Optimal Fiscal Policy under Preference Heterogeneity

Marek Kapička
CERGE-EI *

November 11, 2024

Abstract

This paper studies optimal fiscal policy in a dynamic economy with aggregate shocks, where agents are heterogeneous in their incomes, preferences toward risk, and the intertemporal elasticity of substitution. I show that i) if low-income people have lower intertemporal elasticity of substitution or higher risk aversion, the optimal marginal tax rates are procyclical, and transfers are countercyclical; ii) optimal marginal tax rates increase by about 2.2 percentage points for each percentage increase in productivity, and iii) US marginal tax rate fluctuations were close to optimal before 1980, but not much after that because they turned countercyclical.

The results are due to preference heterogeneity, which leads to procyclical fluctuations in the implicit price of the informational rent of the agents across time and states and, in turn, to fluctuations in marginal taxes and transfers.

1 Introduction

Substantial empirical evidence documents preference heterogeneity across people. The literature reviewed below robustly shows that richer people tend to have a higher intertemporal elasticity of substitution and a lower coefficient of risk aversion than poorer

*Email: marek.kapicka@cerge-ei.cz. I would like to thank Laurence Ales, Árpád Ábrahám, Mike Golosov, and Ctirad Slavík, participants of the CERGE-EI reading group as well as participants in many conferences, for their comments. All errors are mine.

people. This paper shows that preference heterogeneity plays a key role in shaping the optimal design of marginal and average taxes in response to aggregate shocks. I consider the optimal policy design in a dynamic Mirrlees economy with aggregate shocks, where agents are heterogeneous in their productivity and preferences, and their type is private information.

I show that heterogeneity in preferences substantially alters the optimal policy prescriptions relative to an economy where all agents have identical preferences. Suppose that, consistently with the empirical evidence, low-productivity agents have lower intertemporal elasticity of substitution, higher risk aversion, or both. In that case, decreasing marginal taxes in a recession and increasing them during an expansion is optimal. I relate fluctuations in the marginal taxes to a novel channel, the fluctuations of the shadow price of the informational rent. I show they are positively related: when the shadow price of the informational rent is high, the social planner reduces the informational rent more aggressively because distortions of low types become relatively less important than the redistribution concerns. To put it differently, it is optimal to increase the tax burden on high-productivity agents by increasing marginal taxes on the low-productivity agents.

The shadow price of the informational rent measures the degree to which different types of agents value the instantaneous utility differently and differs from the informational rent itself. The distinction gains importance with heterogeneous preferences: if both types have identical preferences, then the shadow price of the informational rent is constant over time and states. Under relatively standard assumptions about the instantaneous utility function, I show an even stronger result: if both types have identical preferences, then the informational rent, and so the marginal tax rates, are constant over time and states as well.

To isolate the role of risk aversion and intertemporal elasticity of substitution, I present three stylized examples. In all examples, high-productivity agents are standard expected utility maximizers, while low-productivity agents differ in each of the three examples. They have infinite risk aversion in the first example, zero elasticity of substitution in the second one, and both in the third one. In all cases, marginal taxes decrease following an adverse productivity shock, but the implications for the future dynamics of marginal taxes differ markedly. Infinite risk aversion produces positive persistency in the marginal tax rates. Decreases in marginal tax rates will thus be followed by low marginal tax rates in the future. On the other hand, zero intertemporal elasticity of substitution

prescribes the opposite pattern: low marginal tax rates are followed by higher marginal tax rates in the future; that is, they overshoot their long-run level. Transfers to low-productivity individuals are generally indeterminate because of Ricardian equivalence. Under the assumption that low-productivity individuals are hand-to-mouth, transfers follow a similar pattern with the opposite sign. When both infinite risk aversion and zero intertemporal elasticity of substitution are present, the effects cancel each other out, and the optimal pattern is remarkably simple in that marginal taxes and transfers are only a function of the current state, bearing similarity to the complete markets solution to the Ramsey problem in a representative agent economy (Lucas and Stokey (1983)).

I compute a continuous-time version of the model to assess the United States post-war tax experience. The model is formulated using the stochastic variational utility approach of Geoffard (1996) and Dumas et al. (2000). It assumes that labor productivity, which follows an Ornstein-Uhlenbeck process, is the only source of aggregate fluctuations. I calibrate the productivity process to match the standard deviation and autocorrelation of US labor productivity for the 1947-2012 time period. The low-productivity agents represent the bottom 80 percent of the population, and I calibrate relative productivity to reflect the relative earnings of both types. To parameterize the preference parameters, I exogenously set the elasticity of intertemporal substitution and risk aversion coefficient for the low-productivity agents to 0.05 and 10 and calibrate them for the high-productivity agents to match the annual risk-free rate and the ratio of the standard deviation of log consumption for both types. The procedure yields the elasticity of intertemporal substitution of 0.214 and the coefficient of risk aversion of 2.773 for the high-productivity agents.

I show that, in the optimum, the marginal tax rates respond strongly to the underlying labor productivity shock: a 1 percent increase in labor productivity leads to an approximately 2.2 percentage point increase in the marginal tax rates. The calibrated model is qualitatively similar to the second example where low-productivity agents have a lower elasticity of substitution, and the marginal tax rates overshoot their long-run level. Quantitatively, however, the subsequent dynamics is relatively weak, similarly to the third example where low-productivity agents have both zero intertemporal elasticity of substitution and infinite risk aversion. Finally, transfers to the low-productivity agents are countercyclical but fluctuate relatively little despite being well-insured against aggregate fluctuations. The reason is that fluctuations in the marginal tax rates already help provide insurance to the low-productivity agents.

I compare the optimal marginal tax rates on the bottom 90 percent of the population to the average marginal tax rates on the corresponding population percentile, as documented by Mertens and Olea (2018). I show that changes in the average US marginal tax rates have been approximately optimal before 1980: both standard deviation and autocorrelation are about right, and they were strongly procyclical. On the other hand, after 1980, US marginal tax rates diverged from the optimal ones, namely because they turned countercyclical. I also show that the variations in the shadow price of the informational rent are quantitatively substantially more important than variations in the informational rent. In their absence, the optimal marginal tax rates would be approximately flat.

1.1 Related Literature

There is substantial empirical evidence that documents heterogeneity in the intertemporal elasticity of substitution across people (see Guvenen (2006) for a review). The literature typically investigates the difference between shareholders and nonshareholders (Vissing-Jørgensen (2002) for US households using the Consumption Expenditure Survey data and Attanasio et al. (2002) for UK households using the Family Expenditure Survey data) and finds that shareholders have a higher intertemporal elasticity of substitution. Their estimates imply a low intertemporal elasticity of substitution for nonstockholders, which roughly corresponds to the low productivity agents in the model. Blundell et al. (1994) directly examine differences across people of different income levels and find that higher-income people exhibit higher intertemporal elasticity of substitution. Furthermore, Attanasio and Browning (1995) find, using the Family Expenditure Survey data, that the intertemporal elasticity of substitution increases with consumption, while Atkeson and Ogaki (1996) and Ogaki and Atkeson (1997) find, using Indian panel data, that the intertemporal elasticity of substitution is larger for the rich than it is for the poor. There is also empirical evidence that wealthier people exhibit lower risk aversion, as shown by Mankiw and Zeldes (1991), and Brav et al. (2002). Both types of findings are consistent with the assumptions in this paper.

The assumption that there are only two types of preference types in the economy is standard in finance and macro-finance. Guvenen (2009) and Gomes and Michaelides (2007) use a similar setup in discrete time to address asset pricing puzzles with limited risk-sharing opportunities, while Gârleanu and Panageas (2015)) use it to address similar issues in a complete market overlapping-generations framework. Anderson (2005) and

Colacito et al. (2019) study a discrete-time version of the model to characterize the conditions under which long-run stationary distribution exists. Backus et al. (2008) study equilibrium asset trades in economies with two types exhibiting extreme preferences similar to the examples in this paper. Borovička (2019) studies the long-run survival and asset pricing in a two-type economy where both types have common recursive preferences but differ in their beliefs about the distribution of aggregate endowment. My paper is technically similar to Gârleanu and Panageas (2015)) and Borovička (2019) in that it studies the problem in continuous time but adds endogenous labor supply and looks at second-best rather than first-best allocations.

This paper uses the Mirrlees approach to the optimal fiscal policy but focuses on issues traditionally studied in the dynamic Ramsey tradition, following the seminal approach of Lucas and Stokey (1983) and Chari et al. (1994). A representative agent Ramsey framework with no scope for redistribution or insurance across different types reduces the scope of the policy prescriptions - one cannot talk about redistribution across people, for example - and can potentially give misleading answers about other dimensions of the policy design. Werning (2007) studies similar issues in a Mirrlees framework but without preference heterogeneity, thus shutting down the key mechanisms in this paper. This paper maintains the assumption of complete markets and differs thus from a more recent Ramsey literature under incomplete markets (Aiyagari et al. (2002), Bhandari et al. (2016)). Such considerations are beyond the scope of this paper.

Almost all the Ramsey literature assumes that the agents have standard time-additive preferences. Karantounias (2018) is an important exception. He considers a dynamic representative agent Ramsey economy with recursive preferences of the Duffie-Epstein-Zin variety (Epstein and Zin (1989), Duffie and Epstein (1992b)), also considered in this paper. His results are, however, conceptually very different from the results in this paper. The key mechanism in a representative agent Ramsey economy is that marginal taxes are tightly connected to the market price of risk and debt. However, the link between the price of debt and marginal tax rates is broken in the Mirrlees framework because marginal taxes no longer have a one-on-one relationship with the average tax rate and level of consumption. One manifestation of this property has already been mentioned: with identical preferences and a relatively standard instantaneous utility function, the marginal tax rates are constant across time and states, in contrast to Karantounias (2018).

2 Three Examples

I will start by characterizing the optimal tax policy in three examples to illustrate how different dimensions of preference heterogeneity manifest in optimal tax and transfer policies. In the first example, low-productivity agents will exhibit infinite risk aversion; in the second, they will exhibit zero intertemporal elasticity of substitution; in the third, they will exhibit both. High-productivity agents will always have standard time-additive expected utility. The examples are presented in discrete time for an easier understanding, but that modeling choice is not essential. A formal treatment of the three examples and proofs of the results of this section are relegated to Appendix A.

There is uncertainty regarding aggregate productivity and government spending. The aggregate shock in period t is denoted by $s_t \in S$, where S is a finite set. The initial shock s_0 is known. The probability of a history of shocks $s^t = (s_0, s_1, \ldots, s_t)$ is $\pi_t(s^t)$. Then aggregate productivity in period t is $Z_t(s^t)$, and government consumption is $G_t(s^t)$.

There are two types of agents, A and B, with population shares α and $1-\alpha$. The types differ in their productivity and their preferences. The productivity of type A agents is equal to the aggregate productivity $Z(s^t)$. The productivity of type B agents is $Z(s^t)\theta$, where $\theta > 1$ is their productivity relative to type A agents. Type B agents are thus always more productive than type A agents.

The allocation consists of consumption and hours worked for each type, $\{c_t^i(s^t), n_t^i(s^t)\}$ for $i = A, B$. The allocation must be feasible,

$$\alpha c_t^A(s^t) + (1-\alpha)c_t^B(s^t) + G_t(s^t) \leq Z_t(s^t)\left[\alpha n_t^A(s^t) + (1-\alpha)\theta n_t^B(s^t)\right] \quad \forall s^t \in S^{t+1}. \quad (1)$$

The preferences of type A agents will differ in the three examples, and the formal representation of type A's lifetime utility $V_0^A(s_0)$ is presented in the respective sections below. Preferences of type B agents are, on the other hand, standard and common in all three examples:

$$V_0^B(s_0) = (1-\beta)\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\pi_t(s^t)U\left(c_t^B(s^t), n_t^B(s^t)\right).$$

The preferences of type B agents are a special case of the Epstein-Zin preferences with zero aversion toward utility risk and infinite elasticity of intertemporal substitution

in utilities. The period utility is

$$U(c,n) = \ln c - \frac{n^{1+\eta}}{1+\eta}, \quad \eta > 0.$$

I follow a standard Mirrlees information structure where the agent's type and hours worked are the agents' private information, while consumption and income are publicly observable. The incentive constraint requires that both types prefer their allocation to the allocation of the other type. Given that type B agents have higher relative productivity, I will restrict attention to situations where only the incentive constraint on type B agents binds:

$$(1-\beta)\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\pi(s^t)U\left(c_t^B(s^t), n_t^B(s^t)\right) \geq (1-\beta)\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\pi(s^t)U\left(c_t^A(s^t), \frac{n_t^A(s^t)}{\theta}\right). \quad (2)$$

The social planner assigns a relative Pareto weight $\lambda \geq 0$ to the type A agent, sufficiently large so that the planner wants to redistribute from type B agents toward type A agents. The planning problem is

$$\max_{c^A, n^A, c^B, n^B} \lambda V_0^A(s_0) + (1-\lambda)V_0^B(s_0) \quad \text{s.t.} \quad (1) \text{ and } (2). \quad (3)$$

Marginal tax rates. The optimum always features zero distortions of the type B agents for well-known reasons. We characterize the solution in terms of the marginal tax rates faced by type A agents in each period and state,

$$\tau_t^A(s^t) = 1 + \frac{U_{nt}^A(s^t)}{U_{ct}^A(s^t)}\frac{1}{Z_t(s^t)}.$$

It can be shown that in all three examples, the marginal tax rate satisfies a common formula

$$\tau_t^A(s^t) = \frac{\kappa(1-\phi)}{\lambda_t(s^t) - \kappa\phi}, \quad \phi = \theta^{-1-\eta}, \quad (4)$$

where κ is the Lagrange multiplier on the incentive constraint (2) and $\lambda_t(s^t)$ represents the relative Pareto weight of type A agents after a given history s^t. The fact that first-best allocations with recursive preferences can be represented by time-varying Pareto

weights is well-known from the existing literature (Anderson (2005)). Here, the same notion is extended to the second-best economy. If type A agents had standard time-additive preferences as well, $\lambda_t(s^t)$ would be a constant, and the marginal tax rates would likewise be constant across time and states. Any variations below are then due to preference heterogeneity. The three examples differ in how the relative Pareto weight $\lambda_t(s^t)$ responds to the aggregate shocks.

It is clear from (4) that the movement of the marginal tax rates is inversely related to the movement of the Pareto weight $\lambda_t(s^t)$. The intuition is the following. The ratio $p = \kappa/\lambda_t(s^t)$ is the shadow price of the informational rent, which in this case is a constant ϕ. A higher Pareto weight decreases the shadow price p since the incentive constraint becomes less relevant. Paying out the informational rent is less costly, and the need to distort type A's allocation decreases.

Transfers. By putting no restrictions on lump-sum transfers, the environment features Ricardian equivalence and the indeterminacy of the optimal debt management plans. This is in stark contrast to the representative agent Ramsey models. The indeterminacy is also stronger than in the heterogeneous agent Ramsey models (Werning (2007)) because assuming that low-income people are *hand-to-mouth* and are fully excluded from the asset markets is insufficient to determine the optimal government debt. With non-linear income tax, even this assumption does not suffice to uniquely determine the asset holdings of high-income people and the government debt. Therefore, I focus only on the transfers to type A agents under the hand-to-mouth assumption.

2.1 Example 1: Differences in Risk Aversion

Low productivity type A agents have infinite intertemporal elasticity of substitution in utilities but, at the same time, infinite aversion toward utility risk. Their preferences are

$$V_t^A(s^t) = (1-\beta)U\left(c_t^A(s^t), n_t^A(s^t)\right) + \beta \min\left\{V_{t+1}^A(s^{t+1})\right\}.$$

The optimum requires that $V_{t+1}^A(s^t, s_{t+1})$ is independent of the next period shock s_{t+1}, and so the lifetime utility must satisfy

$$V_t^A(s^{t-1}) = (1-\beta)U\left(c_t^A(s^t), n_t^A(s^t)\right) + \beta V_{t+1}^A(s^t) \quad \forall s^t \in S^{t+1}. \tag{5}$$

Equation (5) implies that current and continuation utility changes are inversely related since they must sum to a value independent of the current shock. A reasonable conjecture is that an adverse aggregate shock (either low productivity or high government spending) will decrease the current utility of the A agent, but his continuation utility will increase. Working backwards, equation (5) defines $V_0^A(s_0)$ in the planning problem (3).

The stochastic process for λ_t satisfies two properties. First, its initial value equals the Pareto weight of the type A agents λ. Second, $\lambda_t(s^t)$ follows a random walk:

$$\lambda_t(s^t) = \sum_{s_{t+1}} \pi(s_{t+1}|s^t)\lambda_{t+1}(s^{t+1}). \tag{6}$$

The social planner keeps the expected value of λ_t constant over time, in expectation, because both types of agents have identical intertemporal elasticity of substitution. However, it does not keep the value of λ constant across states. This is because type A agents are infinitely risk averse, and the social planner maintains an identical lifetime utility across states. The random walk property implies that changes in the Pareto weights are permanent. As we shall see, this is the opposite pattern of what's optimal if type A agents have zero IES instead of infinite risk aversion.

Given that the Pareto weights follow a random walk, it follows immediately from Jensen's inequality that taxes on type A agents have an upward drift:

$$\tau_t(s^t) < \sum_{s_{t+1}} \pi_t(s_{t+1}|s^t)\tau_{t+1}(s^{t+1}) \quad \forall s^t.$$

While marginal taxes themselves have an upward drift, it is easy to see that the inverse of marginal taxes follows a random walk. The random walk component has an interesting analog in the dynamic Ramsey literature under incomplete markets. In that case, the random walk comes from the Ramsey planner's inability to issue state-contingent claims and equalize the marginal cost of collecting taxes across states. In contrast, the model in this section features complete markets, but type A agents are unwilling to transfer utility across states.

2.2 Example 2: Differences in IES

Low productivity type A agents have zero intertemporal elasticity of substitution in utilities and zero aversion toward utility risk. Their preferences are

$$V_t^A(s^t) = \min\left\{U\left(c_t^A(s^t), n_t^A(s^t)\right), \sum_{s_{t+1}} \pi(s_{t+1}|s^t) V_{t+1}^A(s^{t+1})\right\}.$$

The optimum requires that the two terms on the right-hand side are equalized. This yields $V_t^A(s^t) = U\left(c_t^A(s^t), n_t^A(s^t)\right)$, and in turn,

$$U\left(c_t^A(s^t), n_t^A(s^t)\right) = \sum_{s_{t+1}} \pi(s_{t+1}|s^t) U\left(c_{t+1}^A(s^{t+1}), n_{t+1}^A(s^{t+1})\right) \quad \forall s^t. \tag{7}$$

It is now the period utility of the type A agents that follows a random walk. Unlike in the previous example with the infinite risk aversion, a decrease in the current utility is not expected to be compensated by an increase in the continuation utility; quite to the contrary, the expected continuation utility decreases just as much.

The relative Pareto weights $\lambda_t(s^t)$ no longer start at λ or follow a random walk. Instead, the implicit Pareto weights exhibit history dependence in the following sense. For any infinite history s^∞, the Pareto weights must sum to the unconditional Pareto weight λ:

$$(1-\beta)\sum_{t=0}^{\infty} \beta^t \lambda_t(s^t) = \lambda \quad \forall s^\infty \in S^\infty. \tag{8}$$

Any change in the implicit Pareto weight today must be met with future changes in the opposite direction. Moreover, given that the future values are discounted, the off-setting change, or the sum of off-setting changes, must be larger in absolute value. To understand, consider an unexpected adverse shock today. It will increase the Pareto weight of agent A because it is easier to decrease agent B's utility: unlike agent A's utility, the decrease can be compensated in the future. But, to compensate agent B, the Pareto weight of agent A must decrease in the future.

The optimal marginal tax rates continue being determined by (4) and so inherit the off-setting property of the implicit Pareto weights. An adverse shock is expected to decrease the marginal taxes on type A today but increase them in the future.

2.3 Example 3: Differences in IES and Risk Aversion

I now modify the preferences of type A agents so that, in addition to infinite aversion toward utility risk, they have zero intertemporal elasticity of substitution in utilities. Their preferences are

$$V_t^A(s^t) = \min\left\{U\left(c_t^A(s^t), n_t^A(s^t)\right), \min\left\{V_{t+1}^A(s^{t+1})\right\}\right\},$$

which is equivalent to requiring that type A agents have a constant period utility not only across states but also across time:

$$U\left(c_t^A(s^t), n_t^A(s^t)\right) = V^A \quad \forall t, s^t \in S^{t+1}$$

for some lifetime utility value V^A. Unlike the two previous examples, the dynamics in the lifetime utility of type A is fully eliminated. This makes the solution similar to the solution to a Ramsey problem under complete markets (Lucas and Stokey (1983)): conditional on the Lagrange multiplier on the incentive constraint and the utility of type A agent, the optimal allocations and taxes are only a function of the current aggregate shocks G and Z, and thus inherit their stochastic properties.

Notably, the stochastic nature of the optimal allocation is inherited by the implicit Pareto weights of type A, which can be written as $\lambda_t(s^t) = \tilde{\lambda}(Z_t(s^t), G_t(s^t))$, for some function $\tilde{\lambda}$, and by the optimal marginal tax rates $\tau_t^A(s^t) = \tilde{\tau}^A\left(Z_t(s^t), G_t(s^t)\right)$ for some function $\tilde{\tau}^A$.

2.4 A numerical example

I compute the optimum in a simple T period economy with aggregate productivity risk. I set $T = 5$. The aggregate productivity is $Z(s_0) = 1$ in the first period, while $Z_t(s^t)$ is an iid random variable taking three possible values: 0.9 (recession), 1.0, and 1.1 (expansion) with equal probability. Both agents have an equal share in the population ($\lambda = 0.5$), type B agents are 40 percent more productive ($\theta = 1.4$), and the Frisch elasticity of labor is 0.5 ($\eta = 2$). The Pareto weight α is set so that the initial period marginal tax rate is 40 percent.

Figure 1 compares the optimal marginal tax rates and transfers in an economy with standard time-additive preferences and in the three extreme economies above. In the

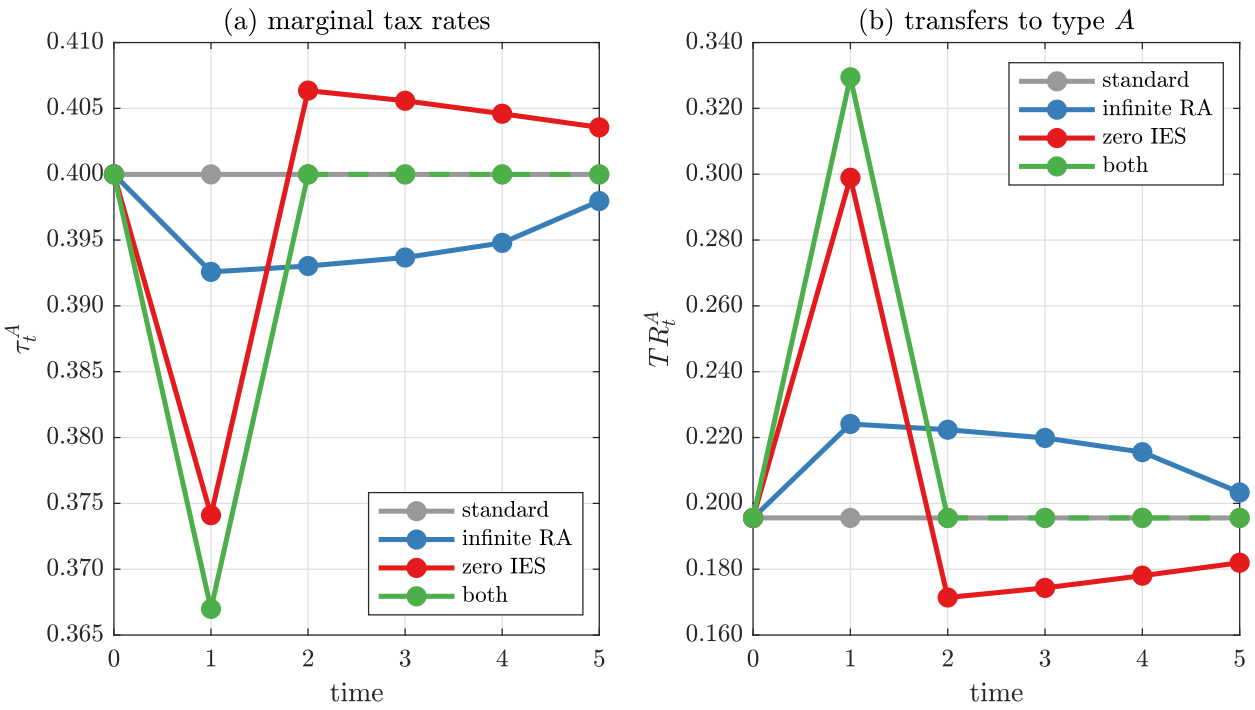


Figure 1: Panel (a): marginal tax rates on type A agents. Panel (b): Transfers to type A agents as a fraction of income. Infinite RA, zero IES, or a combination of both refers to preferences of type A agents.

depicted scenario, there is a temporary low productivity shock in period 1, followed by a return to the average productivity from period two. As the grey line shows, the optimal marginal tax rates and transfers are constant across time and states under standard preferences.

The left panel shows that, in all three examples, marginal tax rates on type A agents decrease on impact as the Pareto weight of type A agent increases. They decrease more under zero IES than under infinite risk aversion, and most under both zero IES and infinite risk aversion, where both forces are combined. However, the three extreme economies differ in their future responses to the current shock. Persistency in the Pareto weights under infinite risk aversion leads to a persistent decline in the marginal tax rates that only slowly increase. With zero IES, future Pareto weights of type A must decrease, leading to lower marginal tax rates. With infinite risk aversion and zero IES, both forces offset each other; marginal taxes are a function of only the current state, and they revert to their initial period level immediately after the shock disappears.

The right panel of Figure 1 shows the optimal transfers to type A under the hand-to-

mouth assumption. The transfers are defined as a ratio of consumption to income minus one. The figure shows that average transfers follow the opposite pattern of the marginal tax rates. This is expected since the high relative Pareto weight of type A agent indicates both low marginal tax rates through the optimal tax formula and, directly, relatively high consumption and high transfers.

Government Spending Shock. A similar analysis can be made concerning an increase in government spending instead of aggregate productivity. Regarding the response of marginal tax rates, the results are qualitatively similar to the ones in Figure 1. A temporary increase in government spending increases the relative Pareto weight of type A agents and reduces their marginal taxes. The subsequent dynamics again has an opposite sign for infinite risk aversion and zero IES, and when working together, they cancel each other. Other scenarios perform similar dynamics. The dynamics of transfers differs, however. A temporary increase in government spending always reduces transfers to type A agents, although the reduction is relatively mild under zero IES and under both zero IES and infinite risk aversion.

2.5 On the importance of commitment

The dynamics of the optimal marginal taxes when type A people have either infinite risk aversion or zero IES (but not both) show that the optimal policy depends on the history of one's income, and the tax function must take the form $T_t = T_t(y^t)$: in both cases the optimal tax and transfer policies cannot be simply represented as a function of the current income of the agent, and cannot be recovered if the government has no commitment. If type A agents have both infinite risk aversion and zero IES, then commitment is still important, but the argument is more subtle because the optimal policy now happens to be a function of only the current aggregate state. To see this, consider policies that do not allow for dependence on past incomes because the government cannot keep track of past shocks. Similarly to Kapička (2006), such a problem can be written as one where the agents are allowed to reconsider their report in each period and state and where the incentive constraint breaks down to a sequence of period incentive constraints

$$u_t^B(s^t) \geq U\left(C\left(V^A, n_t^A(s^t)\right), \frac{n_t^A(s^t)}{\theta}\right) \quad \forall s^t.$$

Since the type A agent must maintain a period utility V^A across state and time, the period utility of type B agent, $u_t^B(s^t)$ must be increasing in the current productivity shock, and decreasing in the current government spending shock. Consequently, the incentive constraints will only bind under the most adverse shock and be slack otherwise. The optimum will then feature zero marginal taxes on type A agents in all states except for the worst one.

3 The Model

I now consider a full version of the problem where both agents have general recursive preferences. I will formulate the problem in a continuous time. The logarithm of the aggregate productivity z_t follows the Ornstein-Uhlenbeck process

$$d \ln z_t = \zeta(\mu - \ln z_t)dt + \sigma dB_t, \tag{9}$$

σ is the volatility of the stochastic process, μ determines its level, and ζ is the mean reversion parameter. The level parameter μ is set so that the mean of the productivity shock is one. In what follows, I will represent the stochastic process for z as a diffusion in the form $dz_t = \mu_z(z)dt + \sigma_z(z)dB_t$, with $\mu_z(z) = \zeta(\mu - \ln z) + \sigma^2/2$ and $\sigma_z(z) = \sigma z$.

As in the examples above, there are two types of agents, A and B, with measures α and $1-\alpha$, and they differ in their relative productivity and preference parameters. The productivity of type A agent is normalized to equal the aggregate productivity z_t. Type B agents are more productive, with productivity $z_t\theta$, where $\theta > 1$ is the relative productivity of the two types. The value of θ is constant across time and states, which rules out relative productivity changes over time and business cycles.[1]

For simplicity, I set the government consumption $G(z)$ as a function of the current state z. The aggregate resource constraint requires that aggregate consumption cannot exceed aggregate production,

$$\alpha c_t^A + (1-\alpha)c_t^B + G(z) \leq z_t \left[\alpha n_t^A + (1-\alpha)\theta n_t^B\right]. \tag{10}$$

The two types of agents A and B have an instantaneous utility function $U(c_t, n_t)$, where $c_t \geq 0$ is consumption and $n_t \geq 0$ are hours worked. The utility function U is

[1] I do so to isolate the role of preference heterogeneity. It is easy to relax this assumption.

common to both types of agents, increasing and concave in the first argument, decreasing and concave in the second argument, and twice differentiable.

The two types of agents differ in their recursive utility aggregator $f^i(U_t, V_t)$, $i = A, B$, that aggregates the instantaneous utility $U_t = U(c_t, n_t)$ and the lifetime utility V_t of the agent. The function f^i takes the Duffie-Epstein-Zin constant intertemporal elasticity form (Duffie and Epstein (1992a,b)),

$$f^i(U, V) = \frac{\delta}{1-\rho^i}\left[(1-\gamma^i)V\right]^{\frac{\rho^i-\gamma^i}{1-\gamma^i}}\left[U^{1-\rho^i} - \left[(1-\gamma^i)V\right]^{\frac{1-\rho^i}{1-\gamma^i}}\right].$$

The parameter γ^i represents the relative risk aversion of type i, while $1/\rho^i$ is the intertemporal elasticity of substitution in utilities.[2] Both agents have a common discount rate δ. Expected utility is recovered if $\gamma^i = \rho^i$, in which case $f^i = \delta c^{1-\rho^i}/(1-\rho^i) - \delta V$. The lifetime utility of the agents is represented by a recursion

$$V_t^i = V_0^i + \mathbb{E}_t\left[\int_{s=t}^{\infty} f^i(U_s^i, V_s^i)\, ds\right], \quad i = A, B, \tag{11}$$

where $U_t^i = U(c_t^i, n_t^i)$ is the instantaneous utility of type i.

For this paper, it will be easier to work with an alternative representation of the recursive utility, the stochastic variational utility, due to Geoffard (1996) and Dumas et al. (2000). Unlike recursive utility, maintains additive separability in preferences but has a future discounted by a stochastic discount rate that is endogenous. As shown by Dumas et al. (2000), the stochastic variational utility represents the lifetime utility V_t^i as a solution to the following maximization problem:

$$\lambda_t^i V_t^i = \min_{\nu^i \in D} \mathbb{E}_t\left[\int_{s=t}^{\infty} \lambda_s^i F^i\left(U_s^i, \nu_s^i\right) ds\right] \quad \text{s.t.} \quad d\lambda_s^i = -\nu_s^i \lambda_s^i ds, \tag{12}$$

for some instantaneous felicity function F^i, where ν_s^i is the discount rate, D is a space of all stochastic discount rate processes, and $\lambda_0^i = 1$. The stochastic process λ_t^i is the discount factor and is related to the discount rate ν_t^i through the differential equation in the constraint.

[2] One needs to distinguish between an intertemporal elasticity of substitution in utilities and an intertemporal elasticity of substitution in consumption, which depends on the function U and will, in general, be different.

The felicity function F^i calculates the contribution of the instantaneous utility U_t to the lifetime utility given the instantaneous discount factor ν_t. For the felicity function for the Duffie-Epstein-Zin preference class, the felicity function has a closed-form expression

$$F^i(U,\nu) = \delta \frac{U^{1-\gamma^i}}{1-\gamma^i}\left(\frac{1-\gamma^i-(1-\rho^i)\frac{\nu}{\delta}}{\rho^i-\gamma^i}\right)^{\frac{\gamma^i-\rho^i}{1-\rho^i}}. \tag{13}$$

If $\gamma^i = \rho^i$, the felicity function becomes a CRRA utility $F^i = \delta U^{1-\gamma^i}/(1-\gamma^i)$ and, in addition, the discount rate that solves (12) becomes a constant equal to δ.

As in the introductory examples, the agent's type is defined by their productivity, relative risk aversion, and intertemporal elasticity of substitution. The agent's type and hours worked are the private information of the agents, while consumption and income are publicly observable. The incentive compatibility constraint requires that both types prefer their allocation to the allocation of the other type. Since type B agents have higher relative productivity, I will again restrict attention to situations where only the incentive constraint on type B agents binds. To that end, define the utility of a type B agent from an allocation of a type A agent as

$$\hat{\lambda}_t^B \hat{V}_t^B = \min_{\hat{\nu}^B \in D} \mathbb{E}_t\left[\int_{s=t}^{\infty} \hat{\lambda}_s^B F^B\left(\hat{U}_s^B, \hat{\nu}_s^B\right) ds\right] \quad \text{s.t.} \quad d\hat{\lambda}_s^B = -\hat{\nu}_s^B \hat{\lambda}_s^B ds, \tag{14}$$

where $\hat{\lambda}_0^B = 1$ and $\hat{U}_t^B = U\left(c_t^A, n_t^A/\theta\right)$, is the instantaneous utility of the deviating agent. The expression for the instantaneous utility reflects that type B agents must only work n_t^A/θ hours to produce what type A agent produces. The incentive compatibility constraint requires

$$V_0^B \geq \hat{V}_0^B. \tag{15}$$

I will characterize the social planner's problem analogously to Dumas et al. (2000), with the main difference being that my problem is a second-best problem and includes the incentive constraint, while theirs is a first-best problem. The social planner assigns a Pareto weight $\lambda^i \geq 0$ to the type i agent. The Pareto weight is assumed to be sufficiently large that the planner wants to redistribute from type B agents toward type A agents.

The planning problem then chooses an allocation $x = \{c^A, n^A, c^B, n^B\}$ to solve

$$W_0^* = \max_x \ \lambda^A V_0^A + \lambda^B V_0^B \quad \text{s.t.} \quad (10) \text{ and } (15).$$

I separate the social planner's problem into two stages. In the first stage, the value of the Lagrange multiplier on the incentive constraint (15) is taken as given. Let κ be that value. After substituting the lifetime utilities from (12), the social planner's value function $W(\lambda^A, \lambda^B, \kappa, z_0)$ is

$$W_0(\lambda^A, \lambda^B, \kappa, z_0) = \max_x \min_{\nu^A, \nu^B, \hat{\nu^B}} \mathbb{E}_0 \int_{s=0}^{\infty} \left[\lambda_s^A F\left(U_s^A, \nu_s^A\right) + \lambda_s^B F\left(U_s^B, \nu_s^B\right) - \hat{\lambda}_s^B F\left(\hat{U}_s^B, \hat{\nu}_s^B\right) \right] ds$$

subject to the resource constraint (10), the laws of motion for the discount factors,

$$\begin{aligned} d\lambda_s^i &= -\nu_s^i \lambda_s^i ds, \quad i = A, B \\ d\hat{\lambda}_s^B &= -\hat{\nu}_s^B \hat{\lambda}_s^B ds, \end{aligned}$$

the initial conditions $\lambda_0^A = \lambda^A$, $\lambda_0^B = \lambda^B + \kappa$, and $\hat{\lambda}_0^B = \kappa$, and the discount rate processes belonging to D. Note that the initial conditions for the stochastic discount factors are conveniently scaled by the initial Pareto weights λ^A and λ^B, as well as by the Lagrange multiplier on the incentive constraint κ and, in addition, that λ_0^B is a sum of the Pareto weight and the Lagrange multiplier on the incentive constraint.

In the second stage, the social planner looks for the value of the Lagrange multiplier κ such that the incentive constraint (15) holds with equality. In what follows, I will characterize the first-stage problem recursively and derive its implications for an arbitrary value of κ. The second-stage problem will be solved numerically in the quantitative section.

3.1 Recursive formulation

Given that preferences are recursive and the productivity process is a diffusion, the planning problem can be characterized recursively. The value function $W(\lambda^A, \lambda^B, \kappa, z)$

satisfies the following Hamilton-Jacobi-Bellman equation:

$$0 = \max_{x} \min_{\nu^A, \nu^B, \hat{\nu}^B} \left\{ \lambda^A \left[F^A \left(U^A, \nu^A \right) - \nu^A W_{\lambda^A} \right] + \lambda^B \left[F^B \left(U^B, \nu^B \right) - \nu^B W_{\lambda^B} \right] \right.$$
$$\left. - \hat{\lambda}^B \left[F^B \left(\hat{U}^B, \hat{\nu}^B \right) + \hat{\nu} W_{\hat{\lambda}^B} \right] + \mu_z(z) W_z + \frac{\sigma_z(z)^2}{2} W_{zz} \right\} \tag{16}$$

subject to (10), where W_{λ^A}, W_{λ^B}, $W_{\hat{\lambda}^B}$, W_z and W_{zz} are partial derivatives of the value function. The recursive formulation is similar to Borovička (2019), again with the main technical difference being that between a first-best and second-best problem.

The first three arguments of the value function represent the weights of the two agents and the deviating agent. This suggests that the value function is homogeneous of degree one in the three weights. This is indeed the case:

Proposition 1. *The value function satisfies for any* $\lambda^A \geq 0$, $\lambda^B \geq 0$, $\hat{\lambda}^B \geq 0$ *and* $z \geq 0$,

$$W(\lambda^A, \lambda^B, \hat{\lambda}^B, z) = (\lambda^A + \lambda^B - \hat{\lambda}^B) W \left(\frac{\lambda^A}{\lambda^A + \lambda^B - \hat{\lambda}^B}, \frac{\lambda^B}{\lambda^A + \lambda^B - \hat{\lambda}^B}, \frac{\hat{\lambda}^B}{\lambda^A + \lambda^B - \hat{\lambda}^B}, z \right).$$

Proof. See Appendix A. □

Using the homogeneity property, we define the value function in its intensive form as

$$w(\lambda, \kappa, z) = W(\lambda, 1 - \lambda + \kappa, \kappa, z),$$

where the normalized state variables are

$$\lambda = \frac{\lambda^A}{\lambda^A + \lambda^B - \hat{\lambda}^B}, \quad \kappa = \frac{\hat{\lambda}^B}{\lambda^A + \lambda^B - \hat{\lambda}^B}.$$

The intensive form value function is a function of two endogenous state variables, λ, which has the usual interpretation as the time-varying normalized Pareto weight of type A, and κ, which has a novel interpretation as a time-varying Lagrange multiplier on the incentive constraint. The state variables have a law of motion

$$d\lambda = \lambda \left[(1 - \lambda)(\nu^B - \nu^A) + \kappa(\nu^B - \hat{\nu}^B) \right] \tag{17}$$
$$d\kappa = \kappa \left[\lambda(\nu^A - \nu^B) + (1 + \kappa)(\nu^B - \hat{\nu}^B) \right]. \tag{18}$$

The Pareto weight of type A will increase not only if type A is momentarily more patient than type B ($\nu^A < \nu^B$) but also if the deviating type B is more patient than the truthtelling type B ($\hat{\nu^B} < \nu^B$). The latter inequality signals that the allocation provided to the type A agent tends to tighten the incentive constraint, and the social planner responds by lowering the Pareto weight for the type A agent.

The lifetime utilities of the agents are not explicit in this formulation of the social planner's problem but can be inferred indirectly once the value function is known. The lifetime utilities of the two types and of the deviating B agent are $V^A = w + (1-\lambda)w_\lambda - \kappa w_\kappa$, $V^B = w - \lambda w_\lambda - \kappa w_\kappa$, and $\hat{V}^B = w - \lambda w_\lambda - (1+\kappa)w_\kappa$.

Rewriting (16) in terms of w and regrouping terms yields the following intensive-form HJB equation:

$$\begin{aligned} 0 = \max_x \min_{\nu^A, \nu^B, \hat{\nu}^B} & \left\{ \lambda F^A\left(U^A, \nu^A\right) + (1-\lambda+\kappa)F^B\left(U^B, \nu^B\right) - \kappa F^B\left(\hat{U}^B, \hat{\nu}^B\right) \right. \\ & - \left[\lambda\nu^A + (1-\lambda+\kappa)\nu^B - \kappa\hat{\nu}^B\right] w + \lambda\left[-(1-\lambda)\nu^A + (1-\lambda+\kappa)\nu^B - \kappa\hat{\nu}\right] w_\lambda \\ & \left. + \kappa\left[\lambda\nu^A + (1-\lambda+\kappa)\nu^B - (1+\kappa)\hat{\nu}\right] w_\kappa + \mu_z(z)w_z + \frac{\sigma_z(z)^2}{2} w_{zz} \right\} \end{aligned} \tag{19}$$

subject to the resource constraint (10). From now on, the HJB equation (19) is the main object of interest. It will now be characterized qualitatively and quantitatively, and its implications for the optimal marginal tax rates will be studied.

3.2 Efficient Allocations and Taxes

The planner's decision consists of three choices: allocating consumption across types, allocating consumption and hours for a given type, and choosing the discount rates for each type. For the sake of brevity, I focus on the first two choices.[3] Consumption sharing across types is determined by the following conditions:

$$\frac{\alpha}{1-\alpha}(1-\lambda+\kappa)\frac{F_u^B}{F_u^A}\frac{U_c^B}{U_c^A} = \lambda - \kappa\frac{\hat{F}_u^B}{F_u^A}\frac{\hat{U}_c^B}{U_c^A}. \tag{20}$$

In a first-best allocation, κ would be zero, and the sharing rule would equate Pareto-weighted marginal utility of consumption $F_u^i U_c^i$ across types. With private information,

[3] Discounting is important because it determines how the Pareto weight λ and Lagrange multiplier on the incentive constraint κ changes, according to their law of motion (17) and (18).

the left-hand side of (20) is depressed, and the marginal rate of substitution for type A is increased relative to type B; that is, consumption of type A is decreased relative to type B to relax the incentive constraint.

To characterize the intratemporal allocation for each type, and the distortions generated in the planning problem, define the intratemporal wedge for each type by[4]

$$\tau^A = 1 + \frac{U_n^A}{U_c^A}\frac{1}{z}, \quad \tau^B = 1 + \frac{U_n^B}{U_c^B}\frac{1}{z\theta}.$$

Then, in the optimum, τ^A and τ^B are equal to

$$\tau^A = \frac{\frac{\kappa}{\lambda}\frac{\hat{F}_u^B}{F_u^A}\left(\frac{\hat{U}_c^B}{U_c^A} - \frac{\hat{U}_n^B}{U_n^A}\frac{1}{\theta}\right)}{1 - \frac{\kappa}{\lambda}\frac{\hat{F}_u^B}{F_u^A}\frac{\hat{U}_n^B}{U_n^A}\frac{1}{\theta}}, \quad \tau^B = 0. \tag{21}$$

As usual, type B is undistorted (no distortion at the top), and type A is distorted downwards. The expression for τ^A shows two potential sources of dynamics. The first one is the information rent. The informational rent is represented in (21) by two terms: $\phi_n = \hat{U}_n^B/U_n^A/\theta$, and $\phi_c = \hat{U}_c^B/U_c^A$. The first term ϕ_n is standard. Keeping other things equal, higher values of ϕ_n will be associated with lower marginal tax rates because the informational rent is larger, so redistributing is more costly. The distinct appearance of the term ϕ_c reflects the nonseparability of preferences between consumption and hours worked, inherent with the Duffie-Epstein-Zin preference class: With separable preferences, ϕ_c is always equal to one. A higher value of ϕ_c will increase the marginal tax rate because it increases the marginal utility of consumption of the deviating agent and tightens the incentive constraint. In the optimal tax formula (4) from the introductory examples, ϕ_n was equal to a constant ϕ, while ϕ_c was equal to one.

The second source of dynamics in τ^A comes from the term

$$p = \frac{\kappa}{\lambda} \times \frac{\hat{F}_u^B}{F_u^A}.$$

While ϕ_n and ϕ_c represent the informational rent in current utils, one can think of p as a time-varying shadow price for the planner's informational rent. If the shadow price of the informational rent increases, it is more beneficial for the social planner to reduce it by

[4] I will loosely refer to the intratemporal wedges as marginal taxes later on.

distorting the allocation for type A, and the marginal tax rate increases. With standard time-additive preferences, κ and λ are constant, while $\hat{F}_u^B$ and F_u^A are equal to one, so p is a constant.

The shadow price p consists of two terms. Changes in the ratio κ/λ reflect the time-varying importance of the incentive constraint relative to type A's Pareto weight. It follows from the laws of motion (17) and (18) that changes in the ratio reduce to a very simple expression:

$$d\left(\frac{\kappa}{\lambda}\right) = \frac{\kappa}{\lambda}\left(\nu^A - \hat{\nu}^B\right).$$

Thus, only the difference between the discount rate of the deviating type B and type A matters for the dynamics of κ/λ. If the instantaneous discount rate of the deviator is higher than that of the type A, then the shadow price of the informational rent will decrease. To put it differently, the distortions will be gradually relaxed in times when the deviator discounts more than type A because they will be less helpful in alleviating the pressure from the incentive constraint (15).[5]

At first sight, the idea that an increase in λ decreases distortions is counterintuitive. In a static Mirrlees model, an increase in the Pareto weight of the low type typically tightens the incentive constraint due to an increased desire to redistribute and leads to a corresponding increase in its Lagrange multiplier. In this model, the link between the instantaneous values of λ and κ is weaker because one can intertemporally rearrange their values without affecting how the incentive constraint binds ex-ante (i.e., for a given initial value κ_0).

Since the term κ/λ has no instantaneous volatility, it will be characterized by lower frequency movements. In contrast, the second term $\hat{F}_u^B/F_u^A$ will impact the marginal tax rate instantaneously. The movements in this ratio come from the changes in the marginal values of the deviator's instantaneous utility $\hat{U}^B$ relative to the truthteller's utility U^A. A low value of the deviator's utility represents times when the informational rent is not very important and can be paid without affecting the incentive constraint too much.

The next proposition shows conditions under which none of the dynamic forces in (21) is active, and the intratemporal wedge is constant across time and states:

Proposition 2. *Suppose that $U(c,n) = c\phi(n)$ and $G(z) = \bar{G}z$. If $\gamma^A = \gamma^B$ and $\rho^A = \rho^B$, then the intratemporal wedge τ^A is constant across time and states.*

[5]To see an extreme example of this principle, consider a two-period discrete-time economy, where type B has a discount factor β equal to zero. Then, it makes no sense to distort the allocation in the second period because the second-period allocation does not enter the incentive constraint at all.

Proof. See Appendix A. □

The joint assumptions about the instantaneous utility function, the government consumption function, and homogeneous preferences ensure that the optimal consumption allocation for both types is proportional to the aggregate productivity shock and that hours worked are independent of the productivity shock. This, in turn, ensures two things: first, the instantaneous discount rates of both types and the deviator are always equal (although state-dependent), and so both λ and κ are constant. Second, the assumptions ensure that the informational rent is constant as well. Taken together, the intratemporal wedge must be constant as well.

4 Quantitative Results

I calibrate the model to the postwar United States, specifically the 1947-2012 period. One unit of time corresponds to a quarter. All parameters are summarized in Table 1.

The parameters σ and ζ of the Ornstein-Uhlenbeck (9) are calibrated to match quarterly standard deviation and autocorrelation of detrended US log labor productivity for the 1947-2012 time period. I measure labor productivity by the output per hour in the nonfarm business sector for all workers, as reported by the BLS, and linearly detrend log productivity. This detrended log labor productivity has standard deviation Std $= 0.018$ and autocorrelation Corr $= 0.975$, which translates to Ornestein-Uhlenbeck parameters $\zeta = -\ln(\text{Corr}) = 0.025$ and $\sigma = \sqrt{2\zeta}\,\text{Std} = 0.004$. The parameter $\mu = -\sigma^2/(4\zeta)$ then normalizes the level of productivity to one. Government consumption is set to be 16% of steady-state output in the calibrated economy, independently of the current productivity, based on NIPA data.[6]

I set α, the fraction of type A agents in the population, to be 0.8. While this fraction is somewhat arbitrary, I chose this value so that the fraction high productivity agents $1-\alpha$ loosely corresponds to the fraction of shareholders in the population. The relative productivity of type B, θ, is then chosen to match the ratio of labor earnings of top-20% and bottom-80% of the population. The ratio of labor earnings is computed from Table I of Kopczuk et al. (2010) and equals 4.232.

[6]Because government consumption is not proportional to labor productivity, Proposition 2 does not apply and marginal taxes are not constant even under preference homogeneity. However, the fluctuations in the optimal marginal taxes coming from G being constant are quantitatively negligible.

Table 1: Model Parameters

parameter	value	target
ζ	0.025	autocorrelation of US log labor productivity, 1947-2012
σ	0.004	standard deviation of US log labor productivity, 1947-2012
G	0.24	gov't consumption to output ratio of 0.16, NIPA
α	0.8	exogenously set
λ	0.932	steady-state average marginal tax rate of 0.294
θ	4.232	20-80 relative earnings, Kopczuk et al. (2010)
η	2	Frisch elasticity of 0.5
δ	0.005	2% annual interest rate
ρ^A	20	exogenously set
γ^A	10	exogenously set
ρ^B	4.667	annual risk-free rate 1.52%, Gârleanu and Panageas (2015)
γ^B	2.773	B-to-A ratio of the standard deviation of log consumption of 2.57

The Pareto weight λ is set as follows. I compute a steady-state version of the planning problem for the economy and choose λ so that it is *optimal* to have the average marginal tax rate in the economy equal to the 1947-2012 average, which is 0.294 based on the data in Mertens and Olea (2018). This produces $\lambda = 0.932$. Since $\lambda > \alpha$, the planner puts more weight on the type A agents than their population share would dictate. The procedure for finding the Pareto weight follows Heathcote and Tsujiyama (2021), and it allows the analysis to separate the cyclical fluctuations in the marginal tax rates from changes in their level that would otherwise be present.

The utility function is assumed to be of the type specified in Proposition 2:

$$U(c,n) = ce^{-\frac{n^{1+\eta}}{1+\eta}}.$$

I set $\eta = 2$, which corresponds to the Frisch elasticity of labor one half. The instantaneous discount rate δ is set to 0.005, which yields a 2% annual discount rate.

The preference parameters of the A type are set exogenously: the elasticity of intertemporal substitution is set to 0.05, which yields $\rho^A = 20$. This is consistent with the empirical evidence that shows that the intertemporal elasticity of substitution is close to zero for nonstockholders (Vissing-Jørgensen (2002), Guvenen (2006)). I set the risk aversion of type A agents to $\gamma^A = 10$. This value is a standard upper bound used in

the literature; see Mehra and Prescott (1985). Both parameters are roughly in line with the existing literature: Guvenen (2009) for example sets $\rho^A = 10$ and $\gamma^A = 6$ (common to both types) for the bottom 80 % of the population, while Gârleanu and Panageas (2015) set precisely the same parameters as this paper, although for the bottom 99% of the population.

For type B, the preference parameters are calibrated. To do so, I compute a competitive equilibrium for an otherwise identical economy, where all agents face a flat marginal tax rate of 0.294, based on the calibration above. The details are provided in Appendix B. I then choose ρ^B and γ^B to match two targets. The first one is the average annual risk-free rate of 1.52%, a 1947-2012 average real return on 6-month commercial paper, as reported by Shiller (1992).[7] The second target is the ratio of the volatility of log consumption for type B and type A agents. Based on the Consumer Expenditure Survey, Malloy et al. (2009) estimate the volatility of stockholders and nonstockholders to be 3.6 and 1.4. Thus, the target ratio is set at 2.57.[8] I compute the respective statistics in the model as a population average over the first 66 years, corresponding to the 1947-2012 period. The procedure yields $\rho^B = 4.667$ and $\gamma^B = 2.773$. That is, type B agents have intertemporal elasticity of substitution of 0.214.

To compute the social planner's problem (19), I discretize the state space in all three dimensions, with 70 equispaced gridpoints for $\lambda \in [0,1]$, 40 equispaced gridpoints for $\kappa \in [0,\bar{\kappa}]$ for $\bar{\kappa}$ sufficiently large, and 21 equispaced gridpoints for $\ln z \in [\mu - 3\sigma/\sqrt{2\zeta}, \mu + 3\sigma/\sqrt{2\zeta}]$. Following Achdou et al. (2021), I compute the HJB equation (19) using the upwind finite difference scheme. The distorted first-best problem used in the calibration is computed similarly. I also use the Kolmogorov forward equation to compute the evolution of the distribution and, subsequently, the average risk-free rate and volatility of log consumption. In all computations, I choose the time step $dt = 0.01$; one time step thus represents less than one day.

4.1 What drives the optimal marginal tax rate?

The optimal allocation exhibits cyclical fluctuations in response to the productivity shock, but also a nonzero trend since both κ and λ are not necessarily stationary. The left panel

[7] Updated data are at, http://www.econ.yale.edu/~shiller/data/chapt26.xls.

[8] The literature typically estimates the volatility of consumption growth. However, the volatility ratio of log consumption and consumption growth in thisd model are almost identical.

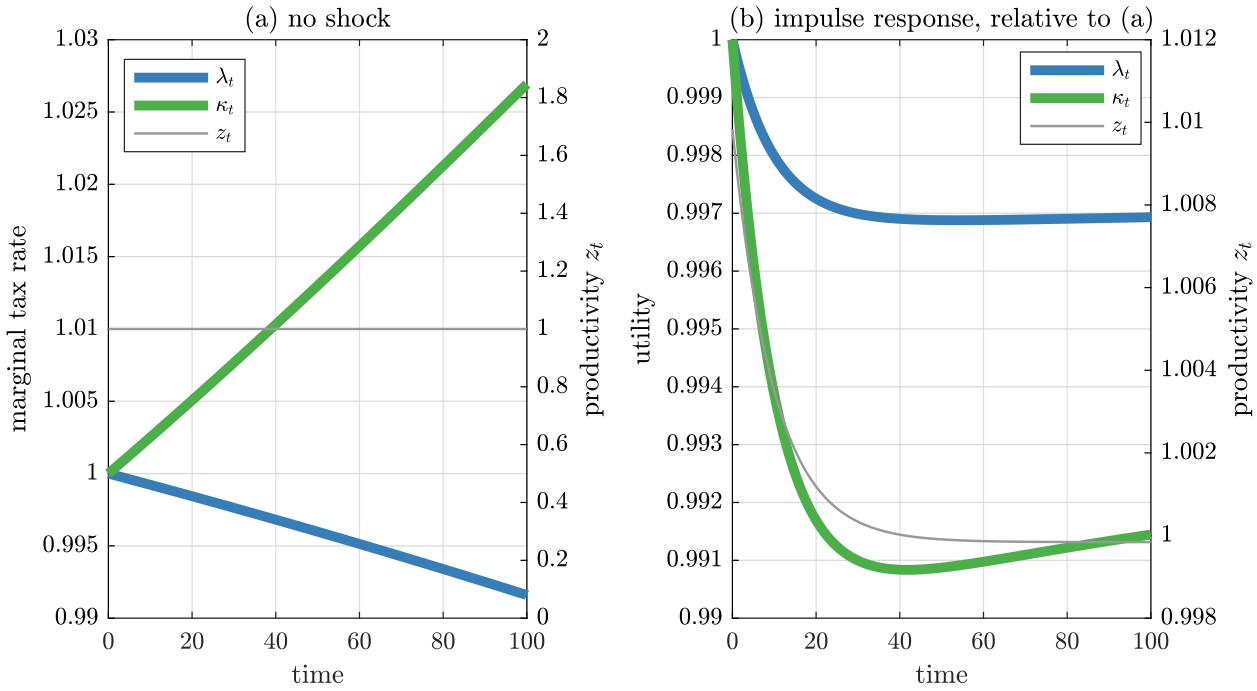


Figure 2: Panel (a) shows the dynamics of λ_t and κ_t in the absence of TFP shocks. Panel (b) shows the response to one standard percent TFP increase relative to the trajectory in (a).

of Figure 2 shows, however, that the trend is quantitatively negligible: in the absence of any TFP shock, λ decreases by only about 1% over 100 years, while κ increases by about 2.5 % over the same period. As the left panel of Figure 3 shows, this secular increase in the shadow price of the informational rent translates into a negligible increase of the marginal tax rate τ^A from 0.307 to less than 0.31. Thus, while theoretically possible, the trend movements are unimportant for the quantitative results.

The right panels of Figures 2 and 3 show the optimal response to a one percent increase in productivity z. The right panel of Figure 2 shows that κ and λ decrease following the shock relative to the no-shock trend. The value of κ, however, decreases more. As a result, the ratio κ/λ, which is relevant for the optimal tax, also decreases. One can infer from Figure 2 that κ/λ is U-shaped, reaches a through after about 25 years, and then very gradually increases back to the original trend from the left panel of Figure 3. The magnitude of its change is minor: at its trough, the value of the ratio declines only 0.6 percent.

The marginal tax rate increases substantially on impact by about 7 percent. Given that the marginal tax rate was 30.8 percent in the absence of the shock, a one percent

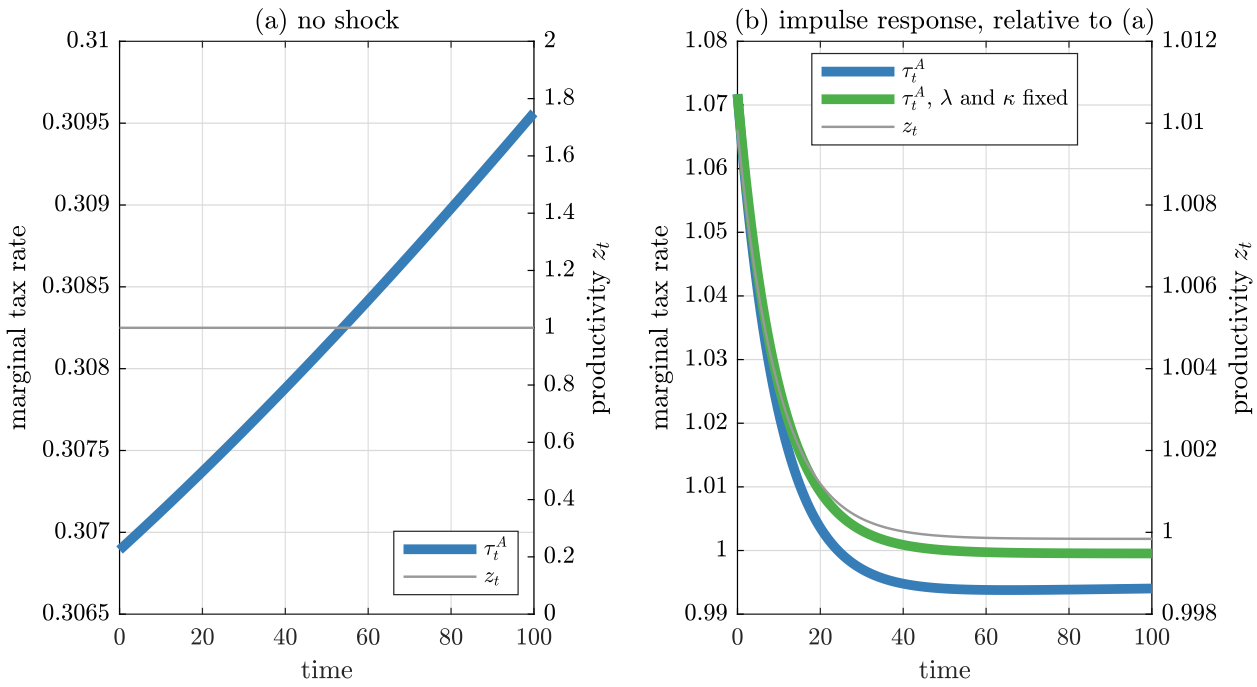


Figure 3: Panel (a) shows the dynamics of τ_t^A in the absence of TFP shocks. Panel (b) shows the response to a one percent TFP increase relative to the trajectory in (a).

productivity shock increases the marginal income tax by 2.2 percentage points. The marginal tax rate declines immediately after the shock, but the response is not monotone. The right panel of Figure 3 shows that the marginal tax rate "overshoots" the trend path and declines below after about 25 years, with a very gradual return to the trend afterward from below (barely visible in the figure spanning 100 years). The overshooting is quantitatively small, by about half a percent at its maximum.

To understand the response of the marginal tax rate τ_t^A, it is useful to separate the instantaneous response to the shock when κ and λ are fixed and a dynamic response where κ and λ adjust. Consider first the instantaneous response. In a first-best economy, an increase in productivity would increase consumption of both types to keep the marginal rate of substitution in consumption unchanged, although it would increase it more for the type B agent. Importantly, an increase in consumption of agent A increases the shadow price of paying the informational rent $p \propto \hat{F}_u^B / F_u^A$, because the marginal utility of the truthtelling agent A decreases faster than the marginal utility of the deviating agent.[9] The social planner responds by reducing the informational rent and its

[9]To see this, note that it follows from (13) that $\hat{F}_u^B / F_u^A \propto (c^A)^{\gamma^A - \gamma^B}$, and that $\gamma^A > \gamma^B$.

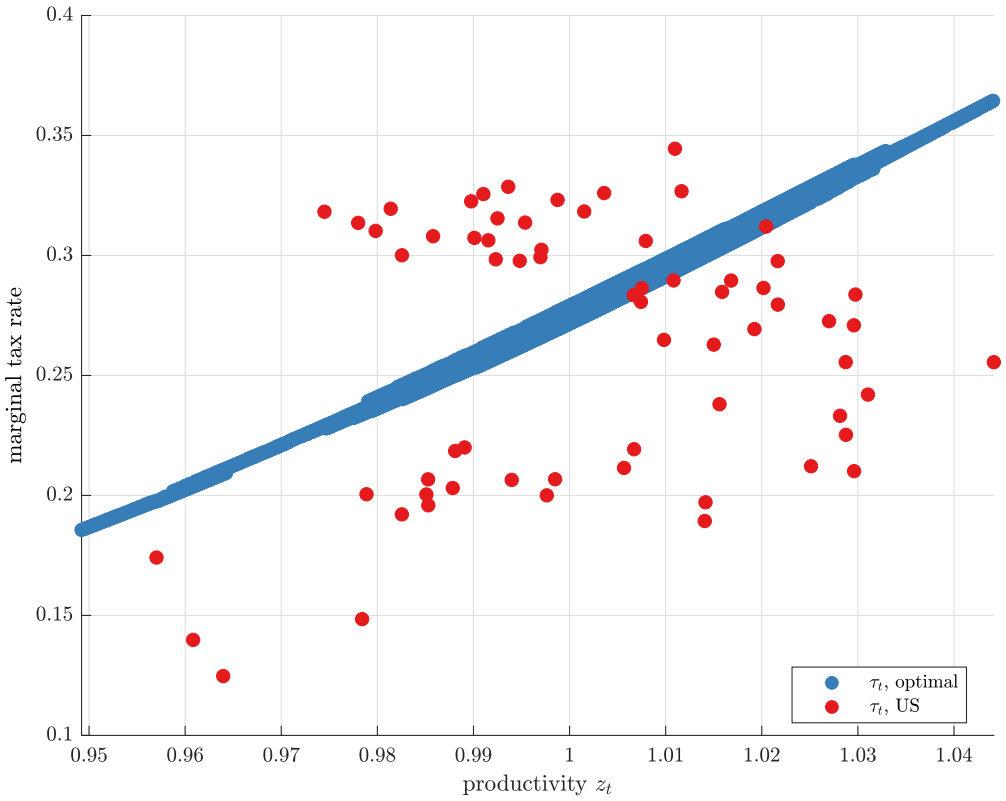


Figure 4: Average marginal tax rates on bottom 90 % of the population versus US productivity realizations. US average marginal tax rates are from Mertens and Olea (2018).

shadow price by reducing both the consumption and labor supply of type A agent. That is, the social planner increases the marginal tax rate on type A agent. In fact, in the optimum, consumption of type A agent turns slightly countercyclical since fluctuations in the marginal tax rates dominate fluctuations in productivity.

The dynamic response of τ_t^A depends on the response of κ/λ. Since κ/λ is smaller than one and U-shaped, the dynamic response contributes to a decline in the marginal tax rates and is responsible for overshooting the trend path. As shown in Figure 3, if κ/λ is kept fixed, there is no overshooting, and the marginal tax rate returns relatively quickly to the trend path. The calibrated model is thus qualitatively similar to Example 2, where type A agents have zero intertemporal elasticity of substitution. However, the magnitude of the dynamic response is relatively small since the marginal tax rates overshoot the trend path only by about 0.5 percent. From this perspective, Example 3, with no subsequent dynamics, is a decent quantitative approximation of the optimum.

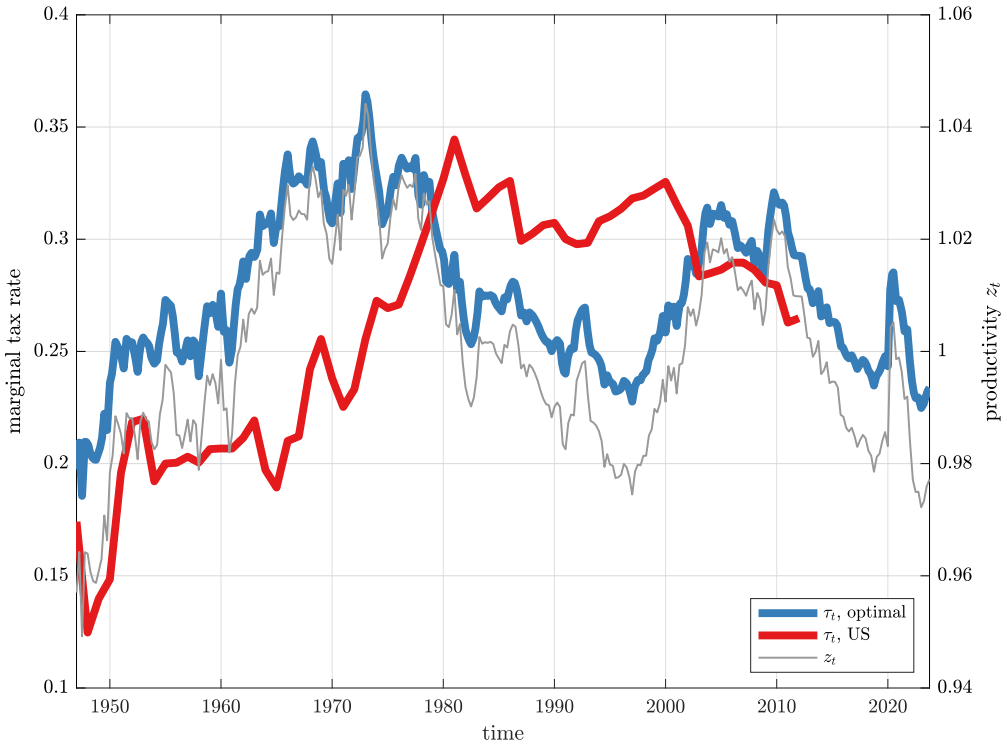


Figure 5: Time series of average marginal tax rates on bottom 90 % of the population for a sequence of US labor productivity realizations. US average marginal tax rates from Mertens and Olea (2018).

4.2 Optimal tax and transfer policy in the US

To characterize the optimal policy for the US economy, I feed the measured US labor productivity for the 1947-2023 period into the model.[10] The productivity time series is shown as a thin black line in Figure 5. I compare the optimal marginal tax rates to the US average marginal tax rates, as estimated by Mertens and Olea (2018) for the 1947-2012 period. Mertens and Olea (2018) estimate the average marginal tax rates for the bottom 90 percent of the population. To match their estimates to the model counterparts where type A agents represent the bottom 80 percent of the population, I use the result that the top 20 percent of the population faces zero marginal tax rate. The average marginal tax rate on the bottom 90 percent of the population in the model is then simply 8/9 of τ_t^A. This value is represented in Figures 4 to 6.

Figure 4 shows the scatterplot of the optimal marginal tax rates and compares them

[10]Since the data are available on a quarterly frequency while the model period dt is smaller, I assume that the productivity shocks are constant within the quarter, and select the shock so that the overall change in productivity over the quarter matches the data.

Table 2: Marginal Tax Rates: Statistical Moments

	1947-2012		1947-1979		1980-2012	
marginal tax rates	US	optimal	US	optimal	US	optimal
standard deviation	0.053	0.037	0.042	0.044	0.019	0.024
autocorrelation	0.964	0.891	0.907	0.901	0.874	0.850
correlation with z_t	0.196	0.997	0.767	0.999	-0.485	0.997

to the US average marginal tax rates. The behavior of the optimal marginal tax rates is consistent with the findings of the previous section: the marginal tax rate increases strongly with the current productivity by approximately 2.2 percentage points for a one percentage point increase in productivity. Overall, the optimal marginal tax rates lie in a broad range between 18 and 35 percent. The long-run dynamics is also present (otherwise, the blue set would be a straight line). However, it is quantitatively small: the optimal marginal tax rate can be within a 2 percent range for a given productivity level, but not more. The US average marginal tax rates exhibit much higher dispersion for a given productivity level and a smaller correlation with productivity overall, but the overall range of 13-35 percent is comparable to the range of optimal taxes.

A time series comparison in Figure 5 gives another perspective on the comparison. The upward trajectory of the US average marginal tax rates after World War II that continues to around 1980 is approximately optimal until the early 1970s. The increase in the average marginal tax rates in the late 1970's is due to inflation and bracket creep (Mertens and Olea (2018)) and, in the absence of a strong labor productivity growth, runs counter to the optimal tax pattern. A decrease in the average marginal tax rates in the 1980s because of a sequence of tax reforms has been approximately optimal; the increase in the early 1990s and a subsequent decline in the 21st century were not.

Table 2 decomposes the post-war era into two subperiods, before and after 1980, and computes statistical moments of the average marginal tax rates. For the whole period, the US average marginal tax rates exhibit higher volatility, substantially weaker correlation with productivity, and a somewhat higher autocorrelation; the first two facts were already visually evident from the scatter plot Figure 4. The pre-1980 subperiod is, on the other hand, almost perfect from the perspective of the standard deviation and autocorrelation. The correlation with productivity still lags behind the optimum,

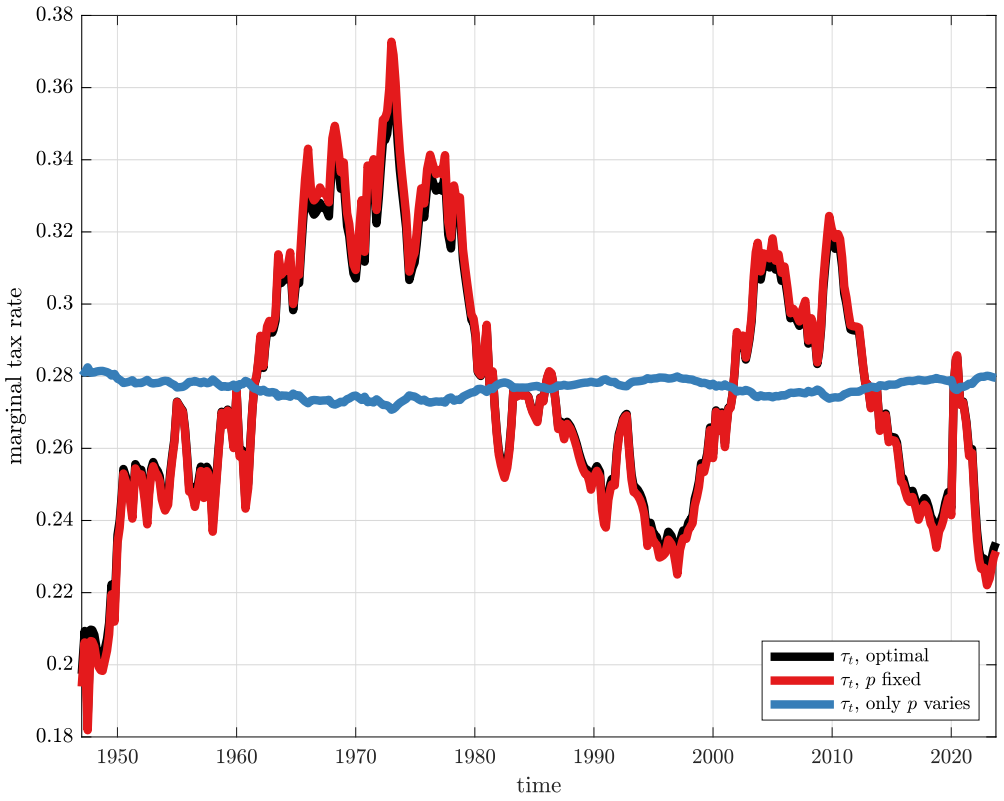


Figure 6: Decomposition of the marginal tax rate fluctuations. The line "p fixed" keeps the shadow price of informational rent $p = \kappa/\lambda \times \hat{F}_u^B / F_u^A$ fixed at average levels. The line "only p varies" keeps all other components fixed at average levels.

but at 0.767, it is substantially higher than 0.196 throughout the whole post-war period. The main difference between what's optimal and the US policy comes in the post-1980 period (and especially from changes after 1990). The US average marginal tax rates become strongly countercyclical, and the correlation with productivity is -0.485.

Figure 6 decomposes the optimal marginal tax rate into two components: changes in the informational rent and changes in the shadow price of the informational rent. The decomposition is striking, although not unexpected, in light of the previous discussion of the impulse response: the shadow price of the informational rent p is almost entirely responsible for the fluctuations of the marginal tax rate. Shutting down changes in p isolates the role of the informational rent, and the fluctuations in the marginal tax rate are an order of magnitude smaller. On the other hand, if one shuts down changes in the informational rent, the resulting marginal tax rates are almost indistinguishable from the optimal ones. The decomposition shows how important preference heterogeneity,

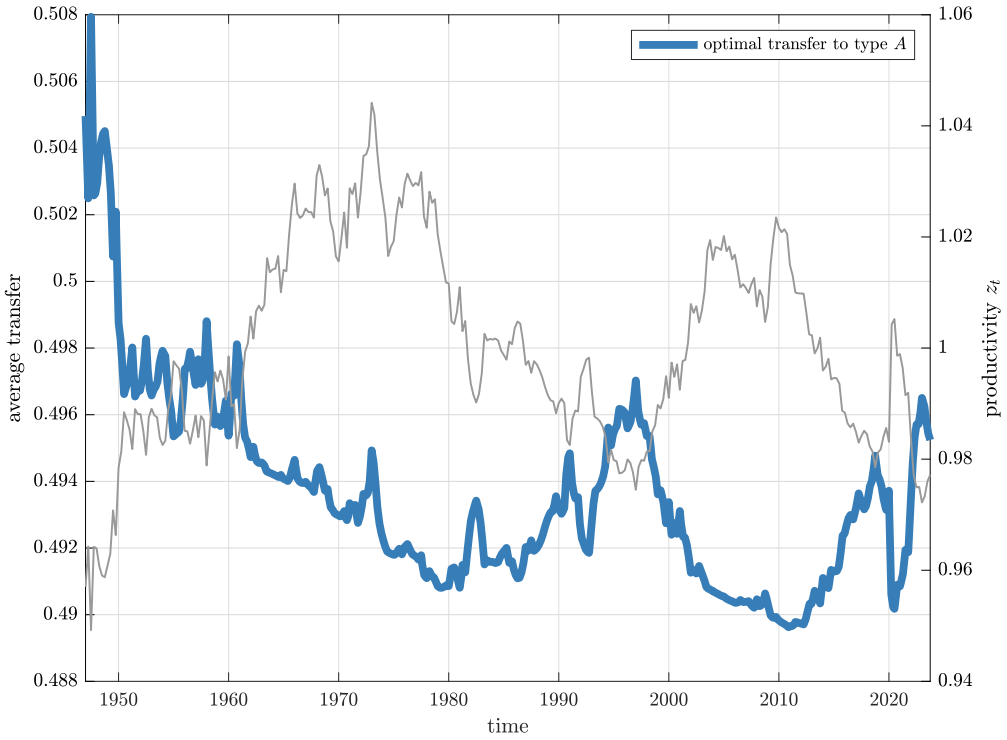


Figure 7: Optimal transfers to type A, $TR_t^A = (c_t^A - y_t^A)/y_t^A$.

responsible for the fluctuations in p, is for the overall results.

It was already discussed in Section 2.4 that transfers are generally indeterminate but can be uniquely determined under the assumption that type A agents are hand-to-mouth. Figure 7 shows the transfers to type A agents $TR_t^A = (c_t^A - y_t^A)/y_t^A$ under this assuption. The optimal transfers are countercyclical, and their correlation with labor productivity is -0.642. Surprisingly, their fluctuations are quantitatively small and range only from 49% to 50.8%. This does not mean that type A agents face large fluctuations in consumption or instantaneous utility; their fluctuations in both are very small. Instead, the fluctuations in transfers are relatively unimportant because the procyclical changes in the marginal tax rates already help reduce the fluctuations in period utility. To see this, note that the intratemporal optimality condition on type A agent is $c_t^A (n_t^A)^\eta = z_t(1 - \tau_t^A)$. Low values of z_t are more than offset by high values of the marginal retention rate $1 - \tau_t^A$, and the left-hand side of the expression is countercyclical. This helps the planner to increase consumption or reduce hours (or both) in recessions.

5 Conclusions

The paper presents several novel results, both quantitative and qualitative. Accounting for preference heterogeneity is of first-order quantitative importance for the procyclical fiscal policy; marginal taxes are strongly procyclical, and a one percent increase in labor productivity above trend should increase the marginal tax rates by about 2.2 percentage points; and the main channel through which the increase operates are fluctuations in the shadow price of the informational rent, a novel concept that is unimportant in the absence of preference heterogeneity. Those results obviously rest on the underlying assumptions of the paper. This section discusses the assumptions and potential extensions.

The heterogeneity in productivity is deliberately restricted to two types to retain tractability. To the extent that adding types with "intermediate" productivity preserves the monotonicity of preferences, with lower productivity types having lower intertemporal elasticity of substitution and higher risk aversion, one would expect the results to be similar, with perhaps the most significant fluctuations in the marginal tax rates at the bottom end of the distribution. Doing so would require introducing additional state variables, one for each type. Allowing instead for an imperfect correlation between productivity and preferences would complicate the analysis substantially due to multidimensional private information. In addition, it would lead to somewhat unnatural optimal policies, where the marginal income tax rates would depend on the declared preference type. In both cases, however, properly calibrating the model would become substantially more challenging because one would need to calibrate the correlation between productivity and preference parameters. Such an exercise seems infeasible without a deeper understanding of how preference parameters are distributed in the population.

It is well known that detrending the aggregate productivity series is not an innocuous procedure when the agents are heterogeneous in the intertemporal elasticity of substitution because the agents exhibit varying willingness to accept a growing consumption. The non-stationary trends, evident in Figures 2 and 3, would likely be strengthened and would likely lead to a nonnegligible drift in the marginal tax rates. However, such forces can be counteracted by, for example, introducing additional heterogeneity in the discount rates and calibrating it to match the existing trends in consumption distribution or simply directly counteracting the underlying trend in productivity. It is unclear how, if at all, the resulting optimum would change.

The assumption that low-productivity agents are hand-to-mouth pins down their transfers but also allows us to sidestep a related important issue of how the risk-free and risky assets should be traded among the two types. In fact, the consumption sharing equation (20) implies that, in the optimum, both agents face a different stochastic discount factor, suggesting that it would be optimal to tax or subsidize the risk-free and risky assets in such case. The analysis of the optimal asset taxes is left for a separate exercise.

Finally, the quantitative analysis has restricted attention to only one source of shock: labor productivity. A more realistic exercise would allow for multiple aggregate shocks, particularly for government consumption fluctuations. That would produce a more nuanced view of the optimal tax policy and would likely reduce the high correlation between the optimal marginal tax rates and productivity, possibly bringing the optimum closer to the data. Solving such a model would require introducing a multidimensional diffusion process for shocks and adding additional state variables, but the model's basic structure would be unchanged. This exercise is left for future research as well.

References

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2021, 04). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The Review of Economic Studies 89*(1), 45–86. 24

Aiyagari, S., A. Marcet, T. Sargent, and J. Seppala (2002). Optimal taxation without state-contingent debt. *Journal of Political Economy 110*(6), 1220–1254. 5

Anderson, E. W. (2005). The dynamics of risk-sensitive allocations. *Journal of Economic Theory 125*(2), 93–150. 4, 8

Atkeson, A. and M. Ogaki (1996). Wealth-varying intertemporal elasticities of substitution: Evidence from panel and aggregate data. *Journal of Monetary Economics 38*(3), 507–534. 4

Attanasio, O., J. Banks, and S. Tanner (2002). Asset holding and consumption volatility. *Journal of Political Economy 110*(4), 771–792. 4

Attanasio, O. P. and M. Browning (1995). Consumption over the life cycle and over the business cycle. *The American Economic Review 85*(5), 1118–1137. 4

Backus, D. K., B. R. Routledge, and S. E. Zin (2008). Who holds risky assets? Technical report, New York University. 5

Bhandari, A., D. Evans, M. Golosov, and T. J. Sargent (2016, 11). Fiscal Policy and Debt Management with Incomplete Markets*. *The Quarterly Journal of Economics 132*(2), 617–663. 5

Blundell, R., M. Browning, and C. Meghir (1994). Consumer demand and the life-cycle allocation of household expenditures. *The Review of Economic Studies 61*(1), 57–80. 4

Borovička, J. (2019). Survival and long-run dynamics with heterogeneous beliefs under recursive preferences. *Journal of Political Economy 128*(1), 206–251. 5, 18

Brav, A., G. M. Constantinides, and C. C. Geczy (2002). Asset pricing with heterogeneous consumers and limited participation: Empirical evidence. *Journal of Political Economy 110*(4), 793–824. 4

Chari, V. V., L. J. Christiano, and P. J. Kehoe (1994). Optimal fiscal policy in a business cycle model. *Journal of Political Economy 102*(4), 617–652. 5

Colacito, R., M. M. Croce, and Z. Liu (2019). Recursive allocations and wealth distribution with multiple goods: Existence, survivorship, and dynamics. *Quantitative Economics 10*(1), 311–351. 5

Duffie, D. and L. G. Epstein (1992a, 05). Asset Pricing with Stochastic Differential Utility. *The Review of Financial Studies 5*(3), 411–436. 15

Duffie, D. and L. G. Epstein (1992b, 02). Stochastic differential utility. *Econometrica 60*(2), 353–394. 5, 15

Dumas, B., R. Uppal, and T. Wang (2000). Efficient intertemporal allocations with recursive utility. *Journal of Economic Theory 93*, 240–259. 3, 15, 16

Epstein, L. G. and S. E. Zin (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica 57*(4), 937–969. 5

Gârleanu, N. and S. Panageas (2015). Young, old, conservative, and bold: The implications of heterogeneity and finite lives for asset pricing. *Journal of Political Economy 123*(3), 670–685. 4, 5, 23, 24

Geoffard, P.-Y. (1996, April). Discounting and Optimizing: Capital Accumulation Problems as Variational Minmax Problems. *Journal of Economic Theory 69*(1), 53–70. 3, 15

Gomes, F. and A. Michaelides (2007, 11). Asset Pricing with Limited Risk Sharing and Heterogeneous Agents. *The Review of Financial Studies 21*(1), 415–448. 4

Guvenen, F. (2006). Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective. *Journal of Monetary Economics 53*(7), 1451–1472. 4, 23

Guvenen, F. (2009). A parsimonious macroeconomic model for asset pricing. *Econometrica 77*(6), 1711–1750. 4, 24

Heathcote, J. and H. Tsujiyama (2021). Optimal income taxation: Mirrlees meets ramsey. *Journal of Political Economy 129*(11), 3141–3184. 23

Kapička, M. (2006). Optimal taxation and human capital accumulation. *Review of Economic Dynamics 9*(4), 612–639. 13

Karantounias, A. G. (2018, 01). Optimal Fiscal Policy with Recursive Preferences. *The Review of Economic Studies 85*(4), 2283–2317. 5

Kopczuk, W., E. Saez, and J. Song (2010, 02). Earnings Inequality and Mobility in the United States: Evidence from Social Security Data Since 1937*. *The Quarterly Journal of Economics 125*(1), 91–128. 22, 23

Lucas, Jr., R. E. and N. L. Stokey (1983). Optimal fiscal and monetary policy in an economy without capital. *Journal of Monetary Economics 12*, 55–93. 3, 5, 11

Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jørgensen (2009). Long-run stockholder consumption risk and asset returns. *The Journal of Finance 64*(6), 2427–2479. 24

Mankiw, N. and S. P. Zeldes (1991). The consumption of stockholders and nonstockholders. *Journal of Financial Economics 29*(1), 97–112. 4

Mehra, R. and E. C. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics 15*, 335–359. 24

Mertens, K. and J. L. M. Olea (2018). Marginal tax rates and income: New time series evidence. *Quarterly Journal of Economics 133*, 1803–1834. 4, 23, 27, 28, 29

Ogaki, M. and A. Atkeson (1997). Rate of time preference, intertemporal elasticity of substitution, and level of wealth. *The Review of Economics and Statistics 79*(4), 564–572. 4

Shiller, R. (1992). *Market Volatility* (1 ed.), Volume 1. The MIT Press. 24

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of Political Economy 110*(4), 825–853. 4, 23

Werning, I. (2007). Optimal fiscal policy with redistribution. *Quarterly Journal of Economics 22*(3), 925967. 5, 8

Appendix

A. Proofs

A1. Solution to the Examples in Section 2

Each of the following subsections formulates the planning problem (3) specifically for the conditions of the example and derives the results mentioned in Section 2.

Example 1. Let $\beta^t \pi(s^t)\lambda_t(s^t)$ be the Lagrange multiplier on the promise-keeping constraint (5), and $\beta^t \pi(s^t)\chi_t(s^t)$ be the Lagrange multiplier on the resource constraint. The Lagrangian is

$$\begin{aligned}
\mathcal{L} &= \lambda V_0^A(s_0) + (1-\lambda+\kappa)(1-\beta)\textstyle\sum_{t=0}^{\infty}\sum_{s^t}\beta^t U\left(c_t^B(s^t), n_t^B(s^t)\right)\pi(s^t) \\
&- \kappa(1-\beta)\textstyle\sum_{t=0}^{\infty}\sum_{s^t}\beta^t U\left(c_t^A(s^t), \dfrac{n_t^A(s^t)}{\theta}\right)\pi(s^t) \\
&- \sum_{t=0}^{\infty}\sum_{s^t}\beta^t\lambda_t(s^t)\left[V_t^A(s^{t-1}) - (1-\beta)U\left(c_t^A(s^t), n_t^A(s^t)\right) - \beta V_{t+1}^A(s^t)\right]\pi(s^t) \\
&- (1-\beta)\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\chi_t(s^t)\left[\alpha\left(c_t^A(s^t) - Z_t(s^t)n_t^A(s^t)\right) + (1-\alpha)\left(c_t^B(s^t) - Z_t(s^t)\theta n_t^B(s^t)\right)\right].
\end{aligned}$$

The first-order condition in V_0^A gives $\lambda_0(s_0) = \lambda$, the first property in the text. The first-order condition in V_t^A gives directly the second property (6). The no-distortion-at-the-top property is derived by combining the first-order conditions in c_t^B and n_t^B. The first-order conditions in c_t^A and n_t^A combined give

$$\left(\lambda_t(s^t) - \kappa\right) U_{ct}^A = \alpha\chi_t(s^t) \tag{22}$$

$$-U_{nt}^A(\lambda_t(s^t) - \kappa\phi) = \alpha\chi_t(s^t)Z_t(s^t), \tag{23}$$

where the second equation uses the functional form for the period utility function to simplify. Rearranging yields the optimal tax formula (4).

Example 2. Denote the Lagrange multiplier on the promise-keeping constraints (7) as $\beta^{t+1}\pi(s^t)\zeta_t(s^t)$. After rearranging the terms involving the promise-keeping constraints,

write the Lagrangian as

$$\begin{aligned}\mathcal{L} &= (1-\lambda+\kappa)(1-\beta)\sum_{t=0}^{\infty}\sum_{s^t}\beta^t U\left(c_t^B(s^t), n_t^B(s^t)\right)\pi(s^t)\\ &+ [\lambda-\beta\zeta_0(s_0)]U\left(c_0^B(s_0), n_0^B(s_0)\right) + \sum_{t=0}^{\infty}\sum_{s^t}\beta^t\left[\zeta_{t-1}(s^{t-1})-\beta\zeta_t(s^t)\right]U\left(c_t^A(s^t), n_t^A(s^t)\right)\pi(s^t)\\ &- \kappa(1-\beta)\sum_{t=0}^{\infty}\sum_{s^t}\beta^t U\left(c_t^A(s^t), \frac{n_t^A(s^t)}{\theta}\right)\pi(s^t)\\ &- (1-\beta)\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\chi_t(s^t)\left[\alpha\left(c_t^A(s^t)-Z_t(s^t)n_t^A(s^t)\right)+(1-\alpha)\left(c_t^B(s^t)-Z_t(s^t)\theta n_t^B(s^t)\right)\right].\end{aligned}$$

Define the implicit Pareto weights by

$$\begin{aligned}\lambda_0(s_0) &= \frac{\lambda-\beta\zeta_0(s_0)}{1-\beta}\\ \lambda_t(s^t) &= \frac{\zeta_{t-1}(s^{t-1})-\beta\zeta_t(s^t)}{1-\beta}\quad \forall s^t.\end{aligned}$$

Summing for each path and canceling terms yields (8). The first-order conditions in c_t^B and n_t^B again produce the no-distortion-at-the-top result. Using the definition of $\lambda_t(s^t)$, the first-order conditions in c_t^A and n_t^A are identical to (22) and (23). Thus, the optimal tax formula (4) applies again.

Example 3. The fact that V^A is constant across time and states simplifies the planning problem considerably. Define the consumption function $C(u,n) = U^{-1}(u,n)$. Using the fact that V^A is also a period utility of type A, write the Lagrangian as

$$\begin{aligned}\mathcal{L} &= \lambda V^A + (1-\lambda+\kappa)(1-\beta)\mathbb{E}\sum_{t=0}^{\infty}\beta^t U\left(c_t^B(s^t), n_t^B(s^t)\right)\\ &- \kappa(1-\beta)\mathbb{E}\sum_{t=0}^{\infty}\beta^t U\left(C\left(V^A, n_t^A(s^t)\right), \frac{n_t^A(s^t)}{\theta}\right)\\ &- (1-\beta)\sum_{t=0}^{\infty}\sum_{s^t}\beta^t\chi_t(s^t)\left[\alpha\left(C(V^A, n_t^A(s^t))-Z_t(s^t)n_t^A(s^t)\right)+(1-\alpha)\left(c_t^B(s^t)-Z_t(s^t)\theta n_t^B(s^t)\right)\right].\end{aligned}$$

The first-order conditions in n_t^A, c_t^B and n_t^B are

$$\begin{aligned} -U_{nt}^A \kappa (1-\phi) &= \alpha \chi_t \left(\frac{U_{nt}^A}{U_{ct}^A} + Z_t \right) \\ (1-\lambda+\kappa) U_{ct}^B &= (1-\alpha)\chi_t \\ -(1-\lambda+\kappa) U_{nt}^B &= (1-\alpha)\chi_t Z_t \theta. \end{aligned}$$

Together with the resource constraint, they determine the optimal allocation. Since nothing in these four equations depends directly on s^t, the solution takes the form $n_t^i(s^t) = \bar{n}^i[Z_t(s^t), G_t(s^t)]$ and $c_t^i(s^t) = \bar{c}^i[Z_t(s^t), G_t(s^t)]$ for some functions $\bar{n}^i$ and $\bar{c}^i$, with the dependence on κ and V^A being implicit. Define the implicit Pareto weight by[11]

$$\lambda_t(s^t) = (1-\lambda+\kappa)\frac{\alpha}{1-\alpha}\frac{U_{ct}^B(s^t)}{U_{ct}^A(s^t)} + \kappa. \tag{24}$$

Combining the first two first-order conditions and rearranging, one obtains that the optimal tax formula (4) again holds for λ defined by (24).

A2. Remaining Proofs

Proof of Proposition 1. To simplify notation, define

$$\lambda = \frac{\lambda^A}{\lambda^A + \lambda^B - \hat{\lambda}^B}, \quad \kappa = \frac{\hat{\lambda}^B}{\lambda^A + \lambda^B - \hat{\lambda}^B}.$$

Then, the proposition can be restated as

$$W(\lambda^A, \lambda^B, \hat{\lambda}^B, z) = (\lambda^A + \lambda^B - \hat{\lambda}^B) W\left(\lambda, 1-\lambda+\kappa, \kappa, z\right).$$

[11] It can be shown that $\lambda_t(s^t)$ from Examples 1 and 2 also satisfies this relationship.

Differentiating the equality with respect to the first three arguments yields

$$\begin{aligned}
W_{\lambda^A}(\lambda^A,\lambda^B,\kappa,z) &= W + (1-\lambda)W_{\lambda^A} - (1-\lambda+\kappa)W_{\lambda^B} - \kappa W_{\hat{\lambda}^B}\\
W_{\lambda^B}(\lambda^A,\lambda^B,\kappa,z) &= W - \lambda W_{\lambda^A} + (\lambda-\kappa)W_{\lambda^B} - \kappa W_{\hat{\lambda}^B}\\
W_{\hat{\lambda}^B}(\lambda^A,\lambda^B,\kappa,z) &= -W + \lambda W_{\lambda^A} + (1-\lambda+\kappa)W_{\lambda^B} + (1+\kappa)W_{\hat{\lambda}^B}\\
W_z(\lambda^A,\lambda^B,\kappa,z) &= (\lambda^A+\lambda^B-\hat{\lambda}^B)W_z\\
W_{zz}(\lambda^A,\lambda^B,\kappa,z) &= (\lambda^A+\lambda^B-\hat{\lambda}^B)W_{zz},
\end{aligned}$$

where the arguments on the right-hand side, λ, $1-\lambda+\kappa$, κ and z, are omitted for brevity.

Let $R(\cdot)$ ber the expression defined by the right-hand side of the HJB equation (16). Substituting in the above expressions and rearranging yields

$$\begin{aligned}
R = (\lambda^A+\lambda^B-\hat{\lambda}^B)\Bigg\{ & \lambda\left[F^A\left(U^A,\nu^A\right) - \nu^A\left(W + (1-\lambda)W_{\lambda^A} - (1-\lambda+\kappa)W_{\lambda^B} - \kappa W_{\hat{\lambda}^B}\right)\right]\\
& + (1-\lambda+\kappa)\left[F^B\left(U^B,\nu^B\right) - \nu^B\left(W - \lambda W_{\lambda^A} + (\lambda-\kappa)W_{\lambda^B} - \kappa W_{\hat{\lambda}^B}\right)\right]\\
& - \kappa\left[F^B\left(\hat{U}^B,\hat{\nu}^B\right) + \hat{\nu}\left(-W + \lambda W_{\lambda^A} + (1-\lambda+\kappa)W_{\lambda^B} + (1+\kappa)W_{\hat{\lambda}^B}\right)\right]\\
& + \mu_z(z)W_z + \frac{\sigma_z(z)^2}{2}W_{zz}W_{zz}\Bigg\},
\end{aligned}$$

and $\lambda^A+\lambda^B-\hat{\lambda}^B$ cancels out, leaving R, and the HJB equation, expressed solely in terms of λ and κ. This verifies the guess. □

Proof of Proposition 2. The proof will use the guess and verify method. Guess that the discount rate is identical to both types and the deviating agent, $\nu^A = \nu^B = \hat{\nu}^B$. Then $d\lambda = d\kappa = 0$, and both λ and κ are constant over time. We will keep them implicit for notational simplicity and write the common discount rate as $\nu(z)$. Guess also that $c^i(z) = \bar{c}^i z$ and $n^i(z) = \bar{n}^i$ for some constants $\bar{c}^i$ and $\bar{n}^i$.

First, I verify that the discount rates are all equal. The first-order condition in ν^i can

be written as

$$F_\nu\left(U_t^i, \nu_t^i\right) = \frac{(U_t^i)^{1-\gamma}}{1-\gamma}\left(\frac{1-\gamma-(1-\rho)\frac{\nu_t^i}{\delta}}{\rho-\gamma}\right)^{-\frac{1-\gamma}{1-\rho}}$$
$$= V_t^i = \mathbb{E}_t\left[\int_{s=t}^{\infty} e^{-\int_t^s \nu_\epsilon^i d\epsilon}\delta\frac{(U_s^i)^{1-\gamma}}{1-\gamma}\left(\frac{1-\gamma-(1-\rho)\frac{\nu_s^i}{\delta}}{\rho-\gamma}\right)^{\frac{\gamma-\rho}{1-\rho}} ds\right].$$

for $i = A, B$, where the first equality replaces F_ν^i with its expression, the second one follows from the first-order condition, and the third one uses the definition of V_t^i. Given that $U_t^i = z_t \bar{c}^i \phi(\bar{n}^i)$, $\bar{c}^i\phi(\bar{n}^i)$ cancels from both sides to obtain

$$\frac{z_t^{1-\gamma}}{1-\gamma}\left(\frac{1-\gamma-(1-\rho)\frac{\nu_t^i}{\delta}}{\rho-\gamma}\right)^{-\frac{1-\gamma}{1-\rho}} = \delta\mathbb{E}_t\left[\int_{s=t}^{\infty} e^{-\int_t^s \nu_\epsilon^i d\epsilon}\frac{z_s^{1-\gamma}}{1-\gamma}\left(\frac{1-\gamma-(1-\rho)\frac{\nu_s^i}{\delta}}{\rho-\gamma}\right)^{\frac{\gamma-\rho}{1-\rho}} ds\right]$$

The condition is independent of i, implying that $\nu^A = \nu^B$. Analogous reasoning establishes that $\nu^B = \hat{\nu}^B$.

Given that discount rates are equal across types, the instantaneous-to-lifetime utility ratio must also be equal across types. That is,

$$\frac{(U^i)^{1-\gamma}}{(1-\gamma)V^i} = \frac{(\hat{U}^B)^{1-\gamma}}{(1-\gamma)\hat{V}^B} = X(z), \quad i = A, B,$$

for some common ratio $X(z)$.

The rest of the proof shows that, given those results, the assumed policy functions indeed solve the system of the first-order conditions. Using $F_u^i = \delta X(z)^{\frac{\gamma-\rho}{1-\gamma}}(U^i)^{-\gamma}$ and $\hat{F}_u^B = \delta X(z)^{\frac{\gamma-\rho}{1-\gamma}}(\hat{U}^B)^{-\gamma}$, the first-order conditions (20) and (21) can be written, after elim-

inating τ^A and τ^B and canceling terms that involve X, as

$$1+\frac{\phi'(\bar{n}^A)}{\phi(\bar{n}^A)}\bar{c}^A=\frac{\kappa\left(\frac{\phi(\bar{n}^A/\theta)}{\phi(\bar{n}^A)}\right)^{-\gamma}\left(\frac{\phi(\bar{n}^A/\theta)}{\phi(\bar{n}^A)}-\frac{\phi'(\bar{n}^A/\theta)}{\phi'(\bar{n}^A)}\frac{1}{\theta}\right)}{\lambda-\kappa\left(\frac{\phi(\bar{n}^A/\theta)}{\phi(\bar{n}^A)}\right)^{-\gamma}\frac{\phi'(\bar{n}^A/\theta)}{\phi'(\bar{n}^A)}\frac{1}{\theta}}$$

$$1+\frac{\phi'(\bar{n}^B)}{\phi(\bar{n}^B)}\frac{\bar{c}^B}{\theta}=0$$

$$\lambda-\kappa\left(\frac{\phi(\bar{n}^A/\theta)}{\phi(\bar{n}^A)}\right)^{1-\gamma}=\frac{\alpha}{1-\alpha}(1-\lambda)\left(\frac{\bar{c}^B}{\bar{c}^A}\right)^{-\gamma}\left(\frac{\phi(\bar{n}^B)}{\phi(\bar{n}^A)}\right)^{1-\gamma}$$

In addition, after dividing both sides bt z_t, the resource constraint can be written in terms of $\bar{c}^i$ and $\bar{n}^i$ as well:

$$\alpha\bar{c}^A+(1-\alpha)\bar{c}^B+\bar{G}\leq\alpha\bar{n}^A+(1-\alpha)\theta\bar{n}^B,$$

where the assumption that $G(z)=\bar{G}z$ was used. The four equations depend only on $\bar{c}^i$ and $\bar{n}^i$ for $i=A,B$, and can be solved for those values. This verifies the guess.

□

B. Distorted First Best

I compute a modification of the first best allocation, where the consumption-labor margin is distorted by a constant intertemporal wedge τ^A. The optimum is charazterized by the value function $w(\lambda,z)$, consumption and labor allocation $c^i(\lambda,z)$ and $n^i(\lambda,z)$, and discount factors $v^i(\lambda,z)$ for $i=A,B$. These functions satisfy the HJB equation, which is just (19) with $\kappa=0$.

$$0=\lambda F^A\left(U(c^A,n^A),v^A\right)+(1-\lambda)F^B\left(U(c^B,n^B),v^B\right)-\left[\lambda v^A+(1-\lambda)v^B\right]w$$
$$+\lambda(1-\lambda)\left(v^B-v^A\right)w_\lambda+\mu_z(z)w_z+\frac{\sigma_z(z)^2}{2}w_{zz}.$$

They also satisfy the resource constraint:

$$\alpha c^A+(1-\alpha)c^B=z\left[\alpha n^A+(1-\alpha)\theta n^B\right].$$

The functions also satisfy the first-order conditions in consumption and hours, dis-

torted by the presence of the intratemporal wedge:

$$\frac{\lambda}{\alpha}F_u^A\left(U(c^A,n^A),\nu^A\right)U_c(c^A,n^A)=\frac{1-\lambda}{1-\alpha}F_u^B\left(U(c^B,n^B),\nu^B\right)U_c(c^B,n^B) \tag{25}$$

$$-U_n(c^A,n^A)=(1-\tau^A)zU_c(c^A,n^A) \tag{26}$$

$$-U_n(c^B,n^B)=(1-\tau^B)z\theta U_c(c^B,n^B). \tag{27}$$

Finally, the functions satisfy the first-order conditions in the discount factor:

$$F_\nu^A\left(U(c^A,n^A),\nu^A\right)=w+(1-\lambda)\partial_\lambda w$$
$$F_\nu^B\left(U(c^B,n^B),\nu^B\right)=w-\lambda\partial_\lambda w.$$

In addition, λ follows the differential equation $d\lambda=\lambda\mu_\lambda$, where the drift is $\mu_\lambda=(1-\lambda)(\nu^B-\nu^A)$.

Decentralization with Complete Markets. The equilibrium stochastic discount factor ζ is a diffusion dependent on the states λ and z. It is written as

$$\frac{d\zeta}{\zeta}=-r(\lambda,z)dt-\beta(\lambda,z)dB$$

for some functions $r(\lambda,z)$ and $\beta(\lambda,z)$. I will define the agent's problem in the complete-markets equilibrium and relate the risk-free rate r and the market price of risk β to the solution of the distorted planning problem above.

Agent problem. Type i agent starts with initial wealth W_0^i and maximizes the expected utility

$$\max_{c_t,n_t}\min_{\nu_t}\mathbb{E}_0\left[\int_0^\infty e^{-\int_0^t\nu_\epsilon d\epsilon}F^i\left(U(c_t,n_t),\nu_t\right)\,dt\right]$$

subject to the present value budget constraint

$$\mathbb{E}_0\int_t^\infty\zeta_t c_t dt=\zeta_0W_0^i+(1-\tau^i)\mathbb{E}_0\int_0^\infty\zeta_t\theta^i n_t\,dt,$$

where $\theta^A=1$ and $\theta^B=\theta$, and the labor income is taxed at a constant rate τ^i.

Let $m_t^i=F_{ut}^iU_{ct}^i$ be the marginal utility of consumption. It follows a diffusion process

$dm_t^i = \mu_m^i m_t^i dt + \sigma_m^i m_t^i dB_t$. The first-order condition in consumption is

$$e^{-\int_0^t \nu_\epsilon^i d\epsilon} m_t = v^i \zeta_t,$$

where v^i is the Lagrange multiplier on the budget constraint. Differentiating the first-order condition in consumption and using Ito's lemma yields

$$d \ln m_t^i = \nu_t^i dt + d \ln \zeta_t = \left(\nu_t^i - r_t - \frac{1}{2} \beta_t^2 \right) dt - \beta_t B_t. \tag{28}$$

Furthermore, since $m_t = m(\lambda_t, z_t)$, differentiate it to obtain

$$d \ln m_t^i = \left(\mu_m^i - \frac{1}{2} (\sigma_m^i)^2 \right) dt + \sigma_m^i dB, \quad \mu_m^i = \frac{m_\lambda^i}{m^i} \lambda \mu_\lambda + \frac{m_z^i}{m^i} \mu_z + \frac{1}{2} \frac{m_{zz}^i}{m^i} \sigma_z^2, \quad \sigma_m^i = \frac{m_z^i}{m^i} \sigma_z.$$

Equating coefficients, the agent's problem implies

$$r = \nu^A - \mu_m^A = \nu^B - \mu_m^B, \quad \beta = -\sigma_m^A = -\sigma_m^B.$$

Since the values of μ_m^A, μ_m^B, σ_m^A, and σ_m^B are computed directly from the planning problem (which can also be used to verify that the second equality in each of the two terms holds), the equations above can be used directly to compute the risk-free rate r and the market price of risk β.