R}$ . The first-order change in a functional  $\Psi(R)$  in response to the reform  $\hat{R}$  is given by the Gateaux derivative  $\hat{\Psi}(R, \hat{R}) \equiv \lim_{\delta \rightarrow 0} (\Psi(R + \delta \hat{R}) - \Psi(R))/\delta$ .

**Earnings.** We start by studying the effect of tax reforms on the worker's earnings. We obtain the following characterization:

**Lemma 5** *The first-order impact of a tax reform  $\hat{R}$  on earnings  $\underline{z}, \bar{z}$  is given by*

$$\hat{\underline{z}} = \left[ -\frac{\hat{R}(\underline{z})}{R'(\underline{z})} + \frac{\hat{U}}{v'(\underline{z})} \right] - \frac{\ell h''(\ell)}{v'(\underline{z})} \hat{\ell} \quad (19)$$

$$\hat{\bar{z}} = \left[ -\frac{\hat{R}(\bar{z})}{R'(\bar{z})} + \frac{\hat{U}}{v'(\bar{z})} \right] + \frac{(1 - \ell)h''(\ell)}{v'(\bar{z})} \hat{\ell}, \quad (20)$$

where  $\hat{U}$  and  $\hat{\ell}$  denote the first-order changes in the reservation utility and labor effort due to the reform. The terms in square brackets constitute the crowd-out, while the terms multiplied by  $\hat{\ell}$  constitute the crowd-in.

Lemma 5 shows that the tax reform affects earnings  $z \in \{\underline{z}, \bar{z}\}$  via three channels, the first two of which jointly constitute the crowd-out, while the last is the crowd-in.

The first term in equations (19) and (20),  $-\hat{R}(z)/R'(z)$ , implies that, *ceteris paribus*, the agent's consumption  $c = R(z)$  remains unchanged despite the tax change. Indeed, this term implies

$$\hat{c} \equiv \hat{R}(z) + R'(z)\hat{z} = \hat{R}(z) - R'(z)\frac{\hat{R}(z)}{R'(z)} = 0.$$

Thus, absent any change in the reservation utility and optimal labor effort, the firm would simply adjust pre-tax earnings so as to keep the worker's disposable income levels  $\underline{c}$  and  $\bar{c}$  fixed. In other words, any attempt by the government to affect consumption insurance would be fully offset by the firm in order to preserve incentives.

Second, the tax reform affects the earnings contract via its impact on the equilibrium reservation utility. The increase in income  $z$  resulting from an increase in the reservation value  $\hat{U} > 0$  is inversely proportional to the marginal utility  $v'(z)$ . Thus, the earnings of the high-performers increase by a larger amount than those of the low-performers with the same ability. This ensures that the utility gain  $\hat{U}$  is distributed uniformly across agents regardless of their performance, thus preserving incentive compatibility (15).

Third, the tax reform modifies the desired effort level. Recall that, by equation (15), eliciting higher effort  $\hat{\ell} > 0$  requires widening the gap between the utility of high- and low-performers by  $\Delta h'(\ell) = h''(\ell)\hat{\ell}$ . The implied change in earnings—the crowd-in effect—is given by the third term in equations (19) and (20).

**Expected Utility.** Next, we evaluate the effect of tax reforms on the worker's expected utility.

**Lemma 6** *The first-order impact of a tax reform  $\hat{R}$  on expected utility  $U(\theta)$  is given by*

$$\hat{U} = \mathbb{E} \left[ \mu(z)u'(R(z))\hat{R}(z) \right] \text{ where } \mu(z) = \frac{1}{\mathbb{E} \left[ \frac{1}{v'} \right]} \text{ for } z \in \{\underline{z}, \bar{z}\}, \quad (21)$$

*Note in particular that, to a first order, the worker's expected utility is unaffected by changes in labor effort (envelope theorem).*

In the standard (full-insurance) model, a tax cut  $\hat{R}(z) > 0$  affects the worker's utility in proportion to their marginal utility of consumption  $u'(R(z))$ , and the envelope theorem ensures that the endogenous behavioral responses have no first-order

impact on utility. This is no longer true in the model with performance pay. While the envelope theorem still applies to the endogenous effort (i.e., crowd-in) responses, the earnings adjustments caused by the crowd-out of private insurance have a first-order impact on welfare. The additional factor  $\mu$  present in equation (21) accounts for these welfare effects.<sup>28</sup>

The intuition is analogous to that laid out in Section 1: A tax cut initially generates a rent to the firm—the first term in equations (19) and (20)—which is then transferred to the worker via free entry and the resulting adjustment of the reservation value. This transfer needs to keep the original effort level incentive-compatible and, hence, the utility difference  $v(\bar{z}) - v(\underline{z})$  unchanged. Thus, earnings  $z$  must change in proportion to  $\frac{1}{v'(z)}$ , leading to the factor  $\mu(z)$  in Lemma 6. Since the marginal utility  $v'$  is decreasing, this means that the high-level pay rises more than the base pay, resulting in a regressive distribution of rents.

To measure the welfare impact of tax reforms, the standard approach is to use marginal social welfare weights, defined as the welfare impact of marginally increasing the consumption of a given agent (Saez & Stantcheva, 2016). In the standard Mirrlees model, the marginal welfare weight of a worker with ability  $\theta$  and earnings  $z$  is given by

$$g(z | \theta) = \frac{1}{\lambda} \alpha(\theta) u'(R(z)), \quad (22)$$

where  $\alpha(\theta)$  is the Pareto weight of type  $\theta$  in the social objective, and  $\lambda$  is the marginal value of public funds. The social marginal welfare weights are convenient in the standard model because, by the envelope theorem, all endogenous responses to tax changes have at most a second-order impact on the agent's utility.

By contrast, recall that in our model the envelope theorem does not apply to all endogenous responses: The crowd-out of private insurance has a first-order impact on welfare. To account for this effect, we define the *modified* marginal social welfare weights  $\tilde{g}(z | \theta)$  as the welfare impact of marginally reducing the tax liability of an agent with type  $\theta$  and realized income  $z$ .

**Corollary 1** *The modified marginal social welfare weights satisfy*

$$\frac{\tilde{g}(\underline{z} | \theta)}{g(\underline{z} | \theta)} = \mu(\underline{z}) < 1 \quad \text{and} \quad \frac{\tilde{g}(\bar{z} | \theta)}{g(\bar{z} | \theta)} = \mu(\bar{z}) > 1. \quad (23)$$

---

<sup>28</sup>Chang and Park (2017) show a similar partial applicability of the envelope theorem in an Alvarez and Jermann (2000) economy with asset trades and endogenous borrowing limits.



Relative to the standard model, the contribution to social welfare of agents with type  $\theta$  is adjusted upwards when they receive a bonus, and downwards otherwise. These regressive adjustments to the marginal social welfare weights reflect the fact that, as explained above, tax cuts are passed through within the firm primarily to the high-performing workers. It is therefore more difficult for the government to target transfers at unlucky workers, i.e., to directly raise the consumption of those whose marginal utility is relatively large.

**Labor Effort.** Next, we characterize the incidence of tax reforms on labor effort. The general analytical expressions for the elasticities of effort with respect to marginal and average tax rates are given in Lemma 8 in the Appendix. For conciseness, the following Lemma only focuses on their signs.

**Lemma 7** *The first-order impact of a tax reform  $\hat{R}$  on labor effort  $\ell$  can be expressed as*

$$\frac{\hat{\ell}}{\ell} = \varepsilon_{\ell, R'(\underline{z})} \frac{\hat{R}'(\underline{z})}{R'(\underline{z})} + \varepsilon_{\ell, R'(\bar{z})} \frac{\hat{R}'(\bar{z})}{R'(\bar{z})} + \varepsilon_{\ell, R(\underline{z})} \frac{\hat{R}(\underline{z})}{R(\underline{z})} + \varepsilon_{\ell, R(\bar{z})} \frac{\hat{R}(\bar{z})}{R(\bar{z})}. \quad (24)$$

*Suppose that the base pay  $\underline{z}$  and high-level pay  $\bar{z}$  are both located in brackets where the marginal tax rate is locally constant. A higher marginal tax rate on base (resp. high-level) pay raises (resp., reduces) labor effort:  $\varepsilon_{\ell, R'(\underline{z})} < 0 < \varepsilon_{\ell, R'(\bar{z})}$ . Suppose moreover that the utility function is logarithmic. A higher average tax rate on base pay raises labor effort:  $\varepsilon_{\ell, R(\underline{z})} < 0$ . A higher average tax rate on high-level pay reduces labor effort, so that  $\varepsilon_{\ell, R(\bar{z})} > 0$ , if and only if  $R'(\underline{z}) < R(\underline{z})/\underline{z}$ .*

In the standard model, effort responds negatively to the marginal tax rate (MTR) due to the substitution effect and positively to the average tax rate (ATR) due to the income effect. Our model is more complex, as there are two relevant levels of pay (base and high-level) and, thus, four tax rates that impact behavior. Focusing first on the average rates, we find that a higher ATR levied on the base pay increases effort, but a higher ATR levied on the high-level pay reduces effort (when the tax schedule is progressive). Intuitively, adjusting effort in this way allows the worker to escape some of the increased tax burden by reducing the probability of receiving the income that is taxed more heavily.

Moving on to marginal tax rates, we find that a higher MTR levied on high-level pay reduces effort, but a higher MTR on base pay, perhaps surprisingly, *increases* effort. Note that a change in marginal rates alone does not modify workers' incentives for effort, since the incentive constraint (15) is unaffected. What are modified, however, are the firm's incentives to offer different pay structures. Intuitively, a higher MTR gives firms incentives to reduce gross pay and save on payroll, since the implied reduction in after-tax income—which matters for workers—is now smaller. A higher MTR at the high-level pay  $\bar{z}$  then leads to a lower value of  $\bar{z}$  and, via the incentive constraint, to a lower effort level. Conversely, a higher MTR on base pay  $\underline{z}$  leads to a lower value of  $\underline{z}$  and, hence, a higher effort level.

Summing up, we find that labor effort responds very differently to taxes on base pay and on high-level pay. In particular, increasing the marginal tax rate on base pay actually encourages effort. These findings underlie our later results on the efficiency gains from taxing bonuses at a lower rate than base pay.

**Taking Stock: Crowd-Out and Crowd-In.** We can now gather our previous results to obtain a full characterization of the incidence of tax reforms on the earnings contract. The following result, obtained by substituting expression (21) into (19) and (20), generalizes the crowding-in and crowding-out forces (highlighted in Section 1) to arbitrary tax systems and reforms.

**Proposition 3** *The first-order impact of a tax reform  $\hat{R}$  on earnings  $\underline{z}, \bar{z}$  is given by*

$$(1 - \ell)\hat{\underline{z}} = - (1 - \varepsilon_{out})(1 - \ell)\frac{\hat{R}(\underline{z})}{R'(\underline{z})} + \varepsilon_{out}\ell\frac{\hat{R}(\bar{z})}{R'(\bar{z})} - \ell\bar{z}\varepsilon_{in}\frac{\hat{\ell}}{\ell} \quad (25)$$

$$\ell\hat{\bar{z}} = (1 - \varepsilon_{out})(1 - \ell)\frac{\hat{R}(\underline{z})}{R'(\underline{z})} - \varepsilon_{out}\ell\frac{\hat{R}(\bar{z})}{R'(\bar{z})} + [\ell\bar{z}\varepsilon_{in} + \underline{z}]\frac{\hat{\ell}}{\ell}, \quad (26)$$

where the crowd-out parameter  $\varepsilon_{out} \in (0, 1)$  and the crowd-in parameter  $\varepsilon_{in} > 0$  are respectively given by

$$\varepsilon_{out} = \frac{1}{\mathbb{E}\left[\frac{1}{v'(z)}\right]}\frac{1 - \ell}{v'(\underline{z})} \quad \text{and} \quad \varepsilon_{in} = \frac{\ell(1 - \ell)h''(\ell)}{\bar{z}v'(\underline{z})},$$

and where the labor supply response  $\hat{\ell}/\ell$  is given by Lemma 7 (and Lemma 8 in the Appendix).

Proposition 3 shows that earnings  $\underline{z}$ ,  $\bar{z}$  respond to changes in tax payments at both levels of compensation. A tax hike on high-level earnings  $\hat{R}(\bar{z}) < 0$ , for instance, leads the firm to lower the base pay  $\underline{z}$  and raise the high-level pay  $\bar{z}$ . The magnitude of this crowding-out of private insurance is determined by the parameter  $\varepsilon_{out} \in (0, 1)$ , for which our model gives a simple analytic expression.<sup>29</sup> Moreover, a tax reform that leads to a lower optimal effort level,  $\hat{\ell} < 0$ , triggers an increase in the base pay and a decrease in the high-level pay. The magnitude of this crowding-in of private insurance is determined by the product of  $\varepsilon_{in}$  and the percentage change in labor effort  $\hat{\ell}/\ell$ .

### 2.3 Optimal Top Tax Rate on Total Earnings

We now proceed to characterizing the optimal non-linear taxes in our setting. In the main body of the paper, we focus on the problem of finding the optimal tax rate in the highest bracket, and we relegate the analysis of the full optimal taxation problem to Appendix B.2. Determining the optimal top tax rate is an especially salient policy question: A large share of the rise in inequality since the 1980s has been driven by the explosion of performance-based forms of compensation for the highest income earners, such as bankers’ bonuses and CEOs’ stock options (see, e.g., Bell and Van Reenen 2013; Lemieux et al. 2009; Piketty and Saez 2003).<sup>30</sup>

We consider the following policy experiment. Fixing the threshold  $z^*$  of the top tax bracket, we optimally set the marginal tax rate  $\tau$  that applies to this bracket. Before stating our main result, we introduce some notation.

**Top and Intermediate Earners.** There are two kinds of workers who react to changes in the top-bracket tax rate. Some (the most productive) of them receive an income in the top tax bracket regardless of their performance: We refer to them as “top earners.” Others earn base pay that is below the top bracket threshold  $z^*$ , but high-level pay that is above that threshold. Because they have a positive probability of landing in the top bracket, their entire contract—including their base salary—adjusts in response to the tax change. We refer to this group as “intermediate earners.” We

---

<sup>29</sup>For instance, if the utility of consumption is CRRA with risk aversion  $\sigma$  and the tax schedule is CRP, we have  $\varepsilon_{out} = (1 - \ell)/[1 - \ell + \ell e^{(p+(1-p)\sigma)\beta}]$ , where  $\beta$  is the bonus rate defined in Section 1.

<sup>30</sup>Bell and Van Reenen (2014) show that in the UK, 83 percent of workers in the top percentile received a bonus in 2008, and that bonuses represented 35 percent of their total compensation (44 percent in the financial sector).

denote by  $\sigma^I$  the fraction of intermediate workers among all (top and intermediate) workers, and by  $s^I$  the income share of intermediate earners in the top bracket.

**Income distribution.** Since earnings of top earners always fall into the top bracket, it is sufficient to keep track of their mean earnings, which we denote by  $Z^T$ . By contrast, for intermediate earners, we need to keep track of the level of their base pay  $\underline{z}$ , which falls short of  $z^*$ , as well as the endogenous frequency with which it (rather than the high-level pay  $\underline{z} + b$ ) is realized. To do so, it suffices to define the mean frequency-adjusted base pay  $(1 - \ell)\underline{z}$  among intermediate earners, which we denote by  $\underline{Z}^I$ . Responses of  $\underline{Z}^I$  to changes in the top tax rate account for both the level responses  $(1 - \ell)\hat{\underline{z}}$  and the frequency responses  $-\hat{\ell}\underline{z}$ . Analogously, we denote the mean frequency-adjusted high-level pay  $\ell\bar{z}$  of intermediate earners by  $\bar{Z}^I$ , and their mean labor effort  $\ell$  by  $L^I$ . Aggregating over intermediate and top earners in the top bracket yields  $Z^* \equiv (1 - \sigma^I)Z^T + \sigma^I\bar{Z}^I$ . Finally, we introduce the Pareto coefficient of the income distribution in the top bracket  $\rho = \frac{E[z|z \geq z^*]}{E[z|z \geq z^*] - z^*}$ , where  $E$  denotes the expectation with respect to the realized (i.e., observed) earnings distribution.

**Elasticities.** Consider the tax reform that raises the marginal tax rate  $\tau$  by  $\hat{\tau} > 0$  above the income level  $z^*$ . To keep track of the impact of this tax reform on various aggregates, we define the elasticity of any variable  $x$  (e.g.,  $L^I$ ,  $\underline{Z}^I$ ,  $Z^*$ ) with respect to the top net-of-tax rate  $1 - \tau$  as  $e_x \equiv -\frac{1-\tau}{\hat{\tau}} \frac{\hat{x}}{x}$ . This shorthand notation measures the total percentage change in  $x$  caused by the reform  $\hat{\tau}$ .<sup>31</sup> These elasticities can be easily expressed in terms of primitives using Lemma 7 and Proposition 5.

**Welfare Weights.** We defined the modified marginal welfare weights  $\tilde{g}$  in Corollary 1. We denote the income-weighted average modified welfare weight in the top bracket by  $\tilde{G}$ .

**Optimum Top Tax Rate.** We now proceed to characterizing the optimal tax rate on total earnings in the top bracket  $\tau$ . Suppose for simplicity that the second highest tax bracket is large enough so that all the intermediate workers pay the same marginal tax rate  $t$  on their base pay.

---

<sup>31</sup>The variable  $e_x$  is thus defined as a “policy elasticity,” a concept proposed by [Hendren \(2016\)](#).

**Theorem 2** *The optimal tax rate on total earnings in the top bracket is given by*

$$\frac{\tau}{1-\tau} = \frac{1 - \tilde{\mathcal{G}} - s^I \left( \frac{t}{1-\tau} \kappa_1 e_{\underline{Z}^I} + \frac{t-\tau}{1-\tau} \kappa_2 e_{L^I} \right)}{\rho e_{Z^*}}, \quad (27)$$

where  $\kappa_1 = (\underline{Z}^I / \bar{Z}^I) \rho^I > 0$ ,  $\kappa_2 = \rho^I - 1 > 0$ , and  $\rho^I = \frac{\bar{Z}^I}{\bar{Z}^I - L^I z^*}$ . The elasticity  $e_{\underline{Z}^I}$  can in turn be decomposed as

$$\kappa_1 e_{\underline{Z}^I} = \overline{\varepsilon_{out}^I} - \rho^I \overline{\varepsilon_{in} e_{\ell}^I} + \kappa_1 \overline{e_{1-\ell}^I}, \quad (28)$$

where  $\overline{\varepsilon_{out}^I}$ ,  $\overline{\varepsilon_{in} e_{\ell}^I}$ ,  $\overline{e_{1-\ell}^I}$  are income-weighted averages (formally defined in the Appendix) of the crowd-out, crowd-in and effort elasticities among intermediate earners, respectively.

The remainder of this section is devoted to analyzing this formula.

**Full-Insurance Benchmark.** Consider first the setting where firms provide full insurance against output risk. The optimal top tax rate is then determined by the following formula (Saez, 2001):

$$\frac{\tau}{1-\tau} = \frac{1 - \mathcal{G}}{\rho e_{Z^*}}, \quad (29)$$

where  $\mathcal{G}$  is the aggregate marginal social welfare weights (see equation (22)) in the top bracket. This expression shows that the optimal tax rate is an increasing function of income inequality measured by the inverse Pareto coefficient  $1/\rho$ , a decreasing function of the average elasticity of top earnings  $e_{Z^*}$  that captures the efficiency cost of raising the tax rate, and a decreasing function of the average social welfare weight  $\mathcal{G}$  that the social objective assigns to agents with earnings  $z \geq z^*$ .

**Welfare Effect of Crowd-Out.** The first novelty introduced by the endogeneity of private insurance is the adjustment to the marginal social welfare weights, which accounts for the first-order welfare effects of crowd-out. As shown in Corollary 1, these responses are captured by placing lower weights on the (unlucky) workers who do not receive a bonus and higher weights on the (lucky) workers who do. Suppose for simplicity that the top bracket is only composed of the lucky intermediate earners. In this case, the average social marginal welfare weight is unambiguously higher than in

the standard model, i.e.,  $\tilde{\mathcal{G}} > \mathcal{G}$ .<sup>32</sup> Thus, the crowd-out reduces the welfare benefits of raising taxes on top incomes and contributes to a strictly lower optimal top tax rate. In the Appendix, we show that if the utility function is logarithmic, this result also holds in the presence of top earners who receive both their base pay and their high-level pay in the top bracket.<sup>33</sup>

**Fiscal Externality: Base-Pay of Intermediate Workers.** Second, applying the benchmark full-insurance formula (29) in our setting would mistakenly ignore that perturbing the top tax rate spills over to workers who could have, but haven't, made it to the top. Indeed, some agents—the unlucky intermediate earners—react to the top tax rate change even though their income is lower than the threshold  $z^*$ . This fiscal externality is captured by the elasticity  $e_{z^I}$  in formula (27). This elasticity can in turn be decomposed into the average crowd-out parameter  $\overline{\varepsilon}_{out}^I$ , the average product of the crowd-in parameter and the individual labor effort elasticity  $\overline{\varepsilon}_{in} e_{\ell}^I$ , and the average elasticity of the base pay frequency  $\overline{e}_{1-\ell}^I$ , as shown in expression (28). Intuitively, the crowd-out and crowd-in forces determine the level response of base pay, while the effort elasticity captures the change in the frequency of base pay realization. If the resulting elasticity  $e_{z^I}$  is negative, say, then a higher top tax rate leads to a higher average base pay for intermediate workers, which increases the tax revenue collected from the lower tax bracket. In that case, this fiscal externality contributes to a higher optimal top tax rate.

**Fiscal Externality: Frequency Margin with a Nonlinear Tax Schedule.** Finally, we need to correct the calculation of the fiscal impact of frequency responses of intermediate earners—i.e., changes in their probability of receiving a bonus due to a change in the top tax rate—when the tax schedule is nonlinear at the top (i.e.,  $\tau \neq t$ ). The elasticities  $e_{z^*}$  and  $e_{z^I}$  already incorporate the frequency responses, but the corresponding fiscal externality terms in formula 27 implicitly assume that both income levels are taxed at the same rate ( $\tau$  or  $t$ , respectively). However, the fiscal impact of these responses is proportional to the difference in tax liabilities between the two states, which is equal to neither  $\tau$  nor  $t$  but to a convex combination of the

<sup>32</sup>This argument takes as a given the marginal value of public funds  $\lambda$ .

<sup>33</sup>More generally, this result holds as long as the utility losses caused by raising the top tax rate,  $(z - z^*)u'(R(z))$ , are increasing with income.

two.<sup>34</sup> The last term in the numerator of (27) makes this necessary correction; This is the same adjustment that led to the term  $\kappa_3$  in formula (13) in Section 1.3.

**Taking stock.** To sum up, the endogeneity of private insurance modifies the standard formula (29) in three ways: The marginal social welfare weights are adjusted to account for the welfare consequences of crowd-out; Two new fiscal externalities arise from the endogenous responses of base pay in the lower bracket and from the frequency responses when the tax schedule is nonlinear at the top. We explore the importance of these effects quantitatively in Section 2.5.

## 2.4 Separate Taxation of Fixed and Variable Pay

So far, we have focused on tax schedules that are functions of total earnings, i.e., the sum of the base pay and, if any, the bonus. This corresponds to the tax treatment of bonuses in most countries. It is conceivable, however, that the tax system could treat the two types of pay differently. In the U.S., for instance, CEO compensation in excess of \$1 million used to be deductible from the company’s corporate income tax only if it was performance-based (Hall & Liebman, 2000).<sup>35</sup> Thus, top performance-based earnings used to have a tax advantage.<sup>36</sup>

In this section we investigate whether there are gains from separating the tax treatment of base earnings and bonuses. We show that it is indeed optimal to tax bonuses at a lower rate than base pay. This provides an explanation for the past tax practice in the U.S. We then characterize analytically the optimal top tax rate on bonuses.

**Local Gains From Separating Taxes on Fixed and Variable Pay.** Suppose that base earnings and bonuses are initially taxed jointly, as in the previous section. For clarity of exposition, we assume that the utility function is logarithmic and the tax schedule is initially linear with rate  $\tau > 0$ , but the main insight holds more

---

<sup>34</sup>To be precise, for intermediate workers we have  $\hat{\ell}(T(\bar{z}) - T(\underline{z})) = \hat{\ell}(\tau(\bar{z} - z^*) + t(z^* - \underline{z}))$ .

<sup>35</sup>Preferential treatment of performance-based compensation ended with the 2017 Tax Cuts and Jobs Act.

<sup>36</sup>One may be concerned that taxing bonuses at a lower effective rate will lead to an effort to disguise base pay as bonuses. This is precisely the reason why performance pay of top CEOs needed to be qualified (i.e., based on objective targets and administered by an external committee) to remain deductible (Rose & Wolfram, 2002). Such a qualification is costly and, arguably, not practical beyond top earners.

generally. Consider the following policy experiment. Raise the marginal tax rate on base pay from  $\tau$  to  $\tau + \delta_z$  and simultaneously lower the marginal tax rate on bonuses from  $\tau$  to  $\tau - \delta_b$ , with  $\delta_z, \delta_b > 0$ , thus effectively decoupling their tax treatment. We engineer this perturbation in such a way that every worker’s expected utility remains constant, which is the case when  $\delta_z = \frac{\ell b}{z} \delta_b$  for all agents.<sup>37</sup>

**Proposition 4** *Suppose that the utility is logarithmic and the tax schedule is initially a linear function of total earnings with a positive tax rate. Lowering the marginal tax rate on bonuses by  $-\delta_b < 0$  and raising the marginal tax rate on base pay by  $\delta_z = \frac{\ell b}{z} \delta_b > 0$  yields a strict improvement in social welfare.*

To understand this result, note first that by Lemma 9 in the Appendix, the tax reform keeps expected utility unchanged,  $\hat{U}(\theta) = 0$  for all  $\theta$ , so that the social objective (11) remains constant. As a result, the tax reform improves social welfare if and only if it yields a strict gain in government revenue.<sup>38</sup> We then show that the tax reform indeed generates higher tax revenue, by raising the effort level of all agents. As we show in Lemma 7, both the reduction in the bonus tax rate and the increase in the base pay tax rate boost labor effort. Therefore, the planner achieves a Pareto improvement by inducing individuals to work harder while keeping their expected utility unchanged, thereby freeing up resources available for redistribution.

The intuition behind this result is as follows: Compensating workers with bonuses for high performance is a *constrained efficient* way of dealing with agency frictions in the labor market. Taxing bonuses removes the only instrument at the firm’s disposal to maintain workers’ incentives for effort. It is, thus, highly distortive. On the other hand, taxing base pay does not have such a large efficiency cost—on the margin it even encourages effort. It is therefore optimal to tax bonuses at a lower rate to minimize distortions, and to tax base pay at a higher rate to provide redistribution.

**Optimal Top Tax Rate on Bonuses.** Our last theoretical result is the characterization of the optimal top tax rate on bonuses for a given tax schedule on base pay.<sup>39</sup>

---

<sup>37</sup>Recall from our analysis of Section 1 that with log utility and a linear tax schedule,  $\ell b/z$  is independent of  $\theta$ .

<sup>38</sup>This tax revenue can, in turn, be rebated to workers to provide a higher level of redistribution.

<sup>39</sup>We refrain from optimizing simultaneously over the top tax rate on bonuses and the top tax rate on base pay for the following reason: Suppose for simplicity that neither bonuses nor base pay



Let  $[b^*, \infty)$  denote the highest bonus tax bracket, and  $\tau_b$  be the corresponding marginal tax rate. Let  $[\underline{z}^*, \infty)$  denote the interval of base earnings of workers whose bonus falls in the top bracket (“top earners”). We suppose for simplicity that  $\underline{z}^*$  is itself in the top bracket of the base pay tax schedule and denote by  $\tau_z$  the corresponding marginal tax rate. Our policy experiment consists of choosing optimally the top bonus tax rate  $\tau_b$ , while keeping the bonus tax schedule  $T_b(\cdot)$  below  $b^*$  and the base pay tax schedule fixed.

Denote the mean frequency-adjusted bonus  $\ell b$  among top earners by  $B$ . Thus, tracking changes in  $B$  accounts for responses to top tax rate changes along two margins: the level of bonuses ( $\hat{\ell}b$ ) and their frequency ( $\hat{\ell}$ ). Second, denote by  $L$  the average probability with which top earners receive a bonus and by  $\underline{Z}$  their average base pay. Finally, the Pareto coefficient of the bonus distribution  $\rho_b$  is given by  $\rho_b = \frac{E[b|b \geq b^*]}{E[b|b \geq b^*] - b^*} = \frac{B/L}{B/L - b^*}$ , where  $E$  denotes the expectation over the (observed) distribution of realized bonuses.

**Theorem 3** *The optimal top tax rate on bonuses is given by*

$$\frac{\tau_b}{1 - \tau_b} = \frac{1 - \tilde{\mathcal{G}} - \frac{\tau_z}{1 - \tau_b} \kappa_1 e_{\underline{Z}} - \frac{t_b - \tau_b}{1 - \tau_b} \kappa_2 e_L}{\rho_b e_B}, \quad (30)$$

where  $t_b = T_b(b^*)/b^*$  is the average bonus tax rate at  $b^*$ ,  $\kappa_1 = (\underline{Z}/B)\rho_b$  and  $\kappa_2 = 1 - \rho_b$ , and the modified marginal social welfare weight  $\tilde{\mathcal{G}}$  is derived formally in the Appendix. The elasticity of mean base earnings  $e_{\underline{Z}}$  can in turn be decomposed as

$$\kappa_1 e_{\underline{Z}} = \overline{\varepsilon_{out}} - \rho_b \overline{\varepsilon_{in} e_{\ell}}, \quad (31)$$

where  $\overline{\varepsilon_{out}}$ ,  $\overline{\varepsilon_{in} e_{\ell}}$  are income-weighted averages of the crowd-in and crowd-out param-

are taxed below the respective thresholds  $b^*$  and  $\underline{z}^*$ , but above these thresholds both forms of pay are taxed with a rate close to 100 percent. In the standard model, taxing top incomes close to 100 percent is never optimal, as it severely distorts the effort of high-ability workers. In the moral-hazard model, however, the effort incentives of high-ability types are driven by  $u(R(\underline{z}, b)) - u(R(\underline{z}, 0)) \approx u(\underline{z}^* + b^*) - u(\underline{z}^*)$ . Thus, effort at the top does not converge to zero, but rather remains at the constant level implied by  $\underline{z}^*$  and  $b^*$ , and the government essentially extracts all the surplus from high-ability workers with almost no distortions. This is reminiscent of [Stantcheva \(2014\)](#), where agency frictions in the labor market (in her case, adverse selection) also allowed the government to redistribute more than in the standard model. We believe that our result, while technically correct, is not policy-relevant, since such heavy taxation of performance-pay contracts would motivate firms to find different ways of incentivizing workers’ effort (e.g., via monitoring), which we do not model explicitly. Fixing the tax schedule on base pay, on the other hand, yields a well-defined and non-corner optimum top tax rate on bonuses, as shown in [Theorem 3](#).

eters, and are defined formally in the Appendix.

It is instructive to compare the optimal bonus tax rate to the benchmark Saez (2001) formula applied to top bonuses:

$$\frac{\tau_b}{1 - \tau_b} = \frac{1 - \mathcal{G}}{\rho_b e_B}. \quad (32)$$

The optimum formula (30) differs from this benchmark in three ways, which are analogous to those we found when studying the optimal joint tax rate on top earnings in Theorem 2.

First, the average marginal social welfare weight at the top is adjusted upwards due to the first-order welfare consequences of crowd-out: We show in the Appendix that, as long as  $t$  is not too high relative to  $\tau_b$ ,  $\tilde{\mathcal{G}} > \mathcal{G}$ . Intuitively, in response to an increase in the bonus tax rate, a share of the extra tax burden is endogenously transmitted to the unlucky workers who have not received a bonus. This lowers the welfare benefits of this tax increase.


Second, a higher tax on bonuses affects the level of base pay according to the elasticity  $e_{\underline{z}}$ . This response can in turn be expressed as the difference between average crowd-out and crowd-in parameters, as shown in equation (31). If, for instance, the crowd-out dominates the crowd-in on average, then  $e_{\underline{z}} > 0$  and a higher tax on bonuses results in lower base pay. In this case, the spillover to the base earnings tax base erodes the tax revenue and reduces the optimal bonus tax rate.

Third, the last term in the numerator of formula (30) is a corrective adjustment to the efficiency cost of taxation due to the frequency margin responses. A reduction in aggregate labor effort by  $e_L > 0$ , and hence in the frequency with which the worker receives a bonus, does not reduce government revenue as much if, rather than being taxed uniformly at rate  $\tau_b$ , bonuses are taxed at a lower average rate up to  $b^*$ . Thus, if  $t_b < \tau_b$ , this correction tends to raise the optimal top tax rate.

## 2.5 Quantitative Analysis

We now adapt the quantitative framework with performance-pay and fixed-pay jobs of Section 1.4 to study the optimal taxation of top incomes and bonuses. We proceed in two steps. First, we evaluate the optimal top tax rate on total earnings, characterized theoretically in Section 2.3. Second, we evaluate the optimal top tax

Figure 4: Optimal Top Tax Rate on Total Income and SCPE



rate on bonuses for a given tax schedule on base earnings, as analyzed in Section 2.4.

**Optimal Top Tax Rate on Total Earnings.** We approximate the empirical income tax schedule with a piecewise linear function with two brackets and a lump-sum transfer. The top bracket starts at  $z^* = \$500,000$  and applies to approximately the top 1 percent of earnings. The top tax rate in the status quo is equal to the marginal tax rate implied by the calibrated CRP tax function at  $z^*$ . We choose the bottom tax rate and the lump-sum transfer by matching the aggregate tax revenue and the average tax rate at  $z^*$  implied by the calibrated CRP tax. We obtain a top tax rate of 49 percent, a bottom tax rate of 41 percent, and a lump-sum transfer of \$16,460. This gives us a well-defined notion of top tax rate that we will choose optimally.

We compute the optimal top tax rate under two social objectives: Rawlsian and utilitarian. The former maximizes the extraction of revenue from top earners, i.e., it sets the average marginal welfare weight at the top  $\tilde{\mathcal{G}}$  to zero in (27). We compare the optimal top rates with the rates obtained if the policymaker ignored the endogeneity of earnings risk (Self-Confirming Policy Equilibrium, SCPE). More specifically, a policymaker in the SCPE ignores the impact of changing the top tax rate on earnings in the lower bracket and calculates the welfare impact of reforms without accounting for the crowd-out of private insurance.<sup>40</sup> The results are depicted in Figure 4.

<sup>40</sup>Here we ignore the third discrepancy between the standard and the optimal tax formula, which comes from the fact that the frequency responses are taxed at a different average tax rate than the level responses (see Theorem 2). Such a policy mistake would arise if the policymaker in the SCPE tracked only the aggregate earnings in the top bracket. Instead, we are implicitly assuming that the policymaker can observe individual earnings and, hence, correctly calculate individual tax burdens.

We find that the top tax rates chosen in the SCPE and in the optimum are almost identical: They differ by less than 1 percentage point. The welfare cost of such a policy mistake is minuscule. Furthermore, and unlike the case of the optimal rate of progressivity analyzed in Section 1.4, we verify that the welfare cost of ignoring endogenous earnings risk would not be much different if all workers had performance-pay contracts. This is because performance-pay contracts are already widely adopted at the top: They account for 76 percent of earnings in the top bracket.

The right panel of Figure 4 shows that the planner who attempts to maximize tax revenue (Rawlsian), but does not account for the impact of the top tax rate on earnings in the lower bracket (SCPE), chooses a top rate that is slightly too low. That is because a higher top rate leads to higher earnings (and, hence, higher tax revenue) in the lower bracket. To understand why this is the case, focus on the intermediate earners whose base pay is in the lower bracket and whose high-level pay is in the top bracket; this group accounts for 90 percent of the performance-pay workers who are affected by the top tax rate. By Proposition 3, a higher top tax rate affects their base pay via two channels: the crowd-out, which reduces it, and the crowd-in, which works in the opposite direction. We quantify these effects in Table 2. While the crowd-out is sizable, and on its own would imply a 4 percent fall in base pay, the crowd-in has an almost equal magnitude and offsets most of this change. As a result, the base pay level falls only by 0.5 percent on average.<sup>41</sup> In addition, the mean frequency at which solely the base pay is paid out increases by 1.7 percent (second column of Table 2). Taking into account these level and frequency changes, we obtain that the overall base earnings of intermediate earners increase, leading to higher tax revenue from spillovers to the lower bracket.

In contrast to the Rawlsian case, the utilitarian policymaker who ignores the endogeneity of private insurance (SCPE) would choose a top tax rate that is marginally too high. Effectively, this planner makes two mistakes: It underestimates the tax revenue from the lower tax bracket (similarly to the Rawlsian planner), and it also underestimates the negative welfare effect on high-income individuals—this second mistake is new. Focus again on the intermediate earners. A naive planner would believe that the tax hike affects only the lucky intermediate workers who receive a

---

<sup>41</sup>There is a third effect at play: The tax reform also affects the relative probabilities with which different intermediate earners receive their base pay; this in turn modifies how the average base pay among these workers is computed in the first column of Table 2. This effect, however, is small relative to the crowd-out and crowd-in (0.1 percent).

Table 2: Impact of top tax rate on base pay of intermediate earners

	level $\underline{z}$	frequency $1 - \pi(\ell)$	overall $\underline{z} \cdot (1 - \pi(\ell))$
Status quo (initial)	\$264,360	0.76	\$202,210
+ crowd-out	-4.0%		
+ crowd-in	+3.4%		
+ change of relative freq.	+0.1%		
Rawlsian optimum (final)	\$263,040 (-0.5%)	0.78 (+1.7%)	\$204,630 (+1.2%)

Note: Top tax rate is increased from 49.3% (status quo) to 63.1% (Rawlsian optimum). All statistics are averages over intermediate earners.

Table 3: Impact of top tax rate on earnings risk

	intermediate earners		top earners	
	bonus rate $\beta$	$Var(\log z   \theta)$	bonus rate $\beta$	$Var(\log z   \theta)$
Status quo (initial)	1.05	0.2	1.1	0.22
+ crowd-out	+9%	+18%	+16%	+35%
+ crowd-in	-18%	-38%	-10%	-26%
Rawlsian optimum (final)	0.95 (-9%)	0.16 (-20%)	1.17 (+6%)	0.24 (+9%)

Note: Top tax rate is increased from 49.3% (status quo) to 63.1% (Rawlsian optimum). All statistics are averages over either intermediate or top earners.

bonus. However, the resulting crowd-out of private insurance contributes to a lower level of base pay and, thus, also negatively affects the unlucky intermediate workers. Although this change in base earnings is mostly offset by the crowd-in effect, its welfare impact remains significant due to the envelope theorem. The negative welfare impact on top earners, whose base pay is in the top bracket, is similarly underestimated. In our quantitative model, the mistake in the welfare effect dominates that in the revenue from the lower bracket, leading to a tax rate that is too high relative to the optimum.

Finally, [Table 3](#) shows how increasing the top tax rate influences the earnings risk of affected workers. Starting from top earners, who earn both their base pay and their high-level pay in the top bracket, we find that the crowd-in offsets most (63–75 percent) of the crowd-out when measured either by the mean bonus rate  $\beta = \log(\bar{z}/\underline{z})$  or by the mean variance of log earnings conditional on ability. As a result, a strong increase in the top tax rate raises the earnings risk of top earners, but much less than the crowd-out alone would suggest.

The picture is quite different for intermediate workers: for them the crowd-in is twice as large as the crowd-out. As a result, earnings risk actually falls after an increase in the top tax rate. To understand this, recall that the crowd-in effect is driven by the endogenous labor effort adjustment. Lemma 7 implies that intermediate earners adjust their labor effort more than do top earners. This is because intermediate earners experience an increase in the marginal and the average tax rate on their high-level pay, both of which reduce effort. In addition, top earners face an increase in the marginal and average tax rates on their base pay, which mitigates their labor effort responses. Quantitatively, the mean earnings of intermediate workers fall on average by 5.6 percent, while those of top workers fall by 3.6 percent, leading to a much more powerful crowd-in for the former group.

These results highlight that (i) the crowd-in is an important force that we need to take into account; and (ii) the relative size of the crowd-out and crowd-in depends on the specifics of the tax reform. In particular, strongly progressive reforms (such as the one experienced by intermediate earners) lead to a stronger effort response and magnify the crowd-in effect.

**Separate Taxation of Top Bonuses.** In Section 2.4 we showed that there are efficiency gains from taxing bonuses at a lower rate than base pay and characterized the optimal top bonus tax. In this section we quantify these results. In particular, we compare our economy in which all earnings are taxed jointly with a single tax schedule to the economy in which top bonuses are taxed with a separate rate that is chosen optimally, taking as given the tax on base pay. For simplicity, we assume that the original joint tax schedule is affine, i.e., it consists of a lump-sum transfer and a constant marginal tax rate.<sup>42</sup> We then suppose that a policymaker can adjust the tax rate on the top 5 percent of bonuses, taking the rest of the tax schedule as given. The results are shown in Table 4.

We find that the optimal bonus tax rate can be substantially lower than the tax on base earnings, and that such a separation can bring sizable welfare benefits, particularly when the initial joint tax rate is high. For instance, consider an initial joint tax rate that is chosen optimally according to a utilitarian or Rawlsian social welfare criterion. In the case of utilitarian social preferences, the optimal tax rate

---

<sup>42</sup>The assumption that the initial joint tax schedule is affine allows us to represent the tax schedule, both before and after the reform, in the additively separable form:  $T(z, b) = T_z(z) + T_b(b)$ , which significantly simplifies the analysis.

Table 4: Optimal Tax on Top Bonuses for a Given Base Pay Tax

joint tax rate (initial)	Utilitarian		Rawlsian	
	bonus tax rate	welfare gain	bonus tax rate	welfare gain
40%	36.8%	0.0%	38.9%	0.0%
60%	42.1%	0.3%	45.1%	0.4%
62%*	42.8%	0.4%	45.8%	0.5%
80%	51.3%	1.6%	54.5%	1.8%
84%**	54.5%	2.2%	57.5%	2.3%

Note: Welfare gain from the optimal separation of the tax on top 5% of bonuses (taking as given the tax on base pay and on bonuses below the top) is expressed as an equivalent percentage change of consumption. \* Optimal Utilitarian joint tax rate. \*\* Optimal Rawlsian joint tax rate.

on top bonuses would then be 19 percentage points lower than that on the base pay (42.9 percent vs. 62 percent), with a welfare gain from separating the tax treatment of the two forms of pay equal to 0.4 percent of consumption. In the case of Rawlsian social preferences, it is optimal to cut the top bonus tax rate relative to the base pay tax rate by 27 percentage points (57.5 percent vs. 84 percent), with a welfare gain of 2.4 percent of consumption. Note that the welfare gains from separating the tax rates increase with the size of the initial tax rate. This is because, for a (sub-optimally) low initial joint tax rate, the efficiency gains from lowering the tax on bonuses are countervailed by the redistributive losses from smaller taxes on the bonuses of the most productive agents.

## Conclusion

We have set up and analyzed a tractable moral-hazard environment in which firms design labor contracts that trade off effort incentives with insurance against performance shocks. The government uses the tax-and-transfer system to redistribute income across workers who differ in uninsurable ex ante ability. The key feature of our model is that earnings risk is endogenous and has a productive role as it motivates labor effort. We found that standard models that ignore the endogeneity of pre-tax earnings risk come very close to accurately evaluating the impact of taxes on labor contracts. We derived optimal tax formulas for the overall rate of tax progressivity, the top tax rate on total earnings, and the top tax rate on bonuses. Perhaps surprisingly,

we showed that it is optimal to tax bonuses at a strictly lower rate than base earnings.

It would be interesting to extend our analysis in several directions. First, we only considered the impact of taxes on compensation for already existing performance-pay jobs. One could also model the incentives of firms to create such performance-pay jobs (rather than monitored jobs) in the first place. Second, in our model, private markets are perfectly competitive and constrained efficient. In other words, we gave private markets their best chance of making government policy redundant. Introducing market power and frictions such as adverse selection in private markets, whereby firms cannot perfectly observe workers' ability are natural next steps. Third, our theoretical analysis delivers predictions regarding the impact of various types of tax reforms on the structure of incentive-based compensation. Testing these predictions empirically should be particularly fruitful.



## References

- Abrahám, A., Alvarez-Parra, F., & Forstner, S. (2016). The effects of moral hazard on wage inequality in a frictional labor market. *Unpublished working paper*.
- Ábrahám, Á., Koehne, S., & Pavoni, N. (2016). Optimal income taxation when asset taxation is limited. *Journal of Public Economics*, 136, 14–29.
- Alvarez, F., & Jermann, U. J. (2000). Efficiency, equilibrium, and asset pricing with risk of default. *Econometrica*, 68(4), 775–797.
- Attanasio, O., & Ríos-Rull, J.-V. (2000). Consumption smoothing in island economies: Can public insurance reduce welfare? *European Economic Review*, 44(7), 1225–1258.
- Bandiera, O., Barankay, I., & Rasul, I. (2005). Social preferences and the response to incentives: Evidence from personnel data. *The Quarterly Journal of Economics*, 120(3), 917–962.
- Bell, B., & Van Reenen, J. (2013). Extreme wage inequality: pay at the very top. *American economic review*, 103(3), 153–57.
- Bell, B., & Van Reenen, J. (2014). Bankers and their bonuses. *The Economic Journal*, 124(574), F1–F21.
- Bénabou, R., & Tirole, J. (2016). Bonus culture: Competitive pay, screening, and multitasking. *Journal of Political Economy*, 124(2), 305–370.
- Besley, T., & Ghatak, M. (2013). Bailouts and the optimal taxation of bonus pay. *American Economic Review*, 103(3), 163–67.
- Boadway, R., & Sato, M. (2015). Optimal income taxation with risky earnings: A synthesis. *Journal of Public Economic Theory*, 17(6), 773–801.
- Brown, C., & Andrabi, T. (2020). Inducing positive sorting through performance pay: Experimental evidence from pakistani schools. *University of California at Berkeley Working Paper*.
- Chang, Y., & Park, Y. (2017). *Optimal taxation with private insurance* (Tech. Rep.).
- Chetty, R., Guren, A., Manoli, D., & Weber, A. (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review*, 101(3), 471–75.
- Chetty, R., & Saez, E. (2010). Optimal taxation and social insurance with endogenous private insurance. *American Economic Journal. Economic Policy*, 2(2), 85.
- Colombi, R. (1990). A new model of income distribution: The pareto-lognormal

- distribution. In *Income and wealth distribution, inequality and poverty* (pp. 18–32). Springer.
- Craig, A. C. (2019). Optimal income taxation with spillovers from employer learning.
- Cremer, H., & Pestieau, P. (1996). Redistributive taxation and social insurance. *International Tax and Public Finance*, 3(3), 281–295.
- Cullen, J. B., & Gruber, J. (2000). Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Economics*, 18(3), 546–572.
- Cutler, D. M., & Gruber, J. (1996a). Does public insurance crowd out private insurance? *The Quarterly Journal of Economics*, 111(2), 391–430.
- Cutler, D. M., & Gruber, J. (1996b). The effect of medicaid expansions on public insurance, private insurance, and redistribution. *The American Economic Review*, 86(2), 378–383.
- Doligalski, P. (2019). Optimal income taxation and commitment on the labor market.
- Doligalski, P., Ndiaye, A., & Werquin, N. (2020). Redistribution with performance pay. Available at SSRN 3581882.
- Eaton, J., & Rosen, H. S. (1980). Optimal redistributive taxation and uncertainty. *The Quarterly Journal of Economics*, 95(2), 357–364.
- Edmans, A., & Gabaix, X. (2011). Tractability in incentive contracting. *The Review of Financial Studies*, 24(9), 2865–2894.
- Edmans, A., Gabaix, X., Sadzik, T., & Sannikov, Y. (2012). Dynamic ceo compensation. *The Journal of Finance*, 67(5), 1603–1647.
- Ferey, A., Haufler, A., & Perroni, C. (2022). Incentives, globalization, and redistribution. *Mimeo*.
- Findeisen, S., & Sachs, D. (2016). Education and optimal dynamic taxation: The role of income-contingent student loans. *Journal of Public Economics*, 138, 1–21.
- Frydman, C., & Molloy, R. S. (2011). Does tax policy affect executive compensation? evidence from postwar tax reforms. *Journal of Public Economics*, 95(11-12), 1425–1437.
- Gietl, D., & Haufler, A. (2018). Bonus taxes and international competition for bank managers. *European Economic Review*, 110, 41–60.
- Gittleman, M., & Pierce, B. (2013). How prevalent is performance-related pay in the united states? current incidence and recent trends. *National Institute Economic Review*, 226, R4–R16.
- Golosov, M., Kocherlakota, N., & Tsyvinski, A. (2003). Optimal indirect and capital

- taxation. *Review of Economic Studies*, 569–587.
- Golosov, M., & Tsyvinski, A. (2007). Optimal taxation with endogenous insurance markets. *The Quarterly Journal of Economics*, 122(2), 487–534.
- Grigsby, J., Hurst, E., & Yildirmaz, A. (2019). *Aggregate nominal wage adjustments: New evidence from administrative payroll data* (Tech. Rep.). National Bureau of Economic Research.
- Hall, B. J., & Liebman, J. B. (2000). The taxation of executive compensation. *Tax policy and the economy*, 14, 1–44.
- Haufler, A., & Nishimura, Y. (2022). Taxing mobile and overconfident top earners. *International Tax and Public Finance*, Forthcoming.
- Heathcote, J., Storesletten, K., & Violante, G. L. (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4), 1693–1754.
- Heathcote, J., & Tsujiyama, H. (2021). Optimal income taxation: Mirrlees meets ramsey. *Journal of Political Economy*, 129(11), 3141–3184.
- Hendren, N. (2016). The policy elasticity. *Tax Policy and the Economy*, 30(1), 51–89.
- Holmstrom, B. (1984). *A note on the first-order approach in principal agent models* (mimeo).
- Holmstrom, B., & Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica: Journal of the Econometric Society*, 303–328.
- Kapička, M., & Neira, J. (2019). Optimal taxation with risky human capital. *American Economic Journal: Macroeconomics*.
- Kaplow, L. (1991). Incentives and government relief for risk. *Journal of risk and Uncertainty*, 4(2), 167–175.
- Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature*, 49(4), 961–1075.
- Krueger, D., & Perri, F. (2011). Public versus private risk sharing. *Journal of Economic Theory*, 146(3), 920–956.
- Lazear, E. P. (2000). Performance pay and productivity. *American Economic Review*, 90(5), 1346–1361.
- Lazear, E. P. (2018). Compensation and incentives in the workplace. *Journal of Economic Perspectives*, 32(3), 195–214.

- Leaver, C., Ozier, O., Serneels, P., & Zeitlin, A. (2021). Recruitment, effort, and retention effects of performance contracts for civil servants: Experimental evidence from rwandan primary schools. *American economic review*, *111*(7), 2213–46.
- Lehmann, E., Parmentier, A., & Van der Linden, B. (2011). Optimal income taxation with endogenous participation and search unemployment. *Journal of Public Economics*, *95*(11-12), 1523–1537.
- Lemieux, T., MacLeod, W. B., & Parent, D. (2009). Performance pay and wage inequality. *The Quarterly Journal of Economics*, *124*(1), 1–49.
- Levitt, S. D., & Syverson, C. (2008). Market distortions when agents are better informed: The value of information in real estate transactions. *The Review of Economics and Statistics*, *90*(4), 599–611.
- Makris, M., & Pavan, A. (2021). Taxation under learning by doing. *Journal of Political Economy*, *129*(6), 1878–1944.
- Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The review of economic studies*, 175–208.
- Miyazaki, H. (1977). The rat race and internal labor markets. *The Bell Journal of Economics*, 394–418.
- Netzer, N., & Scheuer, F. (2007). Taxation, insurance, and precautionary labor. *Journal of Public Economics*, *91*(7-8), 1519–1531.
- Park, Y. (2014). Optimal taxation in a limited commitment economy. *Review of Economic Studies*, *81*(2), 884–918.
- Piketty, T., & Saez, E. (2003). Income inequality in the united states, 1913–1998. *The Quarterly journal of economics*, *118*(1), 1–41.
- Raj, A. (2019). *Optimal income taxation in the presence of networks of altruism* (Tech. Rep.). Working Paper.
- Rogerson, W. P. (1985). The first-order approach to principal-agent problems. *Econometrica: Journal of the Econometric Society*, 1357–1367.
- Rose, N. L., & Wolfram, C. (2002). Regulating executive pay: Using the tax code to influence chief executive officer compensation. *Journal of Labor Economics*, *20*(S2), S138–S175.
- Rothschild, C., & Scheuer, F. (2016). Optimal taxation with rent-seeking. *The Review of Economic Studies*, *83*(3), 1225–1262.
- Rothschild, M., & Stiglitz, J. (1976). Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of*

- Economics*, 90(4), 629–649.
- Sachs, D., Tsyvinski, A., & Werquin, N. (2020). Nonlinear tax incidence and optimal taxation in general equilibrium. *Econometrica*, 88(2), 469–493.
- Saez, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies*, 68(1), 205–29.
- Saez, E., & Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(1), 24–45.
- Sannikov, Y. (2008). A continuous-time version of the principal-agent problem. *The Review of Economic Studies*, 75(3), 957–984.
- Scheuer, F. (2013). Adverse selection in credit markets and regressive profit taxation. *Journal of Economic Theory*, 148(4), 1333–1360.
- Schoeni, R. F. (2002). Does unemployment insurance displace familial assistance? *Public Choice*, 110(1), 99–119.
- Shearer, B. (2004). Piece rates, fixed wages and incentives: Evidence from a field experiment. *The Review of Economic Studies*, 71(2), 513–534.
- Shourideh, A. (2014). Optimal taxation of wealthy individuals. *Working paper*.
- Sleet, C., & Yazici, H. (2017). Taxation, redistribution and frictional labor supply. In *2017 meeting papers 298 (society for economic dynamics, 2017)*.
- Spence, M. (1978). Product differentiation and performance in insurance markets. *Journal of Public Economics*, 10(3), 427–447.
- Stantcheva, S. (2014). Optimal income taxation with adverse selection in the labour market. *The Review of Economic Studies*, rdu005.
- Stantcheva, S. (2017). Optimal taxation and human capital policies over the life cycle. *Journal of Political Economy*, 125(6), 1931–1990.
- Thanassoulis, J. (2012). The case for intervening in bankers’ pay. *The Journal of Finance*, 67(3), 849–895.
- Varian, H. R. (1980). Redistributive taxation as social insurance. *Journal of public Economics*, 14(1), 49–68.
- Wilson, C. (1977). A model of insurance markets with incomplete information. *Journal of Economic Theory*, 16(2), 167–207.