

Redistribution with Performance Pay

Paweł Doligalski

University of Bristol and Group for Research in Applied Economics

Abdoulaye Ndiaye

New York University, Stern School of Business, and Centre for Economic Policy Research

Nicolas Werquin

Federal Reserve Bank of Chicago, Toulouse School of Economics, and Centre for Economic Policy Research

Half of the jobs in the United States feature pay for performance. We derive incidence and optimum formulas for the rate of tax progressivity and the top income tax rate when such labor contracts arise from moral hazard frictions within firms. Our first main result is that the sensitivity of the worker's compensation to performance is roughly invariant to tax progressivity. Second, the optimal tax schedule is strictly less progressive than in standard models that treat pretax earnings risk as exogenous. Quantitatively, the welfare cost of not accounting for performance pay when choosing tax progressivity is 0.3% of consumption.

Introduction

The considerable increase in income inequality observed since the 1980s is in large part due to the explosion of performance-based forms of

Opinions expressed in this article are ours and do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System. Nicolas Werquin acknowledges support from Agence Nationale de la Recherche (ANR) under grant ANR-17-EURE-0010 (Investissements d'Avenir program). We thank Árpád Ábrahám, Anmol Bhandari,

Electronically published May 10, 2023

Journal of Political Economy Macroeconomics, volume 1, number 2, June 2023.

© 2023 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/724511>

remuneration at the top of the income distribution, such as the rise in bankers' bonuses or CEOs' stock options. 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. Lemieux, MacLeod, and Parent (2009) estimate that in the United States almost half of all jobs—and three-quarters within the top percentile of the income distribution—involve performance-based 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-pay contracts? In particular, how do taxes affect the pay-performance sensitivity—that is, the relationship between a worker's performance and their compensation? How should the overall progressivity of the income tax schedule and the tax rate on top earners be designed to account for the existence of performance pay?

This paper provides answers to these questions. 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: *ex ante* ability differences, which cannot be insured by firms, and *ex post* performance (or output) shocks. Specifically, as in Rogerson (1985), the output of each worker is stochastic and either high or low. Workers can increase the probability of high output by exerting more (unobservable and costly) effort. Firms can observe workers' ability and their realized output but not their labor effort. As a result, they design a contract that offers a base salary plus a bonus when realized output is high. Such a contract provides only partial insurance against output risk: while providing better insurance lowers the firm's cost of giving workers their reservation utility, it also reduces their incentives to exert effort.¹ The government

Joydeep Bhattacharya, Katherine Carey, Ashley Craig, Alex Edmans, Emmanuel Farhi, Antoine Ferey, Xavier Gabaix, Daniel Garrett, Mike Golosov, Andreas Haufler, Greg Kaplan, 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 el ene 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. Code replicating the tables and figures in this article can be found in the Harvard Dataverse: <https://doi.org/10.7910/DVN/CMBSPF>. This paper was edited by Anmol Bhandari.

¹ 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 microfounded by models of adverse selection, in which incentive pay is offered to attract workers with higher unobserved ability. Lazear (2000), Brown and Andrabi (2020), and Leaver et al. (2021) find empirical support for both moral hazard and adverse selection, the former accounting for at least half of the overall effect of performance pay on productivity.

uses the income tax to redistribute between workers with different *ex ante* abilities, taking into account the fact that the amount of earnings risk embedded in the labor contracts—the size and the frequency of receiving the bonus—is endogenous to the tax system.

We first take a positive standpoint and evaluate the incidence of nonlinear taxes on the design of performance-pay contracts. An increase in the rate of tax progressivity causes two offsetting effects on the sensitivity of earnings to performance. First, it leads to a standard crowding out of the private insurance provided by the firm: in response to an improvement in social insurance via tax progressivity, firms endogenously raise the pay-performance sensitivity (i.e., spread out pretax earnings), so that the worker's incentives for effort remain unchanged.² 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 turn, eliciting a lower effort level allows the firm to improve insurance by lowering the sensitivity of pretax earnings to performance. This earnings compression counteracts the direct crowd-out mechanism.

The first main result of our paper is that while taken separately these effects are both significant, they are of the same order of magnitude. Thus, summing them implies that taxes barely affect the sensitivity of (pretax) earnings to performance. Furthermore, we show that this insight is robust to alternative forms of performance pay: it continues to hold when optimal contracts are a linear function of performance, such as piece rates or commissions (Holmstrom and Milgrom 1987); when they are convex, such as stock options (Edmans and Gabaix 2011); or when incentives are provided over time (Edmans et al. 2012). This 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).

Next, we take a normative standpoint. The second main result of our paper is that even though the combined effect of the crowd out and crowd in leaves the labor contract almost unaffected by taxes, the endogeneity of earnings risk reduces the optimal level of progressivity of the tax schedule. We derive two variations of this insight: we provide analytical formulas for (i) the optimal rate of progressivity in a parametric class of tax functions and (ii) the optimal tax rate on top income earners. In both cases, the optimum is strictly lower than in the benchmark setting with exogenous private insurance—as long as the social welfare weight placed on top earners is positive in the second case.

² Theoretically, the crowd out of private insurance has been shown to severely limit the ability of governments to provide social insurance—see Attanasio and Ríos-Rull (2000), Golosov and Tsyvinski (2007), and Krueger and Perri (2011). Empirically, such crowding out has been observed in health insurance (Cutler and Gruber 1996a, 1996b; Schoeni 2002) and unemployment insurance (Cullen and Gruber 2000).

These results follow from the fact that, by the envelope theorem, the crowd-in channel affects workers' utility only up to a second order. The crowd out, on the other hand, has first-order negative welfare consequences. As a result, any insurance gain from higher progressivity comes at the welfare cost of crowding out private insurance. This eliminates the role of the income tax as an insurance device against output shocks. Moreover, we show that tax cuts are less accurately targeted toward high-marginal-utility agents than in models where private insurance is exogenous. Indeed, the need to preserve incentives requires an adjustment to gross earnings in favor of the high performers of a given ability type, who therefore command a disproportionately large share of the tax cut. This regressive distribution of rents within firms further dampens the welfare benefits of redistributive taxation.

To complement this theoretical analysis, we calibrate the model to US data. Our model features both performance-pay and fixed-pay jobs and accurately matches the first and second moments of the earnings distribution within each job category, as well as 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% of the crowd out in terms of the ratio of bonus to base pay and more than 100% 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. Focusing first on the optimal rate of tax progressivity, we find that a utilitarian planner would choose a progressivity rate of 0.38, which is more than twice as high as the current progressivity in the United States. Furthermore, a “naive” planner who ignores the endogeneity of earnings risk would choose an even higher progressivity rate of 0.43. The welfare cost of such a policy mistake is equivalent to a 0.3% fall in consumption. We show that these results crucially depend on the share of performance-pay jobs in the economy: if all jobs in the United States featured performance pay—rather than only half, as is currently the case—then the welfare cost of ignoring endogenous earnings risk would more than quadruple. Second, we find that the optimal tax rate on top earners is lower than the one chosen by the naive planner, who ignores these workers' utility losses due to the crowd out of private insurance, and that the difference between the two increases in their weight in the social welfare objective. If top earners have a relatively low marginal social welfare weight (relative to the average worker), then the policy mistake of the naive planner will be small as well. Performance pay can justify much lower tax rates at the very top but only when society places a substantial weight on the welfare of the highest earners.

Finally, we examine the sensitivity of our main results to the chosen parameter values and data moments, including the Frisch elasticity of labor effort, the variance of log earnings in performance-pay jobs, and the empirical probability with which performance-pay workers receive bonuses. Our main results are robust to these alternative parameter values. In all the alternative calibrations we consider, the crowd in of private insurance always offsets most (85%–98%) of the crowd out, and the rate of progressivity chosen when ignoring the endogeneity of earnings risk is higher than the optimal rate by 0.03–0.075, relative to the baseline value of 0.05.

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. Stantcheva (2014) focuses on the adverse selection model of the labor market of Miyazaki (1977). Scheuer (2013) analyzes a choice between payroll employment and entrepreneurship that 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. Empirically, there is growing reduced-form and structural evidence that moral hazard in labor markets is pervasive (Foster and Rosenzweig 1994; Prendergast 1999; Shearer 2004; Lazear and Oyer 2010; Bandiera, Barankay, and Rasul 2011; Ábrahám, Alvarez-Parra, and Forstner 2016), that employers are important providers of insurance for their employees (Guiso, Pistaferri, and Schivardi 2005; Lamadon 2016; Friedrich et al. 2019; Lamadon, Mogstad, and Setzler 2019), and that the fraction of jobs with explicit pay for performance is high and rising (Lemieux, MacLeod, and Parent 2009; Bloom and Van Reenen 2010; Bell and Van Reenen 2014; Grigsby, Hurst, and Yildirmaz 2019; Makridis and Gittleman 2022).

Ferey, Haufler, and Perroni (2022) also study the effects of tax policy in a setting with performance-pay contracts due to moral hazard. While we allow for a continuous effort choice, they 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.³ Chetty and Saez (2010) derive a reduced-form sufficient 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 analytically characterize the effects of nonlinear government policy on private insurance contracts via crowd-out and crowd-in responses and

³ 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.

derive explicit theoretical formulas for the tax incidence and optimal taxes in terms of underlying structural parameters. Kaplow (1991) shows that the optimal linear tax in a setting with endogenous private insurance constrained by moral hazard is zero. This result holds in our setting as well, but we also allow for ex ante heterogeneous agents, which gives the government a redistributive role.

An important strand of papers studies income taxation in the presence of endogenous consumption insurance, which can take the form of private insurance markets (Cremer and Pestieau 1996; Netzer and Scheuer 2007), asset trades (Park 2014; Ábrahám, Koehne, and Pavoni 2016; Chang and Park 2017), or informal exchanges in family networks (Attanasio and Ríos-Rull 2000; Krueger and Perri 2011; Heathcote, Storesletten, and Violante 2017; Raj 2019). In contrast to these papers, it is pretax 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 nonlinear, while changes in consumption risk with exogenous wages do not have a direct fiscal impact.⁴ Other papers endogenize earnings risk by focusing on human capital accumulation (Findeisen and Sachs 2016; Stantcheva 2017; Craig 2019; Kapička and Neira 2019; Makris and Pavan 2021), 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., Eaton and Rosen 1980; Varian 1980; Boadway and Sato 2015). 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.

The paper is organized as follows. In section I, we set up our baseline environment. Section II contains our analysis of the incidence of tax progressivity on the design of labor contracts. In section III, we characterize the optimal rate of progressivity and the optimal tax rate of top earners. Section IV evaluates our findings quantitatively. The proofs and additional results are collected in the appendix.

I. Environment

There is a continuum of mass one of agents indexed by their exogenous ability $\theta \in \mathbb{R}_+$ distributed according to the cumulative distribution function F and density f . Preferences over consumption c and labor effort ℓ are represented by the utility function $\log(c) - h(\ell)$, where h is twice continuously differentiable and strictly convex.

⁴ Naturally, there can be indirect effects from consumption insurance to tax revenue through precautionary labor supply, as in Netzer and Scheuer (2007).

A worker with ability θ who provides effort ℓ can produce two levels of output: either θ with probability $\pi(\ell)$ or zero with probability $1 - \pi(\ell)$, 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 $\pi(\ell) = \ell \in [0, 1]$.⁵

Firms observe both the agent’s ability and their realized output but not their effort. The labor contract thus specifies two values for earnings z : a base pay \underline{z} if output is low and a high-level pay \bar{z} if output is high, with $\bar{z} \geq \underline{z}$. We denote the bonus by $b = \bar{z} - \underline{z}$ and the pass-through of output risk to log earnings by $\beta \equiv \log(\bar{z}/\underline{z})$. The parameter β is our measure of the performance sensitivity of pay.

The government levies nonlinear income taxes. A worker with earnings z consumes their after-tax earnings $c = R(z)$, where $R : \mathbb{R}_+ \rightarrow \mathbb{R}$ is the retention function. Throughout the paper, we denote the utility over pretax earnings by $v(z) \equiv \log(R(z))$ and assume that v is concave. This restriction imposes that the tax schedule is not too regressive;⁶ in particular, it always holds under the following class of tax schedules, on which some of our results rely.

DEFINITION 1 (CRP tax schedule). The tax schedule has a constant rate of progressivity (CRP) if there exist $\tau \in \mathbb{R}$ and $p \in (-\infty, 1)$ such that $R(z) = ((1 - \tau)/(1 - p))z^{1-p}$.⁷

Labor contract.—A firm that hires a worker with ability θ 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 maximize its expected profit:

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

where $\mathbb{E}[z|\theta] = (1 - \ell)\underline{z} + \ell\bar{z}$ denotes agent θ ’s expected earnings under the contract, subject to the incentive constraint

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

and the participation constraint

⁵ The disutility of effort must then be renormalized 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 and π .

⁶ See the appendix for details. 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.

⁷ The CRP tax code is a good approximation of the US tax system; see, e.g., Heathcote, Storesletten, and Violante (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 $R(z)/z$.

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

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

We assume that there is free entry of firms in the labor market. Thus, the equilibrium reservation value $U(\theta)$ is such that

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

As long as the equilibrium effort level is interior, the first-order condition of the maximization problem (2) is necessary and sufficient to ensure global incentive compatibility:⁸

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

or equivalently, $R(\bar{z})/R(\underline{z}) = \exp(h'(\ell))$. This equation plays an important role in our analysis. Intuitively, inducing a worker to provide a given level of labor supply requires promising a larger reward for performance if the marginal disutility of effort is higher. Since h is convex, this implies that eliciting a higher effort from workers in the presence of moral hazard can be achieved only by raising their exposure to output risk. In particular, when the tax schedule is CRP, it immediately implies that the pass-through of output risk to log earnings is given by $\beta = h'(\ell)/(1 - p)$.

In the following lemma, we derive the equilibrium labor contract between a firm and a worker with ability θ . For simplicity, it focuses on the case of a CRP tax schedule, which leads to a particularly sharp characterization. We generalize it to arbitrary tax schedules in the appendix.

LEMMA 1 (Equilibrium contracts). Suppose that the tax schedule is CRP. Equilibrium earnings are given by

$$\underline{z}(\theta) = \frac{1}{1 + \ell(e^\beta - 1)} \ell \theta \quad \text{and} \quad \bar{z}(\theta) = \frac{e^\beta}{1 + \ell(e^\beta - 1)} \ell \theta, \quad (6)$$

where the pass-through parameter β is given by

$$\beta = \frac{h'(\ell)}{1 - p}. \quad (7)$$

The effort level ℓ is independent of θ and satisfies

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

Expected utility is given by

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

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

⁸ Rogerson (1985) credits an unpublished paper by Holmstrom (1984) for the first proof of validity of the first-order approach in such a setting.

To interpret lemma 1, consider first the benchmark setting where the firm can perfectly monitor the worker's effort. In this case, the firm provides full insurance against output risk, so that $\beta = 0$. Workers with ability θ who provide effort ℓ 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, in which the relationship $z = \theta\ell$ between labor effort and income also holds.

Earnings.—Equation (6) characterizes the equilibrium level of base pay and high-level pay in the presence of moral hazard frictions. Notice first that the previous relationship continues to hold on average: the free-entry condition (4) imposes that workers' expected earnings $\mathbb{E}[z|\theta]$ remain equal to their expected output $\ell\theta$. However, providing effort incentives implies that realized earnings must now be dispersed around their mean; that is, the firm provides only partial insurance against output risk: $0 < \underline{z} < \theta\ell < \bar{z} < \theta$. At the heart of our paper lies the observation that the optimal degree of within-firm insurance (in particular, the pass-through parameter β) is endogenous to the tax system.

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 θ . Keeping the earnings structure \underline{z} , \bar{z} , b fixed, the firm incurs an additional cost b , since it must pay a bonus to the worker more frequently.⁹ In addition, the firm needs to raise the earnings spread to incentivize the worker to actually exert this extra effort. Specifically, we saw that the pass-through parameter β must increase proportionately to the rise in the marginal disutility of effort, $h'(\ell)$. Exposing risk-averse workers to more earnings risk creates an additional cost for the firm, because keeping their participation constraint satisfied requires increasing their mean earnings. We call the corresponding term in (8) (second term on the right-hand side) the marginal cost of incentives (MCI).

Utility.—Equation (9) decomposes the worker's expected utility into three components. The first is the utility they would attain under full insurance, $\mathcal{U}(\theta) = v(\theta\ell) - h(\ell)$. The second term on the right-hand side (in curly brackets) captures that the incompleteness of private insurance makes risk-averse workers worse off. The utility loss associated with a given earnings lottery z is equal to the utility difference between expected earnings $\mathbb{E}[z|\theta]$ and the certainty equivalent.¹⁰ Third, this utility loss is weighted by $(1 - p)$: all else 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.

⁹ Note that, by eq. (5), this is just sufficient to compensate the worker for the higher effort level and keep the participation constraint satisfied.

¹⁰ Recall that the certainty equivalent z^{CE} is defined by $\log(z^{CE}) = \mathbb{E}[\log(z)|\theta]$.

II. Tax Progressivity and Performance Sensitivity of Pay

In this section, we study the incidence of taxes on the equilibrium labor contracts—in particular, on the amount of earnings risk to which the worker is exposed. For clarity of exposition, we focus on the case of a CRP tax schedule and relegate the general incidence analysis with arbitrary tax systems to the appendix. We show that a higher rate of tax progressivity has two opposite effects—crowd out and crowd in—on pretax earnings risk. These two effects almost exactly offset each other, so that private insurance remains roughly invariant to tax policy.

Crowd out.—Equation (7) shows that the pass-through β is inversely proportional to $1 - p$, so that, ceteris paribus, a higher rate of progressivity p leads to a higher earnings spread. This direct effect is a standard crowding out of private 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 pretax earnings to preserve the worker's incentives for effort. The elasticity of the pass-through with respect to tax progressivity is given by $\varepsilon_{\beta,1-p} \equiv \partial \log \beta / \partial \log(1 - p) = -1$. This value implies that, absent effort responses, the firm adjusts the contract so as to keep consumption insurance (measured by the variance of log consumption) fixed.

Crowd in.—Second, β is proportional to $h'(\ell)$. As a result, tax progressivity affects earnings risk indirectly via the endogenous choice of labor effort. Denote by $\varepsilon_{\ell,1-p} \equiv \partial \log \ell / \partial \log(1 - p)$ the elasticity of labor effort (or, equivalently, the frequency of receiving a bonus) with respect to tax progressivity.

LEMMA 2 (Labor supply elasticity). In the model with exogenous earnings risk, we have $\varepsilon_{\ell,1-p} = \varepsilon_{\ell}^f / (1 + \varepsilon_{\ell}^f)$, where $\varepsilon_{\ell}^f = h'(\ell) / (\ell h''(\ell))$ gives the Frisch elasticity of labor supply. In our setting, we have $\varepsilon_{\ell,1-p} > \varepsilon_{\ell}^f / (1 + \varepsilon_{\ell}^f) > 0$. In particular, higher tax progressivity reduces optimal labor effort.

This result shows that in our setting, tax progressivity disincentivizes labor effort—as in standard models of taxation. This is because greater progressivity makes high-powered incentives more costly for the firm to provide: when high levels of income are taxed away more heavily, eliciting marginally higher effort from the worker requires a larger increase in the dispersion of pretax earnings and therefore a higher cost for the firm.¹¹ Notice that raising tax progressivity causes a stronger labor supply response than in the exogenous-risk environment, because the MCI channel is muted when effort is observable.

¹¹ This argument echoes the wage-cum-labor demand effect of Lehmann, Parmentier, and Van der Linden (2011), whereby tax progressivity tends to make high pretax wages less attractive (which in their setting leads to lower unemployment).

Now, since it is optimal to reduce the worker’s level of labor supply in response to a rise in tax progressivity, the firm is able to provide a higher level of insurance against output risk. This leads to a compression of the pretax earnings distribution—that is, a crowding in of private insurance. Let $\varepsilon_{\beta,\ell} \equiv \partial \log \beta / \partial \log \ell$ denote the elasticity of the pass-through parameter β with respect to the desired level of labor supply. Importantly, an immediate consequence of equation (7) is that $\varepsilon_{\beta,\ell} = 1/\varepsilon_\ell^F$. The crowd-in effect is then given by the product $\varepsilon_{\beta,\ell} \cdot \varepsilon_{\ell,1-p}$.

The following proposition summarizes these effects. It is the first main result of this paper.

PROPOSITION 1 (Incidence of tax progressivity). The total impact of tax progressivity on earnings risk is given by

$$\frac{d \log \beta}{d \log(1-p)} = \varepsilon_{\beta,1-p} + \varepsilon_{\beta,\ell} \varepsilon_{\ell,1-p} = -1 + \frac{\varepsilon_{\ell,1-p}}{\varepsilon_\ell^F} > -\frac{\varepsilon_\ell^F}{1 + \varepsilon_\ell^F}. \quad (10)$$

It is negative (net crowding out) if $\varepsilon_{\ell,1-p} < \varepsilon_\ell^F$ and positive (net crowding in) otherwise.¹² The crowd in offsets at least a share $1/(1 + \varepsilon_\ell^F)$ of the crowd out.

Crowd out and crowd in approximately offset each other.—Proposition 1 implies that the crowding out and crowding in of private insurance have comparable orders of magnitude. The key insight is that the strength of the moral hazard friction is inversely proportional to the Frisch elasticity of labor supply: $\varepsilon_{\beta,\ell} = 1/\varepsilon_\ell^F$. If the elasticity of labor effort $\varepsilon_{\ell,1-p}$ is roughly equal to the Frisch elasticity, then the crowding in is equal to $\varepsilon_{\ell,1-p}/\varepsilon_\ell^F \approx 1$ —that is, about the same size (and the opposite sign) as the direct crowd out $\varepsilon_{\beta,1-p} = -1$. Thus, on net, tax progressivity has only a small impact on the amount of within-firm insurance: $d \log \beta / d \log(1-p) \approx 0$.

Importantly, this insight is robust to the size of the labor supply responses to taxes. To see this, suppose that effort hardly diminishes in response to an increase in tax progressivity. A naive intuition would suggest that the crowd in, which is driven by labor supply responses, is then negligible. But proposition 1 shows that the crowd-in effect is bounded from below by $1/(1 + \varepsilon_\ell^F)$, which converges to one as $\varepsilon_\ell^F \rightarrow 0$. Thus, not only does the crowding in of private insurance remain large but it offsets (at least as much as) the entire crowd out in this polar case. Intuitively, it is precisely because of the worker’s inelastic behavior that the firm is able to dramatically reduce their exposure to output risk without hindering incentives beyond the small desired reduction in labor effort. In other words, since

¹² We show in the appendix that $\varepsilon_{\ell,1-p} \geq \varepsilon_\ell^F$ if and only if $\ell \geq 1/2$. Thus, social insurance leads to a net crowding out of private insurance if and only if the frequency of receiving the bonus is below 50%.

$\varepsilon_{\beta,\ell} = 1/\varepsilon_\ell^F \rightarrow \infty$, the product $\varepsilon_{\beta,\ell} \varepsilon_{\ell,1-p}$ does not vanish in the limit, and the crowding in 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 strictly positive, this need not be exactly the case. Suppose that the disutility of (unnormalized) 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 ε_ℓ^F in proposition 1, however, is lower than e , since it is that of the function $h \circ \pi^{-1}$, where π is concave (see n. 5). 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 $\varepsilon_{\beta,\ell} \varepsilon_{\ell,1-p} > 1/(1 + 0.2) = 0.83$. That is, the earnings risk adjustment due to labor effort responses offsets at least 83% of the crowd out of private insurance caused by tax progressivity. Even when we raise the Frisch elasticity to $e = 1$, at the high end of empirical estimates, the crowd-in effect offsets at least 75% 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 section C, we study the moral hazard framework of Holmstrom and Milgrom (1987) with preferences $u(c, \ell) = -(1/\gamma) \exp(-\gamma(c - h(\ell)))$, 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 γ . When workers are risk neutral, the crowd in exactly offsets the crowd out. 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.

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. These represent the largest share of all performance-based earnings contracts. Nevertheless, it is important to evaluate whether our main insights carry over to other types of compensation. In appendix section C, we set up several alternative frameworks. The first, based on Holmstrom and Milgrom (1987), provides conditions under which linear contracts are optimal—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, for example, promotions or salary raises. To study such dynamic effects, we use the framework of Edmans et al. (2012).

Despite their differences, all of these models share a common structure: when the utility is logarithmic and the tax schedule is CRP, the slope of the earnings contract, which measures the sensitivity of log earnings to output shocks, is given by $h'(\ell)/(1 - p)$.¹³ This is a direct consequence of the local incentive constraint, which is common to all of these models of moral hazard 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 regardless of these modeling differences: a tax change creates crowd-out and crowd-in effects on the pretax 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 approximately insensitive to tax policy.

III. Optimal Taxation

In this section, we derive formulas for the optimal rate of tax progressivity and the optimal tax rate on top income earners. The government chooses the tax schedule, or the retention function R , to maximize a weighted utilitarian social welfare function

$$\int_0^\infty \alpha(\theta) U(\theta) f(\theta) d\theta \tag{11}$$

subject to a budget constraint

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

where the Pareto weights $\alpha(\theta) \geq 0$ are continuous in θ and satisfy $\int_0^\infty \alpha(\theta) f(\theta) d\theta = 1$ and where $G \geq 0$ is an exogenous expenditure requirement. We denote by s the ratio of public spending to aggregate output, so that $\hat{s} \equiv s/(1 - s)$ is the ratio of public to private consumption.

A. Optimal Tax Progressivity

We first restrict the tax schedule to the CRP class. Proposition 2 characterizes the optimal rate of tax progressivity p under the additional assumptions that the distribution of ability types is lognormal and that the

¹³ In the case of Holmstrom and Milgrom (1987), the theoretical restrictions on agents' preferences allow us to study only a constant absolute risk aversion utility function with affine taxes. In this case, the relevant notion of pass-through becomes $h'(\ell)/(1 - \tau)$, where τ represents the constant marginal tax rate. Our discussion continues to hold once we focus on reforms of the tax rate τ .

social welfare objective is utilitarian. The operator $\mathbb{V}(\cdot|\theta)$ denotes the variance conditional on ability θ .

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

$$\frac{p}{(1-p)^2} = \frac{\sigma_\theta^2 + \kappa_1(1 + \varepsilon_{\beta,1-p})\mathbb{V}(\log z|\theta)}{(1 + \hat{s}/p + \kappa_2)\varepsilon_{\ell,1-p} + \kappa_3 \varepsilon_{\beta,\ell} \varepsilon_{\ell,1-p}\mathbb{V}(\log z|\theta)}, \tag{13}$$

with $\kappa_1 > 0$, $\kappa_2 > 0$ for small enough $p > 0$, and $\kappa_3 > 0$ whenever $p > 0$.¹⁴ The optimal rate of progressivity (13) is strictly lower than with exogenous earnings risk ($\varepsilon_{\beta,1-p} = \varepsilon_{\beta,\ell} = 0$).

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

$$\frac{p}{(1-p)^2} = \frac{\sigma_\theta^2}{(1 + \hat{s}/p)\varepsilon_{\ell,1-p}}.$$

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

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

$$\frac{p}{(1-p)^2} = \frac{\sigma_\theta^2 + \kappa_1\mathbb{V}(\log z|\theta)}{(1 + \hat{s}/p + \kappa_2)\varepsilon_{\ell,1-p}}. \tag{14}$$

Compared with the full-insurance benchmark, the dispersion of earnings in the population is now mechanically larger than that of ability types θ . Tax progressivity thus plays two roles: redistribution across ex ante ability differences (measured by the variance σ_θ^2) and social insurance against ex post earnings risk (measured by the conditional variance of pretax earnings $\mathbb{V}(\log z|\theta)$).¹⁵ This contributes to raising the optimal rate of progressivity. Moreover, a change in taxes now not only affects the income levels \underline{z} , \bar{z}

¹⁴ We have $\kappa_1 = 1/(\beta(1-p))(R(\bar{z}) - R(\underline{z}))/\mathbb{E}[R(z)|\theta]$, $\kappa_2 = \ell((1-p)/p)(b/\mathbb{E}[z|\theta] - (1-p)(R(\bar{z}) - R(\underline{z}))/\mathbb{E}[R(z)|\theta])$, $\kappa_3 = ((1-p)/(\beta p))(b/\mathbb{E}[z|\theta] - (R(\bar{z}) - R(\underline{z}))/\mathbb{E}[R(z)|\theta])$. If $\kappa_2, \kappa_3 > 0$, the optimal p is also strictly lower than in the full-insurance Mirrlees benchmark.

¹⁵ 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.

on the intensive margin but also triggers a response on the frequency margin by altering the probability ℓ with which the high income level \bar{z} , and hence the high tax payment, occurs. Accounting for these additional distortions gives rise to the additional term κ_2 in (14).¹⁶

Fiscal externalities from endogenous insurance.—Finally, consider the general case where earnings risk, captured by the pass-through β , is endogenous to taxes. The optimal tax formula (13) first accounts for the fact that if the tax schedule is strictly progressive, a spread of the pretax earnings distribution caused by the crowding out of private insurance (analyzed in sec. II) generates a positive fiscal externality—that is, 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 $-(p/(1-p)^2)\kappa_3(\varepsilon_{\beta,1-p} + \varepsilon_{\beta,\ell} \varepsilon_{\ell,1-p})\mathbb{V}(\log z|\theta)$. Note that the denominator of formula (13) features only the negative externality from the crowd-in channel. This is because, as we argue next, the positive fiscal externality caused by the crowd out is (more than) compensated for by its negative welfare impact.

Welfare losses from endogenous insurance.—The crowd out of private insurance also causes two welfare losses. Importantly, by the envelope theorem, these losses are *not* counteracted by corresponding welfare gains from the crowd-in responses: recall that the crowding in of private insurance due to an increase in tax progressivity operates via adjustments in optimal labor effort, thus leading to (at most) second-order welfare effects. First, the crowd out of private insurance exactly offsets the welfare benefits of social insurance, since $1 + \varepsilon_{\beta,1-p} = 0$ in the numerator of (13). Intuitively, any attempt by the government to compress the distribution of disposable income leads the firm to raise the dispersion of pretax earnings one for one to preserve the worker’s effort incentives.¹⁷ 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 θ , measured by the variance σ_θ^2 .

¹⁶ The intensive-margin behavioral responses affect government revenue proportionately to the income-weighted marginal tax rates $\mathbb{E}[T'(z)|z]_{\varepsilon_{\ell,1-p}}$, while the frequency-margin responses affect revenue proportionately to the gap in total tax payments (or average tax rates) between the high- and low-performance states $\ell(T(\bar{z}) - T(\underline{z}))_{\varepsilon_{\ell,1-p}}$. If $p = 0$, then $\kappa_2 = 0$; in this case, the marginal and average tax rates coincide, so that no additional correction is necessary.

¹⁷ 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.

The second 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. To build intuition, suppose that the tax liabilities at the base pay and at the high-level pay are lowered by the same amount. Such a reform increases the worker's reservation value $U(\theta)$ by $\Delta U > 0$. The incentive constraint (5) implies that, to maintain effort ℓ , the ex post utility of both low and high performers must increase by the same amount ΔU .¹⁸ 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 nonnegative, lower the base pay. This implies that high performers capture a disproportionately large share of the tax cut. This regressive distribution of rents within the firm—away from individuals whose marginal utility of consumption is the highest—further reduces the welfare benefits of redistribution via progressive taxes and exactly offsets the positive fiscal externality described earlier.

Taking stock.—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.¹⁹ As a result, the optimal rate of progressivity (13) is strictly lower than in a setting with exogenous earnings risk.

B. Optimal Top Tax Rate

In sections II and III.A, we focused on CRP tax schedules. While this functional form closely approximates the marginal and average taxes in the bulk of the income distribution, it does not naturally lend itself to analyzing the optimal taxation of high-income earners, since it imposes that the marginal tax rate converges to 100% at the top. In this section, we relax the CRP restriction and characterize the optimal top marginal tax rate. This 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

¹⁸ 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 eqq. (19) and (20) in the appendix.

¹⁹ Another way to decompose the novel effects is as follows: (i) the sum of a negative fiscal externality from crowd in and a positive fiscal externality from crowd out, which is close to zero since we saw in proposition 1 that the crowd in and the crowd out approximately offset each other; (ii) two negative welfare effects of crowd out. Under the parametric assumptions of proposition 2, the positive fiscal externality and the second negative welfare effect of crowd out exactly cancel out. This is no longer the case in a more general environment: both of these effects appear explicitly in the optimal tax formulas we derived in our working paper, Doligalski, Ndiaye, and Werquin (2020).

forms of compensation for the highest-income earners, such as bankers' bonuses and CEOs' stock options (see, e.g., Piketty and Saez 2003; Lemieux, MacLeod, and Parent 2009; Bell and Van Reenen 2013).²⁰

Recall that in the full-insurance Mirrlees benchmark, the welfare effects of taxes are summarized by the marginal social welfare weights (MSWWs; see, e.g., Saez and Stantcheva 2016). The MSWW of an agent with ability θ and earnings z is defined as the impact on social welfare of marginally increasing their consumption:

$$g(z|\theta) = \frac{1}{\lambda} \alpha(\theta) u'(R(z)), \tag{15}$$

where $\alpha(\theta)$ represents the Pareto weight of type θ in the social objective and λ represents the marginal value of public funds.

PROPOSITION 3 (Optimal top tax rate). Suppose that the distribution of ability types has an upper Pareto tail, and denote by ρ^* the implied Pareto coefficient of the right tail of the income distribution. Suppose moreover that the average elasticity of expected earnings with respect to the retention rate (weighted by expected earnings) among workers who receive pay above z^* with positive probability converges to a limit \mathcal{E}^* as $z^* \rightarrow \infty$. Suppose finally that the marginal social welfare weights converge to a constant g^* at the top and that the optimal tax rate τ^* converges to a constant above some income threshold. We then have

$$\frac{\tau^*}{1 - \tau^*} = \frac{1 - \psi_g \cdot g^*}{\psi_\rho \cdot \rho^* \cdot \mathcal{E}^*}, \tag{16}$$

where $\psi_g > 1$ and $\psi_\rho > 1$. Conditional on the values of the sufficient statistics, the optimal top tax rate is lower than in the full-insurance (Mirrlees) model ($\psi_g = \psi_\rho = 1$) and in the exogenous-risk model ($\psi_g = 1$).

To understand this result, recall from our analysis in section III.A that a key step in designing optimal taxes with performance-pay contracts consists of properly accounting for the welfare consequences of perturbing tax rates. In the full-insurance Mirrlees benchmark, the envelope theorem ensures that all endogenous behavioral responses to tax changes have at most a second-order impact on the agent's utility. Hence, a tax cut raises the worker's utility in proportion to their marginal utility of consumption and therefore raises social welfare by $g(z|\theta)$. This is no longer true in the model with performance pay. Recall that the endogenous responses to a tax change can be decomposed into a direct crowd out and an indirect crowd in of the earnings contract. While the envelope theorem still applies to the endogenous effort responses and hence to the

²⁰ Bell and Van Reenen (2014) show that in the United Kingdom 83% of workers in the top percentile received a bonus in 2008 and that bonuses represented 35% of their total compensation (44% in the financial sector).

crowding in of private insurance, the earnings adjustments caused by the crowding out have a first-order impact on welfare.²¹

To account for this implied welfare implications of the crowd out, we define the *modified* MSWW $\tilde{g}(z|\theta)$ as the welfare impact of marginally reducing the tax liability of an agent with type θ and realized income $z \in \{\underline{z}, \bar{z}\}$. They satisfy

$$\frac{\tilde{g}(\bar{z}|\theta)}{g(\bar{z}|\theta)} = \frac{1/v'(\bar{z})}{\mathbb{E}[1/v'(z)|\theta]} > 1 \quad \text{and} \quad \frac{\tilde{g}(\underline{z}|\theta)}{g(\underline{z}|\theta)} = \frac{1/v'(\underline{z})}{\mathbb{E}[1/v'(z)|\theta]} < 1. \quad (17)$$

These expressions imply that, relative to the standard model, the contribution of agents with type θ to social welfare is adjusted upward when they receive a bonus and downward otherwise. These regressive adjustments to the MSWW reflect the fact that, as explained in section III.A, tax cuts are passed through within the firm primarily to the high-performing workers. Intuitively, a tax cut initially generates a rent to the firm, which is then transferred to the worker via free entry and the resulting adjustment of the reservation value. This distribution of rents to the workers needs to keep the effort level incentive compatible and hence the utility difference $v(\bar{z}) - v(\underline{z})$ unchanged (see eq. [5]). In turn, this implies that earnings $z \in \{\underline{z}, \bar{z}\}$ must change in proportion to the respective inverse marginal utilities $1/v'(z)$. Since the marginal utility is decreasing, this means that the high-level pay ends up rising more than the base pay. Such a regressive distribution of rents makes it more difficult than in the full-insurance environment to target transfers to the workers who need it the most—that is, to directly raise the consumption of the agents whose marginal utility is relatively large (namely, the unlucky individuals of a given ability who do not receive a bonus). This amounts to placing lower (respectively, higher) effective weights on the poorer (respectively, richer) workers, resulting in a higher top MSWW $\tilde{g}^*/g^* \equiv \psi_g > 1$ and a lower optimal top tax rate.²²

Finally, since the labor contract is chosen before output is realized, the relevant labor supply elasticity \mathcal{E}^* in expression (16) is the average effort response (to a tax change above income z^*) of all workers who have a positive probability of earning more than this income threshold—including the unlucky ones whose realized pay ends up below z^* . Taking

²¹ 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.

²² Notice that if the marginal social welfare weight placed on the top earners converges to zero (which would be the case, e.g., under a utilitarian planner if the marginal utility of consumption converges to zero), then the welfare effect of the crowd out is irrelevant, since $\psi_g g^* = g^* = 0$. However, empirical studies based on an “inverse optimum” approach consistently conclude that the effective weight that society places on the marginal consumption of top earners is generally positive (see, e.g., Lockwood and Weinzierl 2016).

\mathcal{E}^* as a sufficient statistic, the Pareto coefficient (which measures the average income above z^* relative to z^*) must be scaled up by a factor $\psi_\rho > 1$ equal to the ratio between the average income of these workers and the average income above z^* . This adjustment further lowers the optimum top tax rate relative to the full-insurance benchmark. Note that \mathcal{E}^* differs from the “naive” elasticity one would obtain by accounting for only the earnings responses within the top bracket $[z^*, \infty)$. While both concepts coincide in the Mirrlees benchmark, \mathcal{E}^* additionally accounts for the fiscal externality of the top tax rate on lower-income workers caused by the fact that these workers *could* have (at the time they chose their effort) fallen in the top tax bracket.²³

IV. Quantitative Analysis

To evaluate our results quantitatively, we extend our baseline model of section II (i.e., with a CRP tax schedule) 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 in the previous sections: output is stochastic, equal to the worker’s ability θ with probability $\pi(\ell_{pp})$ and zero otherwise, where ℓ_{pp} represents their effort level. The optimal contract specifies earnings as a function of the output realization according to the pass-through 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 ℓ_{fp} represents 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. 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 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 $x_E = \exp(\rho_{\theta,j})$.

A. Calibration

We calibrate the model to match empirical evidence on performance-pay jobs, earnings elasticities, and the overall earnings distribution in the United States. The chosen parameter values are summarized in table 1.

²³ This fiscal externality can in turn be decomposed into crowd out, crowd in, and frequency responses of base earnings. In general, \mathcal{E}^* can be smaller or larger than the naive elasticity, depending on the relative importance of these competing forces.

TABLE 1
CALIBRATED PARAMETERS

Parameter	Value	Description	Source or Target
s_{pp}	.45	Share of performance-pay jobs	Lemieux, MacLeod, and Parent 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 performance-pay jobs	Difference in mean earnings between job types
σ_{θ}	.29	Normal variance of log ability	Variance of log earnings in the economy
ρ_{θ}	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
ϕ	.82	Curvature of $\pi(\cdot)$ function	Difference in variance of log earnings between job types
ε_t^f	.5	Frisch elasticity of labor effort	Chetty et al. 2011; Keane 2011
ρ	.181	Tax progressivity	Heathcote, Storesletten, and Violante 2017
G/Y	.188	Share of government spending in GDP	Heathcote and Tsujiyama 2021

NOTE.—All the target moments are matched exactly.

We examine the robustness of our main results to alternative parametrizations at the end of this section.

Lemieux, MacLeod, and Parent (2009) use the Panel Study of Income Dynamics (PSID) 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% higher in performance-pay jobs than in fixed-pay jobs in 1998.²⁴ 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 $\pi(\ell) = \bar{\pi} \ell^{\phi}$, with $\phi \in (0, 1]$. In the data, the average probability of receiving a bonus conditional on having a performance-pay job is 23%, which pins down $\bar{\pi}$.²⁵ The exponent ϕ affects the magnitude of earnings risk due to performance pay. Lemieux, MacLeod, and Parent

²⁴ This finding is consistent with other data sources. Gittleman and Pierce (2013) and Makridis and Gittleman (2022) use the National Compensation Survey to show that performance-pay jobs pay higher hourly compensation than fixed-pay jobs conditional on the work level (a proxy for skill) and that performance pay is more prevalent at higher work levels. Similarly, Grigsby, Hurst, and Yildirmaz (2019) use payroll data to show that the share of compensation paid in bonuses is increasing with earnings.

²⁵ Based on table II and n. 15 in Lemieux, MacLeod, and Parent (2009), we calculate that (i) in the 1990s, 82% 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%. 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%.

(2009) report that the variance of log earnings in performance-pay jobs is 42% higher than in fixed-pay jobs. This excess variance can be explained by either the additional earnings risk generated by stochastic bonuses (controlled by ϕ) or 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% of the cross-sectional variance of log earnings of performance-pay workers. We show below that our main conclusions are robust to this assumption.

A concern is that earnings records in the PSID are top coded. Top coding may influence the measured moments of the earnings distribution, including the relative mean and variance of earnings at the two job types. We address this concern in two ways. First, we verify that our main results are robust to higher mean and variance of earnings in performance-pay jobs relative to fixed-pay jobs (see fig. 5). Second, 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 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 $\rho_{\theta,pp} = \rho_{\theta,fp} = 2.2$ and choose $\sigma_{\theta}^2 = 0.29$ to match the overall variance of log earnings.

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

The implied distribution of earnings and job types is depicted in figure 1. Specifically, in figure 1B, we compare the (untargeted) shares of performance-pay jobs by earnings quartiles in the data and in the model.²⁶ The calibrated model successfully matches the empirical prevalence of performance-pay jobs: in both the data and the model, the share of performance-pay jobs is approximately 40% in the bottom three quartiles

²⁶ The data shares are computed for the year 1998 from PSID using the methodology of Lemieux, MacLeod, and Parent (2009). We are grateful to Karl Schulz for computing these values.

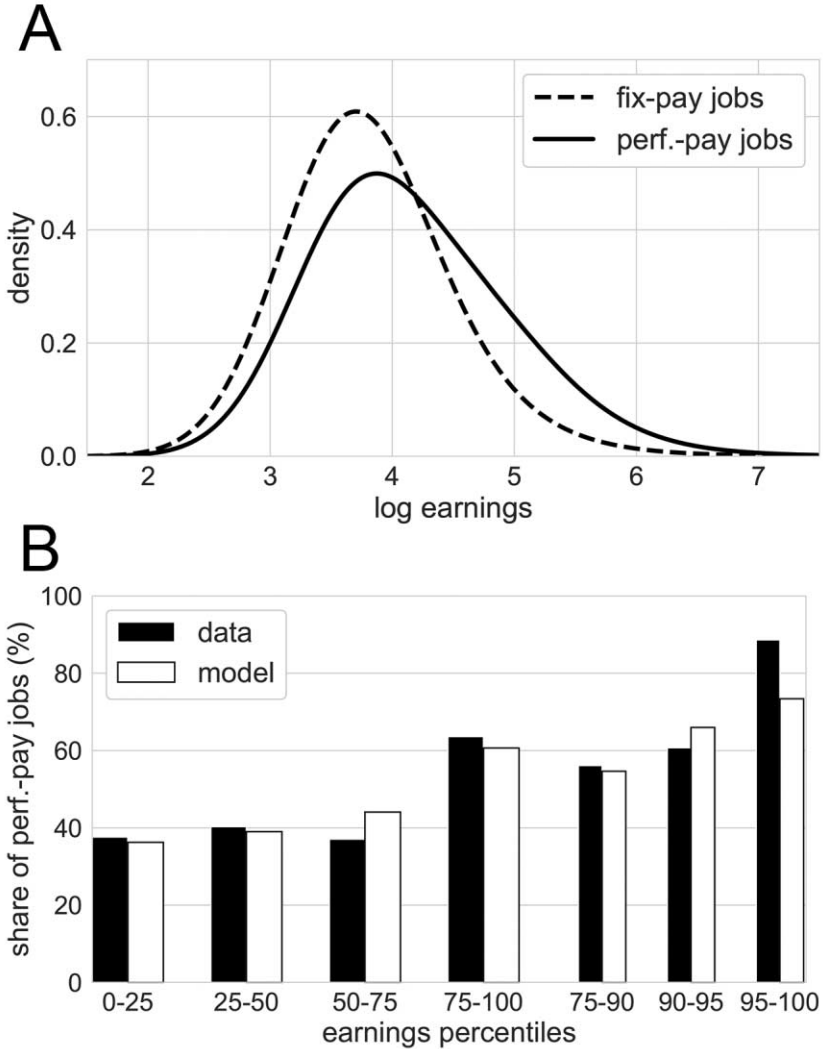


FIG. 1.—Joint distribution of earnings and job types.

and rises to more than 60% in the top quartile. Zooming in on the top quartile, the model matches the data well up until the top 5% of earners, for whom the share of performance-pay jobs is slightly underestimated.²⁷

²⁷ In fig. 5, we consider an alternative calibration in which the ratio of mean earnings in performance-pay vs. fixed-pay jobs is 25% higher. This calibration yields a share of performance-pay jobs in the top 5% that closely matches the empirical value. As evident from the figure, our main results are virtually unchanged.

B. Incidence of a Large Tax Reform: From Status Quo to 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 proposition 2). We find that the utilitarian optimum rate of progressivity is equal to 0.376. This is more than twice as high as the current rate of tax progressivity in the United States, and the implied increase in social welfare is equivalent to a 3.7% increase in consumption (see fig. 2A). In this subsection, we

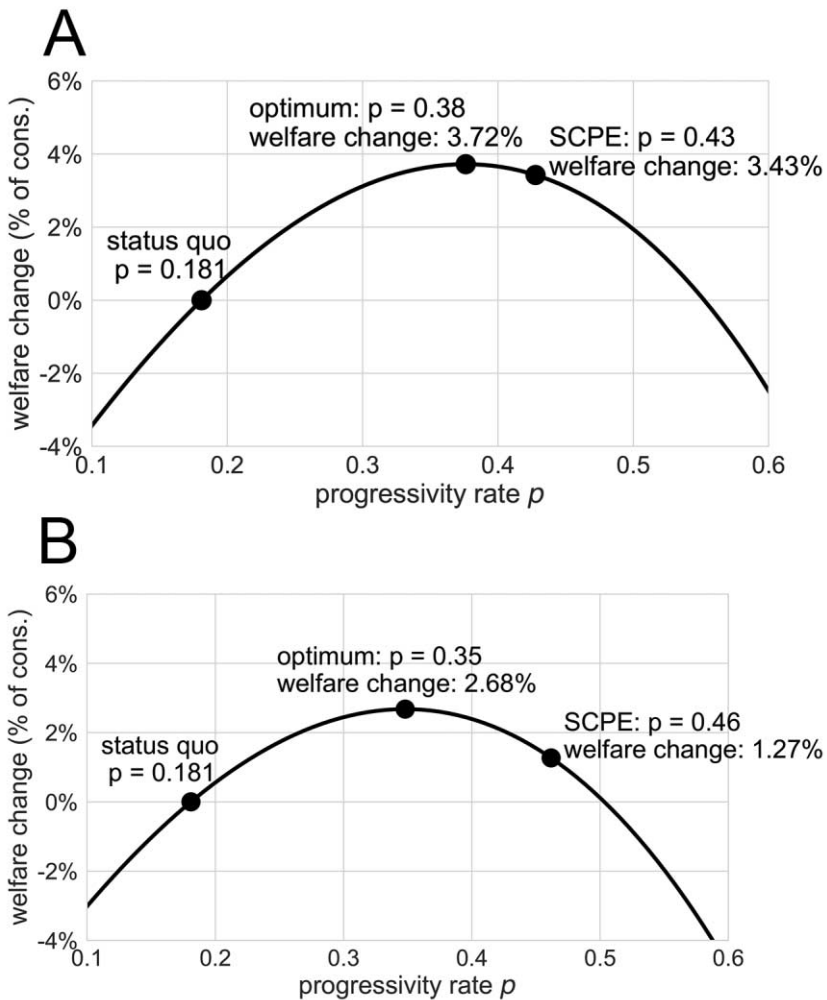


FIG. 2.—Optimal progressivity and self-confirming policy equilibrium.

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 pass-through rate β 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 β is dominated by the impact of a lower effort ℓ_{pp} , distorted downward 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, providing better social insurance via tax progressivity would crowd out private insurance one for one—that is, keep the variance of log consumption fixed. For that to happen, the pass-through β would need to increase from 1.08 to 1.42, raising the log earnings risk of performance-pay workers by 72%. However, in equilibrium firms choose to elicit a 9% 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%.

C. 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

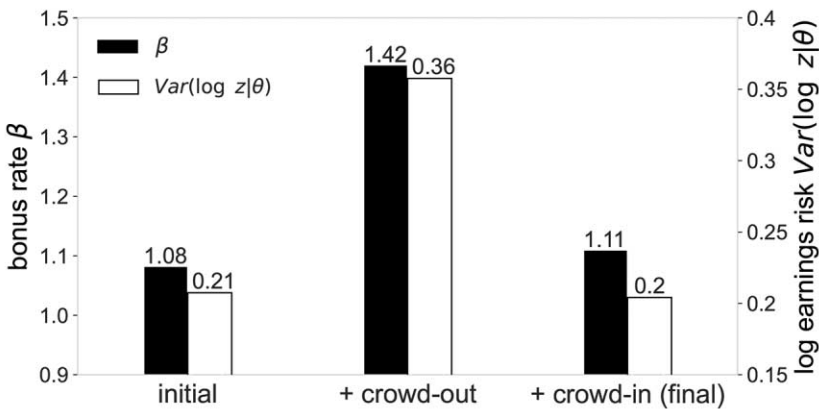


FIG. 3.—Optimal progressivity reform and earnings risk. The tax progressivity rate is increased from 0.181 (status quo) to 0.376 (utilitarian optimum).

of progressivity with a rate chosen by the government that erroneously assumes that all of the earnings risk is exogenous. This naive optimum 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 Rothschild and Scheuer (2016), we call the resulting allocation a self-confirming policy equilibrium (SCPE). The results are depicted in figure 2A.

A policy maker 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% 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 figure 2B. 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% to 1.3%). Thus, the finding that the cost of ignoring performance pay for the design of tax policy is modest is not a theoretical necessity. Rather, it is a quantitative result that is owed to the fact that only half of the jobs in the United States 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.

D. Optimal Top Tax Rate

We showed theoretically that performance pay reduces the optimal top tax rate, understood as the limit of the marginal tax rate in the optimal tax schedule as earnings go to infinity (proposition 3). We illustrate this result in figure 4, with the optimal tax rate shown as a solid line. We also

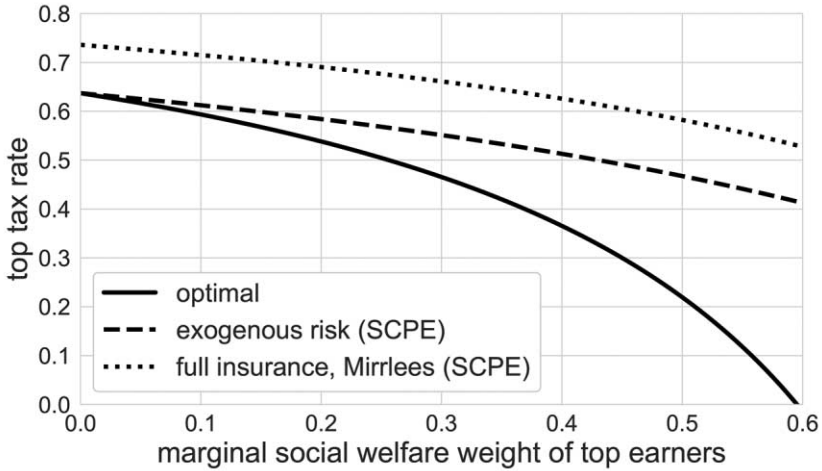


FIG. 4.—Top tax rates and performance pay.

plot the top tax rates chosen by the planner who does not take into account the endogeneity of earnings risk (dotted line) and by the planner who assumes that the underlying economy involves full insurance against performance shocks, as in the standard Mirrlees model (dashed line). In line with proposition 3, the optimal tax rate is lower than the tax rate under exogenous risk, which is in turn lower than the tax rate under full insurance.

We are particularly interested in the gap between the optimal top tax rate and the tax rate chosen by the exogenous-risk planner. Recall from proposition 3 that this difference is driven by the negative welfare effect suffered by top earners due to the crowd out of private insurance—which implies that the standard marginal social welfare weights underestimate the utility cost of raising the top tax rate. Thus, this difference critically depends on the marginal social welfare weight on top earners, g^* . The exogenous-risk planner selects the top tax rate correctly, at 64%, when it aims to maximize the revenue from top earners ($g^* = 0$), which happens, for instance, under standard utilitarian or Rawlsian social preferences. For positive values of g^* , the policy mistake made by the exogenous-risk planner increases in a convex fashion: it remains fairly small for g^* below 0.2 and quickly becomes large for higher values of g^* . For instance, when $g^* = 0.6$ —meaning that the marginal consumption of top earners is valued at 60% of that of the average worker in the population—the exogenous-risk planner would make a 40 percentage point mistake in setting the top tax rate. We conclude that performance pay can justify top tax rates

that are much lower than those predicted by the standard model, but this requires placing relatively high social value on the welfare on top earners.

E. Robustness

In this section, we study the sensitivity of our main results to the chosen parameters. We consider the following alternative calibrations.

First, there is a wide range of empirical estimates of the (Frisch) elasticity of labor effort; see, for example, Chetty et al. (2011). In the alternative calibrations 1 and 2, we consider higher ($\varepsilon_l^f = 1$) and lower ($\varepsilon_l^f = 0.1$) values of the Frisch elasticity of labor effort.

Second, so far, we have assumed that performance pay explains all the excess variance of log earnings in performance-pay jobs relative to fixed-pay jobs. However, other factors such as adverse selection of workers in different jobs types, firm sector specificity (Makridis and Gittleman 2022), the prevalence of overtime relative to fixed hours (Grigsby, Hurst, and Yildirmaz 2019), and retention incentives can explain part of the excess variance of log earnings in performance-pay jobs. Therefore, in the alternative calibration 3, we assume that earnings risk due to performance pay explains only half of the excess variance of log earnings in performance-pay jobs in the cross section. More empirical work using microdata on the different components of earnings is needed to understand the role of performance pay in explaining the variance of log earnings.

Third, the moments of the earnings distribution of performance-pay jobs we obtained from the PSID can be underestimated due to top coding. In calibration 4, we set a ratio of mean earnings at performance-pay versus fixed-pay jobs that is 25% higher than in the baseline calibration. In calibration 5, we set the ratio of the variance of log earnings at performance-pay versus fixed-pay jobs at 25% higher than in the baseline calibration.

Fourth, the probability of receiving performance pay can take a wide range of values depending on the frequency at which it is estimated. We thus set the probability of receiving performance pay, conditional on having a performance-pay job, at the upper bound from Lemieux, MacLeod, and Parent (2009) in calibration 6.

We present the results of these robustness exercises in figure 5 by plotting the values of two numbers that correspond to the two insights of our paper: (i) the share of crowd out that is offset by the crowd-in channel when progressivity increases from the status quo to the optimum and (ii) the difference between the optimal rate of tax progressivity and that chosen by the planner who ignores the endogeneity of earnings risk (SCPE).

Our main results are robust to these parameter values. We find that, depending on the calibration, the crowd in offsets most (85%–98%) of the crowd out and that the rate of progressivity chosen when the endogeneity of earnings risk is ignored is higher than the optimal one by 0.03–0.075.

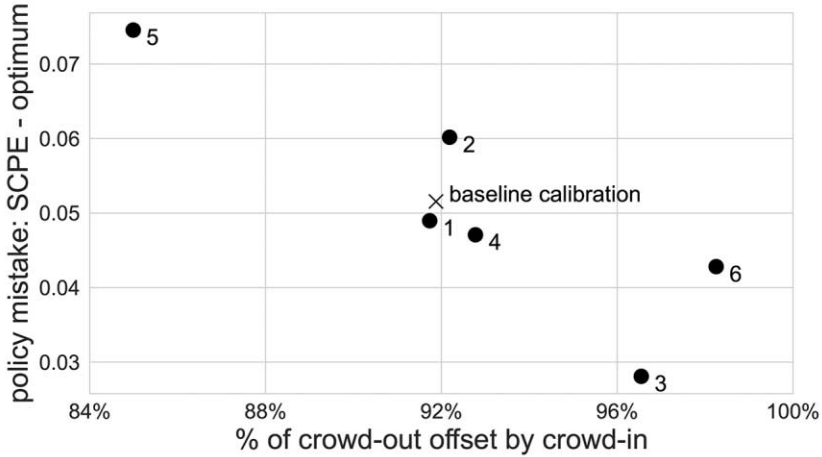


FIG. 5.—Robustness to alternative calibrations. The horizontal axis shows the share of crowd out that is offset by the crowd in when progressivity increases from the status quo to the optimum. The vertical axis shows the difference between the tax progressivity chosen by the planner who ignores the endogeneity of earnings risk (SCPE) and in the optimum. Alternative calibrations considered: (1) higher Frisch elasticity: $\varepsilon_l^f = 1$; (2) lower Frisch elasticity: $\varepsilon_l^f = 0.1$; (3) stochastic bonuses explain 50% (rather than 100%) of the excess variance at performance-pay jobs; (4) ratio of mean earnings at performance-pay versus fixed-pay jobs 25% higher than in baseline; (5) ratio of the variance of log earnings at performance-pay versus fixed-pay jobs 25% higher than in baseline; (6) probability of receiving performance pay, conditional on having a performance-pay job, at the upper bound from Lemieux, MacLeod, and Parent (2009).

Interestingly, varying the Frisch elasticity hardly affects the results. The reason is that what matters is the curvature of the function $h \circ \pi^{-1}$, which depends on both ε_l^f and ϕ (see discussion following proposition 1). Matching a given variance of earnings risk at performance-pay jobs pins down the overall curvature of $h \circ \pi^{-1}$, which means that any variation in ε_l^f is offset by the adjustment of ϕ . Conversely, we should expect that varying the magnitude of earnings risks due to stochastic bonuses has a relatively stronger impact on the results. This is exactly what we see in the alternative calibrations 3 and 5.

V. 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 of motivating labor effort.

We found that standard models that ignore the endogeneity of pretax earnings risk come very close to accurately evaluating the impact of taxes on the sensitivity of pay to performance or earnings risk. We derived optimal tax formulas for the overall rate of tax progressivity and the top income tax rate. Our findings show that it is optimal to reduce the rate of progressivity and the top income tax rate, compared with an economy with exogenous private insurance.

It would be interesting to extend our analysis in several directions. First, we considered the impact of taxes on compensation only 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 fruitful.

References

- Ábrahám, Á., F. Alvarez-Parra, and S. Forstner. 2016. "The Effects of Moral Hazard on Wage Inequality in a Frictional Labor Market." Working paper.
- Ábrahám, Á., S. Koehne, and N. Pavoni. 2016. "Optimal Income Taxation When Asset Taxation Is Limited." *J. Public Econ.* 136:14–29.
- Alvarez, F., and U. J. Jermann. 2000. "Efficiency, Equilibrium, and Asset Pricing with Risk of Default." *Econometrica* 68 (4): 775–97.
- Attanasio, O., and J.-V. Ríos-Rull. 2000. "Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?" *European Econ. Rev.* 44 (7): 1225–58.
- Bandiera, O., I. Barankay, and I. Rasul. 2011. "Field Experiments with Firms." *J. Econ. Perspectives* 25 (3): 63–82.
- Bell, B., and J. Van Reenen. 2013. "Extreme Wage Inequality: Pay at the Very Top." *A.E.R.* 103 (3): 153–57.
- . 2014. "Bankers and Their Bonuses." *Econ. J.* 124 (574): F1–F21.
- Bloom, N., and J. Van Reenen. 2010. "Why Do Management Practices Differ across Firms and Countries?" *J. Econ. Perspectives* 24 (1): 203–24.
- Boadway, R., and M. Sato. 2015. "Optimal Income Taxation with Risky Earnings: A Synthesis." *J. Public Econ. Theory* 17 (6): 773–801.
- Brown, C., and T. Andrabi. 2020. "Inducing Positive Sorting through Performance Pay: Experimental Evidence from Pakistani Schools." Working paper, Univ. California, Berkeley.
- Chang, Y., and Y. Park. 2017. "Optimal Taxation with Private Insurance." Report.
- Chetty, R., A. Guren, D. Manoli, and A. Weber. 2011. "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins." *A.E.R.* 101 (3): 471–75.

- Chetty, R., and E. Saez. 2010. "Optimal Taxation and Social Insurance with Endogenous Private Insurance." *American Econ. J. Econ. Policy* 2 (2): 85–114.
- Colombi, R. 1990. "A New Model of Income Distribution: The Pareto-Lognormal Distribution." In *Income and Wealth Distribution, Inequality and Poverty*, edited by C. Dagum and M. Zenga, 18–32. Berlin: Springer.
- Craig, A. C. 2019. "Optimal Income Taxation with Spillovers from Employer Learning." Working paper.
- Cremer, H., and P. Pestieau. 1996. "Redistributive Taxation and Social Insurance." *Internat. Tax and Public Finance* 3 (3): 281–95.
- Cullen, J. B., and J. Gruber. 2000. "Does Unemployment Insurance Crowd Out Spousal Labor Supply?" *J. Labor Econ.* 18 (3): 546–72.
- Cutler, D. M., and J. Gruber. 1996a. "Does Public Insurance Crowd Out Private Insurance?" *Q.J.E.* 111 (2): 391–430.
- . 1996b. "The Effect of Medicaid Expansions on Public Insurance, Private Insurance, and Redistribution." *A.E.R.* 86 (2): 378–83.
- Doligalski, P. 2019. "Optimal Income Taxation and Commitment on the Labor Market." Working paper.
- Doligalski, P., A. Ndiaye, and N. Werquin. 2020. "Redistribution with Performance Pay." Working paper.
- Eaton, J., and H. S. Rosen. 1980. "Optimal Redistributive Taxation and Uncertainty." *Q.J.E.* 95 (2): 357–64.
- Edmans, A., and X. Gabaix. 2011. "Tractability in Incentive Contracting." *Rev. Financial Studies* 24 (9): 2865–94.
- Edmans, A., X. Gabaix, T. Sadzik, and Y. Sannikov. 2012. "Dynamic CEO Compensation." *J. Finance* 67 (5): 1603–47.
- Ferey, A., A. Haufler, and C. Perroni. 2022. "Incentives, Globalization, and Redistribution." Working paper.
- Findeisen, S., and D. Sachs. 2016. "Education and Optimal Dynamic Taxation: The Role of Income-Contingent Student Loans." *J. Public Econ.* 138:1–21.
- Foster, A. D., and M. R. Rosenzweig. 1994. "A Test for Moral Hazard in the Labor Market: Contractual Arrangements, Effort, and Health." *Rev. Econ. and Statis.* 76 (2): 213–27.
- Friedrich, B., L. Laun, C. Meghir, and L. Pistaferri. 2019. "Earnings Dynamics and Firm-Level Shocks." Working Paper no. 25786, NBER, Cambridge, MA.
- Frydman, C., and R. S. Molloy. 2011. "Does Tax Policy Affect Executive Compensation? Evidence from Postwar Tax Reforms." *J. Public Econ.* 95 (11/12): 1425–37.
- Gittleman, M., and B. Pierce. 2013. "How Prevalent Is Performance-Related Pay in the United States? Current Incidence and Recent Trends." *Nat. Inst. Econ. Rev.* 226:R4–R16.
- Golosov, M., N. Kocherlakota, and A. Tsyvinski. 2003. "Optimal Indirect and Capital Taxation." *Rev. Econ. Studies* 70 (3): 569–87.
- Golosov, M., and A. Tsyvinski. 2007. "Optimal Taxation with Endogenous Insurance Markets." *Q.J.E.* 122 (2): 487–534.
- Grigsby, J., E. Hurst, and A. Yildirmaz. 2019. "Aggregate Nominal Wage Adjustments: New Evidence from Administrative Payroll Data." Working Paper no. 25628, NBER, Cambridge, MA.
- Guiso, L., L. Pistaferri, and F. Schivardi. 2005. "Insurance within the Firm." *J.P.E.* 113 (5): 1054–87.
- Heathcote, J., K. Storesletten, and G. L. Violante. 2017. "Optimal Tax Progressivity: An Analytical Framework." *Q.J.E.* 132 (4): 1693–754.

- Heathcote, J., and H. Tsujiyama. 2021. "Optimal Income Taxation: Mirrlees Meets Ramsey." *J.P.E.* 129 (11): 3141–84.
- Holmstrom, B. 1984. "A Note on the First-Order Approach in Principal Agent Models." Working paper.
- Holmstrom, B., and P. Milgrom. 1987. "Aggregation and Linearity in the Provision of Intertemporal Incentives." *Econometrica* 55 (2): 303–28.
- Kapička, M., and J. Neira. 2019. "Optimal Taxation with Risky Human Capital." *American Econ. J. Macroeconomics* 11 (4): 271–309.
- Kaplow, L. 1991. "Incentives and Government Relief for Risk." *J. Risk and Uncertainty* 4 (2): 167–75.
- Keane, M. P. 2011. "Labor Supply and Taxes: A Survey." *J. Econ. Literature* 49 (4): 961–1075.
- Krueger, D., and F. Perri. 2011. "Public versus Private Risk Sharing." *J. Econ. Theory* 146 (3): 920–56.
- Lamadon, T. 2016. "Productivity Shocks, Long-Term Contracts and Earnings Dynamics." Working paper.
- Lamadon, T., M. Mogstad, and B. Setzler. 2019. "Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market." Working Paper no. 25954, NBER, Cambridge, MA.
- Lazear, E. P. 2000. "Performance Pay and Productivity." *A.E.R.* 90 (5): 1346–61.
- . 2018. "Compensation and Incentives in the Workplace." *J. Econ. Perspectives* 32 (3): 195–214.
- Lazear, E. P., and P. Oyer. 2010. "Personnel Economics." In *Handbook of Organizational Economics*, edited by R. Gibbons and J. Roberts, 479–519. Princeton, NJ: Princeton Univ. Press.
- Leaver, C., O. Ozier, P. Serneels, and A. Zeitlin. 2021. "Recruitment, Effort, and Retention Effects of Performance Contracts for Civil Servants: Experimental Evidence from Rwandan Primary Schools." *A.E.R.* 111 (7): 2213–46.
- Lehmann, E., A. Parmentier, and B. Van der Linden. 2011. "Optimal Income Taxation with Endogenous Participation and Search Unemployment." *J. Public Econ.* 95 (11/12): 1523–37.
- Lemieux, T., W. B. MacLeod, and D. Parent. 2009. "Performance Pay and Wage Inequality." *Q.J.E.* 124 (1): 1–49.
- Lockwood, B. B., and M. Weinzierl. 2016. "Positive and Normative Judgments Implicit in US Tax Policy, and the Costs of Unequal Growth and Recessions." *J. Monetary Econ.* 77:30–47.
- Makridis, C. A., and M. Gittleman. 2022. "On the Cyclicity of Real Wages and Employment: New Evidence and Stylized Facts from Performance Pay and Fixed Wage Jobs." *J. Law, Econ., and Org.* 38 (3): 889–920.
- Makris, M., and A. Pavan. 2021. "Taxation under Learning by Doing." *J.P.E.* 129 (6): 1878–944.
- Mirrlees, J. A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Rev. Econ. Studies* 38 (2): 175–208.
- Miyazaki, H. 1977. "The Rat Race and Internal Labor Markets." *Bell J. Econ.* 8 (2): 394–418.
- Netzer, N., and F. Scheuer. 2007. "Taxation, Insurance, and Precautionary Labor." *J. Public Econ.* 91 (7/8): 1519–31.
- Park, Y. 2014. "Optimal Taxation in a Limited Commitment Economy." *Rev. Econ. Studies* 81 (2): 884–918.
- Piketty, T., and E. Saez. 2003. "Income Inequality in the United States, 1913–1998." *Q.J.E.* 118 (1): 1–41.

- Prendergast, C. 1999. "The Provision of Incentives in Firms." *J. Econ. Literature* 37 (1): 7–63.
- Raj, A. 2019. "Optimal Income Taxation in the Presence of Networks of Altruism." Working paper.
- Rogerson, W. P. 1985. "The First-Order Approach to Principal-Agent Problems." *Econometrica* 53 (6): 1357–67.
- Rose, N. L., and C. Wolfram. 2002. "Regulating Executive Pay: Using the Tax Code to Influence Chief Executive Officer Compensation." *J. Labor Econ.* 20 (S2): S138–S175.
- Rothschild, C., and F. Scheuer. 2016. "Optimal Taxation with Rent-Seeking." *Rev. Econ. Studies* 83 (3): 1225–62.
- Saez, E., and S. Stantcheva. 2016. "Generalized Social Marginal Welfare Weights for Optimal Tax Theory." *A.E.R.* 106 (1): 24–45.
- Scheuer, F. 2013. "Adverse Selection in Credit Markets and Regressive Profit Taxation." *J. Econ. Theory* 148 (4): 1333–60.
- 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." *Rev. Econ. Studies* 71 (2): 513–34.
- Sleet, C., and H. Yazici. 2017. "Taxation, Redistribution and Frictional Labor Supply." Paper presented at the 2017 annual meeting of the Society for Economic Dynamics, Edinburgh, Scotland, June 22–24.
- Stantcheva, S. 2014. "Optimal Income Taxation with Adverse Selection in the Labour Market." *Rev. Econ. Studies* 81 (3): 1296–329.
- . 2017. "Optimal Taxation and Human Capital Policies over the Life Cycle." *J.P.E.* 125 (6): 1931–90.
- Varian, H. R. 1980. "Redistributive Taxation as Social Insurance." *J. Public Econ.* 14 (1): 49–68.