Dynamic Optimal Taxation with Endogenous Skill Premia*

Zongbo Huang† Jason Ravit‡ Michael Sockin§

Princeton University Citigroup UT Austin (McCombs)

October 20, 2016

PRELIMINARY — PLEASE DO NOT CITE WITHOUT PERMISSION

Abstract

We embed imperfect substitutability across skill levels into a dynamic Mirrlees model and uncover a novel intertemporal wage compression channel in optimal labor taxation that can rationalize redistributive programs such as the Earned Income Tax Credit. In contrast to the wage compression channel found in static models, this dynamic channel lowers the optimal tax rate at the bottom because it allows the planner to reduce the cost of providing insurance to unskilled workers while deterring skilled agents from misreporting. The optimal labor tax is progressive in the short-run and our channel is quantitatively significant in comparison to other channels highlighted in the literature.

*We are grateful to Mark Aguiar, Mikhail Golosov, Harrison Hong, Nobuhiro Kiyotaki, Ilyana Kuziemko, Ben Moll, Galo Nuño, Richard Rogerson, Harvey Rosen, and Yuliy Sannikov for many helpful comments and discussions, and participants at the North America Summer Meeting of the Econometric Society, Princeton Student Macro/International Workshop, Microeconomic Theory Student Lunch Seminar, and Public Finance Working Group for feedback. The views expressed in this paper are those of the authors and do not necessarily reflect the position of Citigroup.

†Department of Economics and Bendheim Center for Finance. Princeton, NJ 08544. E-Mail: zongboh@princeton.edu

‡Financial Strategy and Solutions Group. New York, NY 10013. E-Mail: jravit314@gmail.com

§Department of Finance. Austin, TX 78712. E-Mail: michael.sockin@mccombs.utexas.edu

1 Introduction

A primary objective of tax systems in advanced countries is to combat rising income inequality, a phenomenon that has become more pronounced in recent decades.[1] Programs like the Earned Income Tax Credit (EITC) in the United States, for instance, aim to reduce inequality by targeting poverty, subsidizing low- and moderate-income workers through negative effective marginal tax rates.[2] While such redistributive policies are ubiquitous in practice, standard models of optimal taxation have difficulty rationalizing these salient features of the tax code. Negative marginal tax rates in these settings are inefficient because they reduce incentives to work, going against the primary goal of the EITC, which is to incentivize low-income people to work more by providing a transfer only if they actually work. We develop a dynamic Mirrlees model of optimal taxation in which workers with different skills are imperfect substitutes in the production process and uncover a novel "intertemporal wage compression" channel that provides a new rationale for the EITC, even in the absence of idle workers: negative marginal rates for low-income workers relax incentive constraints in earlier periods, making it cheaper for the planner to provide insurance against shocks earlier in life while simultaneously deterring high-income workers from shirking.

Similar to Mirrlees (1971) and its dynamic counterparts in the "new dynamic public finance" (NDPF) literature, we cast the optimal (nonlinear) tax function as the outcome of a mechanism design problem. In our model, agents receive privately-observed skill shocks over time and sell their labor to a representative firm that aggregates their effective labor (labor supplied multiplied by skill) into a final numéraire good. Importantly, we build on Stiglitz (1982) and Ales, Kurnaz and Sleet (2015) and assume that agents with different skills are imperfect substitutes in this production process. As a consequence, changes in relative labor supply change relative wages across workers, giving rise to skill premia. Our model is motivated by the findings in Katz and Murphy (1992), Acemoğlu and Autor (2011), and Autor (2014) that not only have skill premia risen over the last few decades, but this trend is a large contributor to growing inequality, especially in the bottom and middle of the distribution.[3] In static settings, imperfect substitutability gives rise to the "intratemporal wage compression" channel, first introduced in Stiglitz (1982), that reduces skill premia by

[1] See Piketty and Saez (2003) for an in-depth survey of the evidence.

[2] The United Kingdom's version is called the Working Tax Credit, and Austria, Belgium, Canada, Denmark, Finland, France, the Netherlands, New Zealand, and Sweden all have similar programs.

[3] Inequality between the bottom 99% and the top 1% has been driven by "superstar effects" and increases in CEO compensation driven by capital gains; see Gabaix et al. (2015) for a reduced form dynamic model and Ales and Sleet (2016) and Scheuer and Werning (2016) for static Mirrlees models with superstar effects.

encouraging (discouraging) labor that reduces (increases) the wages of high-skill (low-skill) workers. In particular, these wages make it costlier for high-skill workers to mimic low-skill workers. The planner can achieve this by subsidizing, i.e., a negative marginal rate, high-skill labor and taxing low-skill labor. Therefore, imperfect substitutability, in itself, cannot rationalize redistributive programs like the EITC.

Our dynamic model introduces a new, "intertemporal wage compression" channel.[4] Here, the planner relaxes high-skill agents' incentive constraints today by reducing their utility in *future* states that are unlikely to occur. In particular, if skill is persistent, a skilled agent who misreports makes the planner believe he will be unskilled in the future. To deter this agent from misreporting, the planner makes these future states very painful for the lying agent and instead redistributes heavily towards states that low-skill agents are likely to face, i.e., provides more insurance to low-skill agents. The planner does this by subsidizing low-skill labor (the EITC) in the future, which raises the supply of low-skill labor and lowers the low-skill wage (and raises skill premia). The threat of low wages and long hours in the future reduces variation in future utility and ensures that the dynamic loss from mimicking low-skill agents outweighs the static gain. Therefore, in our model, the purpose of the EITC is to lower the cost of providing valuable (dynamic) insurance to unskilled workers by relaxing skilled agents' incentive constraints in earlier periods. Negative rates are possible precisely because wages adjust downward to offset the higher labor supply. We are not making any statements about the EITC's current focus on the extensive margin; rather, by focusing on only the intensive margin, we highlight new, complementary dynamic insurance and incentive motives. Since incentivizing high-skill agents to report truthfully is more important than insuring them, the old "intratemporal" channel dominates our "intertemporal" channel in the right tail, as in Golosov, Troshkin and Tsyvinski (2016*b*). Further, since more skilled agents receive a higher wage, they can more easily mimic less skilled workers through excessive saving. This causes the intertemporal wedge, or implicit savings tax, to be higher for more skilled agents and lower for less skilled agents.

After characterizing the optimal tax system, we perform several numerical exercises to demonstrate that for empirically reasonable parameter values, our intertemporal wage compression channel is both qualitatively and quantitatively significant. While the wage com-

[4] Besides these two channels, optimal labor taxes are shaped by two additional forces. First, the planner must insure agents against shocks in the current period while incentivizing them to work, similar to Diamond (1998). Second, the planner can reduce the cost of insurance against earlier shocks by relaxing earlier incentive constraints. This is the "intertemporal" force from Golosov, Troshkin and Tsyvinski (2016*b*). Decreasing the substitutability between skill types lowers these two terms for low-skill workers and raises them for high-skill workers.

pression term is monotonically decreasing from above to below zero in static models, its dynamic counterpart is hump-shaped and is negative at both ends of the skill distribution. This pattern reflects the dynamic insurance and incentives motives highlighted above that are absent from static models. We illustrate that for a range of reasonable parameter values, the marginal labor tax rate in the left tail is negative, consistent with redistributive programs like the EITC. We also show that the labor tax is progressive in the short-run, in contrast to models with linear production functions in which skill premia are absent. In the cross-section, the optimal labor tax's shape is determined by both the traditional forces from standard Mirrlees models and the wage compression channel and exhibits the hump-shaped pattern described above. Welfare gains relative to simple tax policies are small.

In addition to characterizing the intertemporal wage compression channel, we make two methodological contributions to the literature. First, we are among the first to extend the techniques developed for heterogeneous agent models in continuous time by, for instance, Lasry and Lions (2007) and Nuño and Moll (2015), to a dynamic public finance setting that features imperfect substitutability in the production process. By casting our mechanism design problem in continuous time, we reduce the planner's problem to a coupled system of partial differential equations, the Hamilton-Jacobi-Bellman equation to solve for policy functions and the Kolmogorov Forward Equation to solve for the (stationary) distribution of agents in the economy.[5] Nuño and Moll (2015) develop a method to solve a social planner's problem with heterogeneous agents in settings in which the market clearing conditions are "linear." We adopt and extend their techniques to our setting with nonlinear aggregation.

Second, our model is an example of a dynamic mechanism design problem in which an agent's private information is persistent. Most mechanism design problems (both static and dynamic) derive the optimal contract with the "first-order approach," where the full set of incentive constraints is replaced with a single first-order condition. This first-order condition, however, is only a necessary condition for optimality, and one must often verify sufficiency ex post to check whether the contract is indeed globally optimal. While there are conditions that guarantee sufficiency of the first-order approach in dynamic Mirrlees models, they are often difficult to impose ex ante and must be checked numerically. In contrast, as long as agents cannot overreport their skill level, which is the same assumption as that in Williams (2011), we validate the "first-order approach." This restriction is typically satisfied in Mirrlees models.

[5]In addition to making the planner's problem more tractable, the tools we employ make the planner's problem more computationally manageable. Strong nonlinearities near the corners of the parameter space, however, impose limits on our ability to solve the model numerically.

The remainder of the paper is organized as follows. Section 2 discusses the related literature. Section 3 introduces the model and derives the planner's mechanism design problem. Section 4 highlights the key properties of the optimal tax system while Section 5 provides a quantitative analysis. Section 6 discusses some of the model's main assumptions and Section 7 concludes.

2 Related Literature

The benchmark static Mirrlees models include Mirrlees (1971), Diamond (1998), and Saez (2001). The NDPF literature adds dynamics to these models by allowing skill to be persistent over time. Early examples include Kocherlakota (2005), Albanesi and Sleet (2006), Golosov, Tsyvinski and Werning (2006); Farhi and Werning (2013) focus on time series properties of the tax system while Golosov, Troshkin and Tsyvinski (2016*b*) investigate the cross-sectional properties, including the top tax rate under various skill distributions. All of these models focus on a linear production technology, and consequently abstract from endogenous skill premia. As such, they lack the forces we highlight that give rise to the optimality of programs like the EITC.

Stiglitz (1982) was the first paper to embed imperfect substitutability into a static Mirrlees model, using the simplest possible setting: static with two types of agents. Jacobs (2012) extends this to include endogenous human capital accumulation and Ales, Kurnaz and Sleet (2015) include many types and a skill-to-task assignment framework. All three find that the wage compression channel is a force for a positive marginal rate at the bottom and negative rate at the top, which is the opposite of what is implied by programs like the EITC.[6] Rothschild and Scheuer (2013, 2014) show that task choice partially undoes the wage compression force in a Mirrlees model in which workers have different skills in different tasks and choose how much to work on each task. Our framework focuses on the interaction between insurance in a dynamic setting and this wage compression channel.

To our knowledge, the only dynamic model of nonlinear taxation with imperfect substitutability is Heathcote, Storesletten and Violante (2016). Though they use the same production technology as us and are motivated by the skill premium literature, they solve for the optimal tax system within a particular class of functions that they argue closely

[6]The reason for the negative rate is that they assume skill is bounded, which Diamond and Saez (2011) argue is not practically useful for thinking about top tax rates. Even with unbounded support and a fat-tailed distribution so that the top rate is positive, however, it is still lower than what it would be with a linear production function

approximates the actual U.S. system. This also means that they take the EITC as given. As far as we are aware, ours is the first dynamic Mirrlees model with endogenous skill premia. While our optimal tax functions are qualitatively similar away from the right tail, ours also allows for history-dependence whereas theirs is a restricted function of only current income.

There are several studies, such as Diamond (1980), Saez (2002), Choné and Laroque (2005), Laroque (2005), Beaudry, Blackorby and Szalay (2009), and Choné and Laroque (2010) that show that redistributive programs like the EITC are desirable. Saez (2002) shows that the optimality of the EITC hinges on two key ingredients: labor supply responses must be concentrated on the extensive margin (whether or not to work instead of how much)[7] and the social planner must put disproportionately-large welfare weights on low-skill workers. Even with a powerful extensive margin, without sufficiently strong redistributive preferences there is no EITC in his model.[8] In contrast, our model focuses exclusively on the intensive margin and the planner is utilitarian, which allows us to isolate the dynamic insurance and incentive value. Rothstein (2010) argues that the intensive margin weakens the EITC's effectiveness because it lowers unskilled workers' wages; our model shows that these costs are outweighed by additional dynamic insurance benefits for unskilled workers and incentive benefits for skilled workers. Heckman, Lochner and Cossa (2003) argue that the EITC can be justified on the grounds that it helps boost human capital accumulation; Stantcheva (2015) demonstrates that the planner indeed should subsidize human capital investment for unskilled workers in a dynamic Mirrlees model, but the subsidy acts as a force for intertemporal savings and the labor tax for these workers is still positive; our model does not require human capital to justify the EITC and hence, like the actual EITC, is not earmarked for specific expenditures.

We solve our heterogeneous agent model by using techniques from the theory of mean field games developed by Lasry and Lions (2007), Achdou et al. (2015), and Nuño and Moll (2015). Gabaix et al. (2015) use these techniques to study the dynamics of inequality over time while Kaplan, Moll and Violante (2016) use them to study monetary policy. To our knowledge we are the first to use them to study optimal contracts and tax policy.

An advantage of working with these techniques is that characterizing the economy's stationary equilibrium is straightforward, and we introduce stochastic retirement shocks in a perpetual youth framework to ensure that a stationary distribution exists.[9] This also

[7] See Eissa and Liebman (1996) for more on how the EITC affects labor supply incentives.

[8] Choné and Laroque (2010), unlike the other authors listed, do not require an extensive margin or strong redistributive preferences; they do, however, require multidimensional heterogeneity.

[9] This contrasts with life cycle models, where agents know exactly when they will retire and this date is the same for everyone. Our model is thus a middle ground between the classic life cycle models and

simplifies the tax system because it removes age/time as a state variable.

Finally, our analysis of the first-order approach in dynamic mechanism design problems follows Fernandes and Phelan (2000), Kapička (2013), and Pavan, Segal and Toikka (2014) in discrete time, and Williams (2011) and Sannikov (2014) in continuous time. While it is well-known that there does not exist a simple, robust condition for the first-order approach to be sufficient for global optimality in dynamic settings, Kapička (2013) and Pavan, Segal and Toikka (2014) derive integral monotonicity conditions for sufficiency in their settings that can be checked numerically ex post.[10] Williams (2011) shows that the first-order approach is sufficient as long as agents' reports are not too volatile, but he shows that his condition is unnecessarily restrictive. By using proof techniques developed by Sannikov (2014), however, we show that the first-order approach is sufficient as long as agents cannot overreport their types. Even though our setting is similar to that in Williams (2011), his proof uses a linear approximation to the utility gain from deviating, while we use a quadratic one, which allows us to derive a sharper sufficient condition.

3 Model

3.1 Environment

Time is continuous and is indexed by $t \in [0, \infty)$. Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$ denote a complete, filtered probability space, where the filtration $\{\mathcal{F}_t\}$ satisfies the usual conditions. The economy has a unit measure of agents who work until they retire. When an agent retires, a new agent is born to maintain a constant population size. Let $i \in [0, 1]$ index individual agents.

Agents are born with heterogeneous earning ability/skill θ_0^i. Ability at each date is private information and evolves according to a diffusion process

$$d\theta_t^i = \mu\left(\theta_t^i\right) dt + \sigma\left(\theta_t^i\right) dZ_t^i, \tag{1}$$

where Z_t^i is a standard Wiener process adapted to the filtration for all (i, t) in the economy[11] and $\mu(\theta), \sigma(\theta) \in C^2(\Omega)$ are measurable functions that satisfy the standard Lipschitz and

the endogenous retirement model of Shourideh and Troshkin (2015) in which production is linear and skill evolves deterministically over the life cycle.

[10] Kapička (2013) shows that his setting nests dynamic Mirrlees models and Golosov, Troshkin and Tsyvinski (2016*b*) check that his condition holds in their model.

[11] Going forward, we will omit the i-superscript when doing so does not create confusion.

growth conditions:

$$\|\mu(t,x)-\mu(t,y)\| \leq A|x-y|\,,\ \|\sigma(t,x)-\sigma(t,y)\| \leq A|x-y|\,,$$
$$\|\mu(t,x)\| \leq B(1+|x|)\,,\ \|\sigma(t,x)\| \leq B(1+|x|)\,,$$

for some constants $A, B > 0$. Going forward, let $\mathcal{F}_t$ denote the natural filtration with respect to $\{\theta_t\}$. The specification in (1) is general enough to accommodate many standard functional forms. Unless otherwise stated, however, we take $\sigma(\theta) = \sigma_\theta \theta$. While most of our results go through for more general processes, this specification keeps the algebra simple and nests some special cases that we discuss below. The Markovian structure of the process means that we do not need to keep track of entire histories of shocks and going forward, all functions are such that, when applied to $\{\theta_t\}$, they are $\mathcal{F}$-adapted.

We restrict $\theta_t \in \Theta \subseteq \mathbb{R}_+$, and define reflecting boundaries $\underline{\theta} = \inf \Theta > 0$, $\overline{\theta} = \sup \Theta < \infty$. The lower bound assumption is crucial for generating long-run dynamics since otherwise, all agents will have zero ability with probability one over the course of their lifetimes since zero is an absorbing boundary. Let $\theta^t = \{\theta_s\}_{s\in[0,t]}$ denote a history of abilities up to date t, and Θ^t the set of all histories up to date t.

As is standard in the optimal taxation literature, θ_t determines an agent's stochastic productivity, which in turn determines an agent's labor income. In particular, an agent with ability θ_t has productivity

$$e_t = e_t(\theta_t)\,,$$

where e_t is strictly increasing and concave. In most NDPF models, $e_t(\theta_t) = \theta_t$ and to simplify the math, we make this assumption, too. Agents supply labor $\ell_t \geq 0$ at date t, and hence $\theta_t \ell_t$ units of effective labor. An agent's total labor income at each date is his total units of effective labor supplied multiplied by his wage per unit of effective labor, $w(\theta_t)$:

$$y_t = w(\theta_t)\,\theta_t \ell_t.$$

Agents work and consume, c_t, until they are hit by a Poisson retirement shock, R_t^i, with intensity κ. An agent is hit (or not hit) with the shock at the beginning of each date so that agents who are not hit work and consume as usual, while agents who are hit immediately retire and collect a "social security" payment. The retirement shock is observable. An alternative interpretation is that agents are hit with a "disability shock" and transfers take the form of disability payments.[12] By the weak law of large numbers, exactly a fraction κ of agents of each type retire at each date. After being hit by the shock, $R_t^i = 1$, an agent retires and receives a lump-sum transfer C_t. After receiving this transfer, the agent exits the economy. When a new agent is born, his initial ability is drawn from a lognormal distribution, $\log(\theta_0) \sim \mathcal{N}(\mu, \upsilon^2)$.

An important consequence of stochastic retirement is that the economy will have a stationary equilibrium and the state variables will have a stationary distribution.[13] In fact, several distributions that are exogenously assumed in the literature emerge endogenously in our setting.

Example 1. *Suppose* $\kappa = 0$ *and*

$$d\theta_t = \theta_t \left[-(1-p)(\log(\theta_t) - \log(\theta^*)) + \frac{\sigma_\theta^2}{2} \right] dt + \sigma_\theta \theta_t \, dZ_t.$$

In this case, Itô's Lemma implies

$$d(\log(\theta_t)) = -(1-p)(\log(\theta_t) - \log(\theta^*))\, dt + \sigma_\theta \, dZ_t.$$

This is called an Ornstein-Uhlenbeck process with persistence $p \in [0,1]$ *and is the continuous time analogue of the AR(1) process in discrete time. As long as* $p < 1$ *and* $\overline{\theta} \to \infty$ *then* θ *has a lognormal stationary distribution with mean and variance that depend on the parameters.*

Example 2. *Suppose*

$$d\theta_t = \left(\mu_\theta + \frac{\sigma_\theta^2}{2} \right) \theta_t \, dt + \sigma_\theta \theta_t \, dZ_t$$

and $\overline{\theta} \to \infty$*. Then Itô's Lemma implies*

$$d(\log(\theta_t)) = \mu_\theta \, dt + \sigma_\theta \, dZ_t,$$

so that $\log(\theta_t)$ *is a geometric Brownian motion. In this case, the stationary distribution of*

[12] See Golosov and Tsyvinski (2006) for a model optimal disability insurance in which the retirement/disability shock is unobservable.

[13] Stochastic retirement is not always necessary to guarantee the existence of a stationary distribution, but many times it is sufficient.

θ is a double Pareto-lognormal distribution with probability density function[14]

$$g(\theta)=\frac{ab}{a+b}\Bigg[\theta^{-(1+a)}\exp\left(a\mu+\frac{a^2\upsilon^2}{2}\right)\Phi\left(\frac{\log(\theta)-\mu-a\upsilon^2}{\upsilon}\right)$$
$$+\theta^{b-1}\exp\left(-b\mu+\frac{a^2\upsilon^2}{2}\right)\left(1-\Phi\left(\frac{\ln(\theta)-\mu+b\upsilon^2}{\upsilon}\right)\right)\Bigg],$$

where a and $-b$ $(a,b>0)$ are the roots of the characteristic equation

$$\frac{\sigma_\theta^2}{2}\zeta^2+\left(\mu_\theta-\frac{\sigma_\theta^2}{2}\right)\zeta-\kappa=0.$$

The double Pareto distribution is a special case when $\upsilon\to 0$, the Pareto lognormal distribution[15] *is a special case when $b\to\infty$, and the lognormal distribution is a special case when $a=b\to\infty$.*

Agents have subjective discount factor $\rho>0$. Each agent's per-date utility while working is separable in consumption and effort:

$$\tilde{u}(c_t,y_t;\theta_t)\equiv u(c_t)-\phi\left(\frac{y_t}{w(\theta_t)\,\theta_t}\right).$$

As usual, $u\in C^2(\Omega)$ is increasing and concave while $\phi\in C^2(\Omega)$ is increasing and convex. Let $v_t^R(C_t)$ denote an agent's utility upon receiving the terminal transfer payment C_t. Agents place weight ψ on utility received in retirement so that if an agent retires in date t, he receives utility $\psi u(C_t)$.

An allocation $x_t:\Theta^t\to\mathbb{R}_+^3$ is an $\mathcal{F}_t$-progressively measurable triple that specifies consumption, terminal consumption, and labor supplied at each date t, conditional on the history θ^t: $x_t=\{x(\theta^t)\}_{\Theta^t}=\{c(\theta^t),C(\theta^t),\ell(\theta^t)\}_{\Theta^t}$. Then an agent's expected lifetime utility from an allocation is given by

$$U(x_t)=\mathbb{E}_0\left[\int_0^\tau e^{-\rho t}\tilde{u}\left(c\left(\theta^t\right),\ell\left(\theta^t\right)\right)dt+e^{-\rho\tau}v_\tau^R\left(C\left(\theta^\tau\right)\right)\right], \tag{2}$$

where $\tau=\inf\{t\mid R_t=1\}$ is the date when the retirement shock hits. Since the shock has a Poisson arrival time, we can rewrite (2) as

$$U(x_t)=\mathbb{E}_0\left[\int_0^\infty e^{-(\rho+\kappa)t}\left[\tilde{u}\left(c\left(\theta^t\right),\ell\left(\theta^t\right)\right)+\kappa v_t^R\left(C\left(\theta^t\right)\right)\right]dt\right].$$

[14] See Reed and Jorgensen (2004).
[15] See Colombi (1990).

In particular, the shock acts just like a reduction in the discount factor.

Agents sell their labor to a representative firm and pay taxes on the income they earn. They can also earn net interest rate r on riskfree savings, a_t. Let $T : \mathbb{R}^2_+ \to \mathbb{R}$ denote a (non-stochastic) tax function on labor income and savings. Then the agent solves

$$\max_{\{c_t, C_t, \ell_t\}} U(x_t)$$

subject to the dynamic budget constraint

$$da_t = [w(\theta_t)\theta_t \ell_t + ra_t - T(w_t(\theta_t)\theta_t\ell_t, ra_t) - c_t]\, dt.$$

A representative firm hires all workers and produces output via a constant elasticity of substitution (CES) production function. The firm maximizes output, Y_t, minus labor costs. That is, the firm solves

$$\max_{\{L(\theta_t)\}} \left\{ \left(\int_\Theta L(\theta_t)^{\frac{\alpha-1}{\alpha}}\, d\theta_t \right)^{\frac{\alpha}{\alpha-1}} - \int_\Theta w(\theta_t)\, L(\theta_t)\ d\theta \right\}$$

for all t, where

$$L(\theta_t) = \int_0^1 \theta_t \ell\left(\theta_t^i\right) g\left(\theta_t^i\right) di$$

is the total units of effective labor supplied by workers with skill θ_t and g is the probability density function over all agents.[16] In particular, all the firm cares about is skill at t, not the entire skill history, so that all agents with the same skill level at t are perfect substitutes while agents with different skill levels are imperfect substitutes. This means that an agent's wage is a function of only his skill today, not his entire history. The firm does not directly observe skill but posts a menu of wage contracts $\{w(\theta_t)\}$ to induce agents to choose the contract meant for them. From the firm's problem,

$$w(\theta_t) = \left(\frac{Y_t}{L(\theta_t)} \right)^{\frac{1}{\alpha}}. \tag{3}$$

This functional form is as in the so-called "canonical model" in the skill premium literature but with all skill levels entering the production function symmetrically. However, it can be easily modified to include skill bias, as we do later, so that the technology puts more weight on more productive workers, in which case the technology resembles that of Katz

[16] Integrating over i is shorthand notation for integrating over histories that terminate at θ_t.

and Murphy (1992). Although CES production functions are not widely used in the public finance literature, Heathcote, Storesletten and Violante (2016) and Ales, Kurnaz and Sleet (2015) both use variations of this function.

As $\alpha \to \infty$, then $w(\theta_t) \to 1$ and we are back in the familiar NDPF setting in which agents are perfect substitutes across types. Now, however, each agent's wage depends not only on his own ability and labor supply decision, but on every other agent's ability and labor supply decision (through Y_t). That the wage differs across agents means there are skill premia across types: the skill premium between agents i and j is

$$\pi_t^{i,j} \equiv \frac{w\left(\theta_t^i\right)}{w\left(\theta_t^j\right)} = \left(\frac{L\left(\theta_t^j\right)}{L\left(\theta_t^i\right)}\right)^{\frac{1}{\alpha}}.$$

Again, as $\alpha \to \infty$ then $\pi_t^{i,j} \to 1$ and there are no skill premia.

Our equilibrium concept follows Mirrlees (1971):

Definition 1 (Tax Equilibrium)**.** *A tax equilibrium is a tax function $T : \mathbb{R}_+^2 \to \mathbb{R}$, an allocation $\{x_t\}$, and wage profiles $\{w(\theta_t)\}$ such that, given public spending $\{G_t\}$:*

1. *The allocation solves an agent's problem;*

2. *For each θ the wage is given by (3);*

3. *The goods market clears,*

$$\int_{\Theta^t} \left[c\left(\theta^t\right) + \kappa C\left(\theta^t\right)\right] g\left(\theta^t\right) d\theta^t + G_t \leq Y_t.$$

Let X denote the set of equilibrium allocations. To characterize tax equilibria we solve for the optimal allocation via a mechanism design problem, with wages and taxes set to ensure that the allocation is implemented as part of the equilibrium.

3.2 Planner's Problem

Following Mirrlees (1971), while the planner observes each agent's income, y_t, and consumption, c_t, C_t, he cannot observe ability, θ_t. Therefore, he cannot observe an agent's labor supply decision, ℓ_t. The planner relies on agents to report θ_t at each date.

In this section, we set up the planner's mechanism design problem. A mechanism specifies, given each agent's report, how much an agent should produce and consume, i.e., an allocation. We do this in two steps: first, we characterize the dynamics of the optimal contract by formulating a continuous time version of the first-order approach and deriving a set of necessary and sufficient conditions for the optimal contract under the relaxed problem to coincide with the optimal contract of the full problem. This also delivers the relevant state variables. Then, we combine the state variables' laws of motion derived via the first-order approach with the economy-wide law of motion and the firm's constraints to derive the planner's full recursive problem.

3.2.1 Contracting Problem

By the Revelation Principle, we can restrict attention to direct revelation mechanisms in which agents truthfully report their types. A report $\sigma_t(\theta^t) \in \Sigma_t$ is an $\mathcal{F}_t$-progressively measurable function that specifies some $\theta_t \in \Theta$ that the agent reports to the planner following the history θ^t, and a reporting strategy $\sigma = \{\sigma_t(\theta^t)\} \in \Sigma$ is a history of reports. Let $\sigma^t \in \Sigma^t$ denote the history of reports generated by reporting strategy σ. The planner can specify allocations directly as a function of histories of reports, $x(\theta^t) = x(\sigma^t)$. Without loss of generality, we can focus on reporting strategies with $\Sigma^t \subseteq \Theta^t$ since otherwise, the planner immediately detects a lie.

Given a history θ^t and a reporting strategy σ, let $v_t^\sigma(\theta^t)$ denote an agent's promised utility/continuation value:

$$v_t^\sigma\left(\theta^t\right) \equiv \mathbb{E}_t\left[\int_t^\tau e^{-\rho(s-t)}\tilde{u}\left(\sigma_s\left(\theta^s\right)\right)ds + e^{-\rho\tau}v_\tau^R\left(\sigma_\tau\left(\theta^\tau\right)\right)\right].$$

The truth-telling strategy specifies $\sigma_t(\theta^t) = \theta_t$ after all histories, for all t. Let $v_t \in \mathcal{V}$ denote the value of v_t^σ under truth-telling. We say that an allocation is *incentive-compatible* if truth-telling yields weakly higher continuation utility than any other reporting strategy σ:

$$v_t\left(\theta^t\right) \geq v_t^\sigma\left(\theta^t\right) \;\; \forall\sigma \in \Sigma,\;\; \theta^t \in \Theta^t. \tag{4}$$

Let X^{IC} denote the set of incentive-compatible allocations, i.e., those allocations that satisfy the incentive-compatibility constraint (4). Of course, for an agent to participate in the game, his utility from an allocation must exceed that of some outside option:

$$U(x) \geq \underline{U}. \tag{5}$$

This is an agent's participation constraint.

While agents do not have access to a savings market, the planner can save on their behalf, earning net interest rate $r > 0$. Then the cost to the planner of delivering an allocation $x = \{c, C, \ell\}$ is

$$\Psi(x) = \mathbb{E}_0 \left[\int_0^\infty e^{-rt} \left[c\left(\theta^t\right) - y\left(\theta^t\right) + \kappa C\left(\theta^t\right) \right] dt \right].$$

The expectation $\mathbb{E}_0$ is over the distribution of agents in the economy, $c - y$ is the planner's net cost to the $1-\kappa$ agents who do not retire, while κC is the cost of social security payments to the κ agents who do retire. The planner selects a tax equilibrium, i.e., an allocation x and tax function T, that is individually rational and incentive compatible for all agents, and is as cheap as possible.

The planner has a utilitarian welfare function so that all agents are equally-weighted. Then the planner's contracting problem, i.e., ignoring the aggregate law of motion and firm's problem, written sequentially, is

$$K\left(\underline{U}\right) \equiv \min_{x \in X} \Psi(x)$$

subject to (4), (5), and $c\left(\theta^t\right), C\left(\theta^t\right), \ell\left(\theta^t\right) \geq 0$. Let x^* denote an optimal equilibrium allocation and T^* an optimal equilibrium tax function. The planner has both an insurance and redistributive motive: the planner insures agents against fluctuations over time in their ability and redistributes, within a date, from higher-ability agents to lower-ability ones. At the same time, the tax system must induce agents to supply labor efficiently. Therefore, the planner must balance insurance/redistribution and incentives.

Unfortunately, this problem is intractable as currently written. The reason is that (4) is a hopelessly complicated set of constraints because we need to rule out every possible deviation following every possible history. To get around this, we use the first-order approach. Instead of dealing with global deviations, we will instead focus on only small, "local" deviations. That is, agents can tell only "small" lies. As we will show, this reduces (4) to a simple first-order condition, which makes the problem much easier to solve. At that point, we will formulate the planner's relaxed problem recursively. Finally, we have to go back and check ex post that the solution to the relaxed problem also solves the full problem above. However, we will derive a simple sufficient condition under which this holds.

Our problem is an example of a dynamic mechanism design problem because agents' private information, θ_t, persists through time. In ordinary mechanism design problems, θ_t and v_t are the only state variables, the latter in order to track promise-keeping through time. Here, however, following Fernandes and Phelan (2000), Kapička (2013), and Farhi and Werning (2013) in discrete time and Williams (2011) and Sannikov (2014) in continuous time, we need a third state variable to track threat-keeping.

To formulate the planner's problem recursively, we need each state variable's law of motion. We already have the law of motion for θ_t, which is simply (1). Our first result, which is standard in the continuous time optimal contracting literature, establishes the law of motion for v_t.

Proposition 1. *Fix a contract and a reporting strategy $\{\widehat{\theta}_t\}$ with finite expected payoff to the agent. Then the process $\{v_t\}$ corresponds to the agent's continuation utility if and only if there exists a process $\widehat{\Delta}_t \in H^2$ such that*[17]

$$dv_t = (\rho v_t - \tilde{u}_t)\, dt + \sigma_\theta \theta_t \widehat{\Delta}_t \, dZ_t + \left(v_t^R - v_t\right)(dR_t - \kappa \, dt) \tag{6}$$

and the transversality condition $\mathbb{E}_t \left[e^{-\rho(T\wedge\tau)} v_{T\wedge\tau}\right] \to 0$ as $T \to \infty$ holds.

The drift of (6) tracks promise-keeping: $\tilde{u}_t$ is just-delivered utility and ρv_t is everything owed going forward. The process $\widehat{\Delta}_t$ is the sensitivity of v_t to shocks in the underlying process θ_t. This also has a natural interpretation in terms of options. If we think of the contract between the principal and the agent as a package of call options, then v_t represents the value of the options and $\widehat{\Delta}_t$ is the "delta" of the options, i.e., the sensitivity of the value to changes in the underlying. In addition, note that promised utility can jump at retirement: if, for example, v_t jumps down when the agent retires, then it must drift up when the agent works to compensate him.

If an agent misreports today, this has long-run effects because these misreports distort the planner's beliefs of how θ_t will evolve going forward. That is, a misreport leads the planner to believe that θ_t is following a different trajectory than it actually is, potentially with large consequences down the line. To deter such deviations, the planner has to provide the agent with some information rent. Formally, an agent's information rent is the sensitivity of his continuation value to his report, evaluated at truth-telling:

$$\Delta_t \equiv \left.\frac{\partial v_t}{\partial \widehat{\theta}_t}\right|_{\widehat{\theta}_t = \theta_t} = \frac{\partial v_t}{\partial \theta_t}. \tag{7}$$

Let $\Delta_t \in \mathcal{D}$. Intuitively, this is simply the impulse response function of v_t to a small change in $\widehat{\theta}_t$: small lies can have potentially large effects because they distort the planner's beliefs

[17] A process $\{X_t\}$ is in the space $H^2[0,T]$ if it is $\mathcal{F}_t$-progressively measurable and is square-integrable, $\mathbb{E}_t\left[\int_0^T X_t^2 \, dt\right] < \infty$.

about θ_t in the future. As a technical aside, this is not an ordinary derivative, but what is called a Malliavin derivative.[18] In discrete time models, we can test the effects of misreports via the one-shot deviation principle but such a result does not carry over to continuous time. The Malliavin derivative is a way of formalizing impulse response functions in continuous time by defining what it means to "differentiate" a Brownian motion. That said, the reader will not lose anything by thinking of Δ_t as an ordinary derivative. We are now ready to state the first-order condition.

Proposition 2. *A necessary condition for truth-telling to be optimal is that for all t,*

$$\widehat{\Delta}_t = \Delta_t, \tag{8}$$

where $\widehat{\Delta}_t$ is the process in Proposition 1.

This result is intuitive: the sensitivity of continuation utility must equal the sensitivity evaluated at truth-telling. Let X^{FOA} denote the set of allocations that satisfy (8), which necessarily satisfies $X^{\text{IC}} \subseteq X^{\text{FOA}}$. Using Proposition 2, we can derive the law of motion.

Proposition 3. *The finite process $\{\Delta_t\}$ is characterized by (7) if and only if, for some $\sigma_{\Delta,t} \in H^2$,*

$$d\Delta_t = \left(\left(\rho - \mu'\left(\theta_t\right)\right)\Delta_t - \tilde{u}_{\theta,t} - \sigma_\theta^2 \sigma_{\Delta,t}\right) dt + \sigma_\theta \sigma_{\Delta,t}\, dZ_t + \left(\Delta_t^R - \Delta_t\right)\left(dR_t - \kappa\, dt\right) \tag{9}$$

and the transversality condition $\mathbb{E}_t\left[e^{-\rho(T\wedge\tau)}\Delta_{T\wedge\tau}\right] \to 0$ as $T \to \infty$ holds.

The planner uses the process $\sigma_{\Delta,t}$ to control the flow of information rent over time. In the drift of (9), $-\mu'\left(\theta_t\right)\Delta_t$ is the decay in information rent while $\tilde{u}_{\theta,t} + \sigma_\theta^2 \sigma_{\Delta,t}$ is delivered information rent. The term Δ^R is the sensitivity of continuation utility at retirement, where Δ_t can jump: if it jumps down then the principal must provide more information rent while the agent works. Continuing with the options analogy, $\sigma_{\Delta,t}$ is related to the "gamma" of the option, or the sensitivity of the delta to changes in the underlying process.

The planner's relaxed problem is

$$K\left(\underline{U}\right) \equiv \min_x \Psi\left(x\right)$$

subject to

$$v_0 = \mathbb{E}_0\left[\int_0^\infty e^{-(\rho+\kappa)t}\left(\tilde{u}_t + \kappa v_t^R\right) dt\right],$$

[18] For an introduction to Malliavin calculus, see Di Nunno, Øksendal and Proske (2009).

the first-order incentive constraint (8) and $c, C, \ell \geq 0$. Consider the state variables $(\theta_t, v_t, \Delta_t)$ driven by the controls $\{c_t, C_t, \ell_t, \sigma_{\Delta,t}\}$ according to the above laws of motion and the appropriate transversality conditions. Then Propositions 1 and 3 imply that optimizing the above objective function subject to these laws of motion is equivalent to the relaxed problem.

An important issue is whether or not the solution to the relaxed problem is necessarily a solution to the full problem. It is well-known that in static problems the Spence-Mirrlees single-crossing condition and a simple monotonicity condition are together sufficient for the solutions to coincide. However, in dynamic mechanism design problems, the conditions (see Kapička (2013) and Pavan, Segal and Toikka (2014)) are often extremely complicated and unintuitive. In these cases, one must solve the relaxed problem and then verify numerically ex post that it indeed solves the whole problem. Fortunately, by working in continuous time, we can provide a simple, intuitive condition. We first need the following assumption.

Assumption 1. *The drift of an agent's skill process satisfies* $2\mu'(\theta_t) + \sigma_\theta^2 < \rho$.

This assumption rules out processes that are "too explosive." If $\log(\theta_t)$ is an Ornstein-Uhlenbeck process, then this condition is

$$\frac{2\sigma_\theta^2 - \rho}{2(1-p)} + \log(\theta^*) < 1 + \log(\theta_t),$$

or $\sigma_\theta < \rho$ if $p = 1$ and $\log(\theta_t)$ is a random walk. While the literature usually sets $p = 1$ so that our condition trivially holds in our calibration, the condition still holds under common estimates of $(\sigma_\theta, \theta^*)$ if $p < 1$ as long as θ_t is not too small. If $\log(\theta_t)$ is a Brownian motion with drift, then the condition is

$$\mu_\theta + \sigma_\theta^2 < \frac{\rho}{2},$$

which holds for reasonable values of $(\mu_\theta, \sigma_\theta, \rho)$ (typically, $\rho \gg \sigma_\theta^2 > \mu_\theta$).

Theorem 1. *Suppose agents cannot overreport their types,* $\widehat{\theta}_t \leq \theta_t$ *for almost all* t. *Then under Assumption 1 and the transversality condition of Proposition 1, the optimal contract under the first-order approach is globally optimal.*

The proof of this result uses a technique from Sannikov (2014): we construct an upper bound for utility following a deviation and verify that an agent's deflated gains process is a supermartingale under any deviation, but a true martingale under truth-telling.[19] The

[19]Sannikov (2008) jointly derives a first-order necessary condition and proves it is sufficient to characterize the optimal contract by considering strategies that play arbitrarily until date t and optimally thereafter, and showing that the implied deflated gains process is a supermartingale. This approach does not work here because we need strategies to be absolutely continuous with respect to $\{\theta_t\}$.

condition in the Theorem that agents cannot overreport is fairly ubiquitous in the optimal taxation literature. Most static models, such as Stiglitz (1982), Jacobs (2012), and Ales, Kurnaz and Sleet (2015) assume it outright. Ales and Maziero (2009) prove that agents never overreport in a two-type dynamic model, while Kapička (2013) numerically solves the model in the region where agents cannot overreport; Williams (2011) explicitly restricts strategies so that agents cannot overreport. One way to micro-found this restriction is to assume that an agent must be able to show that he has any income that he reports.

3.2.2 Full Problem

Now we formulate the planner's full recursive problem. In addition to the contracting constraints, the planner must account for the firm's problem and the labor market clearing conditions. Rearranging (3),

$$w\left(\theta_t\right)^{\alpha} L\left(\theta_t\right) = Y_t.$$

Combining this with the labor market clearing condition for each θ yields

$$w\left(\theta_t\right)^{\alpha} \underbrace{\int_{\mathcal{V}\times\mathcal{D}} \theta_t \ell\left(\theta_t, v_t, \Delta_t\right) g\left(\theta_t, v_t, \Delta_t\right) d\left(v_t, \Delta_t\right)}_{=L(\theta_t)} = \underbrace{\int_{\Theta\times\mathcal{V}\times\mathcal{D}} w\left(\theta_t\right) \theta_t \ell\left(\theta_t, v_t, \Delta_t\right) g\left(\theta_t, v_t, \Delta_t\right) d\left(v_t, \Delta_t, \theta_t\right)}_{=Y_t} \tag{10}$$

for all θ. The first integral is over histories up to but not including θ_t while the second integral is over the entire history θ^t. The wage $w\left(\theta_t\right)$ acts like an additional state variable in the planner's problem so we need its law of motion in order to apply standard stochastic optimal control methods. Define

$$w'\left(\theta_t\right) = z\left(\theta_t\right). \tag{11}$$

This ensures that the wage function is differentiable and the tax equilibrium is well-behaved. Here, z_t is a control so the planner controls the shape of the wage function. Recall that since all the firm cares about is skill at t, the wage is a function of θ_t alone and the planner must respect this measurability constraint. Finally, we need a law of motion for the economy-wide distribution, $g\left(\theta, v, \Delta, t\right)$, which is given by the Kolmogorov Forward Equation (KFE)

$$\frac{\partial g}{\partial t} = \mathcal{A}^* g + \kappa \widetilde{f}\left(\theta_0, v_0, \Delta_0\right), \tag{12}$$

$$1 = \int_{\Theta\times\mathcal{V}\times\mathcal{D}} g\left(\theta, v, \Delta, t\right) d\left(\theta, v, \Delta\right), \tag{13}$$

where $\widetilde{f}(\theta_0, v_0, \Delta_0) = \widetilde{f}_0$ is the density of reborn agents[20] and $\mathcal{A}^* g$ is the adjoint of the infinitesimal generator

$$\mathcal{A}^* g = -\sum_{i=1}^{n} \frac{\partial}{\partial s_i} \left[\mu_i(\boldsymbol{s}, \boldsymbol{m}, \boldsymbol{A}) \, g(\boldsymbol{s}, t)\right] + \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \frac{\partial^2}{\partial s_i \partial s_j} \left[\left(\sigma(\boldsymbol{s}) \, \sigma(\boldsymbol{s})^{\mathbf{T}}\right)_{i,k} g(\boldsymbol{s}, t)\right] - \kappa g(\boldsymbol{s}, t).$$

Here, $\boldsymbol{s} = (s_1, \ldots, s_n)$ is the vector of state variables, $\boldsymbol{m}$ is the vector of controls, and $\boldsymbol{A}$ is the vector of aggregate variables. In our case, these are, respectively, (θ, v, Δ), $\{c, C, \ell, \sigma_\Delta, z\}$, and the wages $w(\theta)$. We will solve for the stationary equilibrium so that the left side of (12) is zero and g does not depend on t.

Summarizing, the planner's objective is to find the cheapest possible allocation subject to the participation constraint (6), the incentive constraint (9), the market clearing conditions (10), the wage law of motion (11), and the law of motion of the distribution as captured by the KFE, (12) and (13). We call an allocation x *efficient* if it solves this problem.

Let $k = k(\theta, v, \Delta)$ denote the planner's cost of insuring an agent in state (v, Δ, θ) and let $K(\theta) = \int_{\mathcal{V}\times\mathcal{D}} k(\theta, v, \Delta) \, g(\theta, v, \Delta) \, d(v, \Delta)$ denote the social impact function for each θ. The total cost is then $K = \int k(\theta, v, \Delta) \, g(\theta, v, \Delta) \, d(\theta, v, \Delta)$. We follow Nuño and Moll (2015) to derive the planner's Hamilton-Jacobi-Bellman (HJB) equation. The planner's objective function, written as an inner product,[21] is

$$K_t = \left\langle e^{-rt} \left(c(\theta_t, v_t, \Delta_t) + \kappa C(\theta_t, v_t, \Delta_t) - w(\theta_t) \, \theta_t \ell(\theta_t, v_t, \Delta_t)\right), g(\theta_t, v_t, \Delta_t)\right\rangle_{\Theta\times\mathcal{V}\times\mathcal{D}\times[0,\infty)}.$$

Let $\chi(\theta_t)$ denote the Lagrange multiplier on (10) evaluated at θ_t and $\eta(\theta_t)$ the multiplier on (11). Then the Lagrangian functional for the planner's problem at each t is

$$\begin{aligned}\mathcal{L}(g_t) = {} & \left\langle e^{-rt} \left(c_t + \kappa C_t - w(\theta_t) \, \theta_t \ell_t + \mathcal{A} k_t - r k_t + \partial_t k_t + \kappa k_t \widetilde{f}_0\right), g(\theta_t, v_t, \Delta_t)\right\rangle_{\Theta\times\mathcal{V}\times\mathcal{D}\times[0,\infty)} \\ & + \left\langle e^{-rt} \chi(\theta_t), w(\theta_t)^{\alpha} L(\theta_t) - Y_t\right\rangle_{\Theta\times\mathcal{V}\times\mathcal{D}\times[0,\infty)} + \left\langle e^{-rt} \eta(\theta_t), z(\theta_t)\right\rangle_{\Theta\times\mathcal{V}\times\mathcal{D}\times[0,\infty)},\end{aligned}$$

where $L(\theta_t)$ and Y_t are as in (10) and $\mathcal{A}k$ is the infinitesimal generator

$$\mathcal{A}k = \sum_{i=1}^{n} \frac{\partial k}{\partial s_i} \mu_i(\boldsymbol{s}, \boldsymbol{m}, \boldsymbol{A}) + \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \left(\sigma(\boldsymbol{s}) \, \sigma(\boldsymbol{s})^{\mathbf{T}}\right)_{i,j} \frac{\partial^2 k}{\partial s_i \partial s_j} - \kappa k.$$

[20]Earlier, we specified that the marginal distribution of θ_0 is a lognormal. If all agents are reborn at a single point $(\theta_0, v_0, \Delta_0)$, this becomes the Dirac delta function evaluated there, $\delta_{(\theta_0, v_0, \Delta_0)}$.

[21]For arbitrary functions $u, g \in L^2(\Omega)$, we define the inner product $\langle u, g\rangle = \int_\Omega ug \, d\boldsymbol{\omega}$, where $\boldsymbol{\omega} \in \Omega$.

Note that we can write the integral over (10) as

$$\int_{\Theta}\chi(\theta)\left[w(\theta)^{\alpha}\theta\int_{\mathcal{V}\times\mathcal{D}}\ell(\theta,v,\Delta)g(\theta,v,\Delta)d(v,\Delta)\right.$$
$$\left.-\int_{\Theta\times\mathcal{V}\times\mathcal{D}}w(\theta')\theta'\ell(\theta',v,\Delta)g(\theta',v,\Delta)d(v,\Delta,\theta')\right]d\theta$$
$$=\int_{\Theta\times\mathcal{V}\times\mathcal{D}}\chi(\theta)\left[w(\theta)^{\alpha}\theta\ell(\theta,v,\Delta)g(\theta,v,\Delta)-\int_{\Theta}w(\theta')\theta'\ell(\theta',v,\Delta)g(\theta',v,\Delta)d\theta'\right]d(v,\Delta,\theta)$$
$$=\int_{\Theta\times\mathcal{V}\times\mathcal{D}}\left[\chi(\theta)w(\theta)^{\alpha}\theta\ell(\theta,v,\Delta)-\left(\int_{\Theta}\chi(\theta')d\theta'\right)w(\theta)\theta\ell(\theta,v,\Delta)\right]g(\theta,v,\Delta)d(v,\Delta,\theta),$$

where the last step follows from integration by parts. Let $\mathcal{X}\equiv\int_{\Theta}\chi(\theta')d\theta'$. Taking the Gateaux derivative of the Lagrangian with respect to k along an arbitrary variation and setting it to zero yields the planner's HJB equation in the stationary equilibrium given θ:

$$(r+\kappa)K(\theta)=\inf_{\{c,C,\ell,\sigma_{\Delta},z\}}\left\{\int_{\mathcal{V}\times\mathcal{D}}\left(c+\kappa C-\theta\ell w(\theta)+k_{\theta}\mu(\theta)+k_{v}\left[(\rho+\kappa)v-\tilde{u}-\kappa v^{R}\right]\right.\right.$$
$$+k_{\Delta}\left[(\rho+\kappa-\mu'(\theta))\Delta-\tilde{u}_{\theta}-\kappa\Delta^{R}-\sigma_{\theta}^{2}\sigma_{\Delta}\right]+\frac{\sigma_{\theta}^{2}}{2}\left[k_{vv}\theta^{2}\Delta^{2}+k_{\Delta\Delta}\sigma_{\Delta}^{2}+k_{\theta\theta}\theta^{2}\right]$$
$$+\sigma_{\theta}^{2}\left[k_{v\Delta}\theta\Delta\sigma_{\Delta}+k_{v\theta}\theta^{2}\Delta+k_{\Delta\theta}\sigma_{\Delta}\right]+\kappa k\widetilde{f_{0}}$$
$$\left.\left.+\chi(\theta)w(\theta)^{\alpha}\theta\ell-\mathcal{X}w(\theta)\theta\ell\right)g(\theta,v,\Delta)d(v,\Delta)+\eta(\theta)z(\theta)\right\}. \quad (14)$$

Since (12) and (14) are both three-dimensional second-order partial differential equations, we need six boundary conditions for each one to pin down a unique solution; we discuss these boundary conditions in Appendix C. This system plus the boundary conditions, what Lasry and Lions (2007) call a "mean field game," together characterize the stationary equilibrium. Intuitively, agents care about only the "mean" of all other agents' actions, not each one individually. The policy functions implied by this system yield the efficient allocation.

This mean field game illustrates a key distinction between familiar models with linear production functions and our model. Agents' policy functions depend on w and with a linear production function, this is simply one so that $z(\theta)=0$ for all θ and g does not matter for solving HJB equation. Therefore, each agent's optimal policies are independent of the distribution. This means we can solve the system *sequentially*: derive the optimal policies via (14), plug these into the KFE, and then solve for the stationary distribution via (12). On the other hand, with a CES function then w depends on g so that the system is coupled and we have to solve the system *simultaneously*: each agent's optimal policies depend on the distribution, which itself depends on each agent's policies. In particular,

(14) depends on the aggregate quantity $\mathcal{X}$, not just $\chi(\theta)$. This generalizes the framework in Nuño and Moll (2015), where all of the market clearing conditions can be aggregated linearly. In their setting, capital is the mean field and the planner's HJB equation can be solved state-by-state while here, the mean fields are the wage functions and we need to solve for everything at once.

It turns out that the HJB equation (14) is difficult to work with because several derivatives depend on the choice variable σ_Δ. To get around this, we use duality to work in an alternative state space parametrized by the Lagrange multipliers in the above problem. We briefly explore this approach in the next section to characterize optimal tax policies.

We solve our model using a finite difference method to approximate the partial differential equations. We use an "upwind scheme" to avoid dealing with boundary conditions, although it turns out that in the reparametrized state space, we do not need boundary conditions anyway to recover a unique solution; more detail can be found in Appendix E.

4 Optimal Tax Policies

This section characterizes the allocations, obtained as solutions to the relaxed problem above.

Marginal distortions in agents' choices can be characterized with "wedges" that represent marginal tax rates.[22] Given a history θ^t, the two key wedges are the labor wedge, $\tau_L(\theta^t)$, and the intertemporal, or savings wedge, $\tau_K(\theta^t)$, both of which are standard in the literature:

$$\tau_L\left(\theta^t\right) = 1 + \frac{\tilde{u}_{y,t}}{\tilde{u}_{c,t}} = 1 - \frac{\phi'\left(\ell_t\right)}{w\left(\theta_t\right)\theta_t u_{c,t}}, \tag{15}$$

$$\tau_K\left(\theta^t\right) = r - \rho + \frac{1}{dt}\frac{\mathbb{E}_t\left[du_{c,t}\right]}{u_{c,t}}. \tag{16}$$

The labor wedge is simply the gap between the marginal rate of substitution and the marginal rate of transformation between consumption and labor while the savings wedge is the difference between the marginal rate of intertemporal substitution and the return on savings. In the first-best allocation with perfect information, both of these wedges are zero. Also, both of these wedges are pre-retirement because at retirement there is no more asymmetric information. As long as the tax function T is differentiable, τ_L is the derivative of T with

[22] We use "wedge" and "tax rate" interchangeably, sidestepping the issue of how to implement the constrained efficient allocation with a simple tax system. Kocherlakota (2005) and Albanesi and Sleet (2006) show how to use tax systems to implement the optimal allocation when skill is not too persistent, while Golosov, Troshkin and Tsyvinski (2016*a*) show that a tax system based on consolidated income accounts can generally implement the optimal allocation.

respect to income (and similarly for τ_K).

Now we will derive and characterize the optimal tax formula. To characterize the wedges we will solve the recursive planner's problem. The first-order conditions of (14) are

$$\begin{aligned}
[c] &: 1 = k_v u_c, \\
[C] &: 1 = k_v v_C^R + k_\Delta \Delta_C^R, \\
[\ell] &: -\theta w(\theta) = k_v \tilde{u}_\ell + k_\Delta \tilde{u}_{\theta\ell} - \chi(\theta) w(\theta)^\alpha \theta + \mathcal{X} w(\theta)\theta, \\
[\sigma_\Delta] &: k_\Delta = k_{\Delta\Delta}\sigma_\Delta + k_{v\Delta}\theta\Delta + k_{\Delta\theta}\theta, \\
[z(\theta)] &: \eta(\theta) = \int_{\mathcal{V}\times\mathcal{D}} k_\Delta \tilde{u}_{\theta z} g(\theta, v, \Delta)\, d(v, \Delta).
\end{aligned}$$

The first two equations are easy to interpret: increasing consumption by one unit increases the planner's cost by one unit. However, this lowers an agent's marginal utility of consumption and makes it easier to satisfy the promise-keeping constraint going forward. The third equation balances the effects of asking a given agent to work more by accounting for the effect on everyone's wages. The fourth equation weighs the cost of delivering more information rent today against the benefit of relaxing the incentive constraint. The final equation weighs the cost of increasing the slope of the wage function against the benefit of loosening incentive constraints. This last equation also highlights that an individual agent's policy functions depend on everyone else in the economy.

Instead of working with the state space as given, it is simpler to work with "dual" variables. In particular, following Farhi and Werning (2013), define $\lambda \equiv k_v$ and $\gamma \equiv k_\Delta$. This reparametrizes the state space as $(\lambda, \gamma, \theta) \in \Lambda \times \Gamma \times \Theta$. This new state space has two advantages over the original. First, as we show here, it is quite simple to characterize the optimal allocations, wedges, and their dynamics. The second advantage, which we explore in Appendices D and E, is that it makes numerically solving for the policy functions simpler.

Using the new state space, the first-order conditions are

$$[c] : 1 = \lambda u_c, \tag{17}$$

$$[C] : 1 = \lambda v_C^R + \gamma \Delta_C^R, \tag{18}$$

$$[\ell] : -\theta w(\theta) = \lambda \tilde{u}_\ell + \gamma \tilde{u}_{\theta\ell} - \chi(\theta) w(\theta)^\alpha \theta + \mathcal{X} w(\theta)\theta, \tag{19}$$

$$[\sigma_\Delta] : \gamma = k_{\Delta\Delta}\sigma_\Delta + k_{v\Delta}\theta\Delta + k_{\Delta\theta}\theta, \tag{20}$$

$$[z(\theta)] : \eta(\theta) = \int_{\Lambda\times\Gamma} \gamma \tilde{u}_{\theta z} g(\lambda, \gamma, \theta)\, d(\lambda, \gamma). \tag{21}$$

There is also an envelope condition from the state variable $w(\theta)$,

$$[\text{Env}]: \; -\eta'(\theta) = L(\theta)\left[1 + \alpha\chi(\theta)\, w(\theta)^{\alpha-1} - \mathcal{X}\right] + \int_{\Lambda\times\Gamma} \gamma \tilde{u}_{\theta w} g(\lambda,\gamma,\theta)\, d(\lambda,\gamma).$$

Since we are working with new state variables, we need their equations of motion. The next result, which is similar to Proposition 5 in Farhi and Werning (2013), does just that.

Proposition 4. *The solution to the planner's problem is characterized by the following: there exists a function $\sigma_{\lambda,t}$ such that the stochastic processes for $\{\lambda_t\}$ and $\{\gamma_t\}$ satisfy the following stochastic differential equations:*

$$d\lambda_t = (r-\rho)\,\lambda_t\, dt + \sigma_\theta \sigma_{\lambda,t}\lambda_t\, dZ_t + \left(\lambda_t^R - \lambda_t\right) dR_t, \tag{22}$$

$$d\gamma_t = \left((r-\rho+\mu'(\theta_t))\,\gamma_t - \sigma_\theta^2 \theta_t \sigma_{\lambda,t}\lambda_t\right) dt + \sigma_\theta \gamma_t\, dZ_t + \left(\gamma_t^R - \gamma_t\right) dR_t \tag{23}$$

with γ_0 free.

The law of motion for λ_t is the familiar "inverse Euler equation" and shows that the inverse of the marginal utility of consumption is a submartingale if $r > \rho$, a supermartingale if $r < \rho$, and a martingale if $r = \rho$. Note that these equations are much simpler and more explicit than those of $\{v_t\}$ and $\{\Delta_t\}$. Finally, it turns out that Δ_t^R, and hence γ_t^R, is zero. This is because once the retirement shock hits, there is no more asymmetric information so agents do not receive any information rent at retirement, and since all process are càdlàg information rent must equal zero at the jump. If the retirement shock is private as in Golosov and Tsyvinski (2006) then $\Delta_t^R \neq 0$ since agents can use the shock to extract more information rent.[23] We set γ_0 so that agents have zero information rent at birth.

Using the equations of motion of the dual variables, we can compute the wedges. To derive an expression for the labor wedge, first note that

$$\tilde{u}_\ell = -\phi_\ell \Longrightarrow \tilde{u}_{\theta\ell} = \left(\frac{1}{\theta} + \frac{w'(\theta)}{w(\theta)}\right)(\phi_\ell + \ell\phi_{\ell\ell}) = \left(\frac{1}{\theta} + \frac{w'(\theta)}{w(\theta)}\right)\left(1 + \frac{1}{\varepsilon}\right)\phi_\ell,$$

where $\varepsilon \equiv \frac{\phi_\ell}{\ell\phi_{\ell\ell}}$ is the Frisch elasticity of labor supply.[24] Plugging these into (19) and recalling

[23] They show that agents do not receive a final transfer unless their assets are sufficiently low.

[24] More generally, $\frac{1+\varepsilon^u}{\varepsilon^c}$, where $\varepsilon^u, \varepsilon^c$ are the uncompensated and compensated elasticities of demand; with separable preferences, this reduces to $1 + \frac{1}{\varepsilon}$.

the definition of $\tau_{L,t}$ from (15),

$$-w(\theta_t)\theta_t + \frac{\phi_{\ell,t}}{u_{c,t}} = \gamma_t\left(1+\frac{1}{\varepsilon}\right)\left(\frac{1}{\theta_t}+\frac{w'(\theta_t)}{w(\theta_t)}\right)\phi_{\ell,t} - w(\theta_t)\left[\chi(\theta_t)w(\theta_t)^{\alpha-1}\theta_t - \mathcal{X}_t\theta_t\right] \Longrightarrow$$

$$\frac{\tau_{L,t}}{1-\tau_{L,t}} = \underbrace{-\frac{\gamma_t}{\lambda_t}\left(1+\frac{1}{\varepsilon}\right)\left(\frac{1}{\theta_t}+\frac{w'(\theta_t)}{w(\theta_t)}\right)}_{\text{Mirrlees}} + \underbrace{\frac{w(\theta_t)\theta_t}{\lambda_t\phi_{\ell,t}}\left[\chi(\theta_t)w(\theta_t)^{\alpha-1} - \mathcal{X}_t\right]}_{\text{Wage Compression}}. \tag{24}$$

This expression generalizes the ones in Stiglitz (1982) and Ales, Kurnaz and Sleet (2015) to include dynamics and the ones in Farhi and Werning (2013) and Golosov, Troshkin and Tsyvinski (2016*b*) to include wage compression forces.

The Mirrlees term is standard in all Mirrlees models and captures the planner's insurance motive against shocks. It is the tax rate the planner would implement if he set agents' wages optimally but ignored any additional effects of the nonlinear production function. The only difference between our expression and the one in Farhi and Werning (2013) is the $\frac{w'(\theta_t)}{w(\theta_t)}$-term, which reflects that wages are not constant across types in our model. Under standard distributions, this ratio is positive for large θ and negative for small θ. This means that more skilled workers also receive a larger wage, making it easier for them to mimic less skilled workers and tightening their incentive constraint. This is a force for a higher marginal tax rate on more skilled agents and a lower rate on less skilled agents. However, the Mirrlees term is always nonnegative and so on its own is inconsistent with the EITC.

The wage compression term was first identified by Stiglitz (1982) and further analyzed by Ales, Kurnaz and Sleet (2015), both in static models. In their contexts, the planner should subsidize labor that decreases the wages of skilled agents and tax labor that decreases the wages of unskilled agents. When skill has bounded support, this means the most skilled agents should face a negative marginal tax rate while the least skilled agents should face a positive rate.[25] Indeed, the negative rate at the top increases the supply of skilled labor, lowering the wages of the most skilled workers, thus reducing skill premia. The negative rate at the top also loosens incentive constraints and makes skilled agents less likely to mimic unskilled ones. Conversely, a higher rate on unskilled workers reduces the supply of unskilled labor and raises their wages.

To uncover more insight, note that χ_t is also a dynamic variable with a law of motion. This means that the Mirrlees and wage compression terms have both intra- and intertemporal

[25] When skill has unbounded support with a fat-tailed distribution, the rate at the top is no longer negative but it is still lower than what it would be if the wage compression channel was not active.

components.[26] The intra- and intertemporal Mirrlees forces have the same interpretations as their analogues in Diamond (1998), Saez (2002), Farhi and Werning (2013), and Golosov, Troshkin and Tsyvinski (2016*b*). The intratemporal force is what drives the planner to insure agents against θ_t conditional on the history θ^{t-}. The intertemporal force is the dynamic insurance motive and allows the planner to relax incentive constraints in earlier periods to reduce the cost of insurance against earlier shocks. The main difference is that our expressions are (unsurprisingly) more complicated because other agents' skill shocks spill over via the wage function.

The most important insight comes from deconstructing the wage compression force into intra- and intertemporal components. In our dynamic model, the planner uses a negative rate at the bottom to deter skilled agents from deviating in the future. While in the short-run a skilled agent benefits from deviating, in the long-run he will receive a lower wage and work more to reach the same income level. The planner also redistributes more to those states that unskilled agents are more likely to face. In other words, the planner uses negative rates at the bottom to deter skilled agents from mimicking unskilled ones while providing more insurance to unskilled agents. That wages are free to adjust in response to labor supply changes gives the planner an additional channel he can use to affect insurance and incentives.

To see how lowering rates allows the planner to further front-load incentives and provide more insurance, recall that a lower rate at the bottom increases the labor supply and lowers wages, and so raises $\tilde{u}_{\theta,t}$. From (9), the drift of Δ_t declines so Δ_{t+dt} declines, the planner front-loads incentives and provides less information rent in later periods. But from (6), this lowers the volatility of v_t and unskilled agents receive more insurance. Since insurance is very valuable for these agents, the intertemporal force dominates the intratemporal one and the total wage compression term is negative. If the Mirrlees term is sufficiently low here then the intertemporal wage compression force makes the overall rate negative, consistent with the EITC. For skilled agents, insurance is less valuable than being incentivized to tell the truth so the intratemporal force dominates. Thus, in the right tail the static and dynamic models agree. In the middle and near the top, the planner sets a (slightly) higher rate because skill premia are less responsive and incentives matter. We conclude that the EITC accomplishes two goals: it reduces the cost of insuring unskilled workers while making misreporting costly in the long-run for skilled agents.

[26] One could use Itô's Lemma on (24) to derive the full expressions.

As for the intertemporal wedge, using that $\tilde{u}_{c,t} = \lambda_t^{-1}$, by Itô's Lemma,

$$d\lambda_t^{-1} = \lambda_t^{-1}\left(\left(\rho - r + \sigma_\theta^2\sigma_{\lambda,t}^2\right)dt - \sigma_\theta\sigma_{\lambda,t}\,dZ_t\right) + \left(\lambda_t^{R^{-1}} - \lambda_t^{-1}\right)dR_t \Longrightarrow$$
$$\frac{\mathbb{E}_t\left[d\tilde{u}_{c,t}\right]/\tilde{u}_{c,t}}{dt} = \frac{\mathbb{E}_t\left[d\lambda_t^{-1}\right]/\lambda_t^{-1}}{dt} = \rho - r + \sigma_\theta^2\sigma_{\lambda,t}^2 + \kappa\left(\frac{v_{C,t}^R}{u_{c,t}} - 1\right) = \rho - r + \sigma_\theta^2\sigma_{\lambda,t}^2.$$

It follows from (16) that

$$\tau_{K,t} = \sigma_\theta^2\sigma_{\lambda,t}^2. \tag{25}$$

Recall that $\sigma_{\lambda,t}$ is the volatility of the inverse of the marginal utility of consumption. This is higher when consumption is more volatile so that the planner is providing more incentives at the cost of less insurance. The intertemporal wedge is larger for more skilled agents with imperfect substitutability than it is with a linear production function. To see why, first recall that a positive intertemporal wedge exists in the first place because agents have an incentive to "double-deviate": they can report truthfully today but over-save, and then under-report in the next period and use the extra savings to over-consume. Since very skilled agents receive higher wages they have a stronger incentive to double-deviate, so the planner deters this by setting an even higher tax on savings.

5 Quantitative Analysis

We now parametrize the model and solve it numerically to illustrate our theoretical results.

5.1 Parametrization

Agents have CRRA utility over consumption and isoelastic disutility of labor,

$$u\left(c_t\right) = \frac{c_t^{1-\omega} - 1}{1-\omega} \text{ and } \phi\left(\ell_t\right) = \frac{\beta}{1+\frac{1}{\varepsilon}}\ell_t^{1+\frac{1}{\varepsilon}}.$$

Agents have logarithmic utility so $\omega = 1$. Chetty (2012) finds that the Frisch elasticity is $\varepsilon = 0.5$ and we set $\beta = 1$. We set $\rho = r = 0.05$ so that the planner and the agents have the same discount rate, 5%, and $\psi = 1$ so that agents value working life and retirement equally.

We set

$$\mu\left(\theta\right) = \theta\left[-(1-p)\left(\log\left(\theta\right) - \log\left(\theta^*\right)\right) + \frac{\sigma_\theta^2}{2}\right],$$

for some constant $p \in [0,1]$ and θ^*. This implies that $\log\left(\theta\right) \sim \mathcal{N}\left(\theta^*, \frac{\sigma_\theta^2}{2(1-p)}\right)$. Farhi and Werning (2013) and Stantcheva (2015) both use this specification and set $p = 1$ so that $\log(\theta_t)$ follows a random walk, based on the finding in Storesletten, Telmer and Yaron (2004) that income is very persistent. However, there is no stationary distribution when $p = 1$ unless $\kappa > 0$ and in that case, $\log(\theta)$ has an exponential stationary distribution. We set $\kappa = 0.025$ to ensure an average working life of 40 years, and $p = 0.95$ so that income is still quite persistent but $\log(\theta)$ has a normal distribution. Recall that new agents' skills are drawn from a lognormal distribution, $\log(\theta_0) \sim \mathcal{N}(\mu, \upsilon^2)$, so we set $\mu = \theta^*$ and $\upsilon^2 = \frac{\sigma_\theta^2}{2(1-p)}$ so that in fact they are drawn from the stationary distribution. We normalize $\theta^* = 1$ and we set $\sigma_\theta^2 = 0.0095$ to match the volatility of innovations to skill, as documented by Heathcote, Storesletten and Violante (2005).

The most important parameter in our model is the elasticity of substitution, α. Unfortunately, there are no widely agreed-on estimates. Katz and Murphy (1992) estimate $\alpha = 1.4$ while Acemoğlu and Autor (2011) find $\alpha = 2.9$, both in a two-type production function with high school ("low-skill") and college ("high-skill") graduates. As far as we know, the only other paper that estimates α in our production function with a continuum of skills is Heathcote, Storesletten and Violante (2016). In their model, $\text{Var}(\log(w(\theta))) = \frac{1}{\alpha^2}$ and while they do not directly observe $w(\theta)$, they do observe the variance of the logarithm of the "total wage" $\theta w(\theta)$, which depends (additively) on $\frac{1}{\alpha^2}$ and the variances of several other processes. They estimate $\alpha = 3.124$ with a standard error of 0.115. Since the 95% confidence interval is so narrow, we use this point estimate in the first set of numerical experiments. In an alternative calibration, the same authors show that if income has a Pareto tail, then α coincides with the tail index. We repeat the experiments with $\alpha = 2.5$, which is at the upper end of estimates of the tail index to ensure a thinner tail.

5.2 Properties of Optimal Tax Policies

We solve for the optimal wedges in the stationary equilibrium, so an agent's tax rate does not depend on when he was born but it does depend on his history of shocks.

Before exploring the tax rates themselves, it is instructive to first understand how skill affects income in our model versus in models with a linear production function. Figure 1 plots the average income of agents with a given skill level for three values of α. While in standard Mirrlees models income is approximately linear in skill (dotted line), in our model (solid and dashed lines) income is strictly convex so that income grows more than one-for-one with skill; the lower the elasticity, the more convex is income. Observe that while all agents receive (weakly) higher income, the agents that benefit most relative to in the linear case are those with the most skill, which means they can more easily mimic less skilled agents than they could with a linear production function.

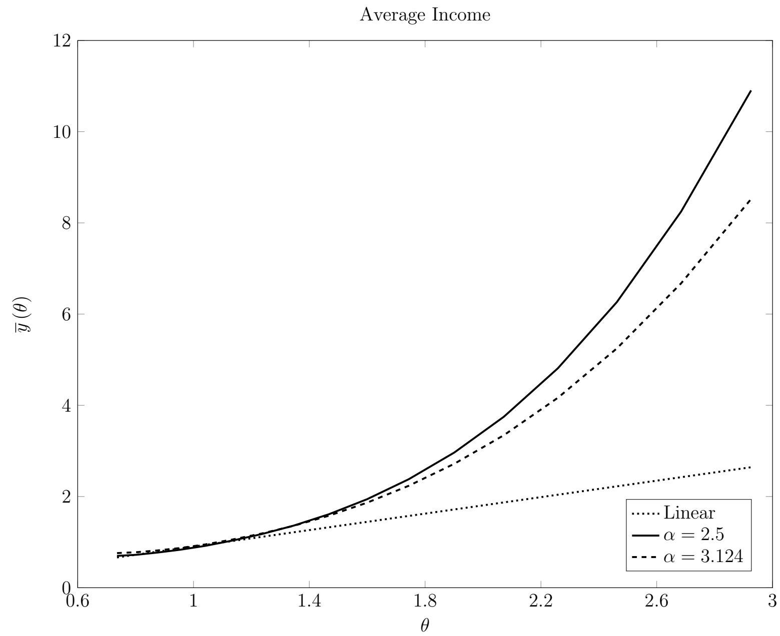

Figure 1: Average income as a function of α.

5.2.1 Decomposing the Labor Wedge

We start by showing the mechanisms at work. Let $\overline{\tau}_j^i(\theta)$ denote the $j \in \{L, K\}$ wedge for an agent with current skill θ, averaged over all histories (of any length) that terminate at θ. Here, $\overline{\tau}_j^M$ is the Mirrlees tax rate described earlier that accounts for only optimal wage-setting and $\overline{\tau}_j^*$ is the overall optimal tax rate that accounts for both optimal wages and the wage compression channel.

Figure 2 plots $\overline{\tau}_L^*(\theta)$ (solid line) vs. $\overline{\tau}_L^M(\theta)$ (dashed line) for $\alpha = 2.5$ (black) and $\alpha = 3.124$ (red). The difference between the solid and dashed lines of the same color is the total wage compression force. In both cases, the Mirrlees rate is approximately hump-shaped, topping out near 80% in both cases, and declining very slowly in the right tail. The overall rate has a more pronounced hump shape, especially in the left tail where it is on average well below the Mirrlees rate. It is also higher than the Mirrlees rate in the middle and near the top in both cases. However, when α is smaller, the overall rate is far lower in the left tail and higher everywhere else. This is because when α is lower, the wage function is "more convex" and lower in the middle of the distribution. This means the planner wants to raise wages here, which he does by increasing the tax rate and discouraging labor. Further, the gap between the two rates closes more quickly when α is larger, meaning that the overall rate will drop below the Mirrlees rate sooner when α is larger. Again, this is because when α is smaller, the wage function is higher in the right tail so that the planner can lower taxes and hence wages more to compress skill premia. Since these are average rates across all histories, it does not rule out negative rates (and hence the EITC) for some agents.

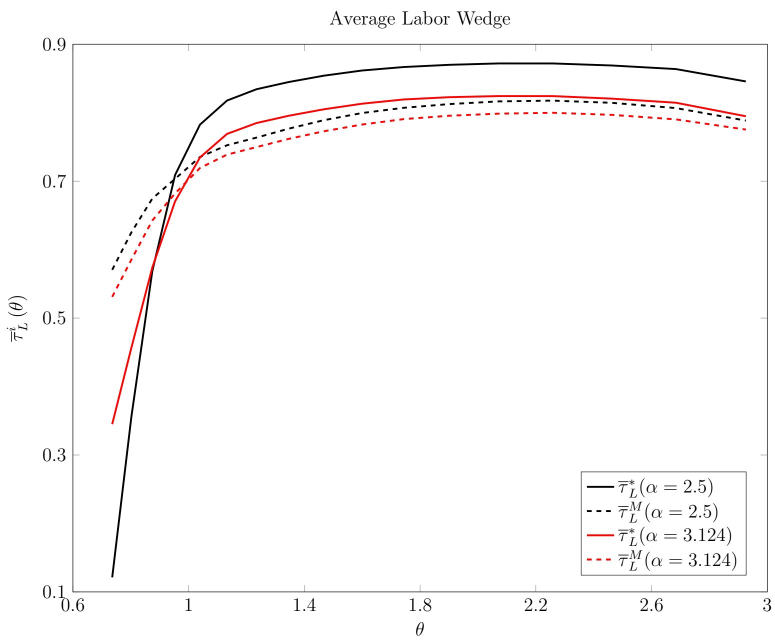

Figure 2: $\overline{\tau}_L^*(\theta)$ vs. $\overline{\tau}_L^M(\theta)$ for $\alpha = 2.5$ and $\alpha = 3.124$.

5.2.2 EITC is Optimal

Now we show that the EITC is consistent with optimal policy. Figure 3 plots the overall rates of an agent who has been at the 25$^{\text{th}}$ percentile for skill for 1 year (solid line) and 50 years (dashed line); the black line is for when $\alpha = 2.5$ and the red line for when $\alpha = 3.124$. First, observe that rates are always positive following a short history in both cases, with a much higher minimum rate when α is larger. On the other hand, following a long history, the minimum rate is well below zero in both cases, consistent with the EITC. The reason the EITC does not kick in immediately is that initially, the planner's problem is essentially static while following long histories, the planner can use negative rates at the bottom to relax more incentive constraints from earlier periods.

A key difference, however, is that the phase-in rate is near 50% when $\alpha = 2.5$ but only around 10% when $\alpha = 3.124$, which suggests that small variations in α can have large effects on the generosity of the EITC. This is consistent with earlier reasoning, as is the observation that the rates in the middle and near the top are higher when α is lower. While at the top the overall rate begins to decline, this occurs so deep in the right tail that our model probably does not accurately describe the labor market for these agents anyway. Therefore, for all intents and purposes, the optimal income tax is progressive. Finally, while our tax system is not age-dependent, for this particular agent, since skill is constant over time the length of the history roughly corresponds to age. Then like in the the age-dependent system in Heathcote, Storesletten and Violante (2016), the EITC in our system is more generous later in life.

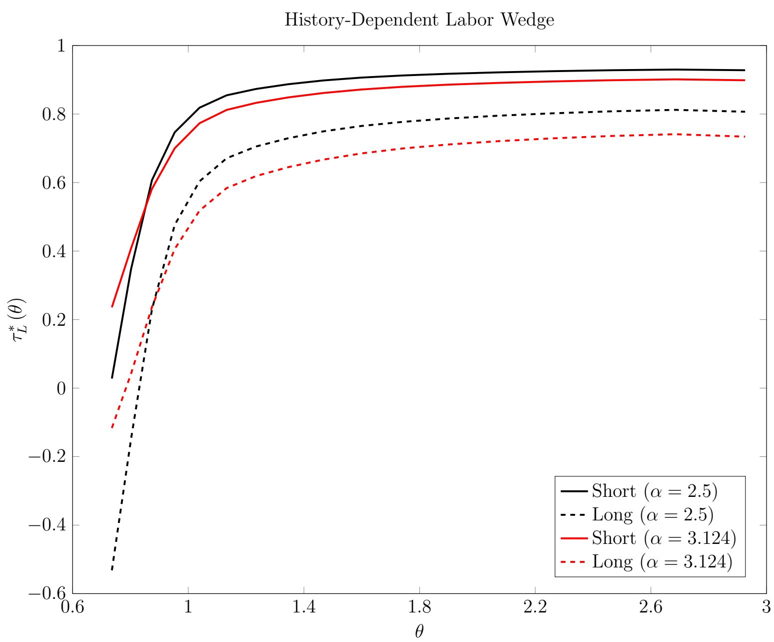

Figure 3: History-dependent labor wedge for low-skill worker for $\alpha = 2.5$ and $\alpha = 3.124$.

5.2.3 Other Properties

Figure 3 shows the long-run behavior of the tax system for an agent whose skill never changes. Consider the impact of a one-time shock to skill on the tax rate, i.e., the impulse response function. This allows us to highlight short-run properties of the tax system. We follow Borovička, Hansen and Scheinkman (2014), who link impulse response functions in continuous time to Malliavin derivatives. We set $\kappa = 0$ and ignore the retirement shock for simplicity. Let $S_t = (\theta_t, \lambda_t, \gamma_t)^{\mathbf{T}}$ denote the vector of state variables with law of motion

$$dS_t = d\begin{bmatrix}\theta_t\\ \lambda_t\\ \gamma_t\end{bmatrix} = \begin{bmatrix}\mu(\theta_t)\\ 0\\ \mu'(\theta_t)\gamma_t - \sigma_\theta^2\theta_t\sigma_{\lambda,t}\lambda_t\end{bmatrix} dt + \begin{bmatrix}\sigma_\theta\theta_t\\ \sigma_\theta\sigma_{\lambda,t}\lambda_t\\ \sigma_\theta\gamma_t\end{bmatrix} dZ_t$$

and initial condition $S_t = s$. Consider an $\mathcal{F}_T$-adapted function $F(S_T)$. By the Clark-Ocone-Haussman Theorem,

$$F(S_T) = \mathbb{E}[F(S_0) \mid \mathcal{F}_0] + \int_0^T \mathbb{E}[D_t(F(S_T)) \mid \mathcal{F}_t]\, dZ_t,$$

where $D_t(F(S_T))$ is the Malliavin derivative of F. Then by the Malliavin chain rule,

$$D_t(F(S_T)) = F'(S_T)^{\mathbf{T}} H_T H_t^{-1}\sigma(S_t)\mathbf{1}_{\{t\leq T\}},$$

where $\{H_t\}$ is the first variation process defined by $H_t = \frac{\partial S_t}{\partial s}$ with $H_0 = \mathbf{I}_3$ and law of motion

$$dH_t = \begin{bmatrix}\mu'(\theta_t) & 0 & 0\\ 0 & 0 & 0\\ \mu''(\theta_t)\gamma_t - \sigma_\theta^2\lambda_t\left(\sigma_{\lambda,t} + \theta_t(\sigma_{\lambda,t})_\theta\right) & -\sigma_\theta^2\theta_t\left(\sigma_{\lambda,t} + \lambda_t(\sigma_{\lambda,t})_\lambda\right) & \mu'(\theta_t) - \sigma_\theta^2\theta_t\lambda_t(\sigma_{\lambda,t})_\gamma\end{bmatrix} H_t\, dt$$
$$+ \begin{bmatrix}\sigma_\theta & 0 & 0\\ \sigma_\theta\lambda_t(\sigma_{\lambda,t})_\theta & \sigma_\theta\left(\sigma_{\lambda,t} + \lambda_t(\sigma_{\lambda,t})_\lambda\right) & \sigma_\theta\lambda_t(\sigma_{\lambda,t})_\gamma\\ 0 & 0 & \sigma_\theta\end{bmatrix} H_t\, dZ_t,$$

where $(\sigma_{\lambda,t})_\theta = \frac{\partial\sigma_{\lambda,t}}{\partial\theta_t}$, and similarly for the other state variables.

Let $\varphi_s(S_T)$ denote the impulse response function of $F(S_T)$ with initial conditions $S_t = s$ and $H_0 = \mathbf{I}_3$. Then

$$\varphi_s(S_T) = \mathbb{E}_s\left[F'(S_T)^{\mathbf{T}} H_T H_t^{-1}\mathbf{1}_{\{t\leq T\}}\right]\sigma(s).$$

Set $F(S_T) = \tau_L(S_T)$ and consider the response to a unit shock to skill. Figure 4 plots the labor wedge's path relative to where it would be had the shock not occurred and shows it is progressive in the short-run.[27] This stands in sharp contrast to models with linear production, as Farhi and Werning (2013) show that the rate is regressive in the short-run. One can show that with a linear production function, the volatility of the tax function is

[27]The plot is for an agent whose skill is at the 83$^{\text{rd}}$ percentile but we have repeated the exercise for agents at very low percentiles and this finding still holds.

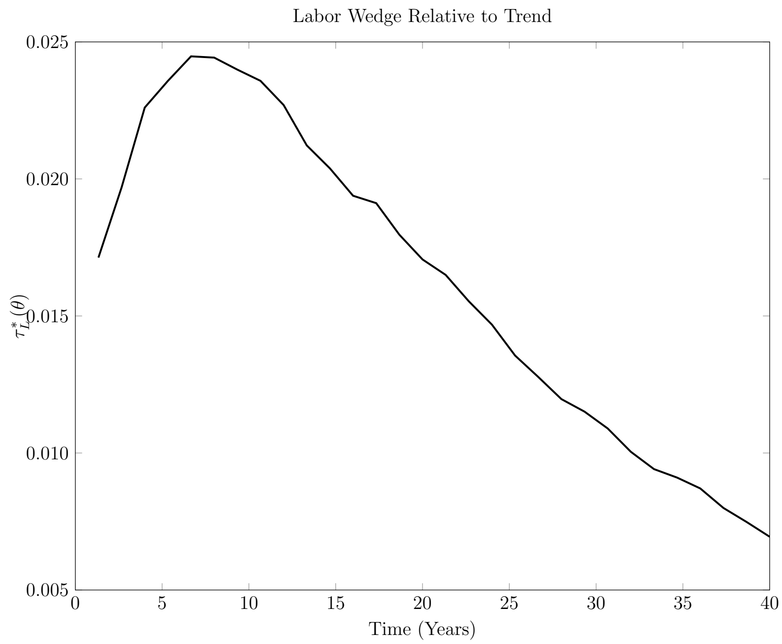

Figure 4: Impulse response function of the labor wedge.

proportional to $-\sigma_{\lambda,t}$ while here there are additional positive terms that come from the variation in $w(\theta)$. This is because a positive shock increases skill premia relative to less skilled agents so the planner undoes this with a higher labor tax.

Finally, Figure 5 plots the intertemporal wedge, or the implicit tax on savings, averaged across histories that terminate at θ. The solid line is the optimal rate with $\alpha = 2.5$, the dashed line is when $\alpha = 3.124$, and the dotted line is the Mirrlees rate. All three functions are increasing over most of the domain and then decrease to zero in the right tail. The reason for the decline is that the most skilled agents can almost perfectly insure themselves against shocks, which means they have almost smooth consumption. This means that $\lambda_t = u_{c,t}^{-1}$ has almost no volatility, which, by (22), means $\sigma_{\lambda,t}$ is close to zero; (25) then implies that $\tau_{K,t}$ is close to zero. However, the optimal rate tops out at just under 1.5% when $\alpha = 2.5$, compared to just under 1% when $\alpha = 3.124$ and under 0.3% for the Mirrlees rate. In addition, the optimal rate is higher for more agents when α is smaller. The reason, as per Figure 1 again, is that skilled agents not only are more productive but also receive a larger wage, which makes it easier for them to mimic less skilled workers and save too much.

5.2.4 A Word on Welfare

Our mechanism design framework puts an upper bound on economy-wide welfare; what is the welfare loss from moving from an optimal, history-dependent tax system to a simpler

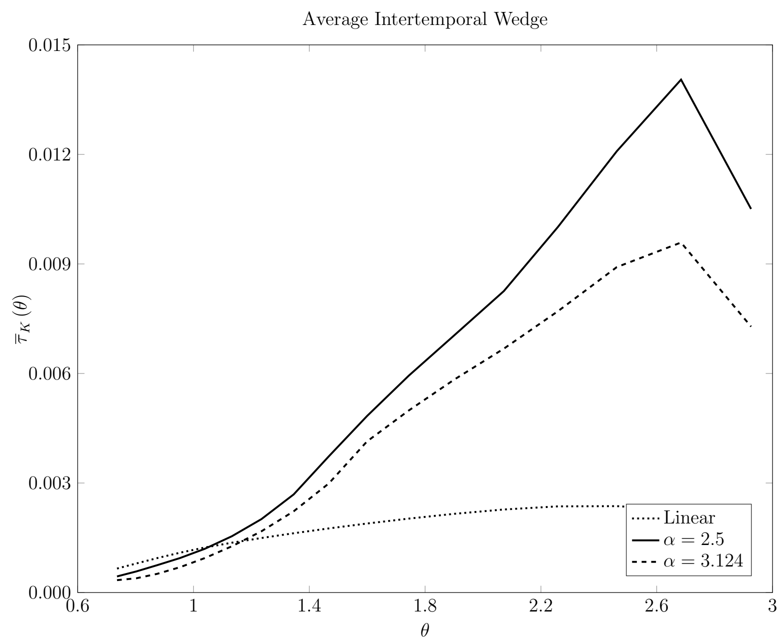

Figure 5: Average intertemporal wedge as a function of α.

one? We compared welfare in our benchmark economy with that in an economy in which agents are unconstrained to borrow and save a riskless asset and face history-independent, linear taxes and found that the welfare gain is essentially a rounding error, well under one percent.[28] This finding that the optimal tax schedule is close to linear is well-documented in both static and dynamic models and suggests that the overall welfare gains from the EITC are not exceptionally large. Further, the shape of our tax function closely resembles the one in Heathcote, Storesletten and Violante (2016),

$$T(y) = y - \tau_0 y^{1-\tau_1}$$

for some parameters τ_0, τ_1, which they argue is a good approximation to the current U.S. system (and thus already has an embedded EITC). Therefore, the gains from modifying the current U.S. are unlikely to be large. However, these statements should not be taken too seriously as there are likely other factors (such as moving away from a lognormal skill distribution) that can increase welfare gains.

[28] We followed Heathcote and Tsujiyama (2016), extended to a dynamic setting, by first solving for the equilibrium allocations given the tax function and then computing the percentage increase in consumption which, if received every period after all histories, would yield the same gain in lifetime utility.

6 Discussion

While our previous analysis highlights our theoretical results, it is worth taking a moment to discuss some of our key assumptions and how relaxing them would affect our model.

6.1 CES Production Function

First, while a few papers (see above) do use CES production functions, its use in the public finance literature is still somewhat limited. The reason, as Salanié (2011) points out, is that, for example, the substitutability between very low- and low-skill workers is probably very different from the substitutability between high- and very high-skill workers, which are both different from the substitutability between low- and high-skill workers. Ales, Kurnaz and Sleet (2015) get around this by micro-founding the production function with a task assignment framework so that the elasticity of substitution differs across pairs of workers. They show that as skilled workers are increasingly locked into complex tasks, the wage compression force becomes more powerful and is thus a force for even lower rates at the top and higher rates at the bottom. However, since their model is static they neglect the opposing dynamic insurance motive. Since the dynamic force is less powerful near the top and more powerful near the bottom, it is not clear how their more complicated production function would affect tax policy in our setting.

6.2 Skill Distribution

Second, in following much of the literature, we assumed that $\log(\theta_t)$ is an Ornstein-Uhlenbeck process. This has two shortcomings: first, the implied stationary distortion is lognormal and hence thin-tailed. It is well-documented[29] that the U.S. income distribution is in fact fat-tailed and Saez (2001) shows that fat tails imply a significantly higher tax rate at the top. However, as the figures above show, convergence in the right tail is so slow that the tax rate near the top as implied by a lognormal stationary distribution is a good enough approximation to the rate implied by a fat-tailed distribution for almost the entire population.[30] On the other hand, a thicker tail would lower the elasticity, α, as argued in Section 5, and this would only strengthen our results.

Second, Guvenen, Ozkan and Song (2014) and Guvenen et al. (2015) show that not only

[29] See Piketty and Saez (2003) and Atkinson, Piketty and Saez (2011), for example.

[30] Indeed, Golosov, Troshkin and Tsyvinski (2016*b*) show that lognormal and Pareto lognormal distributions yield very similar tax functions.

is the skill shock process very persistent, it is also highly leptokurtic, a property that does not hold for lognormal (or Pareto lognormal) distributions. Gabaix et al. (2015) suggest using a skill process with distinct growth regimes,

$$d(\log(\theta_t)) = \mu_{\theta j}\ dt + \sigma_{\theta j}\ dZ_t + \text{Injection} - \text{Death},$$

where $\mu_{\theta j}, \sigma_{\theta j}$ are the growth rate and volatility, respectively, of skill in regime $j \in \{1, \ldots, J\}$, with agents being born into each regime with some probability and randomly transitioning between regimes during life. Not only does this help with fitting the micro data, they also show that it has the potential to explain the dynamics of inequality.

Golosov, Troshkin and Tsyvinski (2016*b*) show that a highly leptokurtic distribution generates a substantially different tax function: instead of being hump-shaped, it is U-shaped and larger in magnitude, including in the left tail. This means that for poor agents, the optimal policy is a guaranteed level of income that is taxed at a very high rate, a so-called "negative income tax." The reason is that the hazard function is U-shaped with leptokurtic distributions, as opposed to monotonic with (Pareto) lognormal ones. However, the hazard function acts through γ_t and as (24) shows, γ_t does not affect the wage compression term. This suggests that while the Mirrlees term will be U-shaped, the wage compression term will have the same pronounced hump-shape as it does with a lognormal distribution.[31] Whether it is powerful enough to force the entire rate negative in the left tail is unclear.

7 Conclusion

We develop a dynamic Mirrlees model of optimal taxation in which agents with different skills are imperfect substitutes in the production process and uncover a novel intertemporal wage compression channel that the planner uses to lower the cost of insurance against shocks. We demonstrate that this insurance motive can drive labor income taxes negative for low-skill agents, helping us to rationalize redistributive programs like the EITC. Since more skilled agents are paid higher wages and can more easily mimic less skilled agents, we further show that the planner also imposes a higher savings tax on these agents to deter this behavior. Finally, in contrast to models with linear production functions, the forces that we introduce cause labor taxes to be progressive in the short-run. This occurs because positive skill shocks

[31] That said, the hazard functions of the lognormal and their leptokurtic distribution are very similar in the left tail anyway. Further, Heathcote and Tsujiyama (2016) show that the tax function's U-shape is actually very sensitive to parameters.

spill over to other agents and raise skill premia.

Finally, we make several methodological contributions to the literature. By working in continuous time, we take advantage of techniques for solving heterogeneous agent models that are applicable to other dynamic mechanism design settings. In addition, while existing dynamic Mirrlees models often must verify the first-order approach ex post numerically, we provide an analytic sufficiency condition, assuming that agent cannot overreport their skill levels, for which the first-order approach is valid.

While our model highlights an important determinant of tax policy relevant for understanding redistributive programs, there are many avenues for future research that would bring our model closer to the policy realm. Adding endogenous retirement decisions, for instance, would allow one to jointly analyze tax and retirement policies with imperfect substitutability. In addition, employing richer production and skill functions, as described in 6, would allow for a more accurate quantitative assessment of the optimal tax schedule relevant for policy recommendations while adding an extensive labor supply margin would allow us to jointly analyze the most important drivers of the EITC.

Appendix

A Section 3 Proofs

A.1 Proof of Proposition 1

Fix a contract and a reporting strategy $\widehat{\theta}_t$, and define the deflated gains process

$$G_t \equiv \mathbb{E}_t\left[\int_0^\infty e^{-\rho s}\tilde{u}_s\left(\widehat{\theta}_s\right) ds\right] = \int_0^t e^{-\rho s}\tilde{u}_s\left(\widehat{\theta}_s\right) ds + e^{-\rho t}v_t.$$

Since G_t is a martingale, by the Martingale Representation Theorem, there exist processes $\widehat{\Delta}_t, Q_t \in H^2$ such that

$$dG_t = e^{-\rho t}\widehat{\Delta}_t\ dZ_t + Q_t\left(dR_t - \kappa\ dt\right).$$

Slightly abusing notation for Q_t, differentiating with respect to t and rearranging,

$$dv_t = \left(\rho v_t - \tilde{u}_t\right) dt + \sigma_\theta \theta_t \widehat{\Delta}_t\ dZ_t + Q_t\left(dR_t - \kappa\ dt\right).$$

Since v_t jumps to v_t^R at retirement, we must have $Q_t = v_t^R - v_t$. To see that the transversality condition holds, note that

$$e^{-\rho t}v_t = \mathbb{E}_t\left[\int_t^{T\wedge\tau} e^{-\rho s}\tilde{u}_s\ ds + e^{-\rho(T\wedge\tau)}v_{T\wedge\tau}\right].$$

Taking the limit as $T \to \infty$ and applying the Monotone Convergence Theorem, the result follows.

For the other direction, if we have a solution to (6) with the desired properties, we can integrate

$$v_t = \mathbb{E}_t\left[\int_t^{T\wedge\tau} e^{-\rho(s-t)}\tilde{u}_s\ ds + e^{-\rho(T\wedge\tau)}v_{T\wedge\tau}\right].$$

Taking the limit as $T \to \infty$, applying the Monotone Convergence Theorem to the first term and the transversality condition to the second term, the result follows.

A.2 Proof of Proposition 2

Rather than working with the actual process, $\{\theta_t\}$, we will work with the logarithm of this process, $\{\log(\theta_t)\}$, which corresponds to the growth rate of θ_t instead of the level. Its law of motion is

$$d\left(\log\left(\theta_t\right)\right) = M\left(\theta_t\right) dt + \sigma_\theta\ dZ_t,$$

where $M(\cdot)$ is a function of $\mu(\cdot)$.

Consider the process defined by

$$d\widehat{Z}_t \equiv dZ_t + \varepsilon a_t\ dt,$$

where $\varepsilon \in \mathbb{R}$ and $a_t \in H^2$.[32] This process is a Brownian motion under another measure

[32] Abusing notation, here and in subsequent proofs, ε is unrelated to the Frisch elasticity of substitution.

$\mathbb{Q} \sim \mathbb{P}$ by the Girsanov Theorem, and this is also induced by some alternative reporting strategy $\widehat{\theta}_t(\varepsilon) \neq \theta_t$. In addition,

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = \xi_t(\varepsilon) = \exp\left(-\int_0^t \varepsilon a_s \, dZ_s - \frac{1}{2}\int_0^t \varepsilon^2 |a_s|^2 \, ds\right)$$

is the Radon-Nikodym derivative. It follows that

$$\left.\frac{d\xi_t}{d\varepsilon}\right|_{\varepsilon=0} = -\int_0^t a_s \, dZ_s.$$

Let $\widehat{\Delta}_t^\ell$ denote the sensitivity of promised utility to changes in $\log\left(\widehat{\theta}_t\right)$, and let W_t denote promised utility as a function of $\log\left(\widehat{\theta}_t\right)$, whose equation of motion is

$$dW_t = (\rho W_t - \tilde{u}_t)\, dt + \sigma_\theta \widehat{\Delta}_{t'}^\ell \, dZ_t + \left(W_t^R - W_t\right)(dR_t - \kappa \, dt).$$

Define two processes,

$$\zeta_t^{t'} \equiv \int_t^{t'} a_s \, dZ_s,$$
$$\Lambda_t \equiv \mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)} \zeta_t^s \tilde{u}_s \, ds\right].$$

We will need the following lemma.

Lemma A1. *An equivalent expression for Λ_t is*

$$\Lambda_t = \mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)} \sigma_\theta a_s \widehat{\Delta}_s^\ell \, ds\right].$$

Proof. Note that

$$\zeta_t^s = \int_t^s a_\tau \, dZ_\tau = \zeta_t^{t'} + \int_{t'}^s a_\tau \, dZ_\tau$$

when $t \leq t' \leq s$. Then

$$\Psi_{t'} \equiv \mathbb{E}_{t'}\left[\int_t^\infty e^{-\rho(s-t)} \zeta_t^s \tilde{u}_s \, ds\right] = \int_t^{t'} e^{-\rho(s-t)} \zeta_t^s \tilde{u}_s \, ds + e^{-\rho(t'-t)} \zeta_t^{t'} W_{t'} + \Lambda_{t'}^t$$

is a martingale, where $\Lambda_{t'}^t$ is the contribution to Λ_t of the compensation after time t',

$$\Lambda_{t'}^t \equiv \mathbb{E}_{t'}\left[\int_{t'}^\infty e^{-\rho(s-t)} \left(\int_{t'}^s a_\tau \, dZ_\tau\right) \tilde{u}_s \, ds\right].$$

Applying Itô's Lemma,

$$\begin{aligned}
d\Psi_{t'} &= e^{-\rho(t'-t)}\left\{\zeta_t^{t'}\left(\tilde{u}_{t'}-\rho W_{t'}\right)\ dt' + a_{t'}W_{t'}\ dZ_{t'} + \zeta_t^{t'}\left[\left(\rho W_{t'}-\tilde{u}_{t'}\right)dt' + \sigma_\theta \widehat{\Delta}_{t'}^{\ell}\ dZ_{t'}\right]\right\} \\
&\quad + e^{-\rho(t'-t)}\zeta_r^{t'}\left(W_{t'}^R - W_{t'}\right)\left(dR_{t'}-\kappa\ dt'\right) + e^{-\rho(t'-t)}\sigma_\theta a_{t'}\widehat{\Delta}_{t'}^{\ell}\ dt' + d\Lambda_{t'}^t \\
&= e^{-\rho(t'-t)}\sigma_\theta a_{t'}\widehat{\Delta}_{t'}^{\ell}\ dt' + e^{-\rho(t'-t)}\left[\left(\sigma_\theta\widehat{\Delta}_{t'}^{\ell}\zeta_t^{t'} + a_{t'}W_{t'}\right)dZ_{t'} + \zeta_t^{t'}\left(W_{t'}^R-W_{t'}\right)\left(dR_{t'}-\kappa\ dt'\right)\right] + d\Lambda_{t'}^t.
\end{aligned}$$

Integrating over $[t,t']$ and taking the expectation at t,

$$0 = \mathbb{E}_t\left[\int_t^s e^{-\rho(t'-t)}\sigma_\theta a_{t'}\widehat{\Delta}_{t'}\ dt'\right] + \underbrace{\mathbb{E}_t\left[\int_t^s d\Lambda_{t'}^t\right]}_{=\mathbb{E}_t[\Lambda_s^t]-\Lambda_t}.$$

Taking the limit as $s\to\infty$ and assuming the transversality condition $\mathbb{E}_t\left[\Lambda_s^t\right]\to 0$, we are done. □

With this lemma, we turn to an agent's problem. An agent's objective function is

$$\mathcal{U}_t(\varepsilon) \equiv \mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\xi_s(\varepsilon)\tilde{u}_s(\varepsilon)ds\right].$$

For truth-telling to be optimal, $\mathcal{U}_t(\varepsilon)$ must be maximized at $\varepsilon=0$. Under the first-order approach, this means its derivative is zero at $\varepsilon=0$. We have

$$\left.\frac{d\mathcal{U}_t}{d\varepsilon}\right|_{\varepsilon=0} = \mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\tilde{u}_s\left(-\int_t^s a_\tau\ dZ_\tau\right)ds\right] + \mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\left.\frac{\partial\tilde{u}_s}{\partial\varepsilon}\right|_{\varepsilon=0}ds\right] = 0.$$

By Lemma A1, we can rewrite the above as

$$\left.\frac{d\mathcal{U}_t}{d\varepsilon}\right|_{\varepsilon=0} = -\underbrace{\mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\sigma_\theta a_s\widehat{\Delta}_s^{\ell}\ ds\right]}_{=\Lambda_t} + \mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\left.\frac{\partial\tilde{u}_s}{\partial\varepsilon}\right|_{\varepsilon=0}ds\right] = 0.$$

Evaluating the derivative above, the entire second term equals

$$\mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\sigma_\theta a_s\int_s^\infty e^{-\rho(\tau-s)}\tilde{u}_\tau' D_s\left(\ln\left(\theta_\tau\right)\right)d\tau\ ds\right] = \mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\sigma_\theta a_s\Delta_s^{\ell}\ ds\right],$$

where we used that $\Delta_t^{\ell}\equiv\mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)}D_t\left(\tilde{u}\left(\ln\left(\theta_s\right)\right)\right)ds\right]$. Putting everything together,

$$\mathbb{E}_t\left[\int_t^\infty e^{-\rho(s-t)}\sigma_\theta a_s\left\{\widehat{\Delta}_s^{\ell}-\Delta_s^{\ell}\right\}ds\right] = 0 \Longrightarrow \widehat{\theta}_s\widehat{\Delta}_s = \widehat{\Delta}_s^{\ell} = \Delta_s^{\ell} = \theta_s\Delta_s$$

at truth-telling. Since $\widehat{\theta}_s=\theta_s$ also holds at truth-telling (by definition), we conclude that $\widehat{\Delta}_s=\Delta_s$.

A.3 Proof of Proposition 3

Let $\{\theta_t\}$ denote the true process and $\{\theta_t + \varepsilon_t\}$ the reported process. Define a process

$$\Phi_t \equiv \int_0^t e^{-\rho s}\tilde{u}_{\theta,s}\varepsilon_s \; ds + e^{-\rho t}\Delta_t\varepsilon_t.$$

The equation of motion for ε_t is

$$d\varepsilon_t = \mu'(\theta_t)\,\varepsilon_t \; dt + \sigma_\theta\theta_t\left(dZ_t - d\widehat{Z}_t\right) + \sigma_\theta\varepsilon_t \; dZ_t,$$

where $\widehat{Z}_t$ is a Brownian motion under the probability measure $\mathbb{Q}$ induced by the report. Since Φ_t is a martingale, its drift is zero:

$$0 = \frac{1}{dt}\mathbb{E}_t\left[d\Phi_t\right] = \tilde{u}_{\theta,t} - \rho\Delta_t + \mu_t^{\Delta} + \mu'(\theta_t)\,\Delta_t + \sigma_\theta^2\sigma_{\Delta,t},$$

where $\sigma_\theta\sigma_{\Delta,t}$ is the volatility of Δ_t. Rearranging this expression gives the desired result. The transversality condition holds because Δ_0 is finite.

Conversely, suppose that $\{\Delta_t\}$ is simply a process that satisfies (9). Then the process

$$\Phi_{t'}^t \equiv \mathbb{E}_{t'}\left[\int_t^{t'\wedge\tau} e^{-\rho(s-t)}\tilde{u}_{\theta,s}\varepsilon_s \; ds\right] + e^{-\rho(t'\wedge\tau-t)}\Delta_{t'}\varepsilon_{t'}$$

is a martingale. Therefore,

$$\Delta_t = \Phi_t^t = \lim_{t'\to\infty}\mathbb{E}_t\left[\Phi_{t'}^t\right] = \lim_{t'\to\infty}\mathbb{E}_t\left[\int_t^{t'\wedge\tau} e^{-\rho(s-t)}\tilde{u}_{\theta,s}\varepsilon_s \; ds\right] + \lim_{t'\to\infty}\mathbb{E}_t\left[e^{-\rho(t'\wedge\tau-t)}\Delta_{t'}\varepsilon_{t'}\right].$$

Then the transversality condition implies that Δ_t satisfies the definition given.

A.4 Proof of Theorem 1

Consider any alternative reporting strategy $\{\widehat{\theta}_t\}$ deviating from truth-telling, $\{\theta_t\}$, where $\widehat{\theta}_t = \theta_t + \varepsilon_t$ and $\varepsilon_t < 0$ for almost all t. To prove that the payoff to the agent from this deviation cannot exceed the equilibrium payoff v_t, we follow Sannikov (2014) by constructing an upper bound for the utility gain from deviating.

To construct the upper bound, we show that it is sufficient to keep track of two additional deviation state variables that matter to the agent's potential deviation value. These variables are the deviation itself, $\varepsilon_t < 0$, and a variable that, in a sense, captures accumulated information rent:

$$\Phi_t \equiv \int_0^t e^{-\rho s}\tilde{u}_{\theta,s}\varepsilon_s \; ds + e^{-\rho t}\Delta_t\varepsilon_t = \mathbb{E}_t\left[\int_0^\infty e^{-\rho t}\tilde{u}_{\theta,t}\varepsilon_t \; dt\right].$$

By applying Itô's Lemma, we see that

$$
\begin{aligned}
d\Phi_t &= e^{-\rho t}\left[\left(\tilde{u}_{\theta,t} + \sigma_\theta^2\sigma_{\Delta,t} - (\rho - \mu'(\theta_t))\,\Delta_t + \mu_t^{\Delta}\right)\varepsilon_t\right]dt + e^{-\rho t}\varepsilon_t\sigma_\theta(1+\sigma_{\Delta,t})\;dZ_t \\
&\quad + e^{-\rho t}\Delta_t\sigma_\theta\theta_t\left(dZ_t - d\hat{Z}_t\right) + e^{-\rho t}\varepsilon_t\left(\Delta_t^R - \Delta_t\right)(dR_t - \kappa\;dt) \\
d\Phi_t\cdot d\varepsilon_t &= e^{-\rho t}\varepsilon_t^2\sigma_\theta^2(1+\sigma_{\Delta,t})dt,
\end{aligned}
$$

where we used results from the proof of Proposition 3. By the Girsanov Theorem, there exists a process $\{\widetilde{X}_t\}$ such that $\widehat{Z}_t \equiv Z_t + \int_0^t \widetilde{X}_s\;ds$ is a Brownian motion under the probability measure induced by a deviation; note that $dZ_t - d\widehat{Z}_t = X_t\;dt \equiv -\widetilde{X}_t\;dt > 0$ since agents cannot overreport.

Given these two variables, we construct a candidate upper bound for an agent's deviation value:

$$\widehat{v}\left(\varepsilon_t, \Phi_t\right) \equiv v_t + \Phi_t\varepsilon_t + L\varepsilon_t^2,$$

where $L > 0$ is a constant to be specified later. If $\widehat{v}\left(\varepsilon_t, \Phi_t\right)$ is indeed an upper bound, then for an agent who has not yet deviated with $\varepsilon_t = \Phi_t = 0$, the upper bound of his deviation value is just v_t.

We will now prove that $\widehat{v}\left(\varepsilon_t, \Phi_t\right)$ is indeed an upper bound on an agent's deviation utility. To simplify notation, we will drop tildes on the utility function, so that $u_t \equiv \tilde{u}\left(\theta_t \mid \theta_t\right)$ and $\widehat{u}_t \equiv \tilde{u}\left(\widehat{\theta}_t \mid \theta_t\right)$. Define the auxiliary deflated gains process $\widehat{G}_t$ associated with any feasible policy $\{\widehat{\theta}_t\}$:

$$\widehat{G}_t \equiv \int_0^{t\wedge\tau} e^{-\rho s}\hat{u}_s\;ds + e^{-\rho(t\wedge\tau)}\widehat{v}\left(\varepsilon_{t\wedge\tau}, \Phi_{t\wedge\tau}\right).$$

To simplify notation, we will omit the retirement time τ unless doing so would be confusing. Clearly,

$$\mathbb{E}_0\left[\widehat{G}_\infty\right] = \mathbb{E}_0\left[\int_0^\infty e^{-\rho t}\hat{u}_t\;dt\right]$$

is the expected payoff under any feasible reporting strategy,[33] given the transversality condition in Proposition 1. On the other hand, $\widehat{G}_0 = \widehat{v}\left(\varepsilon_0, \Phi_0\right)$ is the proposed upper bound of an agent's deviation value given the current relevant deviation states (ε_0, Φ_0). For the upper bound to be valid, we need $\widehat{G}_t$ to be a supermartingale for any deviation strategy from the agent's perspective, and a martingale under truth-telling.

Differentiating $\widehat{G}_t$ with respect to t, we have

$$e^{\rho t}d\widehat{G}_t = \hat{u}_t\;dt - \rho\widehat{v}_t\;dt + d\widehat{v}_t.$$

[33] Throughout the proof, $\mathbb{E}$ refers to the expectation under the measure induced by the deviation strategy.

Applying Itô's Lemma and taking conditional expectations to capture only the drift terms,

$$
\begin{aligned}
e^{\rho t}\mathbb{E}_t\left[\frac{d\widehat{G}_t}{dt}\right] &= \hat{u}_t - \rho\left(v_t + \Phi_t\varepsilon_t + L\varepsilon_t^2\right) + \underbrace{\rho v_t - u_t}_{=\mathbb{E}_t[dv_t]} + \mathbb{E}_t\left[d\left(\Phi_t\varepsilon_t\right)\right] + L\mathbb{E}_t\left[d\varepsilon_t^2\right] \\
&= \hat{u}_t - u_t - \rho\Phi_t\varepsilon_t - \rho L\varepsilon_t^2 + \underbrace{e^{-\rho t}\varepsilon_t\left[\left(u_{\theta,t} - \rho\Delta_t + \mu_t^{\Delta} + \mu'\left(\theta_t\right)\Delta_t + \sigma_\theta^2\sigma_{\Delta,t}\right)\varepsilon_t + \sigma_\theta\theta_t\Delta_t X_t\right]}_{=\varepsilon_t\mathbb{E}_t[d\Phi_t]} \\
&+ \underbrace{\Phi_t\left[\mu'\left(\theta_t\right)\varepsilon_t + \sigma_\theta\theta_t X_t\right]}_{=\Phi_t\mathbb{E}_t[d\varepsilon_t]} + \underbrace{e^{-\rho t}\sigma_\theta^2\left(1 + \sigma_{\Delta,t}\right)\varepsilon_t^2}_{=\mathbb{E}_t[d\Phi_t\cdot d\varepsilon_t]} + \underbrace{2L\mu'\left(\theta_t\right)\varepsilon_t^2 + 2L\sigma_\theta\theta_t X_t\varepsilon_t + L\sigma_\theta^2\varepsilon_t^2}_{=L\mathbb{E}_t\left[d\varepsilon_t^2\right]} \\
&= \hat{u}_t - u_t + \sigma_\theta\theta_t X_t\Phi_t + \varepsilon_t\left(\sigma_\theta\theta_t X_t\left(e^{-\rho t}\Delta_t + 2L\right) + \left(\mu'\left(\theta_t\right) - \rho\right)\Phi_t\right) \\
&+ \varepsilon_t^2\left(e^{-\rho t}\sigma_\theta^2\left(1 + \sigma_{\Delta,t}\right) + L\left(2\mu'\left(\theta_t\right) + \sigma_\theta^2 - \rho\right)\right),
\end{aligned}
$$

where the last step uses the fact that, by Proposition 3, under the first-order approach

$$
\mu_t^{\Delta} = \left(\rho - \mu'\left(\theta_t\right)\right)\Delta_t - u_{\theta,t} - \sigma_\theta^2\sigma_{\Delta,t}.
$$

Now, since $\widehat{\theta}_t \leq \theta_t$ and so $\varepsilon_t < 0$ and $\Delta_t, X_t \geq 0$, we must have $\sigma_\theta\theta_t X_t\Phi_t < 0$. Using the definition of Φ_t, and taking a first-order Taylor expansion around $\varepsilon_t = 0$, the first three terms can be written as

$$
\begin{aligned}
&\widehat{u}_t - u_t + \sigma_\theta\theta_t Y_t\int_0^t e^{-\rho s}u_{\theta,s}\varepsilon_s\, ds + \sigma_\theta\theta_t Y_t e^{-\rho t}\Delta_t\varepsilon_t \\
&\qquad = \sigma_\theta\theta_t Y_t\int_0^t e^{-\rho s}u_{\theta,s}\varepsilon_s\, ds + \left(\sigma_\theta\theta_t Y_t e^{-\rho t}\Delta_t + \widehat{u}_{\widehat{\theta},t}\right)\varepsilon_t + O\left(\varepsilon_t^2\right).
\end{aligned}
$$

It follows that

$$
\begin{aligned}
e^{\rho t}\mathbb{E}_t\left[\frac{d\widehat{G}_t}{dt}\right] &= \sigma_\theta\theta_t Y_t\int_0^t e^{-\rho s}u_{\theta,s}\varepsilon_s\, ds + \varepsilon_t\left(2\sigma_\theta\theta_t Y_t\left(e^{-\rho t}\Delta_t + L\right) + \left(\mu'\left(\theta_t\right) - \rho\right)\Phi_t + \widehat{u}_{\widehat{\theta},t}\right) \\
&+ \varepsilon_t^2\left(e^{-\rho t}\sigma_\theta^2\left(1 + \theta_t\Gamma_t\right) + L\left(2\mu'\left(\theta_t\right) + \sigma_\theta^2 - \rho\right)\right) + O\left(\varepsilon_t^2\right) \\
&< \varepsilon_t\Phi_t\left(\mu'\left(\theta_t\right) - \rho\right) + \varepsilon_t^2\left(e^{-\rho t}\sigma_\theta^2\left(1 + \theta_t\Gamma_t\right) + L\left(2\mu'\left(\theta_t\right) + \sigma_\theta^2 - \rho\right)\right) + O\left(\varepsilon_t^2\right) \\
&< \varepsilon_t^2\left(e^{-\rho t}\sigma_\theta^2\left(1 + \theta_t\Gamma_t\right) + L\left(2\mu'\left(\theta_t\right) + \sigma_\theta^2 - \rho\right)\right) + O\left(\varepsilon_t^2\right).
\end{aligned}
$$

The first inequality uses that $\varepsilon_t < 0 \leq X_t$ and that I can choose L large enough that $2\sigma_\theta\theta_t Y_t + \widehat{u}_{\widehat{\theta},t} > 0$ regardless of the sign of $\widehat{u}_{\widehat{\theta},t}$. The second inequality uses that $\mu'\left(\theta_t\right) < \rho$ (a consequence of Assumption 1) and $\varepsilon_t, \Phi_t < 0$. Now, since $\theta_t \in \Theta_t = \left[\underline{\theta}, \overline{\theta}\right] \subseteq \mathbb{R}_+$ this means that $\sigma_{\Delta,t}$ is bounded above and below. Since $2\mu'\left(\theta_t\right) + \sigma_\theta^2 < \rho$ by Assumption 1, by again choosing L sufficiently large, we can guarantee that the coefficient on ε_t^2 is negative. Therefore, we know that

$$
e^{\rho t}\mathbb{E}_t\left[\frac{d\widehat{G}_t}{dt}\right] < -\beta_t\varepsilon_t^2
$$

for some time-dependent coefficient $\beta_t > 0$. This parabola is always nonpositive and strictly negative if $\varepsilon_t \neq 0$. Therefore, $\widehat{G}_t$ has strictly negative drift for any deviation strategy, and zero drift under truth-telling. To ensure that $\widehat{G}_t$ is a supermartingale, we need to ensure that all of the Brownian and Poisson terms have zero expectation on any time horizon (see Revuz and Yor (1999)). But since everything is bounded, this holds.

Now, given the fact that $\widehat{G}_t$ is a supermartingale, we have

$$\begin{aligned}\widehat{v}\left(\varepsilon_0, \Phi_0\right) = \widehat{G}_0 \geq \mathbb{E}_0\left[\lim_{t\to\infty} \widehat{G}_t\right] &= \mathbb{E}_0\left[\lim_{t\to\infty}\int_0^{t\wedge\tau} e^{-\rho s}\hat{u}_s\, ds + \lim_{t\to\infty} e^{-\rho(t\wedge\tau)} v_{t\wedge\tau}\right] \\ &= \mathbb{E}_0\left[\lim_{t\to\infty}\int_0^{t\wedge\tau} e^{-\rho s}\hat{u}_s\, ds\right],\end{aligned}$$

which is an agent's deviation payoff. The last equality uses the transversality condition on v_t in Proposition 1. This implies that $\widehat{v}_t$ is an upper bound for the agent's deviation value.

Thus, $\widehat{v}\left(\varepsilon_t, \Phi_t\right)$ is an upper bound for the agent's potential deviation value given the deviation states (ε_t, Φ_t). Then, for an agent who has not yet deviated with $\varepsilon_t = \Phi_t = 0$, the upper bound of his deviation value is just v_t. Because the equilibrium strategy achieves this upper bound, the equilibrium strategy is indeed globally optimal. Thus, the equilibrium strategy that achieves v_t is indeed optimal.

B Section 4 Proofs

B.1 Proof of Proposition 4

By the Envelope Theorem with respect to v,

$$(r+\kappa)\, k_v = (\rho+\kappa)\, k_v + \frac{\mathbb{E}\left[dk_v\right]}{dt}$$

while by Itô's Lemma,

$$\begin{aligned}d\lambda_t &= \mathbb{E}_t\left[d\lambda_t\right] + \sigma_\theta\left(k_{vv,t}\theta_t\Delta_t + k_{v\Delta,t}\sigma_{\Delta,t} + k_{v\theta,t}\theta_t\right) dZ_t + \left(\lambda_t^R - \lambda_t\right) dR_t \\ &= (r-\rho)\,\lambda_t\, dt + \sigma_\theta\left(k_{vv,t}\theta_t\Delta_t + k_{v\Delta,t}\sigma_{\Delta,t} + k_{v\theta,t}\theta_t\right) dZ_t + \left(\lambda_t^R - \lambda_t\right) dR_t,\end{aligned}$$

where the second line uses the envelope condition. From the first-order condition for $\sigma_{\Delta,t}$,

$$\sigma_{\Delta,t} = \frac{k_{\Delta,t} - k_{v\Delta,t}\theta_t\Delta_t - k_{\Delta\theta,t}\theta_t}{k_{\Delta\Delta,t}}.$$

Plugging this in yields the volatility:

$$\begin{aligned}\mathrm{Vol}\left(\lambda_t\right) &= \sigma_\theta\left(k_{vv,t}\theta_t\Delta_t + \frac{k_{v\Delta,t}}{k_{\Delta\Delta,t}}\left(K_{\Delta,t} - k_{v\Delta,t}\theta_t\Delta_t - k_{\Delta\theta,t}\theta_t\right) + k_{v\theta,t}\theta_t\right) \\ &= k_{v,t}\cdot\frac{\sigma_\theta}{k_{v,t}}\left(\frac{k_{vv,t}k_{\Delta\Delta,t} - k_{v\Delta,t}^2}{k_{\Delta\Delta,t}}\theta_t\Delta_t + \frac{k_{v\Delta,t}}{k_{\Delta\Delta,t}}\left(k_{\Delta,t} - k_{\Delta\theta,t}\theta_t\right) + k_{v\theta,t}\theta_t\right) \\ &\equiv \sigma_\theta\sigma_{\lambda,t}\lambda_t.\end{aligned}$$

Next, once again applying the Envelope Theorem to (14),

$$(r+\kappa)\,k_\Delta = (\rho+\kappa-\mu'(\theta))\,k_\Delta + \sigma_\theta^2\theta\,(k_{vv}\theta\Delta + k_{v\Delta}\sigma_\Delta + k_{v\theta}\theta) + \frac{\mathbb{E}\,[dk_\Delta]}{dt}.$$

Using Itô's Lemma,

$$\begin{aligned} d\gamma_t &= \mathbb{E}_t\,[d\gamma_t] + \sigma_\theta\,(k_{v\Delta,t}\theta_t\Delta_t + k_{\Delta\Delta,t}\sigma_{\Delta,t} + k_{\Delta\theta,t}\theta_t)\,dZ_t + \left(\gamma_t^R - \gamma_t\right)dR_t \\ &= \left((r-\rho+\mu'(\theta))\,\gamma_t - \sigma_\theta^2\theta_t\sigma_{\lambda,t}\lambda_t\right)dt + \sigma_\theta\gamma_t\,dZ_t + \left(\gamma_t^R - \gamma_t\right)dR_t, \end{aligned}$$

where the second line uses (20).

We give an alternative proof in Appendix D based on the stochastic maximum principle.

Online Appendix to "Dynamic Optimal Taxation with Endogenous Skill Premia"

Zongbo Huang, Jason Ravit, and Michael Sockin

C Mean Field Game Boundary Conditions

C.1 Hamilton-Jacobi-Bellman Equation

The mean field game in our model comprises two coupled, second-order, three-dimensional partial differential equations, so we need six boundary conditions for each one to pin down a unique solution.

We begin with the HJB equation (14). First, both $\underline{\theta}$ and $\overline{\theta}$ are reflecting boundaries:

$$k_\theta\left(v,\Delta,\underline{\theta}\right)=0, \tag{C-1}$$

$$k_\theta\left(v,\Delta,\overline{\theta}\right)=0. \tag{C-2}$$

Let $\underline{v}$ and $\overline{v}$ denote the lowest and highest possible continuation utilities, respectively, and $\underline{\Delta}$ and $\overline{\Delta}$ the lowest and highest possible information rents, respectively. Since an agent can abandon the contract if continuation utility drops below its autarky value, we set $\underline{v}=v^{\text{aut}}$. For v to reflect back inwards at v^{aut}, we need

$$k_v\left(v^{\text{aut}},\Delta,\theta\right). \tag{C-3}$$

Since information rent cannot be negative, $\underline{\Delta}=0$. At this point, the principal does not ask the agent to produce anything so $y=0$ and the cost to the planner is simply the cost of the consumption stream needed to achieve promised utility v. Now, let

$$v=\frac{u\left(c\right)+\kappa v^{R}\left(C\right)}{\rho+\kappa}$$

denote the agent's promised utility as a function of consumption (since there's full insurance during working life and retirement). Since the retirement shock is observable, the marginal utility of consumption just before and just after retirement must be equal:

$$c^{-\omega}=u'\left(c\right)=v^{R\prime}\left(C\right)=\psi u'\left(C\right)=\psi C^{-\omega}\Longrightarrow C=\psi^{\frac{1}{\omega}}c.$$

This provides a link between consumption pre- and post-retirement. Then

$$v=\frac{u\left(c\right)+\kappa\psi u\left(\psi^{\frac{1}{\omega}}c\right)}{\rho+\kappa}=\frac{c^{1-\omega}\left(1+\kappa\psi^{\frac{1}{\omega}}\right)-\left(1+\kappa\psi\right)}{\left(1-\omega\right)\left(\rho+\kappa\right)}\Longrightarrow c=\left(\frac{\left(1-\omega\right)\left(\rho+\kappa\right)v+1+\kappa\psi}{1+\kappa\psi^{\frac{1}{\omega}}}\right)^{\frac{1}{1-\omega}}.$$

It follows that[34]

$$k(v,0,\theta)=\frac{c+\kappa C}{\rho+\kappa}=\left(\frac{(1-\omega)(\rho+\kappa)v+1+\kappa\psi}{1+\kappa\psi^{\frac{1}{\omega}}}\right)^{\frac{1}{1-\omega}}\frac{1+\kappa\psi^{\frac{1}{\omega}}}{\rho+\kappa}. \tag{C-4}$$

For the fifth boundary condition, we impose a reflecting boundary at $\overline{\Delta}$, or

$$k_{\Delta}\left(v,\overline{\Delta},\theta\right)=0. \tag{C-5}$$

This upper reflecting boundary is endogenous and comes from setting $\mathbb{E}_t\left[d\Delta_t\right]=0$ and using the first-order condition for ℓ. Finally, since an agent will not be promised more utility than are not receive in the first best allocation, we set $\overline{v}=v^{\mathrm{FB}}$. Let k^* denote the planner's cost function under perfect information. Then

$$k\left(v^{\mathrm{FB}},\Delta,\theta\right)=k^*\left(v^{\mathrm{FB}},\Delta,\theta\right). \tag{C-6}$$

C.2 Kolmogorov Forward Equation

Note that the FKE can be written as

$$0=-\sum_{j=1}^{3}\frac{\partial}{\partial x_j}\left(a_i(x)\,g\right)+\frac{1}{2}\sum_{i,j=1}^{3}\frac{\partial^2}{\partial x_i\partial x_j}\left(\sigma_{ij}(x)\,g\right).$$

It is useful to write this in divergence form,

$$\nabla\cdot\mathbf{S}=0,$$

where

$$S_i=a_i(x)-\frac{1}{2}\sum_{j=1}^{3}\frac{\partial}{\partial x_j}\left(\sigma_{ij}(x)\,g\right)$$

is the i^{th} element of $\mathbf{S}$. Here, $\mathbf{S}$ represents the "probability flux" representing the flux of particles through the point $x=(x_1,x_2,x_3)=(v,\Delta,\theta)$.

Several of the boundary conditions are reflecting. If variable-i is reflecting, then

$$\mathbf{S}_i\cdot\widehat{\mathbf{n}}_i=0,$$

where $\widehat{\mathbf{n}}_i$ is the outward unit normal vector in the i^{th} direction evaluated at the reflecting boundary. Since θ has reflecting boundaries, this implies that $\widehat{\mathbf{n}}_i=(0,0,1)$ at $\overline{\theta}$ and $\widehat{\mathbf{n}}_i=$

[34] With logarithmic utility, the condition is $k(v,0,\theta)=\exp\left(\frac{(\rho+\kappa)v-\kappa\psi\ln(\psi)}{\rho+\kappa\psi}\right)\frac{1+\kappa\psi^{\frac{1}{\omega}}}{\rho+\kappa}$.

$(0,0,-1)$ at $\underline{\theta}$. Thus, the two boundary conditions are

$$\mu\left(\overline{\theta}\right) = \frac{1}{2}\sum_{j=1}^{3}\frac{\partial}{\partial x_j}\left(\sigma_{3j}\left(x\right)g\right)\Big|_{\theta=\overline{\theta}},$$

$$\mu\left(\underline{\theta}\right) = \frac{1}{2}\sum_{j=1}^{3}\frac{\partial}{\partial x_j}\left(\sigma_{3j}\left(x\right)g\right)\Big|_{\theta=\underline{\theta}}.$$

Similarly, $\{\underline{v},\overline{v}\} = \{v^{\text{aut}}, v^{\text{FB}}\}$ and $\{\underline{\Delta},\overline{\Delta}\} = \{0,\overline{\Delta}\}$ are reflecting. The outward normal unit vector associated with these are $(\mp 1,0,0)$ and $(0,\mp 1,0)$, respectively. The remaining four boundary conditions can then be derived as above.

D Contracts via Optimal Stochastic Control

This section outlines a procedure to construct an equivalent, but simplified planner's problem. We rely on the stochastic maximum principle to solve an unconstrained maximization problem in the space of Lagrange multipliers.

To simplify the problem, we work with $\log\left(\theta_t\right)$ instead of θ_t. In addition, v_t and Δ_t now refer to the state variables in terms of $\log\left(\theta_t\right)$. The laws of motion of the state variables are

$$\begin{aligned}
d\left(\log\left(\theta_t\right)\right) &= \left(\frac{\mu\left(\theta_t\right)}{\theta_t} - \frac{\sigma_\theta^2}{2}\right)dt + \sigma_\theta\, dZ_t = \mu^{\log}\left(\log\left(\theta_t\right)\right)dt + \sigma_\theta\, dZ_t,\\
dv_t &= \left(\rho v_t - \tilde{u}_t\right)dt + \sigma_\theta \Delta_t\, dZ_t + \left(v_t^R - v_t\right)\left(dR_t - \kappa\, dt\right),\\
d\Delta_t &= \left(\left(\rho - \frac{d\mu^{\log}\left(\log\left(\theta_t\right)\right)}{d\log\left(\theta_t\right)}\right)\Delta_t - \tilde{u}_{\log(\theta),t}\right)dt + \sigma_\theta\sigma_{\Delta,t}\, dZ_t + \left(\Delta_t^R - \Delta_t\right)\left(dR_t - \kappa\, dt\right).
\end{aligned}$$

To simplify notation, we will relabel $\log\left(\theta_t\right)$ with θ_t and $\widetilde{\mu}\left(\theta\right) \equiv \mu^{\log}\left(\log\left(\theta\right)\right)$. The planner's HJB equation (14) is very difficult to work with in its current form because the directions of the higher-order derivatives depend on the choice of σ_Δ. We will instead reparametrize the state space in terms of Lagrange multipliers on the state variables.

We will set up and solve the planner's problem via the stochastic maximum principle. Recall that the controls are $\{c, C, \ell, \sigma_\Delta, z\}$. Let λ and γ denote the Lagrange multipliers on the laws of motion of v and Δ, respectively, and $\sigma_\theta\sigma_\lambda, \sigma^\gamma$ their respective volatilities. Then the Lagrangian is

$$\begin{aligned}
\mathcal{L}\left(g\right) = \int_{\Lambda\times\Gamma}\mathscr{L}\left(\lambda,\gamma,\theta\right)g\left(\lambda,\gamma,\theta\right)d\left(\lambda,\gamma\right) &= \left\langle e^{-rt}\left[c + \kappa C - w\left(\theta\right)e^{\theta}\ell\right], g\right\rangle + \left\langle e^{-rt}f, \mathcal{A}^*g + \kappa\widetilde{f}_0\right\rangle\\
&+ \left\langle e^{-rt}\chi\left(\theta\right), w\left(\theta\right)^{\alpha}L\left(\theta\right) - Y\right\rangle + \left\langle e^{-rt}\eta\left(\theta\right), z\left(\theta\right)\right\rangle.
\end{aligned}$$

Integrating by parts,

$$\begin{aligned}
\mathcal{L}\left(g\right) &= \left\langle e^{-rt}\left[c + \kappa C - w\left(\theta\right)e^{\theta}\ell\right], g\right\rangle + \left\langle e^{-rt}\left(\mathcal{A}f + \kappa\widetilde{f}_0 - rf\right), g\right\rangle\\
&+ \left\langle e^{-rt}\chi\left(\theta\right), w\left(\theta\right)^{\alpha}L\left(\theta\right) - Y\right\rangle + \left\langle e^{-rt}\eta\left(\theta\right), z\left(\theta\right)\right\rangle.
\end{aligned}$$

Expanding this out and substituting in λ, γ yields

$$\begin{aligned}\mathcal{L}(g) &= \left\langle e^{-rt}\left[c+\kappa C - w(\theta)e^{\theta}\ell + \kappa\widetilde{f}_0 - (r+\kappa)f\right], g\right\rangle + \left\langle e^{-rt}\left[\lambda\left(\dot{v} - ((\rho+\kappa)v - \tilde{u})\right)\right], g\right\rangle \\ &+ \left\langle e^{-rt}\left[\gamma\left(\dot{\Delta} - (\rho+\kappa-\widetilde{\mu}'(\theta))\Delta + \tilde{u}_\theta\right)\right], g\right\rangle + \left\langle e^{-rt}\left[f_\theta\left(\dot{\theta} - \widetilde{\mu}'(\theta)\right)\right], g\right\rangle \\ &- \left\langle e^{-rt}\left[\sigma_\theta^2\sigma_\lambda\Delta + \sigma_\theta\sigma_\Delta\sigma^\gamma + \sigma_\theta \mathrm{Vol}(f_\theta)\right], g\right\rangle + \left\langle e^{-rt}\chi(\theta), w(\theta)^\alpha L(\theta) - Y\right\rangle + \left\langle e^{-rt}\eta(\theta), z(\theta)\right\rangle.\end{aligned}$$

Integrating by parts one more time,

$$\begin{aligned}\mathcal{L}(g) &= \left\langle e^{-rt}\left[c+\kappa C - w(\theta)e^{\theta}\ell + \kappa\widetilde{f}_0 - (r+\kappa)f\right], g\right\rangle - \left\langle e^{-rt}\left[\lambda\left((\rho+\kappa)v - \tilde{u}\right) + v\dot{\lambda}\right], g\right\rangle \\ &- \left\langle e^{-rt}\left[\gamma\left((\rho+\kappa-\widetilde{\mu}'(\theta))\Delta - \tilde{u}_\theta\right) + \Delta\dot{\gamma}\right], g\right\rangle + \left\langle e^{-rt}\left[f_\theta\left(\dot{\theta} - \widetilde{\mu}'(\theta)\right)\right], g\right\rangle \\ &- \left\langle e^{-rt}\left[\sigma_\theta^2\sigma_\lambda\Delta + \sigma_\theta\sigma_\Delta\sigma^\gamma + \sigma_\theta \mathrm{Vol}(f_\theta)\right], g\right\rangle + \left\langle e^{-rt}\chi(\theta), w(\theta)^\alpha L(\theta) - Y\right\rangle + \left\langle e^{-rt}\eta(\theta), z(\theta)\right\rangle.\end{aligned}$$

Differentiating with respect to σ_Δ yields

$$[\sigma_\Delta]: \ -\sigma_t^\gamma\sigma_\theta = 0.$$

Thus γ_t has zero volatility under the optimal contract and $\lambda_t = \tilde{u}_{c,t}^{-1}$. Next, we can see that the drift of λ_t is $\mathscr{L}_v + (r+\kappa)\lambda_t = 0$ and the drift of γ_t is $\mathscr{L}_\Delta + (r+\kappa)\gamma_t = (r - \rho + \widetilde{\mu}'(\theta_t))\gamma_t - \sigma_\theta^2\sigma_{\lambda,t} = \widetilde{\mu}'(\theta_t)\gamma_t - \sigma_\theta^2\sigma_{\lambda,t}$ since we've assumed $r = \rho$. Thus we can write

$$d\lambda_t = \sigma_\theta\sigma_{\lambda,t}\,dZ_t + \left(\lambda_t^R - \lambda_t\right)dR_t, \tag{D-1}$$

$$d\gamma_t = \left(\widetilde{\mu}'(\theta_t)\gamma_t - \sigma_\theta^2\sigma_{\lambda,t}\right)dt + \left(\gamma_t^R - \gamma_t\right)dR_t. \tag{D-2}$$

In particular, λ_t has no drift and γ_t has no volatility, i.e., it is a slow-moving variable. If we compare these first-order conditions to the ones that we would obtain from the original HJB equation (under the reparametrized state space), then we see that $\lambda_t = k_{v,t}$ and $\gamma_t = k_{\Delta,t}$.

We solve for v, Δ, and the planner's profit function p jointly via a system of partial differential equations over the new state space. Let $\mathrm{P}(\theta) = \int_{\Lambda\times\Gamma} \mathrm{p}(\lambda,\gamma,\theta)\, g(\lambda,\gamma,\theta)\, d(\lambda,\gamma)$. Then these functions can be represented recursively as

$$\begin{aligned}(r+\kappa)v &= \tilde{u} + \kappa v^R + v_\theta\widetilde{\mu}(\theta) + v_\gamma\left(\widetilde{\mu}'(\theta)\gamma - \sigma_\theta^2\sigma_\lambda\right) + \frac{\sigma_\theta^2}{2}\left[v_{\lambda\lambda}\sigma_\lambda^2 + v_{\theta\theta}\right] + v_{\lambda\theta}\sigma_\theta^2\sigma_\lambda, \\ (r+\kappa)\Delta &= \tilde{u}_\theta + \Delta_\theta\widetilde{\mu}(\theta) + \Delta_\gamma\left(\widetilde{\mu}'(\theta)\gamma - \sigma_\theta^2\sigma_\lambda\right) + \frac{\sigma_\theta^2}{2}\left[\Delta_{\lambda\lambda}\sigma_\lambda^2 + \Delta_{\theta\theta}\right] + \Delta_{\lambda\theta}\sigma_\theta^2\sigma_\lambda, \\ (r+\kappa)\mathrm{P}(\theta) &= \int_{\Lambda\times\Gamma}\Bigg\{w(\theta)e^\theta\ell - c - \kappa C + \mathrm{p}_\theta\widetilde{\mu}(\theta) + \mathrm{p}_\gamma\left(\widetilde{\mu}'(\theta)\gamma - \sigma_\theta^2\sigma_\lambda\right) + \frac{\sigma_\theta^2}{2}\left[\mathrm{p}_{\lambda\lambda}\sigma_\lambda^2 + \mathrm{p}_{\theta\theta}\right] \\ &\quad + \mathrm{p}_{\lambda\theta}\sigma_\theta^2\sigma_\lambda + \kappa\mathrm{p}\widetilde{f}_0 + \chi(\theta)w(\theta)^\alpha e^\theta\ell - \mathcal{X}w(\theta)e^\theta\ell\Bigg\}g(\theta,\lambda,\gamma)\,d(\lambda,\gamma) + \eta(\theta)z(\theta).\end{aligned}$$

Alternatively, taking the Gateaux derivative of the Lagrangian with respect to f along an arbitrary variation, all three equations are characterized by the function $f : (\lambda,\gamma,\theta) \to \mathbb{R}$

that solves

$$(r+\kappa) F(\theta) = \max_{\{c,C,\ell,z,\sigma_\lambda\}} \Bigg\{ \int_{\Lambda\times\Gamma} \Big[w(\theta) e^{\theta}\ell - c - \kappa C + \lambda\left(\tilde{u} + \kappa v^R\right) + \gamma \tilde{u}_\theta + f_\theta \widetilde{\mu}(\theta)$$
$$+ f_\gamma \left(\widetilde{\mu}'(\theta_t)\gamma - \sigma_\theta^2 \sigma_\lambda\right) + \sigma_\theta^2 \left(\frac{f_{\lambda\lambda}\sigma_\lambda^2 + f_{\theta\theta}}{2} + \sigma_\lambda f_{\lambda\theta}\right) + \kappa f \widetilde{f_0}$$
$$+ \chi(\theta) w(\theta)^{\alpha} e^{\theta}\ell - \mathcal{X} w(\theta)\ell \Big] g(\lambda,\gamma,\theta) d(\lambda,\gamma) + \eta(\theta) z(\theta) \Bigg\}, \tag{D-3}$$

where for any triple $(\lambda, \gamma, \theta)$, the optimal allocation solves

$$\begin{aligned}
c &\in \underset{c'}{\operatorname{argmax}} \{\lambda\tilde{u} - c'\}, \\
C &\in \underset{C'}{\operatorname{argmax}} \left\{\lambda v^R - C'\right\}, \\
\ell &\in \underset{\ell'}{\operatorname{argmax}} \left\{ w(\theta) e^{\theta}\ell' + \lambda\tilde{u} + \gamma\tilde{u}_\theta + \chi(\theta) w(\theta)^{\alpha} e^{\theta}\ell' - \mathcal{X} w(\theta) e^{\theta}\ell' \right\}, \\
\sigma_\lambda &\in \underset{\sigma_\lambda'}{\operatorname{argmax}} \left\{ \frac{f_{\lambda\lambda}{\sigma_\lambda'}^2}{2} + (f_{\lambda\theta} - f_\gamma)\sigma_\lambda' \right\}, \\
z &\in \underset{z'}{\operatorname{argmax}} \left\{ \int_{\Lambda\times\Gamma} [\lambda\tilde{u} + \gamma\tilde{u}_\theta] g(\lambda,\gamma,\theta) d(\lambda,\gamma) + \eta z' \right\}.
\end{aligned}$$

Notice that p is the objective function in a constrained optimization problem while f is the objective function of an "unconstrained" problem. Therefore, these four equations can be linked via

$$f(\lambda,\gamma,\theta) = \max_{\{v,\Delta\}} \{\mathrm{p}(v,\Delta,\theta) + \lambda v + \gamma\Delta\}.$$

Then the Envelope Theorem implies

$$v = f_\lambda, \qquad \Delta = f_\gamma, \qquad \mathrm{p} = f - \lambda f_\lambda - \gamma f_\gamma.$$

Note that the fourth condition implies that the second-order terms in (D-3) can be simplified to $\frac{\sigma_\theta^2}{2}\left(f_{\theta\theta} - f_{\lambda\lambda}\sigma_\lambda^2\right)$. Crucially, (D-3) is considerably easier to solve than is (14). The boundary conditions are

$$\begin{aligned}
f_\theta(\cdot,\cdot,\underline{\theta}) &= \lambda f_{\lambda\theta}(\cdot,\cdot,\underline{\theta}) + \gamma f_{\gamma\theta}(\cdot,\cdot,\underline{\theta}), \\
f_\theta\left(\cdot,\cdot,\overline{\theta}\right) &= \lambda f_{\lambda\theta}\left(\cdot,\cdot,\overline{\theta}\right) + \gamma f_{\gamma\theta}\left(\cdot,\cdot,\overline{\theta}\right), \\
f_\lambda(\cdot,0,\cdot) &= \frac{u(c) + \kappa v^R(C)}{\rho+\kappa}, \\
f_\gamma(\cdot,0,\cdot) &= 0, \\
\mathrm{p}(\cdot,0,\cdot) &= f(\cdot,0,\cdot) - \lambda f_\lambda(\cdot,0,\cdot) = -\frac{c+\kappa C}{\rho+\kappa}.
\end{aligned}$$

However, since γ_t never actually reaches zero there is a single non-explosive solution to each equation so the boundary conditions are not actually necessary to solve the problem.

The result below provides a set of sufficient conditions for a contract as determined by (D-1) and (D-2) to solve the relaxed problem.

Proposition 5. *Suppose that the function f solves (D-3). Then (D-1) and (D-2) define a contract in which (v_t, Δ_t, p_t) are given by*

$$v_t = f_{\lambda,t}, \qquad \Delta_t = f_{\gamma,t}, \qquad p_t = f_t - \lambda f_{\lambda,t} - \gamma f_{\gamma,t}.$$

The contract is a solution of the relaxed problem if the Hessian $H(f)$ is positive definite.

Proof. The proof follows the one in Sannikov (2014) and has two key steps. First, we establish the mappings between the original state variables and the dual variables. Then, we show that the principal's profit is bounded above under any alternative contract. Since the last two constraints do not affect the proof, we omit them for brevity.

Step 1: Original/Dual Mapping

Lemma D1. *If f solves (D-3) then $v_t = f_\lambda(\lambda_t, \gamma_t, \theta_t)$, $\Delta_t = f_\gamma(\lambda_t, \gamma_t, \theta_t)$, and the planner's continuation payoff is $f(\lambda, \gamma_t, \theta_t) - \lambda_t v_t - \gamma_t \Delta_t$ in the contract defined by (D-1) and (D-2).*

Proof. Differentiating (D-3) with respect to λ and using the Envelope Theorem,

$$(r+\kappa) f_\lambda - \left(\tilde{u} + \kappa v^R\right) = f_{\lambda\theta}\widetilde{\mu}(\theta) + f_{\lambda\gamma}\widetilde{\mu}'(\theta)\gamma + \frac{\sigma_\theta^2}{2}\left[f_{\lambda\lambda\lambda}\sigma_\lambda^2 + f_{\theta\theta\lambda}\right] + \sigma_\theta^2\sigma_\lambda\left(f_{\lambda\lambda\theta} - f_{\lambda\gamma}\right).$$

The right side equals the drift of the process $f_\lambda(\lambda_t, \gamma_t, \theta_t)$ when (λ_t, γ_t) follow (D-1) and (D-2). Thus as long as the transversality condition holds, f_λ is an agent's continuation value v. Similarly, differentiating with respect to γ,

$$(r+\kappa) f_\gamma - \tilde{u}_\theta = f_{\gamma\theta}\widetilde{\mu}(\theta) + f_{\gamma\gamma}\widetilde{\mu}'(\theta)\gamma + f_\gamma\widetilde{\mu}'(\theta) + \frac{\sigma_\theta^2}{2}\left[f_{\lambda\lambda\gamma}\sigma_\lambda^2 + f_{\gamma\theta\theta}\right] + \sigma_\theta^2\sigma_\lambda\left(f_{\lambda\gamma\theta} - f_{\gamma\gamma}\right)$$

and the right side is the drift of f_γ so as long as the transversality condition holds then $f_\gamma = \Delta$. Finally, subtracting λ times $\mathbb{E}[df_\lambda]$ and γ times $\mathbb{E}[df_\gamma]$ from (D-3),

$$\begin{aligned}(r+\kappa)(f - \lambda f_\lambda - \gamma f_\gamma) = w(\theta)e^\theta\ell - c - \kappa C + \underbrace{\left[f_\theta - \lambda f_{\lambda\theta} - \gamma f_{\gamma\theta}\right]}_{=\frac{\partial(f-\lambda f_\lambda - \gamma f_\gamma)}{\partial\theta}}\widetilde{\mu}(\theta) + \underbrace{\left[f_\gamma - \lambda f_{\lambda\gamma} - \gamma f_{\gamma\gamma} - f_\gamma\right]}_{=\frac{\partial(f-\lambda f_\lambda - \gamma f_\gamma)}{\partial\gamma}}\widetilde{\mu}'(\theta)\gamma \\ + \underbrace{\left[f_{\lambda\theta} - f_\gamma - \lambda\left(f_{\lambda\lambda\theta} - f_{\lambda\gamma}\right) - \gamma\left(f_{\lambda\gamma\theta} - f_{\gamma\gamma}\right)\right]}_{=\frac{\partial^2(f-\lambda f_\lambda - \gamma f_\gamma)}{\partial\lambda\partial\theta} - \frac{\partial(f-\lambda f_\lambda - \gamma f_\gamma)}{\partial\gamma}}\sigma_\theta^2\sigma_\lambda \\ + \frac{\sigma_\theta^2}{2}\left[\underbrace{\left(-f_{\lambda\lambda} - \lambda f_{\lambda\lambda\lambda} - \gamma f_{\lambda\lambda\gamma}\right)}_{=\frac{\partial^2(f-\lambda f_\lambda - \gamma f_\gamma)}{\partial\lambda^2}}\sigma_\lambda^2 + \underbrace{f_{\theta\theta} - \lambda f_{\lambda\theta\theta} - \gamma f_{\gamma\theta\theta}}_{=\frac{\partial^2(f-\lambda f_\lambda - \gamma f_\gamma)}{\partial\theta^2}}\right].\end{aligned}$$

Hence, the process

$$\overline{\mathrm{p}}_t = \int_0^t e^{-(r+\kappa)s}\left(w(\theta_s)e^{\theta_s}\ell_s - c_s - \kappa C_s\right)ds + e^{-(r+\kappa)t}\left(f(\lambda_t, \gamma_t, \theta_t) - \lambda_t v_t - \gamma_t\Delta_t\right)$$

is a martingale. Since $\overline{\mathrm{p}}_t = \mathbb{E}_t\left[\overline{\mathrm{p}}_{T\wedge\tau}\right]$ it follows that $f(\lambda_t, \gamma_t, \theta_t) - \lambda_t v_t - \gamma_t \Delta_t$ is the planner's continuation payoff under the contract with a given agent. □

Step 2: Bounding the Value Function

We will now show that under any alternative contract, the principal's value function is bounded above. We do this in two steps.

Lemma D2. *Consider an alternative contract characterized by the controls (c, C, y, σ_Δ) and let (v, Δ) denote the state variables under these controls. If $H(f)$ is positive definite, then there exist processes*

$$d\lambda_t = \mu_t^\lambda \, dt + \sigma_\theta \sigma_{\lambda,t} \, dZ_t + J_t^\lambda \, dR_t \text{ and } d\gamma_t = \mu_t^\gamma \, dt + \sigma_\theta \sigma_{\gamma,t} \, dZ_t + J_t^\gamma \, dR_t$$

such that $v_t = f_{\lambda,t}$ and $\Delta_t = f_{\gamma,t}$.

Proof. We need to construct the processes above such that the drifts and volatilities match. To match volatilities, by Itô's Lemma,

$$\sigma_\theta \underbrace{\begin{bmatrix} f_{\lambda\lambda} & f_{\lambda\gamma} & f_{\lambda\theta} \\ f_{\lambda\gamma} & f_{\gamma\gamma} & f_{\gamma\theta} \end{bmatrix}}_{\equiv M(f)} \begin{bmatrix} \sigma_\lambda \\ \sigma_\gamma \\ 1 \end{bmatrix} = \sigma_\theta \begin{bmatrix} \Delta \\ \sigma_\Delta \end{bmatrix}.$$

Since $H(f)$ is positive definite, so are all of its principal submatrices, and positive definite matrices are invertible, so there exists a unique solution to this system of equations. To match drifts, we set up another system of equations:

$$M(f) \begin{bmatrix} \mu^\lambda \\ \mu^\gamma \\ \widetilde{\mu}(\theta) \end{bmatrix} + \cdots = \begin{bmatrix} (\rho + \kappa) v - \tilde{u} - \kappa v^R \\ (\rho + \kappa - \widetilde{\mu}'(\theta)) \Delta - \tilde{u}_\theta \end{bmatrix}.$$

Similarly, since $H(f)$ is positive definite a solution exists. The same logic applies to the jump term. □

We will now show that $\overline{\mathrm{p}}$ is a supermartingale under alternative contracts. By Itô's Lemma, the drift of $f(\lambda_t, \gamma_t, \theta_t) - \lambda_t v_t - \gamma_t \Delta_t$ is

$$\begin{aligned} \text{Drift} &= f_\lambda \mu_t^\lambda + f_\gamma \mu_t^\gamma + f_\theta \widetilde{\mu}(\theta_t) + \frac{\sigma_\theta^2}{2} \begin{bmatrix} \sigma_{\lambda,t} & \sigma_{\gamma,t} & 1 \end{bmatrix} H(f) \begin{bmatrix} \sigma_{\lambda,t} \\ \sigma_{\gamma,t} \\ 1 \end{bmatrix} - \mu_t^\lambda v_t - \mu_t^\gamma \Delta_t - \sigma_\theta^2 \begin{bmatrix} \sigma_{\lambda,t} & \sigma_{\gamma,t} \end{bmatrix} \begin{bmatrix} \Delta_t \\ \sigma_{\Delta,t} \end{bmatrix} \\ &\quad - \lambda_t \left((\rho + \kappa) v_t - \tilde{u}_t \right) - \gamma_t \left((\rho + \kappa - \widetilde{\mu}'(\theta_t)) \Delta_t - \tilde{u}_{\theta,t} \right) \\ &= f_\theta \widetilde{\mu}(\theta_t) - \frac{\sigma_\theta^2}{2} \begin{bmatrix} \sigma_{\lambda,t} & \sigma_{\gamma,t} & 1 \end{bmatrix} H(f) \begin{bmatrix} \sigma_{\lambda,t} \\ \sigma_{\gamma,t} \\ 1 \end{bmatrix} + \sigma_\theta^2 \left(f_{\lambda\theta} \sigma_{\lambda,t} + f_{\gamma\theta} \sigma_{\gamma,t} + f_{\theta\theta} \right) \\ &\quad - \lambda_t \left((\rho + \kappa) v_t - \tilde{u}_t \right) - \gamma_t \left((\rho + \kappa - \widetilde{\mu}'(\theta_t)) \Delta_t - \tilde{u}_{\theta,t} \right), \end{aligned}$$

where the final equality uses volatility matching. Under (D-1) and (D-2), this drift is

$$f_\theta \widetilde{\mu}(\theta_t) + \frac{\sigma_\theta^2}{2}\left(f_{\theta\theta} - f_{\lambda\lambda}{\sigma_{\lambda,t}^*}^2\right) - \lambda_t\left((\rho+\kappa)v_t - \tilde{u}_t\right) - \gamma_t\left(\left(\rho+\kappa-\widetilde{\mu}'(\theta_t)\right)\Delta_t - \tilde{u}_{\theta,t}\right),$$

where $\sigma_{\lambda,t}^*$ is the value under the optimal contract. It follows that the drift of $\overline{\mathrm{p}}$ is zero, i.e., it is a martingale under the optimal contract. On the other hand, with arbitrary laws of motion, the drift of $\overline{\mathrm{p}}$ changes by ($e^{-(r+\kappa)t}$ times)

$$\begin{aligned} &-\frac{\sigma_\theta^2}{2}\begin{bmatrix}\sigma_{\lambda,t} & \sigma_{\gamma,t} & 1\end{bmatrix} H(f)\begin{bmatrix}\sigma_{\lambda,t}\\ \sigma_{\gamma,t}\\ 1\end{bmatrix} + \sigma_\theta^2\left(f_{\lambda\theta}\sigma_{\lambda,t} + f_{\gamma\theta}\sigma_{\gamma,t} + f_{\theta\theta}\right) - \frac{\sigma_\theta^2}{2}\left(f_{\theta\theta} - f_{\lambda\lambda}{\sigma_{\lambda,t}^*}^2\right) \\ &= -\frac{\sigma_\theta^2}{2}\begin{bmatrix}x_t & \sigma_{\gamma,t} & 0\end{bmatrix} H(f)\begin{bmatrix}x_t\\ \sigma_{\gamma,t}\\ 0\end{bmatrix} \le 0 \end{aligned}$$

for some x_t that depends on the other processes and the inequality holds because $H(f)$ is positive definite. That is, under any alternative contract, $\overline{\mathrm{p}}$ is a supermartingale and a martingale under the optimal contract. □

E Numerical Appendix

This section describes our numerical procedure to solve the full model. Most of our algorithm is based on Nuño and Moll (2015) so we defer the interested reader there for more details on individual steps. We work in the reparametrized state space described in Appendix D.

We use a finite difference method and approximate the planner's value function, f, on a grid of $(\theta, \lambda, \gamma)$. Let $(\theta_i, \lambda_j, \gamma_k)$, with $i = 1, \ldots, I$, $j = 1, \ldots, J$, and $k = 1, \ldots, K$ denote a grid point, and

$$f_{i,j,k} = f(\theta_i, \lambda_j, \gamma_k)$$

the function f evaluated at that point. We approximate the derivatives with either a forward- or backward-difference approximation,

$$\begin{aligned} f_\theta(\theta_i, \lambda_j, \gamma_k) &\approx f^F_{\theta;i,j,k} = \frac{f_{i+1,j,k} - f_{i,j,k}}{\theta_{i+1} - \theta_i}, \\ f_\theta(\theta_i, \lambda_j, \gamma_k) &\approx f^B_{\theta;i,j,k} = \frac{f_{i,j,k} - f_{i-1,j,k}}{\theta_i - \theta_{i-1}}, \\ f_{\theta\theta;i,j,k} &\approx 2\frac{(\theta_i - \theta_{i-1}) f_{i+1,j,k} - (\theta_{i+1} - \theta_{i-1}) f_{i,j,k} + (\theta_{i+1} - \theta_i) f_{i-1,j,k}}{(\theta_i - \theta_{i-1})(\theta_{i+1} - \theta_i)(\theta_{i+1} + \theta_{i-1})}, \end{aligned}$$

and similarly for the other state variables. As for the mixed partial (recall that γ has zero volatility), this is

$$f_{\theta\lambda;i,j,k} = \frac{f_{i+1,j,k+1} - f_{i+1,j,k-1} - f_{i-1,j,k+1} + f_{i-1,j,k-1}}{4(d\theta)(d\lambda)}.$$

Along the boundaries,

$$f_{\theta\lambda;i,j,k}\left(\underline{\theta}\right) = \frac{f_{i+1,j,k+1} - f_{i+1,j,k} - f_{i-1,j,k+1} + f_{i-1,j,k}}{2\left(d\theta\right)\left(d\lambda\right)},$$

$$f_{\theta\lambda;i,j,k}\left(\overline{\theta}\right) = \frac{f_{i+1,j,k} - f_{i-1,j,k} - f_{i+1,j,k-1} + f_{i-1,j,k-1}}{2\left(d\theta\right)\left(d\lambda\right)},$$

$$f_{\theta\lambda;i,j,k}\left(\overline{\lambda}\right) = \frac{f_{i,j,k+1} - f_{i,j,k-1} - f_{i-1,j,k+1} + f_{i-1,j,k-1}}{2\left(d\theta\right)\left(d\lambda\right)},$$

$$f_{\theta\lambda;i,j,k}\left(\underline{\lambda}\right) = \frac{f_{i+1,j,k+1} - f_{i+1,j,k-1} - f_{i,j,k+1} + f_{i,j,k-1}}{2\left(d\theta\right)\left(d\lambda\right)},$$

while at the corners,

$$f_{\theta\lambda;i,j,k}\left(\underline{\theta},\underline{\lambda}\right) = \frac{f_{i+1,j,k+1} - f_{i+1,j,k} - f_{i,j,k+1} + f_{i,j,k}}{\left(d\theta\right)\left(d\lambda\right)},$$

$$f_{\theta\lambda;i,j,k}\left(\underline{\theta},\overline{\lambda}\right) = \frac{f_{i,j,k+1} - f_{i,j,k} - f_{i-1,j,k+1} + f_{i-1,j,k}}{\left(d\theta\right)\left(d\lambda\right)},$$

$$f_{\theta\lambda;i,j,k}\left(\overline{\theta},\underline{\lambda}\right) = \frac{f_{i+1,j,k} - f_{i+1,j,k-1} - f_{i,j,k} + f_{i,j,k-1}}{\left(d\theta\right)\left(d\lambda\right)},$$

$$f_{\theta\lambda;i,j,k}\left(\overline{\theta},\overline{\lambda}\right) = \frac{f_{i,j,k} - f_{i-1,j,k} - f_{i,j,k-1} + f_{i-1,j,k-1}}{\left(d\theta\right)\left(d\lambda\right)}.$$

We use an "upwind scheme" to determine whether to use the forward or backward approximation: use the forward approximation when the drift of the state variable is positive and a backward approximation when the drift is negative. Then, using an implicit method, the discretized version of the HJB equation (setting $\kappa = 0$ for notational simplicity) is[35]

$$\sum_{j,k}\left[\frac{f^{n+1}_{i,j,k} - f^{n}_{i,j,k}}{\Delta} + r f^{n+1}_{i,j,k}\right] g^{n}_{i,j,k} = \sum_{j,k}\Bigg[y^{n}_{i,j,k} - c^{n}_{i,j,k} + \sum_{i',j',k'} \xi_{i',j',k'} f^{n+1}_{i',j',k'} + \left(\chi^{n}_{i} w^{n\alpha-1}_{i} - \mathcal{X}^{n} w^{n}_{i}\right) e^{\theta_i} \ell^{n}_{i,j,k}\Bigg] g^{n}_{i,j,k} + \eta^{n}_{i} z^{n}_{i}, \tag{E-1}$$

where $\xi_{i',j',k'}$ are the (functional) coefficients on $f^{n+1}_{i',j',k'}$ when the forward and backward approximations are written out, and depend on the drift of the state variables and whether this drift is positive or negative; $(i', j', k') \in \{i-1, i, i+1\} \times \{j-1, j, j+1\} \times \{k-1, k, k+1\}$. Since (E-1) is a system of $J \times K$ linear equations for each i, we can rewrite it in matrix form:

$$\frac{\hat{\mathbf{f}}^{n+1}_{i} - \hat{\mathbf{f}}^{n}_{i}}{\Delta} + r\hat{\mathbf{f}}^{n+1}_{i} = \hat{\mathbf{u}}^{n}_{i} + \mathbf{A}^{n}_{i}\hat{\mathbf{f}}^{n+1}_{i} + \boldsymbol{\eta}^{n}_{i}\mathbf{z}^{n}_{i},$$

where $\mathbf{A}^{n}_{i}$ is a discrete approximation of the infinitesimal generator given θ_i and must satisfy the properties of a Poisson matrix: (1) all rows sum to zero; (2) diagonal elements are

[35] Abusing notation, here Δ refers to the step size, not information rent.

nonpositive; and (3) off-diagonal elements are nonnegative. Also,

$$\hat{\mathbf{f}}_i^{n+1} = \mathbf{f}_i^{n+1} \circ \mathbf{g}_i^n = \left(f_{i,1,1}^{n+1}, \ldots, \ldots, \ldots, f_{i,J,K}^{n+1}\right)^{\mathbf{T}} \circ \left(g_{i,1,1}^n, \ldots, g_{i,J,K}^n\right)^{\mathbf{T}},$$
$$\hat{\mathbf{u}}_i^n = \mathbf{u}_i^n \circ \mathbf{g}_i^n = \left(y_{i,1,1}^n - c_{i,1,1}^n, \ldots, y_{i,J,K}^n - c_{i,J,K}^n\right)^{\mathbf{T}} \circ \left(g_{i,1,1}^n, \ldots, g_{i,J,K}^n\right)^{\mathbf{T}}$$

are the pointwise products of the value and objective functions, respectively, given θ_i, with the density function. This can be rearranged into

$$\left[\left(\frac{1}{\Delta} + r\right)\mathbf{I} - \mathbf{A}_i^n\right]\hat{\mathbf{f}}_i^{n+1} = \hat{\mathbf{u}}_i^n + \frac{\hat{\mathbf{f}}_i^n}{\Delta} \equiv \hat{\mathbf{d}}_i^n + \boldsymbol{\eta}_i^n \mathbf{z}_i^n.$$

Nuño and Moll (2015) show that the KFE can then be written

$$\mathbf{A}_i^{\mathbf{T}} \mathbf{g}_i = \mathbf{h}_i,$$

where $\mathbf{A}_i = \lim_{n\to\infty} \mathbf{A}_i^n$ and $\mathbf{h}_i$ is the vector of zeros with -1 as the first entry; we do this for each i.

The relevant first-order conditions, not discretized, are

$$[c] : \lambda u_c - 1 = 0, \tag{E-2}$$
$$[C] : \lambda v_C^R - 1 = 0, \tag{E-3}$$
$$[\ell] : w(\theta) e^\theta + \lambda \tilde{u}_\ell + \gamma \tilde{u}_{\theta\ell} + \chi(\theta) w(\theta)^\alpha e^\theta - \mathcal{X} w(\theta) e^\theta = 0, \tag{E-4}$$
$$[z] : \int_{\Lambda\times\Gamma} \gamma \tilde{u}_{\theta z} g(\lambda, \gamma, \theta) d(\lambda, \gamma) + \eta(\theta) = 0, \tag{E-5}$$

along with the envelope condition

$$[w] : L(\theta) + \int_{\Lambda\times\Gamma} \gamma \tilde{u}_{\theta w} g(\lambda, \gamma, \theta) d(\lambda, \gamma) + \alpha\chi(\theta) w(\theta)^{\alpha-1} L(\theta) - \mathcal{X} L(\theta) + \eta'(\theta) = 0. \tag{E-6}$$

Combining (E-5) and (E-6), plus the boundary condition $\eta(\underline{\theta}) = 0$,

$$\begin{aligned} 0 = &\int_{\Lambda\times\Gamma} \gamma \tilde{u}_{\theta w} g(\lambda, \gamma, \theta) d(\lambda, \gamma) \\ &+ \int_\Theta \left[\int_{\Lambda\times\Gamma} \gamma \tilde{u}_{\theta w} g(\lambda, \gamma, \theta) d(\lambda, \gamma) + \alpha\chi(\theta) w(\theta)^{\alpha-1} L(\theta) - \mathcal{X} L(\theta)\right] d\theta. \end{aligned} \tag{E-7}$$

Since we work with processes that generate unbounded stationary distributions, we need to truncate the domain of the skill process, $\Theta = [\underline{\theta}, \overline{\theta}]$, to solve the model numerically.

Our solution algorithm to solve the full mean field game is as follows:

1. Guess the wage function and its derivative $w^0(\theta_i)$, $z^0(\theta_i)$ for $i = 1, \ldots, I$.

2. Guess the Lagrange multipliers $\chi^0(\theta_i)$ and $\eta^0(\theta_i)$ for $i = 1, \ldots, I$.

3. Guess the value function $f^0(\theta_i, \lambda_j, \gamma_k)$ for $i = 1, \ldots, I$, $j = 1, \ldots, J$, and $k = 1, \ldots, K$.

4. For each $(\theta_i, \lambda_j, \gamma_k)$, use (E-2)–(E-6) and their approximations above to solve for the policy functions $\{c, C, \ell, \sigma_\Delta\}$; iterate on f^n until convergence.

5. Solve for the stationary distribution via $0 = \mathcal{A}^* g + \kappa \widetilde{f}_0$.

6. Use (E-7) to iterate on χ for each θ_i.

7. Based on the resulting stationary distribution, iterate on the wage function via the market clearing condition.

8. If $\|w^{n+1}(\theta) - w^n(\theta)\| > \epsilon$, return to Step 1; otherwise, stop.

F Extensions

F.1 Skill Bias in Production

In our baseline model, differences in wages are entirely due to differences in labor supply across skill. To put our model more in line with the skill premium literature, we can add skill bias in the production function so that the production process favors more skill workers. Following the working paper version of Heathcote, Storesletten and Violante (2016), let

$$Y_t = \left(\int_\Theta \exp(\varrho \theta_t) L(\theta_t)^{\frac{\alpha-1}{\alpha}} d\theta_t \right)^{\frac{\alpha}{\alpha-1}},$$

which implies that the wage per unit of effective labor is

$$w(\theta_t) = \exp(\varrho \theta_t) \left(\frac{Y_t}{L(\theta_t)} \right)^{\frac{1}{\alpha}}.$$

This implies that the skill premium between two skills is

$$\pi_t^{i,j} = \exp\left(\varrho\left(\theta_t^i - \theta_t^j\right)\right) \left(\frac{L(\theta_t^j)}{L(\theta_t^i)} \right)^{\frac{1}{\alpha}}.$$

In particular, the premium includes an amplifying exogenous component. Also, note that now income grows more rapidly as a function of θ.

To compare the model with and without skill bias, we set α and ϱ so as to keep the variance of the total wage the same in both cases; obviously, there are a continuum of possible combinations. Since we do not know of any joint estimates of these two parameters, for each of the cases in the previous figures, we pick two examples. In Figure 6, we choose (α, ϱ) to be comparable with the $(\alpha, \varrho) = (2.5, 0)$ case from before and in Figure 7, we choose them to be comparable with the $(\alpha, \varrho) = (3.124, 0)$ case. The two example parameter vectors are $(2.8, 0.095)$ and $(3, 0.14)$ in the first case and $(3.5, 0.062)$ and $(4, 0.117)$ in the second case. In both cases, panel (a) reproduces Figure 2 while panel (b) reproduces Figure 3.

Figures 6 and 7 make clear that a strong skill bias calls for a smaller EITC, if at all. Indeed, while optimal policy called for the EITC after a long history of low shocks without skill bias, with skill bias the EITC following a long history is consistent with optimal policy in

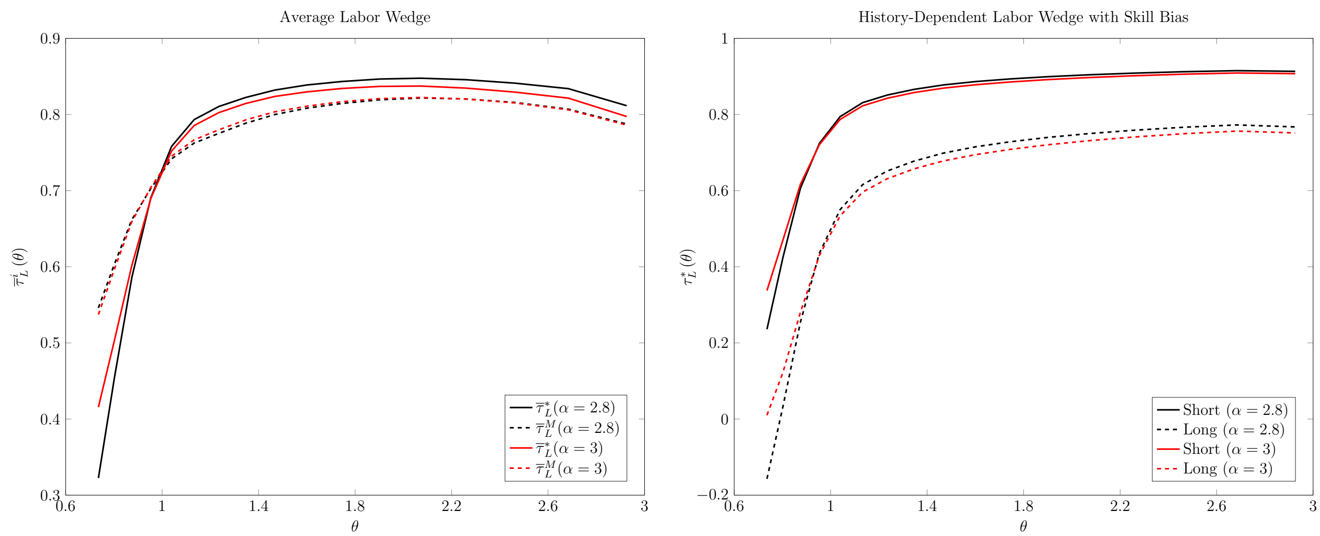

(a) Average labor wedge. (b) Overall rate for low-skill agent.

Figure 6: Labor wedge with skill bias consistent with $(\alpha, \varrho) = (2.5, 0)$.

only one case, $(\alpha, \varrho) = (2.8, 0.095)$. To understand this, recall that an agent's wage is shaped by both the exogenous component and the endogenous labor supply/demand component. When the skill bias is very strong, the wage is shaped more by the exogenous part than the endogenous one. This effectively makes the production function closer to linear and as argued above, weakens the intertemporal wage compression force and thus the EITC. The figures show that the degree of skill bias has the largest effects in the left tail and very small effects elsewhere in the distribution.

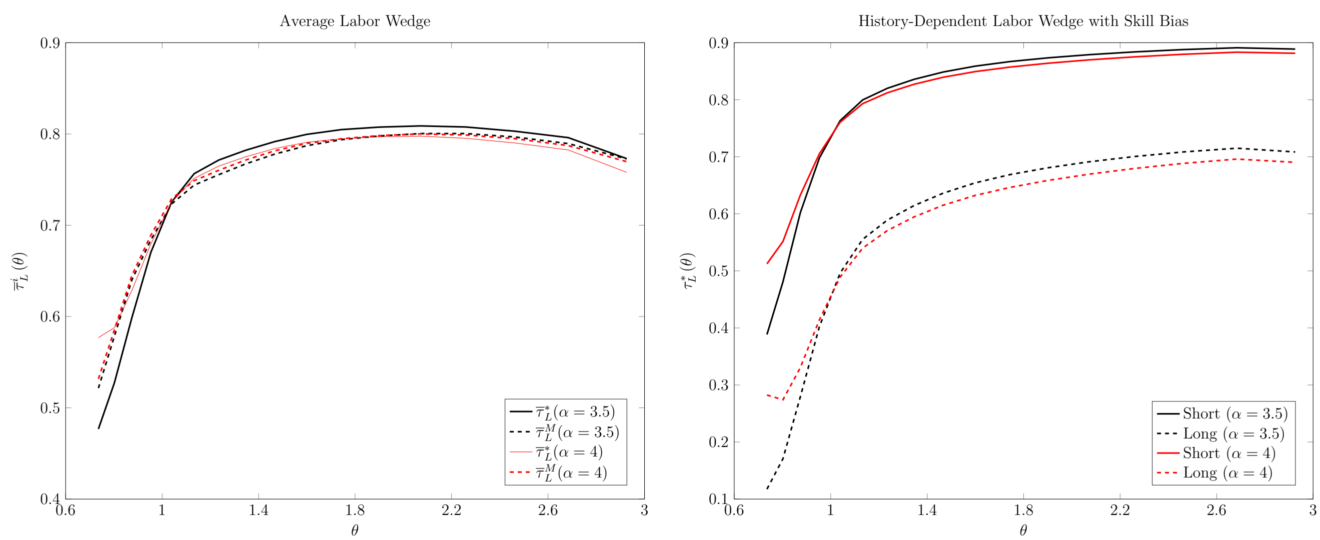

(a) Average labor wedge. (b) Overall rate for low-skill agent.

Figure 7: Labor wedge with skill bias consistent with $(\alpha, \varrho) = (3.124, 0)$.

F.2 Endogenous Human Capital Formation

Recent papers such as Kapička and Neira (2015) and Stantcheva (2015) augment the standard Mirrlees framework to include endogenous human capital formation over the lifecycle. Goldin and Katz (2007) find that rising education wage premia explain 60% to 70% of the increase in U.S. wage inequality between 1980 and 2005 while Heckman, Lochner and Cossa (2003) argue that the EITC is valuable precisely because it boosts human capital investment. By choosing how much to invest in human capital, such as schooling or on-the-job training, an agent can directly affect his earning potential, instead of being totally at the mercy of the exogenous process $\{\theta_t\}$.

This section extends our baseline model à la Stantcheva (2015) to study how our intertemporal wage compression force affects tax policy in the presence of human capital. Let $h_t \in \mathcal{H}$ denote an agent's stock of human capital at date t. Agents can augment this human capital at each date by investing money. Specifically, each agent has access to a technology $H(m_t)$ that allows him to grow his human capital through an observable monetary investment m_t. This technology is best thought of as going to school and exhibits decreasing returns to investment, $H_m > 0$ and $H_{mm} < 0$.[36] Human capital evolves according to

$$dh_t = (H(m_t) - \delta h_t)\, dt \equiv \mu_{h,t}\, dt. \tag{F-1}$$

The δ-term captures the fact that human capital depreciates over time if agents do not invest to grow it. One can imagine that a worker's knowledge and experience in an industry become less useful over time as newer technologies and techniques are developed. Consequently, the worker must continuously invest to maintain the acquired portion of his productivity. Therefore, an agent's type is now $(\theta_t, h_t) \in \Theta \times \mathcal{H}$. When new agents are born, (θ_0, h_0) is drawn from some distribution $\widetilde{f}(\theta_0, h_0)$. To compare with the discrete time model in Stantcheva (2015), to augment human capital s_{t-1} by an amount e_t, $s_t = s_{t-1} + e_t$, the agent must spend $M(e_t)$. Here, s_t corresponds to h_t, m_t corresponds to M_t, $\mu_{h,t}$ corresponds to e_t, and hence $M'(e_t)$ to $m_{\mu_h,t}$.

An agent's productivity depends on exogenous skill and endogenous human capital,

$$e_t = e(\theta_t, h_t),$$

where $e_h > 0$. Stantcheva (2015) takes

$$e_t = \left(\theta_t^{1-\rho} + h_t^{1-\rho}\right)^{\frac{1}{1-\rho}},$$

with $e_t = \theta_t h_t$ as a special case when $\rho = 1$. An agent's labor income at date t is then

$$y_t = w(\theta_t, h_t)\, e(\theta_t, h_t)\, \ell_t,$$

where $w(\theta_t, h_t)$ is still the wage per unit of effective labor, but can now depend on human capital as well. To simplify notation, let $\overline{w}(\theta_t, h_t)$ denote the total wage component of income, i.e., $\overline{w}(\theta_t, h_t) \equiv w(\theta_t, h_t)\, e(\theta_t, h_t)$.

[36] It is straightforward to modify the function so that more skilled agents accumulate human capital more easily, $H_{m\theta} > 0$.

The planner operates the same CES production function as before, the lone difference being that the firm cares about skill and human capital instead of only skill (skill is still unobservable while human capital is observable). Let $L(\theta_t, h_t)$ denote the total labor supply of agents with profile (θ_t, h_t) and L_x the derivative with respect to argument $x \in \{\theta, h\}$. It follows that the wage per unit of effective labor is

$$w(\theta_t, h_t) = \left(\frac{Y_t}{L(\theta_t, h_t)}\right)^{\frac{1}{\alpha}}.$$

As in Stantcheva (2015), an important variable here is the Hicksian coefficient of complementarity between ability and human capital, $\rho_{\theta h,t}$:

$$\rho_{\theta h,t} \equiv \frac{\overline{w}_{\theta h,t}\overline{w}_t}{\overline{w}_{\theta,t}\overline{w}_{h,t}}.$$

There, $\rho_{\theta h} = \rho$ in the CES aggregator above while here that is not the case. For example, if $\rho = 1$ so that $e(\theta, h) = \theta h$, then

$$\rho_{\theta h,t} = 1 - \frac{1}{\alpha}\frac{L_t L_{\theta h,t} - L_{\theta,t}L_{h,t}}{\left(\frac{L_t}{\theta_t} - L_{\theta,t}\right)\left(\frac{L_t}{h_t} - L_{h,t}\right)},$$

which equals 1 as $\alpha \to \infty$ but generally not otherwise. The difference is that here, the derivatives account for both individual and and aggregate wage effects that depend on the distribution and each agent's labor decision, both of which are endogenous. The reason is that while investing in human capital increases an agent's productivity, it also increases the supply of skilled labor, which then feeds back into agents' decisions and hence affects the distribution, too. Thus, complementarity is determined by the direct effect via an agent's productivity and the indirect effect via the wage function. Depending on the magnitudes of these effects, income could increase or decrease. Ability and human capital are complements when $\rho_{\theta h,t} > 1$ and substitutes when $\rho_{\theta h,t} < 1$; see Stantcheva (2015) for a more in-depth discussion of this coefficient.

This highlights the main difference between our model and hers: in her model, although skill has an endogenous component, because the production function is linear, agents' labor supply and human capital investment decisions still do not affect each other. This means that our wage compression channel is absent and hence labor taxes cannot be negative. As we argued with the baseline model, it is precisely the lower insurance cost driven by these labor supply effects that is responsible for our findings.

Since human capital investments are observable, the planner's mechanism design problem is very similar to the one in the baseline model: a reporting strategy is still a report of only θ_t and the laws of motion for the promise- and threat-keeping constraints are the same. There are a few small differences, however: an agent's utility $\tilde{u}_t$ depends on h_t, the planner's HJB equation (and KFE) contains an additional term for the human capital law of motion, an allocation now specifies the incremental human capital change, $\mu_{h,t}$, there is an additional cost from human capital investment, the resource constraints and their Lagrange multipliers depend on h_t, too, and there is a constraint for the h-derivative of the wage function. Thus the planner's problem is

$$
\begin{aligned}
(r+\kappa) K(\theta,h) = \inf_{\{c,C,\ell,\mu_h,\sigma_\Delta,z\}} \Bigg\{ \int_{\mathcal{V}\times\mathcal{D}} \Bigg(& c + \kappa C + m(\mu_h) - \ell \overline{w}(\theta,h) + k_\theta \mu(\theta) + k_h \mu_h + \\
& k_v \left[(\rho+\kappa) v - \tilde{u} - \kappa v^R \right] + k_\Delta \left[(\rho + \kappa - \mu'(\theta)) \Delta - \tilde{u}_\theta - \kappa \Delta^R - \sigma_\theta^2 \sigma_\Delta \right] \\
& + \frac{\sigma_\theta^2}{2} \left[k_{vv} \theta^2 \Delta^2 + k_{\Delta\Delta} \sigma_\Delta^2 + k_{\theta\theta} \theta^2 \right] + \sigma_\theta^2 \left[k_{v\Delta} \theta \Delta \sigma_\Delta + k_{v\theta} \theta^2 \Delta + k_{\Delta\theta} \sigma_\Delta \right] + \kappa k \widetilde{f}_0 \\
& + \chi(\theta,h) w(\theta,h)^\alpha e(\theta,h) \ell - \mathcal{X} \overline{w}(\theta,h) \ell \Bigg) g(\theta,h,v,\Delta)\, d(v,\Delta) \\
& + \eta_\theta(\theta,h) z_\theta(\theta,h) + \eta_h(\theta,h) z_h(\theta,h) \Bigg\},
\end{aligned} \tag{F-2}
$$

where z_s is the derivative of the wage function with respect to the state variable $s \in \{\theta, h\}$ and η_s is its Lagrange multiplier.

Since human capital investment is now a choice variable, it has a first-order condition:

$$
[\mu_h] : m_{\mu_h} = \frac{1}{H^{-1}(\mu_h + \delta h)} = -k_h, \tag{F-3}
$$

where m_{μ_h} is computed using the Inverse Function Theorem. Also, since the HJB equation includes another state (with a first-order derivative), we need one more endogenous boundary condition. Let $\overline{h}$ denote the value of h_t at which $dh_t = 0$. Since agents will never invest in human capital above this point, this is a reflecting boundary, $k_h\left(v, \Delta, \theta, \overline{h}\right) = 0$.

As before, we can characterize marginal distortions using wedges. The labor and intertemporal wedges are defined in the same way as before and have the same interpretations, the lone difference being that h_t is an additional state variable. Let $\tau_H(\theta^t)$ denote the human capital wedge defined by

$$
\tau_H\left(\theta^t\right) = m_{\mu_h,t} - \frac{1}{\rho+\kappa} \left[\ell_t \overline{w}_{h,t} (1 - \tau_{L,t}) + \frac{1}{dt} \frac{\mathbb{E}_t\left[d\left(u_{c,t}\left(m_{\mu_h,t} - \tau_{H,t}\right)\right)\right]}{u_{c,t}} \right]. \tag{F-4}
$$

This is simply the continuous time analogue of $\tau_{S,t}$ in Stantcheva (2015). This wedge represents the implicit subsidy from the planner to an agent for an incremental investment in human capital: an agent receives $\tau_{H,t}$ when human capital increases by $\mu_{h,t}$. This version of the human capital wedge takes as given labor and savings distortions and so acts only to undo these distortions;[37] this is the "Siamese twins" result from Bovenberg and Jacobs (2005).

Because the human capital wedge explicitly and implicitly depends on the other two wedges, we need to adjust it to account for these distortions. We define the net subsidy on human capital expenses, $t_{h,t}$, as

$$
t_{h,t} = \frac{\tau_{H,t}^* - \tau_{L,t} m_{\mu_h,t}^* + p_t}{(1-\tau_{L,t})\left(m_{\mu_h,t}^* - \tau_{H,t}^*\right)} - \frac{\mathcal{X}_t w_{h,t} - \chi(\theta_t, h_t)\left(w_t^\alpha e_t\right)_h}{\overline{w}_{h,t}(1-\tau_{L,t})}, \tag{F-5}
$$

[37] Jacobs (2012) defines a human capital wedge that ignores labor distortions.

where

$$\tau_{H,t}^{*} \equiv \tau_{H,t} - \frac{1}{\rho+\kappa}\frac{1}{dt}\frac{\mathbb{E}_t\left[d\left(u_{c,t}\tau_{H,t}\right)\right]}{u_{c,t}}$$

is the dynamic, risk-adjusted subsidy,

$$m_{\mu_h,t}^{*} \equiv m_{\mu_h,t} - \frac{1}{\rho+\kappa}\frac{1}{dt}\frac{\mathbb{E}_t\left[d\left(u_{c,t}m_{\mu_h,t}\right)\right]}{u_{c,t}}$$

is the dynamic, risk-adjusted cost and

$$p_t \equiv \frac{1-\tau_{L,t}}{\rho+\kappa}\left[m_{\mu_h,t}\tau_{K,t} + \frac{1}{dt}\frac{\mathbb{E}_t\left[du_{c,t}\cdot dm_{\mu_h,t}\right]}{u_{c,t}}\right]$$

is the risk-adjusted savings distortion. This net subsidy ensures that the tax system is neutral with respect to human capital. The second term in (F-5) is new and accounts for the fact that investing in human capital affects other agents' wages and hence their incentive to invest themselves. Given an agent's (distorted) labor and savings decisions, this subsidy ensures that human capital investment is efficient. In effect, it captures the redistributive and insurance properties of the human capital subsidy while filtering out anything that acts only to undo other distortions.

We will now derive an equation that links the labor wedge with the net subsidy. From Itô's Lemma, the Envelope Theorem, and (F-3),

$$\frac{1}{dt}\mathbb{E}_t\left[dm_{\mu_h,t}\right] = \left(r+\kappa\right)m_{\mu_h,t} - \ell_t\overline{w}_{h,t} - \gamma_t\tilde{u}_{\theta h,t} + \left(\chi\left(\theta_t,h_t\right)\left(w_t^{\alpha}e_t\right)_h - \mathcal{X}_t\overline{w}_{h,t}\right)\ell_t.$$

Rearranging and using the definition of $t_{h,t}$ in (F-5), at the optimum the labor and human capital wedges satisfy

$$t_{h,t}^{*} = \left[\frac{\tau_{L,t}^{*}}{1-\tau_{L,t}^{*}} - \mathrm{WC}_t\right]\frac{\varepsilon}{1+\varepsilon}\left(1-\rho_{\theta h,t}\right), \tag{F-6}$$

where WC_t is the wage compression term, now accounting for the effects of human capital. The term in brackets is just the Mirrlees term.

With a linear production function, $t_{h,t}^{*}$ and hence (F-6) reduce to the expression in Stantcheva (2015). The main insight still holds: if $1 < \rho_{\theta h,t}$ then skilled agents do not disproportionately benefit from acquiring human capital so the planner subsidizes its acquisition to make up for higher labor taxes. That said, it may not be the case that $\rho_{\theta h,t} < 1$ for everyone at the same time because $L\left(\theta,h\right)$ varies across types.

Because the wage compression term is negative and hence the Mirrlees term larger than the optimal rate in the two tails (see Figure 2), the subsidy is largest (relative to if the planner ignored the wage compression force) in these regions.[38] However, the reasons for the relatively larger subsidies differ across agents. In the left tail, $\tau_{L,t}^{*}$ is small, if not negative, but $-\mathrm{WC}_t \gg 0$ while in the right tail, $\tau_{L,t}^{*}$ is large and $-\mathrm{WC}_t > 0$. This means that not

[38] The distinction is crucial: the Mirrlees term is lowest for the lowest-skilled agents so the actual subsidy is very small for them. However, absent the wage compression term, the subsidy would be even smaller.

only could low-skill agents receive a transfer via the EITC, they could receive another small subsidy specifically for human capital investment.

Consider an agent with large $e(\theta, h)$. Then the planner should subsidize human capital investment because doing so increases the supply of skilled effective labor, pushing down their wages and lowering skill premia as in static models. For agents at the other extreme, subsidizing human capital investment widens skill premia as before but it also allows the planner to provide dynamic insurance against wage fluctuations. Once again, this insurance is more valuable than offsetting rising skill premia so the optimal policy combines lower labor taxes with human capital subsidies for these agents. However, it is not ex ante clear to what extent the human capital subsidy crowds out the EITC, or vice versa. In other words, policy calls for lower labor taxes but not necessarily negative ones. It is also not clear how this affects the savings tax. While we have focused on the two extremes of the distribution, agents in the middle still receive a human capital subsidy but it is smaller than what they would receive absent the wage compression channel.

As a point of comparison, Jacobs (2012) takes $e(\theta, h) = \theta h$ in a static model. Though his definition of the human capital wedge is different, he finds that high-skill agents should receive positive human capital subsidies (and low-skill agents should pay a tax) for the same reason that they should face negative labor income taxes, namely that both reduce skill premia. Once again, adding dynamics changes the implications because the missing intertemporal wage compression force works in the opposite direction.

References

Acemoğlu, and David Autor. 2011. “Skills, Tasks, and Technologies: Implications for Employment and Earnings.” In *Handbook of Labor Economics.* , ed. David Card and Orley Ashenfelter, 1043–1171. Elsevier.

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2015. “Heterogeneous Agent Models in Continuous Time.” Princeton University Working Paper.

Albanesi, Stefania, and Christopher Sleet. 2006. “Dynamic Optimal Taxation with Private Information.” *Review of Economic Studies*, 73(1): 1–30.

Ales, Laurence, and Christopher Sleet. 2016. “Taxing Top CEO Incomes.” Tepper School of Business Working Paper.

Ales, Laurence, and Pricila Maziero. 2009. “Accounting for Private Information.” Tepper School of Business Working Paper.

Ales, Laurence, Musab Kurnaz, and Christopher Sleet. 2015. “Technical Change, Wage Inequality, and Taxes.” *American Economic Review*, 105(10): 3061–3101.

Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez. 2011. “Top Incomes in the Long Run of History.” *Journal of Economic Literature*, 49(1): 3–71.

Autor, David H. 2014. “Skills, Education, and the Rise of Earnings Inequality Among the ‘Other 99 Percent’.” *Science*, 344(6186): 843–51.

Beaudry, Paul, Charles Blackorby, and Dezsö Szalay. 2009. “Taxes and Employment Subsidies in Optimal Redistribution Programs.” *American Economic Review*, 99(1): 216–42.

Borovička, Jaroslav, Lars Peter Hansen, and José A. Scheinkman. 2014. “Shock Elasticities and Impulse Responses.” *Mathematics and Financial Economics*, 8(4): 333–54.

Bovenberg, A. Lans, and Bas Jacobs. 2005. “Redistribution and Education Subsidies are Siamese Twins.” *Journal of Public Economics*, 89(11-12): 2005–35.

Chetty, Raj. 2012. “Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply.” *Econometrica*, 80(3): 969–1018.

Choné, Philippe, and Guy Laroque. 2005. “Optimal Incentives for Labor Force Participation.” *Journal of Public Economics*, 89(2-3): 395–425.

Choné, Philippe, and Guy Laroque. 2010. “Negative Marginal Tax Rates and Heterogeneity.” *American Economic Review*, 100(5): 2532–47.

Colombi, Roberto. 1990. "A New Model of Income Distribution: The Pareto-Lognormal Distribution." In *Income and Wealth Distribution, Inequality, and Poverty. Studies in Contemporary Economics.* , ed. Camilo Dagum and Michele Zenga, 18–32. Berlin Heidelberg:Springer.

Diamond, Peter A. 1980. "Income Taxation with Fixed Hours of Work." *Journal of Public Economics*, 13(1): 101–10.

Diamond, Peter A. 1998. "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates." *American Economic Review*, 88(1): 83–95.

Diamond, Peter, and Emmanuel Saez. 2011. "The Case for a Progressive Tax: From Basic Research to Policy Recommendations." *Journal of Economic Perspectives*, 25(4): 165–90.

Di Nunno, Giulia, Bernt Øksendal, and Frank Proske. 2009. *Malliavin Calculus for Lévy Processes with Applications to Finance.* Springer.

Eissa, Nada, and Jeffrey B. Liebman. 1996. "Labor Supply Responses to the Earned Income Tax Credit." *Quarterly Journal of Economics*, 111(2): 605–37.

Farhi, Emmanuel, and Iván Werning. 2013. "Insurance and Taxation Over the Life Cycle." *Review of Economic Studies*, 80(2): 596–635.

Fernandes, Ana, and Christopher Phelan. 2000. "A Recursive Formulation for Repeated Agency with History Dependence." *Journal of Economic Theory*, 91(2): 223–47.

Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2015. "The Dynamics of Inequality." NBER Working Paper No. 21363.

Goldin, Claudia, and Lawrence F. Katz. 2007. "Long-Run Changes in the Wage Structure: Narrowing, Widening, Polarizing." *Brookings Papers on Economic Activity*, 38(2): 135–68.

Golosov, Mikhail, Aleh Tsyvinski, and Iván Werning. 2006. "New Dynamic Public Finance: A User's Guide." In *NBER Macroeconomics Annual 2006.* , ed. Daron Acemoğlu, Kenneth Rogoff and Michael Woodford, 317–88. MIT Press.

Golosov, Mikhail, and Aleh Tsyvinski. 2006. "Designing Optimal Disability Insurance: A Case for Asset Testing." *Journal of Political Economy*, 114(2): 257–79.

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski. 2016*a*. "Consolidated Income Accounts." Cornell University Working Paper.

Golosov, Mikhail, Maxim Troshkin, and Aleh Tsyvinski. 2016*b*. "Redistribution and Social Insurance." *American Economic Review*, 106(2): 359–86.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2015. "What Do Data on Millions of U.S. Workers Reveal About Life Cycle Earnings Risk?" University of Minnesota Working Paper.

Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy*, 122(3): 621–60.

Heathcote, Jonathan, and Hitoshi Tsujiyama. 2016. "Optimal Income Taxation: Mirrlees Meets Ramsey." Staff Report 507, Federal Reserve Bank of Minneapolis.

Heathcote, Jonathan, Kjetil Storesletten, and Gianluca L. Violante. 2005. "Two Views of Inequality Over the Life Cycle." *Journal of the European Economic Association*, 3(2–3): 765–75.

Heathcote, Jonathan, Kjetil Storesletten, and Gianluca L. Violante. 2016. "Optimal Tax Progressivity: An Analytical Framework." NBER Working Paper No. 19899.

Heckman, James L., Lance Lochner, and Ricardo Cossa. 2003. "Learning-by-Doing Versus On-the-Job Training: Using Variation Induced by the EITC to Distinguish Between Models of Skill Formation." In *Designing Inclusion: Tools to Raise Low-End Pay and Employment in Private Enterprise.* , ed. Edmund S. Phelps, 74–130. Cambridge University Press.

Jacobs, Bas. 2012. "Optimal Redistributive Tax and Education Policies in General Equilibrium." *International Tax and Public Finance*, 20(2): 312–37.

Kapička, Marek. 2013. "Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First-Order Approach." *Review of Economic Studies*, 80(3): 1027–54.

Kapička, Marek, and Julian Neira. 2015. "Optimal Taxation with Risky Human Capital." U.C. Santa Barbara Working Paper.

Kaplan, Greg, Benjamin Moll, and Gianluca L. Violante. 2016. "Monetary Policy According to HANK." NBER Working Paper No. 21897.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics*, 107(1): 35–78.

Kocherlakota, Narayana R. 2005. "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation." *Econometrica*, 73(5): 1587–1621.

Laroque, Guy. 2005. "Income Maintenance and Labor Force Participation." *Econometrica*, 73(2): 341–76.

Lasry, Jean-Michel, and Pierre-Louis Lions. 2007. "Mean field games." *Japanese Journal of Mathematics*, 2(1): 229–60.

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

Nuño, Galo, and Benjamin Moll. 2015. "Controlling a Distribution of Heterogeneous Agents." Princeton University Working Paper.

Pavan, Alessandro, Ilya Segal, and Juuso Toikka. 2014. "Dynamic Mechanism Design: A Myersonian Approach." *Econometrica*, 82(2): 601–53.

Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913–1998." *Quarterly Journal of Economics*, 118(1): 1–39.

Reed, William J., and Murray Jorgensen. 2004. "The Double Pareto-Lognormal Distribution - A New Parametric Model for Size Distributions." *Communications in Statistics - Theory and Methods*, 33(8): 1733–53.

Revuz, Daniel, and Marc Yor. 1999. *Continuous Martingales and Brownian Motion.* Springer.

Rothschild, Casey, and Florian Scheuer. 2013. "Redistributive Taxation in the Roy Model." *Quarterly Journal of Economics*, 128(2): 623–68.

Rothschild, Casey, and Florian Scheuer. 2014. "A Theory of Income Taxation Under Multidimensional Skill Heterogeneity." NBER Working Paper No. 19822.

Rothstein, Jesse. 2010. "Is the EITC as Good as an NIT? Conditional Cash Transfers and Tax Incidence." *American Economic Journal: Economic Policy*, 2(1): 177–208.

Saez, Emmanuel. 2001. "Using Elasticities to Derive Optimal Income Tax Rates." *Review of Economic Studies*, 68(1): 205–29.

Saez, Emmanuel. 2002. "Optimal Income Transfer Programs: Intensive Versus Extensive Labor Supply Responses." *Quarterly Journal of Economics*, 117(3): 1039–73.

Salanié, Bernard. 2011. *The Economics of Taxation.* MIT Press.

Sannikov, Yuliy. 2008. "A Continuous-Time Version of the Principal-Agent Problem." *Review of Economic Studies*, 75(3): 957–84.

Sannikov, Yuliy. 2014. "Moral Hazard and Long-Run Incentives." Princeton University Working Paper.

Scheuer, Florian, and Iván Werning. 2016. "The Taxation of Superstars." *Quarterly Journal of Economics*, Forthcoming.

Shourideh, Ali, and Maxim Troshkin. 2015. "Incentives and Efficiency of Pension Systems." University of Pennsylvania Working Paper.

Stantcheva, Stefanie. 2015. "Optimal Taxation and Human Capital Policies Over the Life-Cycle." *Journal of Political Economy*, Forthcoming.

Stiglitz, Joseph E. 1982. "Self-Selection and Pareto Efficient Taxation." *Journal of Public Economics*, 17(2): 213–40.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. "Consumption and Risk Sharing Over the Life Cycle." *Journal of Monetary Economics*, 51(3): 609–33.

Williams, Noah. 2011. "Persistent Private Information." *Econometrica*, 79(4): 1233–75.