

Capital Taxation*

Quantitative Explorations of the Inverse Euler Equation

Emmanuel Farhi

Harvard University
efarhi@harvard.edu

Iván Werning

MIT
iwerning@mit.edu

September 25, 2011 (13:16)

Abstract

Economies with private information provide a rationale for capital taxation. In this paper we ask what the welfare gains from following this prescription are. We develop a method to answer this question in standard general equilibrium models with idiosyncratic uncertainty and incomplete markets. We find general equilibrium to be important and greatly reduce the computed welfare gains. Once these effects are taken into account the gains are relatively modest.

* Farhi is grateful for the hospitality of the University of Chicago. Werning is grateful for the hospitality of Harvard University and the Federal Reserve Bank of Minneapolis during the time this draft was completed. We thank Florian Scheuer for outstanding research assistance. We thank comments and suggestions from Daron Acemoglu, Manuel Amador, Marios Angeletos, Patrick Kehoe, Narayana Kocherlakota, Mike Golosov, Ellen McGrattan, Robert Shimer, Aleh Tsyvinski and seminar participants at the Federal Reserve Bank of Minneapolis, Chicago, MIT, Harvard, Urbana-Champaign, Carnegie-Mellon, Wharton and the Conference on the Macroeconomics of Imperfect Risk Sharing at the UC-Santa Barbara, Society for Economic Dynamics in Vancouver and the NBER Summer institute. Special thanks to Mark Aguiar, Pierre-Olivier Gourinchas, Fatih Guvenen, Dirk Krueger, Fabrizio Perri, Luigi Pistaferri and Gianluca Violante for enlightening exchanges on the available empirical evidence for consumption and income. All remaining errors are our own.

1 Introduction

Advanced countries use the tax system as a tool for social insurance. [Friedman \[1962\]](#) proposed a “negative income tax” where individuals with low enough income would owe a negative tax, and thus collect transfers. In particular, he advocated a simple system where a constant marginal tax on income is combined with a lump sum rebate—resulting in a linear tax function with a negative intercept [[Sheshinski, 1972](#)]. For a given distribution of before-tax income, such a tax reduces the dispersion in after-tax income, thus achieving a more equitable outcome. Mirrlees’ seminal work refined this idea by allowing for fully non-linear taxation of income.

These contributions focus on the taxation of labor income to provide social insurance. Is there a similar role to be played by capital taxation? In an important paper [Atkinson and Stiglitz \[1976\]](#), provided a negative answer. Under some conditions, they showed that the optimum taxes labor income only, leaving capital income untaxed. Their analysis rested on a model where individuals work and save, but do not face uncertainty about future earnings. A model that incorporates future uncertainty seems better suited for thinking about social insurance. This turns out to make an important difference. Indeed, when individuals do face uncertainty, the optimal insurance arrangement calls for taxing capital positively: constrained efficient allocations satisfy an Inverse Euler equation, as opposed to the agent’s standard intertemporal Euler equation, implying a positive capital tax [[Diamond and Mirrlees, 1977](#), [Rogerson, 1985](#), [Ligon, 1998](#), [Golosov et al., 2003](#)].

Hence, there is a strong theoretical case for positive capital taxation as part of an optimal social insurance system. The purpose of this paper is to evaluate the quantitative contribution of capital taxation towards welfare. One can view our exercise as a counterpart to Lucas’ well-known “Supply Side Economics”. [[Lucas, 1990](#)] computed the welfare gains from removing a linear tax on capital in a Ramsey model, where it is well known that capital should not be taxed [[Chamley, 1986](#), [Judd, 1985](#)]. In a similar spirit, here, we use a dynamic Mirrleesian model where capital should be taxed and compute the welfare gains of taxing capital optimally, starting from a situation without capital taxes. Our focus, as was Lucas’, is on the welfare implications of the capital tax itself, not those of the entire tax system.

A key input of our model is the uncertainty of earnings. Workers are affected by various shocks that contribute towards this uncertainty. For example, unemployment risk, health shocks, or outright disability. More generally, individuals may not be able to perfectly predict the path of their future skill or earning opportunities. Many papers find it fruitful to isolate a particular source of risk, such as unemployment [e.g. [Hopenhayn and](#)

Nicolini, 1997, Shimer and Werning, 2008] or disability [e.g. Diamond and Mirrlees, 1977, Golosov and Tsyvinski, 2006], and study the optimal design of an insurance program that takes aim at this particular risk. However, for the purposes of evaluating the overall benefits from taxing capital, focusing on a particular risk would be a mistake. Instead, it is important to capture uncertainty more broadly, in a way that reflects estimates for overall earnings processes from panel data, not particular sources of risk. This is consistent with a large incomplete-markets literature that has been quite successful modeling consumption and wealth, taking as an input the estimated earnings process. We follow this tradition. Our method is flexible and allows us to study such model economies. In so doing, we bridge the gap between a more theoretical New Dynamic Public Finance literature and this more applied incomplete-markets literature.

We lay down a Mirrleesian economy with neoclassical technology and idiosyncratic uncertainty. Agents consume and work, experiencing skill shocks that are private information. This means that feasible allocations must be incentive compatible. In this setting, as long as utility is additively separable, constrained efficient allocations satisfy the Inverse Euler equation. Because this condition is incompatible with the agents' standard Euler equation, saving distortions are needed and take the form of a positive tax on capital. Starting from an equilibrium where agents save freely, introducing these optimal taxes increases welfare. We are interested in computing and understanding the determinants of these welfare gains. These measures answer the question posed earlier, regarding the importance of capital taxation as a component of social insurance.

Two approaches are possible to quantify the magnitude of these welfare gains, and we pursue both. The first approach requires modeling the stochastic process for individual consumption directly. The second models consumption indirectly. It uses as a starting point equilibrium allocations from a competitive economy where agents are subject to idiosyncratic skill shocks and can save freely by accumulating a risk-free bond. The advantage of this approach is that direct empirical knowledge on the consumption process is more limited than that for income and wealth.

Following the first approach, we are able to obtain closed-form solutions when utility is logarithmic and consumption follows a geometric random walk. We use the simplicity of our closed-form solutions to derive some insights that illuminate the main determinants for welfare gains. In particular, our formulas highlight how the magnitude of optimal capital taxes, as well as the efficiency gain from these taxes, depend positively on the variance of consumption growth.

With a partial-equilibrium analysis, using a fixed interest rate, the welfare gains span a wide range, going from 0% to 10%, depending primarily on the variance of consump-

tion growth, for which empirical evidence is limited. Indeed, theoretically, the gains are potentially unbounded. This establishes that it is incorrect to simply presume that the welfare gains are necessarily small, underscoring the need for a quantitative assessment that uncovers their main determinants.

The large and wide range of welfare gains found in the partial equilibrium context lead us to consider a general equilibrium setting. We find that, general equilibrium effects can dramatically mitigate welfare gains. With a capital share of $1/3$, we show that welfare gains range from 0% to 0.25%, for plausible values of the variance in consumption growth.

Turning to the second approach, we adopt an Aiyagari-Bewley incomplete-market model. Individuals face idiosyncratic labor income risk. They can save in a risk-free asset, but cannot borrow. At a steady-state equilibrium the interest rate is constant and equal to the marginal product of capital. Although individual consumption fluctuates, the cross-sectional distribution of assets and consumption is invariant. Starting from this steady state, we introduce optimal saving distortions and compute the associated welfare gains, taking into account the resulting transitional dynamics. The solution is no longer closed form, but we show that it is sufficiently tractable to be computed numerically.

There are two important differences in this setting, relative to the geometric random-walk case where closed-form solutions were available. On the one hand, as is well known, agents are able to smooth consumption quite effectively in Aiyagari-Bewley models. This tends to minimize the variance of consumption growth, mitigating welfare gains from saving distortions. On the other hand, equilibrium consumption in these models is not well described by a geometric random walk. Indeed, a steady state, with a stable cross-sectional distribution for consumption, requires a mean reverting consumption process. We show that this tends to magnify potential welfare gains.

We compute the steady state equilibrium for a range of parameters for the income processes and the coefficient of relative risk aversion. For our baseline calibration, we find that welfare gains are relatively small, below 0.2%. Away from this baseline, we find that welfare gains increase with the coefficient of relative risk aversion and with the variance and persistence of the income process. However, unless one allows the planner to relax borrowing constraints, the welfare gains remain relatively modest.

Taken together these results suggest that, once general equilibrium considerations are taken into account, the welfare gains from optimal capital taxation, as prescribed by the Inverse Euler equation, are relatively modest.

On the methodological front, our strategy requires perturbing the consumption assignment and holding the labor assignment unchanged, while preserving incentive compatibility. The new allocation satisfies the Inverse Euler equation and delivers the same

utility while freeing up resources. The reduction in resources is our measure of welfare gains. By leaving the labor assignment unchanged, we focus on the gains from introducing savings distortions, without changing the incentive structure implicit in the labor assignment. In this way, we sidestep resolving the optimal trade-off between insurance and incentives.

There are several advantages to our approach. First, our exercise does not require specifying some components of the economy. In particular, no knowledge of individual labor assignment or the disutility of work function is required. In this way, the degree to which work effort responds to incentives is not needed. This robustness is important, since empirical knowledge of these elasticities remains incomplete. Indeed, our welfare gains depend only on the original consumption assignment, the utility function for consumption and technology. Second, our planning problem has the advantage of being tractable, even for rich specifications of uncertainty. In our view, having this flexibility is important for quantitative work. Finally, by focusing our attention on savings distortions in this way, the welfare gains we compute can be traced back to a few intuitive parameters, such as the variance of consumption growth, the coefficient of relative risk aversion and the concavity of the production function.

Related Literature. The question we address is largely unexplored. One reason is that it is difficult to solve dynamic economies with private information, except for very particular cases, such as shocks that are i.i.d. over time.¹ The optimal taxation literature based on models with private information [see [Golosov, Tsyvinski, and Werning, 2006](#), and the references therein] often characterizes the constrained efficient allocations and taxes, but rarely provides a quantitative analysis of the efficiency gains, especially not focusing on the contribution of savings distortions. Recent progress has been made applying a “first-order approach” to characterize and compute constrained efficient allocations when shocks are persistent [[Farhi and Werning, 2010](#), [Golosov et al., 2010](#)]. However this approach remains computationally challenging and has only been applied to date in partial equilibrium settings. Moreover, the focus in these papers is on characterizing the optimum, not evaluating the contribution from capital taxes. Our approach zeros in on the contribution of capital taxes. It is tractable enough to be carried out in general equilibrium.

Another difference is that much of the literature has proceeded by studying particular aspects of social insurance in isolation. In contrast, in our paper, for an evaluation of the benefits from saving distortions for social insurance, it is important to take a broader

¹ Other special cases that have been extensively explored are unemployment and disability insurance.

perspective. [Golosov and Tsyvinski \[2006\]](#) study disability insurance and [Shimer and Werning \[2008\]](#) study unemployment insurance. In both cases, the nature of the stochastic process for shocks allows for a low dimensional recursive formulation that is numerically tractable. [Golosov and Tsyvinski \[2006\]](#) provide a quantitative analysis of disability insurance. Disability is modeled as an absorbing negative skill shock. They calibrate their model and compute the welfare gains that can be reaped by moving from the most efficient allocation that satisfies free savings to the optimal allocation. They focus on logarithmic utility and report welfare gains of 0.5%. [Shimer and Werning \[2008\]](#) provide a quantitative analysis of unemployment insurance. They consider a sequential job search model, where a risk averse, infinitely lived worker samples wage offers from a known distribution. Regarding savings distortions, they show that with CARA utility allowing agents to save freely is optimal. With CRRA utility, savings distortions are optimal, but they find that the efficiency gains they provide are minuscule. As most quantitative exercises to date, both [Golosov and Tsyvinski \[2006\]](#) and [Shimer and Werning \[2008\]](#) are set in partial equilibrium settings with linear technologies.

Second, there is a vast literature studying and calibrating incomplete-market Aiyagari-Bewley economies to consumption, income and wealth data. These papers emphasize the role of consumers self-smoothing through the precautionary accumulation of risk-free assets. In most positive analyses, government policy is either ignored or else a simple transfer and tax system is included and calibrated to current policies. In some normative analyses, some reforms of the transfer system, such as the income tax or social security, are evaluated numerically [e.g. [Conesa and Krueger, 2005](#)]. Our paper bridges the gap between this literature and the optimal tax literature.

The notion of efficiency used in the present paper is often termed *constrained-efficiency*, because it imposes the incentive-compatibility constraints that arise from the assumed asymmetry of information. Another, very different, notion has exists for exogenously incomplete-market models [see [Geanakoplos and Polemarchakis, 1985](#)]. The idea is to take the available asset structure as given and ask whether individuals can change their trading positions in way that, while not individually optimal, collectively produces a Pareto improvement. This notion has been applied by [Davila, Hong, Krusell, and Rios-Rull \[2005\]](#) in the Aiyagari-Bewley context.

2 A Two-Period Economy

This section presents a simple example economy with two-periods and a linear technology. This allows us to introduce some key concepts quickly and fix some basic ideas that

are useful for the rest of the paper. The next sections generalize the model to an infinite horizon setting with general technologies.

In this example economy there are two periods and a continuum of ex-ante identical agents. They consume in both periods, but work only in the second period.² In the second period, they experience a labor productivity shock, the realization of which is independent across agents. The presence of uncertainty creates a role for insurance.

If productivity were publicly observable, then the first best would be attainable. The first best features perfect insurance in consumption, so that consumption is equalized across workers. It also requires workers with higher productivity to produce more.

Following the seminal work of [Mirrlees \[1971\]](#), we assume instead that productivity is private information and revealed to each worker at the beginning of the second period. With private information, the first best is no longer feasible, because it is not incentive compatible. There is now a tradeoff between insurance and incentives. The constrained efficient allocation, or second best, can be derived as the solution of a planning problem using the tools of mechanism design. The allocation now features distortions in labor and savings. It can be implemented using distortionary nonlinear taxes on labor and capital. Roughly speaking, one can think of labor distortions as arising from the attempts to provide insurance while trading off incentives. Savings distortions help spread these incentive efficiently across periods, by influencing consumption.

If agents were able to save freely, without saving distortions, then the standard Euler equation would hold. Instead, the constrained efficient allocation satisfies an Inverse Euler equation, which is incompatible with the Euler equation. It is this conflict between the Euler and Inverse Euler equations that requires the introduction of savings distortions.

How much do savings distortions contribute towards welfare? That is the question we focus on in this paper. That is, we wish to explore the quantitative importance of satisfying the Inverse Euler equation, instead of the Euler equation.

To make things more concrete, imagine that we start in a competitive equilibrium with only self-insurance, where agents face no taxes on labor or capital. Now consider moving to the constrained efficient allocation, resulting in welfare gains. No doubt, it does not make sense to attribute all these efficiency gains to the introduction of savings distortions. Indeed, the move from autarky to the constrained efficient allocation also involves the introduction of labor distortions and a resulting change in the labor allocation. In this paper, we are not interested in the entire welfare gains of moving from self-insurance to the constrained efficient allocation. Thus, we propose a measure that isolates the contri-

²It is easy to allow for work in the first period, but this feature is not essential. The example economy is meant to be the simplest setting to introduce basic conceptual points.

bution stemming from savings distortions.

Preferences. There are two periods $t = 0, 1$. Agents are ex ante identical. Consumption takes place in both periods, while work occurs only in period $t = 1$. Agents obtain utility

$$v = U(c_0) + \beta \mathbb{E} [U(c_1) - V(n_1; \theta)]. \quad (1)$$

where U is the utility function from consumption, V is the disutility function from effective units of labor (hereafter: labor for short) and \mathbb{E} is the expectations operator. Uncertainty is captured by an individual shock $\theta \in [\underline{\theta}, \bar{\theta}] \equiv \Theta$ that affects the disutility of effective units of labor, where Θ is an interval of \mathbb{R} . We will sometimes refer to θ as a skill shock. To capture the idea that uncertainty is idiosyncratic, we assume that a version of the law of large number holds so that for any function f on Θ , $\mathbb{E}[f]$ corresponds to the average of f across agents.

The utility function U is assumed increasing, concave and continuously differentiable. We assume that the disutility function V is continuously differentiable and that, for any $\theta \in \Theta$, the function $V(\cdot, \theta)$ is increasing and convex. We also assume the single crossing property: $\frac{\partial}{\partial n_1} V(n_1; \theta)$ is strictly decreasing in θ , so that a high shock θ indicates a low disutility from work.

Technology. We assume that technology is linear with and a rate of return on savings equal to q^{-1} and normalize labor productivity to 1. It is convenient to change variables and define an allocation by the triplet $\{u_0, u_1, n_1\}$ with $u_0 \equiv U(c_0)$, $u_1(\theta) \equiv U(c_1(\theta))$. The cost of an allocation is

$$k_0 = c(u_0) + q \mathbb{E} [c(u_1(\theta)) - n_1(\theta)].$$

where $c \equiv U^{-1}$ is the inverse of the utility function.

Incentive-Compatibility. The shock realizations are private information to the agent, so we must ensure that allocations are incentive compatible. By the revelation principle we can consider, without loss of generality, a direct mechanism where agents report their shock realization θ in period $t = 1$ and are assigned consumption and labor as a function of this report. As is well known, the single crossing property implies that incentive

compatibility is equivalent to the condition that n_1 be non-decreasing and

$$u_0 + \beta (u_1(\theta) - V(n_1(\theta), \theta)) = u_0 + \beta (u_1(\underline{\theta}) - V(n_1(\underline{\theta}), \underline{\theta})) + \beta \int_{\underline{\theta}}^{\theta} \frac{\partial V}{\partial \hat{\theta}}(n_1(\hat{\theta}); \hat{\theta}) d\hat{\theta}. \quad (2)$$

Perturbations. Substituting [equation \(2\)](#) into [equation \(1\)](#) implies that

$$u_0 + \beta u_1(\theta) = v + \beta V(n_1(\theta), \theta) + \beta \mathbb{E} \left[\int_{\tilde{\theta}}^{\theta} \frac{\partial V}{\partial \hat{\theta}}(n_1(\hat{\theta}); \hat{\theta}) d\hat{\theta} \right] \quad (3)$$

where the expectation is over $\tilde{\theta} \in \Theta$. For given v and $\{n_1\}$, the right hand side is a fixed function of $\theta \in \Theta$. Given a labor assignment, utility from consumption $u_0 + \beta u_1(\theta)$ is pinned down up to a constant because of the incentives required for $\{n_1\}$.

For any value of u_0 this equation can be seen as determining the value of $u_1(\theta)$, which then varies with u_0 in a parallel fashion, across $\theta \in \Theta$. Another way to see this is to fix a utility level v and labor assignment $\{n_1\}$. Then for a given a utility assignments $\{u_0, u_1\}$ that satisfy [equation \(1\)](#) and [equation \(2\)](#), all other feasible utility assignments $\{\tilde{u}_0, \tilde{u}_1\}$ can be obtained by the following perturbation:

$$\tilde{u}_0 = u_0 - \beta \Delta \text{ and } \tilde{u}_1(\theta) = u_1(\theta) + \Delta \text{ for all } \theta \in \Theta,$$

for some $\Delta \in \mathbb{R}$.

Efficiency and The Inverse Euler Equation. Using this perturbation, it follows that a utility assignment $\{u_0, u_1\}$ satisfying [equation \(1\)](#) and [equation \(2\)](#) minimizes the cost k_0 if and only if $\Delta = 0$ minimizes $c(u_0 - \beta \Delta) + q \mathbb{E} [c(u_1(\theta) + \Delta)]$ over Δ , holding $\{u_0, u_1\}$ fixed. The necessary and sufficient first order condition is

$$\frac{1}{U'(c(u_0))} = \frac{1}{\beta q^{-1}} \mathbb{E} \left[\frac{1}{U'(c(u_1(\theta)))} \right], \quad (4)$$

which is known as the *Inverse Euler equation*.

In contrast, when agents can save freely, the standard Euler equation

$$U'(c(u_0)) = \beta q^{-1} \mathbb{E} [U'(c(u_1(\theta)))] \quad (5)$$

must hold.³ Note that if the Inverse Euler equation holds, Jensen’s inequality implies that

$$U'(c(u_0)) < \beta q^{-1} \mathbb{E}[U'(c(u_1(\theta)))] , \quad (6)$$

as long as consumption at $t + 1$ is uncertain conditional on information at t . This shows that [equation \(4\)](#) is incompatible with the standard Euler equation [equation \(5\)](#). Thus, the optimum cannot be implemented by allowing agents to save freely at the technology’s rate of return.

An Intuition for the Optimality of Positive Savings Distortions. An intuition for the desirability of savings distortions is as follows. Start from a baseline allocation $\{u_0, u_1\}$ with free savings, so that the Euler equation holds. Consider now perturbing this allocation so that $\tilde{u}_0 = u_0 - \beta\Delta$ and $\tilde{u}_1(\theta) = u_1(\theta) + \Delta$. The pertubated allocation is incentive compatibility and delivers the same level of utility. For some $\Delta < 0$, it satisfies the Inverse Euler equation. The corresponding perturbed allocation has a lower cost than the baseline allocation.

Intuitively, although one is used to thinking of a risk-free asset as the natural savings tool, our analysis suggests that, in economies with private information, the proper tool is one that does not upset incentive compatibility. This is accomplished by an intertemporal tradeoff that amounts to risk-free saving, in the amount Δ , in terms of utility, instead of consumption. However, risk free savings in utility is equivalent to a saving in a hypothetical asset that is risky in terms of its consumption payoff $\Delta/U'(c_1(\theta))$ —it must pay out more whenever the marginal utility of consumption is low (i.e. when consumption is high). From a finance perspective, this is the definition of a risky return. As a result, the optimum involves a negative holding of such an asset. By borrowing today against good states of nature tomorrow, the planner increases today’s consumption relative to tomorrow’s. This explains why $\Delta < 0$ is optimal.

Why does this imply a positive tax on capital? Consumers face a conventional risk-free saving technology, instead of the ideal, but hypothetical, risky asset described above (which is risk-free in terms of utility, but risky in terms of consumption). Thus, taxes are needed to induce agents to make the correct consumption and saving choice. In particular, we need a positive tax on capital to induce the agent to raise today’s consumption

³If agents can save freely then an allocation must satisfy the Euler equation, in addition to the incentive compatibility constraint. Indeed, for a given labor assignment $\{n_1\}$ and utility level v , the utility assignments $\{u_0, u_1\}$ are uniquely determined by incentive compatibility, the requirement that the allocation deliver expected utility v and the Euler equation. That is, there exists a unique utility assignment that satisfies the Euler equation supporting $\{n_1\}$ and v . The converse is not generally true, although in practice we have found that this is not an issue for the allocations that we compute.




Figure 1: Efficiency gains as a function of n_1 for given v . The upper (green) curve imposes free savings. The lower (blue) curve implements the Inverse Euler equation with optimal savings distortions.

relative to tomorrow's, to match the planner optimum.

Efficiency Gains from Savings Distortions. We have established that allowing agents to save freely is costly. The goal of this paper is to explore the magnitude of this cost. The lower curve in Figure 1 represents the minimum cost k_0 , for a given v , as a function of the labor assignment $\{n_1\}$. Along this curve, the Inverse Euler equation holds. Similarly, the upper curve represents the cost of the allocation when agents can save freely. Along this curve, the Euler equation holds. The distance between these curves represents the cost of foregoing the optimal taxation of savings and allowing, instead, agents to save freely. Indeed, the vertical distance represents precisely the added cost, for a given v and $\{n_1\}$, of imposing the Euler equation instead of the Inverse Euler equation.

We now consider three distinct allocations: (a) self-insurance, where agents can save freely with the technological return q^{-1} and the output from their labor is untaxed; (c) where the planner is subject to the incentive compatibility constraints and the constraint that agents can save freely; and (b) the constrained optimum, which solves a planning problem subject to the incentive compatibility constraints.

Allocation (b) can be implemented by a combination of nonlinear taxes on labor and capital, allocation (c) requires a nonlinear labor tax, while allocation (a) features no taxes.

These allocations are displayed in Figure 1. Point (b) lies at the minimum of the upper curve, while (c) lies at the minimum of the lower curve. Point (a) is depicted on the far right of the upper curve to convey the idea that we expect labor to be higher since it is not distorted by taxation.

Suppose the economy starts at self-insurance at point (a), and moves towards the constrained efficient allocation, at point (c). In the figure, this can be seen as a jump from the upper curve to lower curve, from point (a) to (a') as well as a movement along the lower curve, from (a') to (c). An alternative is to decompose the change from (a) to (c) as a movement along the upper curve to point (c') and then a vertical switch to the lower curve, from (c') to (c).

The move from (a) to (c) combines the introduction of optimal labor and savings distortions. Therefore, it would be incorrect to attribute the entire gain between points (a) and (c) to the introduction of savings distortions. Instead, we propose interpreting the vertical distances between the two curves, either between (a) and (a') or between (c') and (c) as measuring the gains from savings distortions.

One might also consider the gains of moving from point (b) to point (c). This represents the gains of introducing savings distortions, when labor distortions are fully optimized in both cases. Note that the vertical distance between (c) to (c') provides an upper bound for the gains from (b) to (c). This provides another rationale for focusing on the vertical distance between the two curves.

Furthermore, in practice, the distance between (a) and (a') is greater than the distance between (c') and (c), because it features greater uncertainty in consumption, which is responsible for the difference between the Euler and Inverse Euler equations in the first place. Consumption inequality is greater at self-insurance because labor income is untaxed, resulting in greater after-tax income uncertainty.

A Numerical Example. To illustrate this analysis, we have compute the three allocations for a parametrized version of this economy. In particular we adopte logarithmic utility $U(c) = \log c$, and $V(n; \theta) = \kappa (n/\theta)^{1+\frac{1}{\varepsilon}}$ and set the Frisch elasticity to $\varepsilon = 1/2$. With only two periods, it is natural to interpret a period as half a working lifetime (20 years). Correspondingly, we set $\beta = (0.96)^{20}$. This also motivates our selection for the productivity distribution, which we take to be log normal with variance equal to 20×0.0161 .⁴ This example is meant to provide a rough illustration of our concepts, not a full-blown calibration.

The table collects our results for this economy. The first row shows that the gains of moving from autarky to the constrained efficient allocation are relatively large, equal to 3.08%. Note, however, that at the constrained efficient allocation there are both savings and labor distortions. Thus, it would be incorrect to attribute these gains to savings dis-

⁴The number 0.0161 corresponds to the variance of the innovations in the permanent component of income estimated by Storesletten, Telmer and Yaron (2004)

Policy Experiment	Efficiency gains
$a \rightarrow c$	3.08%
$b \rightarrow c$	0.33%
$c' \rightarrow c$	0.39%
$a \rightarrow a'$	1.38%

Table 1: Efficiency gains for an example economy. The gains are expressed as a percentage of aggregate consumption in autarky.

tortions that implement the Inverse Euler equation. In contrast, the other rows in the table provide measures of efficiency gains that may be attributed to savings distortions. In particular, the second row shows the gains from the optimum with free savings, to the constrained optimum where agents face optimal savings distortions. The gains in this case are significantly smaller, 0.33%. Also shown are the vertical distances, (c') to (c) and (a) to (a') , with gains of 0.39% and 1.38%. These numbers represent the gains, for a given labor allocation, of moving from the Euler to the Inverse Euler equation. As such, they capture one notion of the gains from savings distortions. Also note that, as claimed earlier, (c') to (c) is an upper bound on the distance between (b) and (c) . In this example it is also the case that the distance between (a) and (a') is an upper bound on the other two measures of efficiency gains from savings distortions, (b) to (c) and (c') to (c) .

Discussion. In this simple example economy, the allocations (a) , (a') , (b) , (c) and (c') are easily computed. In the more general infinite horizon setup with general technologies and arbitrary stochastic processes for productivity that we consider in the rest of the paper, this is no longer the case. In particular, computing numerically points (b) , (c) and (c') is out of reach in most cases.

In the rest of the paper, we develop a general method to compute theoretically and quantitatively, for a given level of utility and a given labor assignment, the efficiency gains that can be realized by moving from a utility assignment that satisfies the Euler equation to one that satisfies the inverse Euler equation. That is, we develop a method to analyze theoretically and numerically the vertical distance between the two curves represented in Figure 1 in very general economies with infinite horizon, arbitrary stochastic processes for idiosyncratic productivity shocks, and general technologies.

There are a number of advantages that come from approaching the efficiency gains from savings distortions through the vertical distance between the two curves in Figure 1. First, we have explained how this vertical distance is informative about more encompassing measures of efficiency gains that also allow for changes in the labor assignment $\{n_1\}$.

Second, our measure of efficiency gains can be computed by a simple perturbation

method. Indeed, consider the feasible allocation $\{u_0, u_1, n_1\}$ that satisfies the Euler equation. We can obtain the utility assignment $\{\tilde{u}_0, \tilde{u}_1\}$ that supports the same utility and labor assignment $\{n_1\}$ through a simple perturbation $\{u_0 - \beta\Delta, u_1 + \Delta\}$ of the baseline utility assignment $\{u_0, u_1\}$ while fixing the labor assignment $\{n_1\}$. Indeed, $\{\tilde{u}_0, \tilde{u}_1\}$ is the unique utility assignment within this class of perturbations that satisfies the Inverse Euler equation. The existence of these simple perturbations greatly facilitates the computation of the efficiency gains from moving from $\{u_0, u_1, n_1\}$ and $\{\tilde{u}_0, \tilde{u}_1, n_1\}$.

Third, given the baseline utility assignment $\{u_0, u_1\}$, no knowledge of either the labor assignment $\{n_1\}$ or the disutility of work V is needed in order to compute our measure of efficiency gains. In other words, one does not need to take a stand on how elastic work effort is to changes in incentives. More generally one does not need to take a stand on whether the problem is one of private information regarding skills or of moral hazard regarding effort, etc. This robustness is a crucial advantage since current empirical knowledge of these characteristics and parameters is limited and controversial. The vertical distance between the two curves in Figure 1 addresses precisely the question of whether the intertemporal allocation of consumption is efficient, without taking a stand on how correctly tradeoff between insurance and incentives has been resolved.⁵

Finally, as this paper will show, the magnitude of the efficiency gains that we compute is determined by some rather intuitive parameters: the relative variance of consumption changes, the coefficient of relative risk aversion and the concavity of the production function.

3 Infinite Horizon

In this section, we lay down our general environment. We then describe a class perturbations that preserve incentive compatibility. These perturbations serve as the basis for our method to compute efficiency gains.

⁵This robustness property, can be more formally described as follows. Consider the set $\Omega(\{u_0, u_1, n_1\}, v)$ of disutility functions V with the following properties: V is continuously differentiable; for any $\theta \in \Theta$, the function $V(\cdot, \theta)$ is increasing and convex; V has the single crossing property that $\frac{\partial}{\partial n_1} V(n_1; \theta)$ is strictly decreasing in θ and [equation \(2\)](#) holds. The set of perturbed allocations $\{\{u_0 - \beta\Delta, u_1 + \Delta\} | \Delta \in \mathbb{R}\}$ is the largest of allocations such that for all $V \in \Omega(\{u_0, u_1, n_1\}, v)$, incentive compatibility [\(2\)](#) and promise keeping [\(1\)](#) hold.

3.1 The Environment

We cast our model within a general Mirrleesian dynamic economy. Our formulation is closest to [Golosov, Kocherlakota, and Tsyvinski \[2003\]](#). This paper obtains the Inverse Euler equation in a general dynamic economy, where agents' privately observed skills evolve as a stochastic process.

Preferences. Our economy is populated by a continuum of agent types indexed by $i \in I$ distributed according to the measure ψ . Preferences generalize those used in [Section 2](#) and are summarized by the expected discounted utility

$$\sum_{t=0}^{\infty} \beta^t \mathbb{E}^i [U(c_t^i) - V(n_t^i; \theta_t^i)]$$

where \mathbb{E}^i is the expectations operator for type i .

Additive separability between consumption and leisure is a feature of preferences that we adopt because it is required for the arguments leading to the Inverse Euler equation.⁶

Idiosyncratic uncertainty is captured by an individual specific shock $\theta_t^i \in \Theta$, where as in [Section 2](#), Θ is an interval of the real line. These shocks affect the disutility of effective units of labor. We sometimes refer to them as skill shocks. The stochastic process for each individual θ_t^i is identically distributed within each type $i \in I$ and independently distributed across all agents. We denote the history up to period t by $\theta^{i,t} \equiv (\theta_0^i, \theta_1^i, \dots, \theta_t^i)$, and by π^i the probability measure on Θ^∞ corresponding to the law of the stochastic process θ_t^i for an agent of type i .

Given any function f on Θ^∞ , we denote the integral $\int f(\theta^{i,\infty}) d\pi(\theta^{i,\infty})$ using the expectation notation $\mathbb{E}^i[f(\theta^{i,\infty})]$ or simply $\mathbb{E}^i[f]$. Similarly, we write $\mathbb{E}^i[f(\theta^{i,\infty})|\theta^{i,t-1}]$, or simply $\mathbb{E}_{t-1}^i[f]$, for the conditional expectation of f given history $\theta^{i,t-1} \in \Theta^t$.

As in [Section 2](#), all uncertainty is idiosyncratic and we assume that a version of the law of large number holds so that for any function f on Θ^∞ , $\mathbb{E}^i[f]$ corresponds to the average of f across agents with type i .

To preview the use we will have for types $i \in I$, note that in our numerical implementation we will assume that skills follow a Markov process. We will consider allocations that result from a market equilibrium where agents save in a riskless asset. For this kind of economy, agent types are then initial asset holdings together with initial skill. The measure ψ captures the joint distribution of these two variables.

⁶ The intertemporal additive separability of consumption also plays a role. However, the intertemporal additive separability of work effort is completely immaterial: we could replace $\sum_{t=0}^{\infty} \beta^t \mathbb{E}[V(n_t; \theta_t)]$ with some general disutility function $\tilde{V}(\{n_t\})$.

It is convenient to change variables, translating consumption allocations into utility assignments $\{u_t^i(\theta^{i,t})\}$, where $u_t^i(\theta^{i,t}) \equiv U(c_t^i(\theta^{i,t}))$. This change of variable will make incentive constraints linear and render the planning problem that we will introduce shortly convex. Then, agents of type $i \in I$ with allocation $\{u_t^i\}$ and $\{n_t^i\}$ obtain utility

$$v^i \equiv \sum_{t=0}^{\infty} \beta^t \mathbb{E}^i \left[u_t^i(\theta^{i,t}) - V(n_t^i(\theta^{i,t}); \theta_t^i) \right] \quad (7)$$

Information and Incentives. The shock realizations are private information to the agent. We invoke the revelation principle to derive the incentive constraints by considering a direct mechanism. Agents are allocated consumption and labor as a function of the entire history of reports. The agent's strategy determines a report for each period as a function of the history, $\{\sigma_t^i(\theta^{i,t})\}$. The incentive compatibility constraint requires that truth-telling, $\sigma_t^{i,*}(\theta^{i,t}) = \theta_t^i$, be optimal:

$$\sum_{t=0}^{\infty} \beta^t \mathbb{E}^i [u_t^i(\theta^{i,t}) - V(n_t^i(\theta^{i,t}); \theta_t^i)] \geq \sum_{t=0}^{\infty} \beta^t \mathbb{E}^i [u_t^i(\sigma^{i,t}(\theta^{i,t})) - V(n_t^i(\sigma^{i,t}(\theta^{i,t})); \theta_t^i)] \quad (8)$$

for all reporting strategies $\{\sigma_t^i\}$ and all $i \in I$.

Technology. Let C_t and N_t represent aggregate capital, labor and consumption for period t , respectively. That is, letting $c \equiv U^{-1}$ denote the inverse of the utility function,

$$\begin{aligned} C_t &\equiv \int \mathbb{E}^i [c(u_t^i)] d\psi \\ N_t &\equiv \int \mathbb{E}^i [n_t^i] d\psi \end{aligned}$$

for $t = 0, 1, \dots$. In order to facilitate our efficiency gains calculations, it will prove convenient to index the resource constraints by e_t which represents the aggregate amount of resources that is being economized in every period. The resource constraints are then

$$K_{t+1} + C_t + e_t \leq (1 - \delta)K_t + F(K_t, N_t) \quad t = 0, 1, \dots \quad (9)$$

where K_t denotes aggregate capital. The function $F(K, N)$ is assumed to be homogenous of degree one, concave and continuously differentiable, increasing in K and N .

Two cases are of particular interest. The first is the neoclassical growth model, where $F(K, N)$ is strictly concave and satisfies Inada conditions $F_K(0, N) = \infty$ and $F_K(\infty, N) = 0$. In this case, we also impose $K_t \geq 0$. The second case has linear technology $F(K, N) =$

$N + (q^{-1} - 1)K$, $\delta = 0$ and $0 \leq q < 1$. One interpretation is that output is linear in labor with productivity normalized to one, and a linear storage technology with safe gross rate of return q^{-1} is available. Another interpretation is that this represents the economy-wide budget set for a partial equilibrium analysis, with constant interest rate $1 + r = q^{-1}$ and unit wage. Under either interpretation, we avoid corners by allowing capital to be negative up to the present value of future output, $K_{t+1} \geq -\sum_{s=1}^{\infty} q^s N_{t+s}$. This represents the natural borrowing limit and allows us to summarize the constraints on the economy by the single present-value condition

$$\sum_{t=0}^{\infty} q^t C_t \leq \sum_{t=0}^{\infty} q^t (N_t - e_t) + \frac{1}{q} K_0.$$

Feasibility. An allocation $\{u_t^i, n_t^i, K_t, e_t\}$ and utility profile $\{v^i\}$ is *feasible* if conditions (7)–(9) hold. That is, feasible allocations must deliver utility v^i to agent of type $i \in I$, and must be incentive compatible and resource feasible.

Free-Savings. For the purposes of this paper, an important benchmark is the case where agents can save, and perhaps also borrow, freely. This increases the choices available to agents, which adds further restrictions relative to the incentive compatibility constraints.

In this scenario, the government enforces labor and taxes as a function of the history of reports, but does not control consumption directly. Disposable after-tax income is $w_t n_t^i(\sigma^{i,t}(\theta^{i,t})) - T_t^i(\sigma^{i,t}(\theta^{i,t}))$.⁷ Agents face the following sequence of budget and borrowing constraints:

$$c_t^i(\theta^{i,t}) + a_{t+1}^i(\theta^{i,t}) \leq w_t n_t^i(\sigma^{i,t}(\theta^{i,t})) - T_t^i(\sigma^{i,t}(\theta^{i,t})) + (1 + r_t) a_t^i(\theta^{i,t-1}) \quad (10a)$$

$$a_{t+1}^i(\theta^{i,t}) \geq \underline{a}_{t+1}^i(\sigma^{i,t}(\theta^{i,t})) \quad (10b)$$

with a_0^i given. We allow the borrowing limits $\underline{a}_{t+1}^i(\theta^{i,t})$ to be tighter than the natural borrowing limits.

Agents with type i maximize utility $\sum_{t=0}^{\infty} \beta^t \mathbb{E}^i [u_t^i(\sigma^{i,t}(\theta^{i,t})) - V(n_t^i(\sigma^{i,t}(\theta^{i,t})); \theta_t^i)]$ by choosing a reporting, consumption and saving strategy $\{\sigma_t^i, c_t^i, a_{t+1}^i\}$ subject to the sequence of constraints (10), taking a_0^i and $\{n_t^i, T_t^i, w_t, r_t\}$ as given. A feasible allocation $\{u_t^i, n_t^i, K_t, e_t\}$ is part of a free-savings equilibrium if there exist taxes $\{T_t^i\}$, such that the optimum $\{\sigma_t^i, c_t^i, a_{t+1}^i\}$ for agent of type i with wages and interest rates given by

⁷A special case of interest is where the dependence of the tax on any history of reports $\hat{\theta}^{i,t}$, can be expressed through its effect on the history of labor $n^{i,t}(\hat{\theta}^{i,t})$. That is, when $T_t^i(\hat{\theta}^{i,t}) = T_t^{i,n}(n^{i,t}(\hat{\theta}^{i,t}))$ for some $T_t^{i,n}$ function.

$w_t = F_N(K_t, N_t)$ and $r_t = F_K(K_t, N_t) - \delta$ satisfies truth telling $\sigma_t^i(\theta^{i,t}) = \theta_t^i$ and generates the utility assignment $u_t^i(\theta^{i,t}) = U(c_t^i(\theta^{i,t}))$.⁸

At a free-savings equilibrium, the incentive compatibility constraints (8) are satisfied. The consumption-savings choices of agents impose additional further restrictions. In particular, a necessary condition is the intertemporal Euler condition

$$U'(c(u_t^i)) \geq \beta(1 + r_{t+1})\mathbb{E}_t^i[U'(c(u_{t+1}^i))], \quad (11)$$

with equality if $a_{t+1}^i(\theta^{i,t}) > \underline{a}_{t+1}^i(\theta^{i,t})$. Note that if the borrowing limits $a_{t+1}^i(\theta^{i,t})$ are equal to the natural borrowing limits, then the Euler equation (11) always holds with equality.

Efficiency. We say that the allocation $\{u_t^i, n_t^i, K_t, e_t\}$ and utility profile $\{v^i\}$ is dominated by the alternative $\{\tilde{u}_t^i, \tilde{n}_t^i, \tilde{K}_t, \tilde{e}_t\}$ and $\{\tilde{v}^i\}$, if $\tilde{v}^i \geq v^i$, $\tilde{K}_0 \leq K_0$, $e_t \leq \tilde{e}_t$ for all periods t and either $\tilde{v}^i > v^i$ for a set of agent types of positive measure, $\tilde{K}_0 < K_0$ or $e_t < \tilde{e}_t$ for some period t .

We say that a feasible allocation is *efficient* if it is not dominated by any feasible allocation. We say that an allocation is *conditionally efficient* if it is not dominated by a feasible with the same labor allocation $n_t^i = \tilde{n}_t^i$.

As explained in Section 2, allocations that are part of a free-savings equilibrium are not conditionally efficient. Conditionally efficient allocations satisfy a first order condition, the Inverse Euler equation, which is inconsistent with the Euler equation. Being part of a free-savings equilibrium therefore acts as a constraint on the optimal provision of incentives and insurance. Efficiency gains can be reaped by departing from free-savings.

3.2 Incentive Compatible Perturbations

In this section, we develop a class of perturbations of the allocation of consumption that preserve incentive compatibility. We then introduce a concept of efficiency, Δ -efficiency, that corresponds to the optimal use of these perturbations. Our perturbation set is large enough to ensure that every Δ -efficient allocation satisfies the Inverse Euler equation. Moreover, we show that Δ -efficiency and conditional efficiency are closely related concepts: Δ -efficiency coincides with conditional efficiency on allocations that satisfy some mild regularity conditions.

⁸Note that individual asset holdings are not part of this definition. This is convenient because asset holdings and taxes are indeterminate due to the usual Ricardian equivalence argument.

A Class of Perturbations. For any period t and history $\theta^{i,t}$ a feasible perturbation, of any baseline allocation, is to decrease utility at this node by $\beta\Delta^i$ and compensate by increasing utility by Δ^i in the next period for all realizations of θ_{t+1}^i . Total lifetime utility is unchanged. Moreover, since only parallel shifts in utility are involved, incentive compatibility of the new allocation is preserved. We can represent the new allocation as $\tilde{u}_t^i(\theta^{i,t}) = u_t^i(\theta^{i,t}) - \beta\Delta^i$, $\tilde{u}_{t+1}^i(\theta^{i,t+1}) = u_{t+1}^i(\theta^{i,t+1}) + \Delta^i$, for all θ_{t+1}^i .

This perturbation changes the allocation in periods t and $t+1$ after history $\theta^{i,t}$ only. The full set of variations generalizes this idea by allowing perturbations of this kind at all nodes:

$$\tilde{u}_t^i(\theta^t) \equiv u_t^i(\theta^t) + \Delta^i(\theta^{i,t-1}) - \beta\Delta^i(\theta^{i,t})$$

for all sequences of $\{\Delta^i(\theta^{i,t})\}$ such that $\tilde{u}_t^i(\theta^{i,t}) \in U(\mathbb{R}_+)$ and such that the limiting condition

$$\lim_{T \rightarrow \infty} \beta^T \mathbb{E}^i[\Delta^i(\sigma^{i,T}(\theta^{i,T}))] = 0$$

for all reporting strategies $\{\sigma_t^i\}$. This condition rules out Ponzi-like schemes in utility.⁹ By construction, the agent's expected utility, for *any* strategy $\{\sigma_t^i\}$, is only changed by a constant Δ_{-1}^i :

$$\sum_{t=0}^{\infty} \beta^t \mathbb{E}^i[\tilde{u}_t^i(\sigma^{i,t}(\theta^{i,t}))] = \sum_{t=0}^{\infty} \beta^t \mathbb{E}^i[u_t^i(\sigma^{i,t}(\theta^{i,t}))] + \Delta_{-1}^i. \quad (12)$$

It follows directly from [equation \(8\)](#) that the baseline allocation $\{u_t^i\}$ is incentive compatible if and only if the new allocation $\{\tilde{u}_t^i\}$ is incentive compatible. Note that the value of the initial shifter Δ_{-1}^i determines the lifetime utility of the new allocation relative to its baseline. Indeed, for any fixed infinite history $\bar{\theta}^{i,\infty}$ [equation \(12\)](#) implies that (by substituting the deterministic strategy $\sigma_t^i(\theta^{i,t}) = \bar{\theta}_t^i$)

$$\sum_{t=0}^{\infty} \beta^t \tilde{u}_t^i(\bar{\theta}^{i,t}) = \sum_{t=0}^{\infty} \beta^t u_t^i(\bar{\theta}^{i,t}) + \Delta_{-1}^i \quad \forall \bar{\theta}^{i,\infty} \in \Theta^\infty \quad (13)$$

Thus, ex-post realized utility is the same along all possible realizations for the shocks.¹⁰

Let $Y(\{u_t^i\}, \Delta_{-1}^i)$ denote the set of utility allocations $\{\tilde{u}_t^i\}$ that can be generated by

⁹ Note that the limiting condition is trivially satisfied for all variations with finite horizon: sequences for $\{\Delta_t^i\}$ that are zero after some period T . This was the case in the discussion of a perturbation at a single node and its successors.

¹⁰ The converse is nearly true: by taking appropriate expectations of [equation \(13\)](#) one can deduce [equation \(12\)](#), but for a technical caveat involving the possibility of inverting the order of the expectations operator and the infinite sum (which is always possible in a version with finite horizon and Θ finite). This caveat is the only difference between [equation \(12\)](#) and [equation \(13\)](#).

these perturbations starting from a baseline allocation $\{u_t^i\}$ for a given initial Δ_{-1}^i . This is a convex set.

Below, we show that these perturbation are rich enough to deliver the Inverse Euler equation. In this sense, they fully capture the characterization of optimality stressed by [Goloso et al. \[2003\]](#).

An allocation $\{u_t^i, n_t^i, K_t, e_t\}$ with utility profile $\{v^i\}$ is Δ -efficient if it is feasible and not dominated by another feasible allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{e}_t\}$ such that $\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, \Delta_{-1}^i)$.

Note that conditional efficiency implies Δ -efficiency, since both concepts do not allow for changes in the labor allocation. Under mild regularity conditions, the converse is also true. More precisely, in [Appendix A](#), we define the notion of *regular* utility and labor assignments $\{u_t^i, n_t^i\}$.¹¹ We then show that Δ -efficiency coincides with conditional efficiency on the class of allocations with regular utility and labor assignments. Indeed, given a regular utility and labor assignment $\{u_t^i, n_t^i\}$, the perturbations $Y(\{u_t^i\}, \Delta_{-1}^i)$ characterize all the utility assignments $\{\tilde{u}_t^i\}$ such that $\{\tilde{u}_t^i, n_t^i\}$ is regular and satisfies the incentive compatibility constraints (8).

Inverse Euler Equation. Building on [Section 2](#), we review briefly the Inverse Euler equation which is the optimality condition for any Δ -efficient allocation.

Proposition 1. *A set of necessary and sufficient conditions for an allocation $\{u_t^i, n_t^i, K_t, e_t\}$ to be Δ -efficient is given by*

$$c'(u_t^i) = \frac{q_t}{\beta} \mathbb{E}_t^i [c'(u_{t+1}^i)] \iff \frac{1}{U'(c(u_t^i))} = \frac{q_t}{\beta} \mathbb{E}_t^i \left[\frac{1}{U'(c(u_{t+1}^i))} \right]. \quad (14)$$

where $q_t = 1/(1 + r_t)$ and

$$r_t \equiv F_K(K_{t+1}, N_t) - \delta \quad (15)$$

is the technological rate of return.

A Δ -efficient allocation cannot allow agents to save freely at the technology's rate of return, since then [equation \(11\)](#) would hold as a necessary condition, which is incompatible with the planner's optimality condition, [equation \(14\)](#).

¹¹ Regularity is a mild technical assumption which is necessary to derive an Envelope condition crucial for our proof.

3.3 Efficiency Gains from Optimal Savings Distortions

In this section we consider a baseline allocation and consider an improvement on it that yields a Δ -efficient allocation. We define a metric for the efficiency gains from this improvement. This metric is our measure of the gains from the introduction of optimal savings distortions.

We define a planning problem to compute these gains. We then show how to solve this planning problem. In particular, we work with a relaxed planning problem that allows us to break up the solution into component planning problems. Moreover, we show how these problems may admit a recursive representation. Indeed, in the logarithmic case we find a closed-form solution. More generally, we show that each component planning problem is mathematically isomorphic to solving an income fluctuations problem where Δ_t plays the role of wealth.

If an allocation $\{u_t^i, n_t^i, K_t, e_t\}$ with corresponding utility profile $\{v^i\}$, is not Δ -efficient, then we can always find an alternative allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{e}_t\}$ that leaves utility unchanged, so that $\tilde{v}^i = v^i$ for all $i \in I$, but economizes on resources: $\tilde{K}_0 \leq K_0$ and $e_t \leq \tilde{e}_t$ with at least one strict inequality. In the rest of the paper, we restrict to cases where $e_t = 0$, $\tilde{K}_0 = K_0$ and $\tilde{e}_t = \tilde{\lambda}C_t$ for some $\tilde{\lambda} > 0$. We then take $\tilde{\lambda}$ as our measure of efficiency gains between these allocations. This measure represents the resources that can be saved in all periods in proportion to aggregate consumption.

We now introduce a planning problem that uses this metric to compute the distance of any baseline allocation from the Δ -efficient frontier. For any given baseline allocation $\{u_t^i, n_t^i, K_t, 0\}$, which is feasible with $e_t = 0$ for all $t \geq 0$, we seek to maximize $\tilde{\lambda}$ by finding an alternative allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda}C_t\}$ with $\tilde{K}_0 = K_0$,

$$\tilde{K}_{t+1} + \int \mathbb{E}^i[c(\tilde{u}_t^i)]d\psi + \tilde{\lambda}C_t \leq (1 - \delta)\tilde{K}_t + F(\tilde{K}_t, N_t) \quad t = 0, 1, \dots \quad (16)$$

and

$$\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, 0).$$

Let $\tilde{C}_t \equiv \int \mathbb{E}^i[c(\tilde{u}_t^i)]d\psi$ denote aggregate consumption under the optimized allocation. The optimal allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda}C_t\}$ in this program is Δ -efficient and saves an amount $\tilde{\lambda}C_t$ of aggregate resources in every period. Our measure of the distance of the baseline allocation from the Δ -efficient frontier is $\tilde{\lambda}$.

For future use, we denote the corresponding sequence of interest rates and intertemporal prices by $1 + \tilde{r}_t = 1 + F_K(\tilde{K}_t, N_t) - \delta$ and $\tilde{q}_t = 1/(1 + \tilde{r}_t)$.

In [Appendix B](#), we explain why this planning problem is numerically tractable, and

detail a method to solve it. The basic idea is to introduce a relaxed planning problem which replaces the resource constraints with a single present value condition for a given sequence of intertemporal prices $\{\tilde{q}_t\}$. The relaxed planning problem can then be further decomposed in a series of component planning problems corresponding to the different types $i \in I$.

In most situations, the baseline allocation admits a recursive representation for some endogenous state variable.¹² This is the case whenever θ_t^i is a Markov process and the baseline allocation depends on the history of shocks $\theta^{i,t-1}$ in a way that can be summarized by an endogenous state x_t^i , with law of motion $x_t^i = M(x_{t-1}^i, \theta_t^i)$ and given initial condition x_0^i (types then correspond to different initial values x_0^i). The endogenous state x_t^i is a function of the history of exogenous shocks $\theta^{i,t}$. Defining the state vector $s_t^i = (x_t^i, \theta_t^i)$ there must exist a function \bar{u} such that $u_t^i(\theta^{i,t}) = \bar{u}(s_t^i)$ for all $\theta^{i,t}$. In this case, the component planning problems can be boiled down to a simple Bellman equation with two state variables, s_t^i and Δ_{t-1}^i , which is amenable to numerical simulations.

4 Idiosyncratic and Aggregate Gains with Log Utility

In this section, we focus on the case of logarithmic utility. When utility is logarithmic, our parallel shifts in utility imply proportional shifts in consumption. We first show that the planning problem can be decomposed into an Idiosyncratic planning problem and a simple Aggregate planning problem. It also makes it possible to solve the idiosyncratic efficiency gains and the corresponding allocation in closed form when the baseline allocation is recursive and features constant consumption. In this case, the optimum in the planning problem and the efficiency gains $\tilde{\lambda}$ can be solved out almost explicitly by combining the solution of the Idiosyncratic problem with that of the Aggregate planning problem. We then illustrate our results in the simple benchmark case where the baseline allocation of consumption is a geometric random walk.

¹² The requirement that the baseline allocation be recursive in this way is hardly restrictive. Of course, the endogenous state and its law of motion depend on the particular economic model generating the baseline allocation. A leading example in this paper is the case of incomplete markets Bewley economies in [Huggett \[1993\]](#) and [Aiyagari \[1994\]](#). In these models, described in more detail in [Section 5](#), each individual is subject to an exogenous Markov process for income or productivity and saves using a riskless asset. At a steady state, the interest rate on this asset is constant, so that the agent's solution can be summarized by a stationary savings rule. The baseline allocation can then be summarized using asset wealth as an endogenous state, with law of motion M given by the agent's optimal saving rule. Another example are allocations generated by a dynamic contract. The state variable then includes the promised continuation utility [see [Spear and Srivastava, 1987](#)] along with the exogenous state.

4.1 A Decomposition: Idiosyncratic and Aggregate

Idiosyncratic Efficiency Gains. The full planning problem maximizes over utility assignments and capital. It is useful to also consider a version of the problem that takes the baseline sequence of capital $\{K_t\}$ as given. Thus, define the *idiosyncratic planning problem* as maximizing λ^I subject to

$$\int \mathbb{E}^i[c(\hat{u}_t^i)]d\psi + \tilde{\lambda}^I C_t = C_t \quad t = 0, 1, \dots$$

and $\{\hat{u}_t^i\} \in Y(\{u_t^i\}, 0)$. The idiosyncratic efficiency gains λ^I represent the constant proportional reduction in consumption that is possible without changing the aggregate sequence of capital. Of course, the total efficiency gains are larger than the idiosyncratic ones: $\tilde{\lambda} \geq \tilde{\lambda}^I$.

The solution of the idiosyncratic planning problem improves over the baseline allocation by ensuring that the marginal rates of substitution corresponding to the Inverse Euler equation $\mathbb{E}_t^i[c'(\hat{u}_{t+1}^i) / (\beta c'(\hat{u}_t^i))]$ are equalized across types and histories in every period. These marginal rates of substitution, however, are not necessarily linked to any technological rate of transformation as in [equation \(15\)](#). The idiosyncratic efficiency gains thus correspond to the gains from equalizing the marginal rate of substitution across types and histories in every period, without changing the sequence of capital.

The aggregate efficiency gains then capture the additional benefits from altering the aggregate allocation to equalize these marginal rates of substitution with the marginal rate of transformation in every period.

Aggregate Efficiency Gains. Given $\tilde{\lambda}^I \in [0, 1)$, the Aggregate planning problem seeks to determine the aggregate allocation $\{\tilde{C}_t, N_t, \tilde{K}_t, \tilde{\lambda}^I C_t\}$ that maximizes $\tilde{\lambda}^A$, subject to $\tilde{K}_0 = K_0$,

$$\tilde{K}_{t+1} + \tilde{C}_t + (\tilde{\lambda}^A + \tilde{\lambda}^I)C_t \leq (1 - \delta)\tilde{K}_t + F(\tilde{K}_t, N_t) \quad t = 0, 1, \dots$$

and

$$\sum_{t=0}^{\infty} \beta^t U(\tilde{C}_t) = \sum_{t=0}^{\infty} \beta^t U(C_t(1 - \tilde{\lambda}^I))$$

We refer to $\tilde{\lambda}^A$ as aggregate efficiency gains.

Proposition 2. Consider a baseline allocation $\{u_t^i, n_t^i, K_t, 0\}$ such that $\sum_{t=0}^{\infty} \beta^t U(C_t)$ is well defined and finite. Suppose that the optimum in the Idiosyncratic planning problem are such that the resource constraints hold with equality. Then the total efficiency gains underlying the planning problem are given by the sum of the idiosyncratic and aggregate efficiency gains: $\tilde{\lambda} = \tilde{\lambda}^I + \tilde{\lambda}^A$.

Proof. See [Appendix D](#). ■

The proof of [Proposition 2](#) also establishes that the utility assignment $\{\hat{u}_t^i\}$ that solves the Idiosyncratic planning problem and the utility assignment $\{\tilde{u}_t^i\}$ that solves the original planning problem are related by $\tilde{u}_t^i = \hat{u}_t^i + \delta_t$ where $\delta_t = U(\tilde{C}_t) - U(C_t(1 - \tilde{\lambda}^I))$.

The analysis of the evolution of the aggregate allocation $\{\tilde{C}_t, \tilde{K}_t\}$ only requires the knowledge of $\tilde{\lambda}^I$, but can otherwise be conducted separately from the analysis of the idiosyncratic problem. The aggregate planning problem is simply that of a standard deterministic growth model, which, needless to say, is straightforward to solve. For example, suppose that the baseline allocation represents a steady state with constant aggregates, $C_t = C_{ss}$, $K_t = K_{ss}$ and $N_t = N_{ss}$. Then the optimized aggregate allocation $\{\tilde{C}_t, \tilde{K}_t\}$ converges to a steady state $\{\tilde{C}_{ss}, \tilde{K}_{ss}\}$ such that $1 - \delta + F_K(\tilde{K}_{ss}, N_{ss}) = 1/\beta$ and $\tilde{C}_{ss} = F(\tilde{K}_{ss}) - \delta\tilde{K}_{ss} - \tilde{\lambda}C_{ss}$. We will put that result to use in [Section 5](#).

The aggregate planning problem is as a version of the problem that Lucas studied in his famous “Supply Side Economics” exercise in [[Lucas, 1990](#)]. Both involve computing the welfare gains along the transition to a new steady state in the neoclassical growth model. There is however, an important difference. Indeed, imagine that the Euler equation holds at the baseline allocation. Then whereas Lucas’ exercise involves computing a transition to a steady state with more capital, our exercise, by contrast, involves a transition to a steady state with less capital. The reason is that he considers removing a capital tax in a model where capital should not be taxed, while we consider introducing optimal capital taxes in a model where capital should be taxed.

4.2 Example: Steady States with Geometric Random Walk

Although the main virtue of our approach is that we can flexibly apply it to various baseline allocations, in this section we begin with a simple and instructive case. We maintain the assumption of logarithmic utility throughout. We take the baseline allocation to be a geometric random walk: $c_{t+1} = \varepsilon_t c_t$ with ε_t i.i.d. across agents and over time. In the language of [Section 3.3](#), the baseline allocation is recursive: $s_{t+1} = \varepsilon_t s_t$ with ε_t i.i.d. and $c(s) = s$, so that $u(s) = U(s)$. Moreover, we assume that $\log(\varepsilon)$ is normally distributed with variance σ_ε^2 so that $\mathbb{E}[\varepsilon] \cdot \mathbb{E}[\varepsilon^{-1}] = \exp(\sigma_\varepsilon^2)$. We also assume that the baseline allocation represents a steady state with constant aggregates $C_t = C_{ss}$, $K_t = K_{ss}$ and $N_t = N_{ss}$ which require $\mathbb{E}[\varepsilon] = 1$. We define $r_{ss} = F_K(K_{ss}, N_{ss}) - \delta$ and $q_{ss} = 1/(1 + r_{ss})$. Moreover, we assume that the Euler equation holds at the baseline allocation, which requires $q_{ss} = \beta \mathbb{E}[\varepsilon^{-1}] = \beta \exp(\sigma_\varepsilon^2)$.

Although extremely stylized, a random walk is an important conceptual and em-

pirical benchmark. First, most theories—starting with the simplest permanent income hypothesis—predict that consumption should be close to a random walk. Second, some authors have argued that the empirical evidence on income, which is a major determinant for consumption, and consumption itself shows the importance of a highly persistent component [e.g. [Storesletten, Telmer, and Yaron, 2004a](#)]. For these reasons, a parsimonious statistical specification for consumption may favor a random walk.

The advantage is that we obtain closed-form solutions for the optimized allocation, the intertemporal wedge and the efficiency gains. The transparency of the exercise reveals important determinants for the magnitude of efficiency gains. A geometric random walk however, is special for the following reason. If we apply the decomposition of [Section 4.1](#), then the idiosyncratic efficiency gains are zero. The entirety of the efficiency gains are aggregate. This is because at the baseline allocation, the marginal rates of substitution $\mathbb{E}_t^i [c'(u_{t+1}^i)/(\beta c'(u_t^i))]$ are already equalized across types and histories to β^{-1} . These marginal rates of substitution are not equalized, however, to the marginal rate of transformation $1 - \delta + F_K(K_{ss}, N_{ss})$. Therefore, this section can be seen as an exploration of the determinants of aggregate efficiency gains.

An Example Economy. Indeed, one can construct an example economy where a geometric random walk for consumption arises as a competitive equilibrium with incomplete markets. To see this, suppose that individuals have logarithmic utility over consumption and disutility from labor $V(n; \theta) = v(n/\theta)$ for some convex function $v(n)$ over work effort n , so that θ can be interpreted as productivity. Skills evolve as a geometric random walk, so that $\theta_{t+1} = \varepsilon_{t+1}\theta_t$, where ε_{t+1} is i.i.d. Individuals can only accumulate a riskless asset paying return q^{-1} equal to the rate of return on the economy's linear savings technology. They face the sequence of budget constraints

$$a_{t+1} + c_t \leq q^{-1}a_t + \theta_t n_t \quad t = 0, 1, \dots$$

and the borrowing constraint that $a_t \geq 0$. Suppose the rate of return q^{-1} is such that $1 \geq \beta q^{-1} \mathbb{E}[\varepsilon^{-1}]$. Finally, suppose that individuals have no initial assets, so that $a_0 = 0$.

In the competitive equilibrium of this example individuals exert a constant work effort n_{ss} , satisfying $n_{ss}v'(n_{ss}) = 1$, and consume all their labor income each period, $c_t = \theta_t n_{ss}$; no assets are accumulated, a_t remains at zero. This follows because the construction ensured that the agent's intertemporal Euler equation holds at the proposed equilibrium consumption process. The level of work effort n_{ss} is defined so that the intra-period consumption-leisure optimality condition holds. Then, since the agent's problem is con-

vex, it follows that this allocation is optimal for individuals. Since the resource constraint trivially holds, it is an equilibrium.

Although this is certainly a very special example economy, it illustrates that a geometric random walk for consumption is a possible equilibrium outcome.

Partial Equilibrium: Linear Technology. We first study the case where the technology is linear with a rate of return $q^{-1} > 1$. This imposes $q = q_{ss} = \beta \exp(\sigma_\varepsilon^2)$. Note that $q < 1$ imposes, for a given discount rate β , an upper bound on the variance of the shocks $\exp(\sigma_\varepsilon^2) < \beta^{-1}$.

Since the idiosyncratic efficiency gains $\tilde{\lambda}^l$ are zero, the solution of the planning problem can be derived by studying the Aggregate planning problem, which takes a remarkably simple form. The aggregate consumption sequence $\{\tilde{C}_t\}$ that solves the Aggregate planning problem is given by

$$\tilde{C}_t = C_{ss} \exp\left(\frac{\beta}{1-\beta}\sigma_\varepsilon^2\right) \exp(-t\sigma_\varepsilon^2).$$

The efficiency gains $\tilde{\lambda}$ and the optimal utility assignment $\{\tilde{u}\}$ are then readily computed. We can also derive the intertemporal wedge τ that measures the savings distortions at the optimal allocation $U'(c(\tilde{u}(s^t))) = \beta(1-\tau)q^{-1}\mathbb{E}[U'(c(\tilde{u}(s^{t+1})))|s^t]$.

Proposition 3. *Suppose that utility is logarithmic and that the technology is linear. Suppose that the baseline allocation is a geometric random walk with constant aggregate consumption $C_t = C_{ss}$, that the shocks ε are lognormal with variance σ_ε^2 and that the Euler equation holds at the baseline. Then the consumption assignment of the solution of the planning problem is given by $\tilde{c}(s^t) = \exp\left(\frac{\beta}{1-\beta}\sigma_\varepsilon^2\right) \exp(-t\sigma_\varepsilon^2) c(s_t)$. The intertemporal wedge at the optimal allocation is given by $\tau = 1 - \exp(-\sigma_\varepsilon^2)$. The efficiency gains are given by $\tilde{\lambda} = 1 - \frac{\beta^{-1} - \exp(\sigma_\varepsilon^2)}{\beta^{-1} - 1} \exp\left(\frac{\beta}{1-\beta}\sigma_\varepsilon^2\right)$.*

The optimized allocation has a lower drift than the baseline allocation. Intuitively, our perturbations based on parallel shifts in utility can be understood as allowing consumers to borrow and save with an artificial idiosyncratic asset, the payoff of which is correlated with their baseline idiosyncratic consumption process: they can increase their consumption today by reducing their consumption tomorrow in such a way that they reduce their consumption tomorrow more in states where consumption is high than in states where consumption is low. The desirable insurance properties of these perturbation make them attractive, leading to a front-loading of consumption. In other words, because our perturbations allow for better insurance, they reduce the benefits of engaging in precautionary savings by accumulating a buffer stock of risk free assets. As a result, it is optimal to

front-load consumption, by superimposing a downward drift $\exp(-\sigma_\varepsilon^2)$ on the baseline allocation, where the variance in the growth rate of consumption σ_ε^2 indexes the strength of the precautionary savings motive at the baseline allocation.

In this example, the Inverse Euler equation provides a rationale for a constant and positive wedge $\tau = 1 - \exp(-\sigma_\varepsilon^2)$ in the agent's Euler equation. This is in stark contrast to the Chamley-Judd benchmark result, where no such distortion is optimal in the long run, so that agents are allowed to save freely at the social rate of return.

The efficiency gains are increasing in σ_ε^2 . Note that when $\sigma_\varepsilon^2 = 0$, there are no efficiency gains. For small values of σ_ε^2 , the wedge is given by $\tau \approx \sigma_\varepsilon^2$. The formula for the efficiency gains then takes the form of a simple Ramsey formula $\tilde{\lambda} \approx (\beta/(1-\beta)^2) \tau^2/2$. At the other extreme, as $\sigma_\varepsilon^2 \rightarrow -\log(\beta)$ the efficiency gains asymptote 100%. The reason is that then $q \rightarrow 1$, implying that the present value of the baseline consumption allocation goes to infinity; in contrast, the cost of the optimal allocation remains finite.

General Equilibrium: Concave Technologies. In this section, we maintain the assumption that utility is logarithmic. We also assume that the baseline allocation is a geometric random walk representing a steady state with constant aggregates, $C_t = C_{ss}$, $K_t = K_{ss}$ and $N_t = N_{ss}$, that the shocks ε are lognormally distributed and that the Euler equation holds at the baseline allocation. We depart from the partial equilibrium assumption of a linear technology, and consider instead the case of concave accumulation technologies. We argue that the efficiency effects may be greatly reduced. This point is not specific to the model or forces emphasized here. Indeed, a similar issue arises in the Ramsey literature, the quantitative effects of taxing capital greatly depend on the underlying technology.¹³ It is important to confront this issue to reach meaningful quantitative conclusions.¹⁴

The point that general equilibrium considerations are important can be made most clearly from the following example. We consider the extreme case of a constant endowment: the economy has no savings technology, so that $C_t \leq N_t$ for $t = 0, 1, \dots$ [Huggett, 1993]. Then the baseline allocation is Δ -efficient. This follows since one finds a sequence of intertemporal prices \tilde{q}_t such that the Inverse Euler equation (14) holds. Thus, in this exchange economy there are no efficiency gains from perturbing the allocation—no effi-

¹³ Indeed, Stokey and Rebelo [1995] discuss the effects of capital taxation in representative agent endogenous growth models. They show that the effects on growth depend critically on a number of model specifications. They then argue in favor of specifications with very small growth effects, suggesting that a neoclassical growth model with exogenous growth may provide an accurate approximation.

¹⁴ A similar point is at the heart of Aiyagari's [1994] paper, which quantified the effects on aggregate savings of uncertainty with incomplete markets. He showed that for given interest rates the effects could be enormous, but that the effects were relatively moderate in the resulting equilibrium of the neoclassical growth model.

ciency gains from saving distortions.¹⁵ Certainly the fixed endowment case is an extreme example, but it serves to illustrate that general equilibrium considerations are extremely important.

Consider now a neoclassical production function $F(K, N)$. Applying the results in Section 4.1, we can decompose the planning problem into a idiosyncratic planning problem and an aggregate planning problem, with a corresponding decomposition for efficiency gains. We have already argued that the corresponding efficiency gains are equal to zero $\lambda^I = 0$. Therefore, and just as in the partial equilibrium case, all the efficiency gains are aggregate. The aggregate problem is a simple modification of the neoclassical growth model. The solution involves a transition to a steady state with an interest rate equal to $\tilde{r}_{ss} = 1/\beta - 1$. The efficiency gains, which reflect the strength of the precautionary savings motive, depend how much lower than \tilde{r}_{ss} is the interest rate $r_{ss} = 1/(\beta \exp(\sigma_\epsilon^2))$ at the baseline allocation.¹⁶ This, in turn depends on the variance of consumption growth σ_ϵ^2 .

Empirical Evidence. Suppose one accepts the random walk specification of consumption as a useful approximation. What does the available empirical evidence say about the crucial parameter σ_ϵ^2 ?

Unfortunately, the direct empirical evidence on the variance of consumption growth is very scarce, due to the unavailability of good quality panel data for broad categories of consumption.¹⁷ Moreover, much of the variance of consumption growth in panel data may be measurement error or attributable to transitory taste shocks unrelated to the permanent changes we are interested in here.

There are a few papers that, somewhat tangentially, provide some direct evidence on the variance of consumption growth. We briefly review some of this recent work to provide a sense of what is currently available. Using PSID data, Storesletten, Telmer, and Yaron [2004a] find that the variance in the growth rate of the permanent component of food expenditure lies between 1%–4%.¹⁸ Blundell, Pistaferri, and Preston [2004] use PSID data, but impute total consumption from food expenditure. Their estimates imply a vari-

¹⁵This point holds more generally in an endowment economy with CRRA utility when the baseline allocation is a geometric random walk.

¹⁶ As already noted in Section 4.1, this analysis is akin to the one in [Lucas, 1990]. But while his exercise involves computing a transition to a steady state with more capital, whereas our exercise involves a transition to a steady state with less capital.

¹⁷ For the United States the PSID provides panel data on food expenditure, and is the most widely used source in studies of consumption requiring panel data. However, recent work by Aguiar and Hurst [2005] show that food expenditure is unlikely to be a good proxy for actual consumption.

¹⁸ See their Table 3, pg. 708.

ance of consumption growth of around 1%.¹⁹ Krueger and Perri [2004] use the panel element in the Consumer Expenditure survey to estimate a statistical model of consumption. At face value, their estimates imply enormous amounts of mobility and a very large variance of consumption growth—around 6%–7%—although most of this should be attributed to a transitory, not permanent, component.²⁰ In general, these studies reveal the enormous empirical challenges faced in understanding the statistical properties of household consumption dynamics from available panel data.

An interesting indirect source of information is the cohort study by Deaton and Paxson [1994]. This paper finds that the cross-sectional inequality of consumption rises as the cohort ages. The rate of increase then provides indirect evidence for σ_ε^2 ; their point estimate implies a value of $\sigma_\varepsilon^2 = 0.0069$. However, recent work using a similar methodology finds much lower estimates [Slesnick and Ulker, 2004, Heathcote, Storesletten, and Violante, 2004].

Calibration. We now display the efficiency gains as a function of σ_ε^2 . We choose three possible discount factors $\beta = 0.96, 0.97$ and 0.98 . For the linear technology model, this imposes $q = \beta \exp(\sigma_\varepsilon^2)$. For the neoclassical growth model, we take $F(K, N) = K^\alpha N^{1-\alpha}$ with $\alpha = 1/3$. Figure 2 plots the efficiency gains as a function of $\exp(\sigma_\varepsilon^2)$ for the linear technology model and for the neoclassical growth model. The figure uses an empirically relevant range for $\exp(\sigma_\varepsilon^2)$.

Consider first the linear technology model. For the parameters under consideration, the efficiency gains range from minuscule—less than 0.1% to very large—over 10%. The effect of the discount factor β is nearly equivalent to increasing the variance of shocks; that is, moving from $\beta = .96$ to $\beta = .98$ has the same effect as doubling σ_ε^2 . To understand this, interpret the lower discounting not as a change in the actual subjective discount, but as calibrating the model to a shorter period length. But then holding the variance of the innovation between periods constant implies an increase in uncertainty over any fixed length of time. Clearly, what matters is the amount of uncertainty per unit of discounted time.

Consider now the neoclassical growth model. There again, there is considerable variation in the size of the efficiency gains, depending on the parameters. However, note that the efficiency gains are much smaller than in the linear technology model. Large differ-

¹⁹ See their footnote 19, pg 21, which refers to the estimated autocovariances from their Table VI.

²⁰ They specify a Markov transition matrix with 9 bins (corresponding to 9 quantiles) for consumption. We thank Fabrizio Perri for providing us with their estimated matrix. Using this matrix we computed that the conditional variance of consumption growth had an average across bins of 0.0646 (this is for the year 2000, the last in their sample; but the results are similar for other years).




Figure 2: Efficiency gains in % when baseline consumption is a geometric random walk for consumption and the Euler equation holds as function of σ_ϵ^2 . The left panel shows the gains in partial equilibrium and the right panel shows the gains in general equilibrium. The bottom dotted line corresponds to $\beta = .96$, the middle dashed line to $\beta = .97$ and top plain line to $\beta = .98$.

ences in interest rates $\tilde{r}_{ss} - r_{ss}$ are necessary to generate substantial efficiency gains: it takes a difference of around 2%, to get efficiency gains that are bigger than 1%.

Lessons. Three lessons emerge from our simple exercise. First, efficiency gains are potentially far from trivial. Second, they are quite sensitive to two parameters of our exercise: the variance in the growth rate of consumption and the subjective discount factor. Third, general equilibrium forces can greatly mitigate them.

5 A Numerical Exploration

In the previous section we provided a method to compute efficiency gains as a function of a baseline consumption process. We also explored the potential gains for a simple random-walk consumption process. On the one hand, this had the advantage of allowing for easily interpretable closed-form solutions, delivering some important insights. On the

other hand, the available empirical evidence reviewed above does not provide reliable estimates for the variance of consumption growth. Moreover, consumption may not be a random walk.

A different strategy is to use consumption processes obtained from an incomplete-market income-fluctuations model. This class of models has been relatively successful at matching data on income, consumption and wealth. Consequently, in this section we take as our baseline the steady-state equilibrium allocation of such a model. Since our goal is to explore the main determinants of the gains from optimal capital taxation we simulate model economies for a wide range of parameters. There is no absolute consensus among the empirical literature, so this broad exploration also has the advantage of spanning the type of parameters used in various studies.

The Model Economy. We follow [Aiyagari \[1994\]](#), who considered a Bewley economy, where a continuum of agents each solve an income fluctuations problem, saving in a risk free asset. Efficiency labor is specified as a first-order autoregressive process in logarithms

$$\log(n_t) = \rho \log(n_{t-1}) + (1 - \rho) \log(n_{ss}) + \varepsilon_t$$

where ε_t is an i.i.d. random variable assumed Normally distributed with mean zero and standard deviation σ_ε . With a continuum of mass one of agents the average efficiency labor supply is $N_{ss} = n_{ss}$.

Labor income is given by the product $w_{ss} n_t$ where w_{ss} is the steady-state wage. Agents face the following sequence of budget constraints

$$a_{t+1} + c_t \leq (1 + r_{ss})a_t + w_{ss}n_t$$

for all $t = 0, 1, \dots$. In addition borrowing is not allowed: $a_t \geq 0$.

The equilibrium steady-state wage is given by the marginal product of labor $w_{ss} = F_N(K_{ss}, N_{ss})$ and the interest rate is given by the net marginal product of capital, $r_{ss} = F_K(K_{ss}, N_{ss}) - \delta$.²¹ For any interest rate $r_{ss} < \beta^{-1} - 1$, agent optimization leads to an invariant cross-sectional distribution for s_t , which we denote by ψ . A steady-state equilibrium requires average assets, under ψ , to equal the capital stock K_{ss} .

Individual consumption is a function of the state variable $s_t \equiv (a_t, n_t)$ which evolves as a Markov process. We take this as our baseline allocation with agents distinguished by

²¹ For simplicity this assumes no taxation. It is straightforward to introduce taxation. However, we conjecture that since taxation of labor income acts as insurance, it effectively reduces the variance of shocks to net income. Lower uncertainty will then only lower the efficiency gains we compute.

their initial conditions $i = s_0$, distributed according to the invariant distribution ψ .

Numerical Method. To solve the planning problem we use the result developed in [Appendix B.1](#). See [Appendix E](#) for more details.

Calibration. We simulate the economy for the wide range of parameter values considered in [Aiyagari \[1994\]](#). The discount factor is set to $\beta = .96$, the production function is Cobb-Douglas with a share of capital of 0.36, capital depreciation is 0.08. The utility function is assumed CRRA so that $U(c) = c^{1-\sigma}/(1-\sigma)$ with $\sigma \in \{1, 3, 5\}$. Aiyagari argues, based on various sources of empirical evidence, for a baseline parametrization with a coefficient of autocorrelation of $\rho = 0.6$ and a standard deviation of labor income of 20%. Following Aiyagari, we also consider different values for the coefficient of relative risk aversion, the autocorrelation coefficient $\rho \in \{0, 0.3, 0.6, 0.9\}$ and the standard deviation of log income, $\text{Std}(\log(n_t)) = \sigma_\varepsilon \in \{0.2, 0.4\}$.

These values are based on the following studies. [Kydland \[1984\]](#) finds that the standard deviation σ_ε of annual hours worked from PSID data is around 15%. Using data from the PSID and the NLS, [Abowd and Card \[1987\]](#) and [Abowd and Card \[1989\]](#) find that the standard deviation of percentage changes in real earnings and annual hours are about 40% and 35% respectively. They report a first order serial correlation coefficient ρ of about 0.3, resulting in an estimate of σ_ε of 34%. Using PSID data, [Heaton and Lucas \[1996\]](#) estimate a range of 0.23 to 0.53 for ρ and a range of 27% to 40% for σ_ε .

Some recent studies, within a life cycle context, consider an alternative approach and estimate a process for log earnings indirectly, by using the increase in the observed cross-sectional inequality of earnings over time within a cohort. Two views have been articulated. The first view, dating back to [Deaton and Paxson \[1994\]](#) and developed most recently by [Storesletten et al. \[2004a\]](#) and [Storesletten et al. \[2004b\]](#), posits that the increase in earning inequality over time within a cohort is due to large and persistent income shocks. This leads to estimates of ρ around or above 0.9 and a range of estimates for σ_ε between 0.3 and 0.6, on the high end of the range of parameter values explored by [Aiyagari \[1994\]](#). The second view, dating back to [Lillard and Weiss \[1979\]](#) and [Hause \[1980\]](#), and developed [Guvenen \[2007\]](#) and [Guvenen \[2009\]](#), argues that the increase in cross sectional earnings inequality over time within a cohort is better explained by an alternative model where agents face individual-specific income profiles. This view leads to lower values for ρ and σ_ε , around 0.8 and 25% respectively. Since both approaches are based on a life-cycle framework, the estimates are not directly relevant for our infinite horizon setup. Indeed, as we have shown, a key object for our analysis is the conditional variance




Figure 3: Path of aggregate consumption for the baseline and the optimized allocations for $\sigma = 1$, $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$ and $\rho = 0.9$.

of consumption growth which depends on the conditional variance of permanent income growth. Life-cycle models incorporate a retirement period, generating a longer horizon for consumption than for labor income. This reduces the impact of permanent income shocks on consumption. Taking this into account, estimates for the persistence ρ and the standard deviation of labor income σ_ε derived in the context of a life-cycle model would therefore have to be adjusted downwards in order to be used in our numerical exercises.

Results. We find that the optimized allocation always features aggregates converging to new steady state values: $\tilde{C}_t \rightarrow \tilde{C}_{ss}$, $\tilde{K}_t \rightarrow \tilde{K}_{ss}$, $\tilde{q}_t \rightarrow \tilde{q}_{ss}$ as $t \rightarrow \infty$. The new steady state interest rate $\tilde{r}_{ss} = 1/\tilde{q}_{ss} - 1$ is always higher than $\tilde{r}_{ss} \geq 1/\beta - 1$ as long as $\sigma \geq 1$, with equality when $\sigma = 1$. This is to be contrasted with the initial interest rate $r_{ss} < 1/\beta - 1$.²²²³ Figure 3 plots the path of aggregate consumption for one particular parameter case. Since the baseline allocation represents a steady state, its aggregate consumption is constant. Aggregate consumption for the optimized allocation is initially above this level, but declines monotonically, eventually reaching a new, lower, steady state. For the same parameter values, Figure 4 shows a typical sample path for individual income, cash in hand, consumption and utility. Optimized consumption appears more persistent and displays a downwards trend in the initial periods.

²²By contrast, when $\sigma < 1$, any steady state of a Δ -efficient allocation features an interest rate lower than $1/\beta - 1$.

²³In the case of logarithmic utility functions, we provided a formal proof of this result in Section 4.1. In the more general CRRRA utility function case, we rely solely on our numerical results.

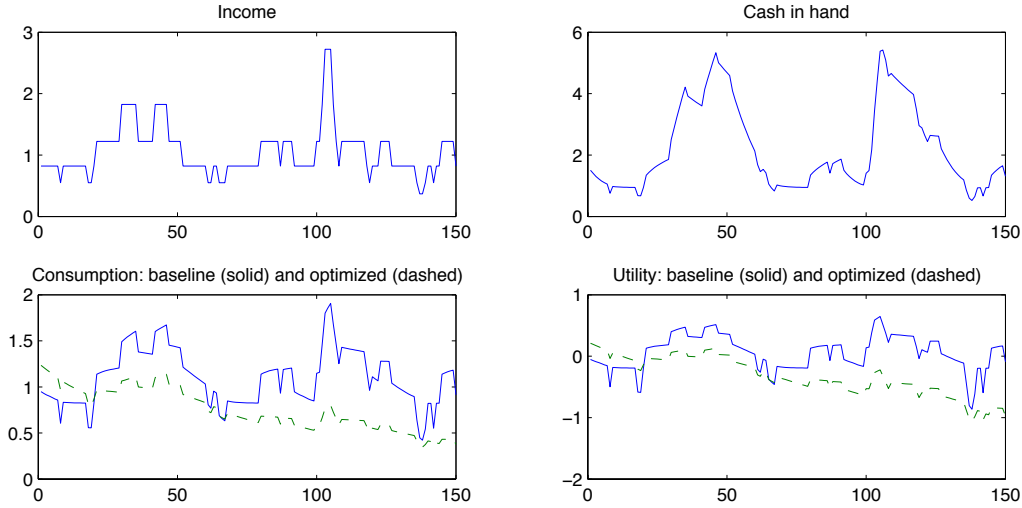


Figure 4: Simulation of a typical individual sample path for $\sigma = 1$, $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$ and $\rho = 0.9$.

Tables 2, 3 and 4 collect the results of our simulations. All tables report our measure of efficiency gains, $\tilde{\lambda}$. In the logarithmic utility case ($\sigma = 1$), Table 2 includes the idiosyncratic and aggregate components $\tilde{\lambda}^I$ and $\tilde{\lambda}^A$. For references, the tables show the baseline and optimized steady state interest rates r_{ss} and \tilde{r}_{ss} . The second column, showing r_{ss} , replicates Aiyagari's Table II (pg. 678).

The baseline's interest rate r_{ss} is decreasing in the size and persistence of the shocks, as well as in the coefficient of relative risk aversion. Precautionary saving motives are stronger and depress the equilibrium interest rate. In particular, $\tilde{r}_{ss} = 1/\beta - 1$ when utility is logarithmic, and $\tilde{r}_{ss} > 1/\beta - 1$ when $\sigma > 1$. Moreover, we find that \tilde{r}_{ss} increases with σ and the size and the persistence of the labor income shocks. These comparative statics for \tilde{r}_{ss} are precisely the reverse of those for r_{ss} . This is perhaps not surprising, given the reversal in the sign of the power coefficients in the Inverse Euler and Euler equations, σ and $-\sigma$, respectively.

Efficiency gains are increasing with the variance of the shocks and with their persistence. They also increase with the coefficient of relative risk aversion. Aiyagari argues, based on various sources of empirical evidence, for a parameterization with a coefficient of autocorrelation of $\rho = 0.6$ and a standard deviation of labor income of 20%. For this specification, we find that efficiency gains are small – below 0.2% for all three values for the coefficient of relative risk aversion.

In the logarithmic case, efficiency gains are moderate – always less than 1.3%. Our decomposition along the lines of Section 4.1 shows that idiosyncratic efficiency gains $\tilde{\lambda}^I$

$\sigma = 1$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .2$						
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains			
			Idiosyncratic	Aggregate	Total	Borrowing
0	4.14%	4.17%	0.0%	0.0%	0.0%	0.0%
0.3	4.13%	4.17%	0.0%	0.0%	0.0%	0.0%
0.6	4.09%	4.17%	0.0%	0.0%	0.0%	0.0%
0.9	3.95%	4.17%	0.2%	0.0%	0.2%	0.1%

$\sigma = 1$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$						
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains			
			Idiosyncratic	Aggregate	Total	Borrowing
0	4.06%	4.17%	0.1%	0.0%	0.1%	0.1%
0.3	3.97%	4.17%	0.2%	0.0%	0.2%	0.2%
0.6	3.79%	4.17%	0.4%	0.0%	0.4%	0.3%
0.9	3.38%	4.17%	1.2%	0.1%	1.3%	1.1%

Table 2: Efficiency Gains for replication of [Aiyagari \[1994\]](#) when $\sigma = 1$.

completely dwarf aggregate efficiency gains $\tilde{\lambda}^A$. Our finding that idiosyncratic gains are modest could have perhaps been anticipated by our illustrative geometric random-walk example, where idiosyncratic gains are zero. Intuitively, efficiency gains from the idiosyncratic component require differences in the expected consumption growth rate across individuals. When individuals smooth their consumption over time effectively the remaining differences are small—as a result, so are the efficiency gains. Efficiency gains from the aggregate component are directly related to the difference between the equilibrium and optimal steady-state capital. With logarithmic utility this is equivalent, to the difference between the equilibrium steady-state interest rate and $\beta^{-1} - 1$, the interest rate that obtains with complete markets. Hence, our finding of low aggregate efficiency gains is directly related to [Aiyagari’s \[1994\]](#) main conclusion: for shocks that are not implausibly large or for moderate risk aversion, precautionary savings are small in the aggregate, in that steady-state capital and interest rate are close their complete-markets levels, as shown in our [Table 2](#).

For the range of parameters that we consider, the efficiency gains range from minuscule—less than 0.1%—to very large—more than 8.4%. However, efficiency gains larger than 1.3% are only reached for combinations of high values of relative risk aversion— σ greater than 3—and both large and highly persistent shocks—a standard deviation of labor income of 40%, and a mean-reversion coefficient ρ greater than 0.6.

$\sigma = 3$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .2$				
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains	
			Total	Borrowing
0	4.09%	4.21%	0.0%	0.0%
0.3	4.02%	4.37%	0.0%	0.0%
0.6	3.88%	4.45%	0.2%	0.2%
0.9	3.36%	4.75%	0.7%	0.7%

$\sigma = 3$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$				
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains	
			Total	Borrowing
0	3.77%	4.62%	0.2%	0.2%
0.3	3.47%	4.77%	0.5%	0.5%
0.6	2.89%	5.03%	1.2%	1.2%
0.9	1.47%	5.56%	4.2%	3.3%

Table 3: Efficiency Gains for replication of [Aiyagari \[1994\]](#) when $\sigma = 3$.

The Role of Borrowing Constraints. The market arrangement in Aiyagari’s economy imposes borrowing constraints that limit the ability to trade consumption intertemporally. At the baseline allocation, the Euler equation holds with equality $U'(c(u_t^i)) = \beta(1 + r_{ss})\mathbb{E}_t^i [U'(c(u_{t+1}^i))]$ holds when assets and current income are high enough. However, for low levels of assets and income, the agent may be borrowing constrained so that the Euler conditions holds with strict inequality $U'(c(u_t^i)) > \beta(1 + r_{ss})\mathbb{E}_t^i [U'(c(u_{t+1}^i))]$.

In contrast, since our planning problem does not impose arbitrary restrictions on the perturbations, it places no such limits on intertemporal reallocation of consumption. Thus, the perturbations can effectively undo limits to borrowing. As a result, part of the efficiency gains we compute can be attributed to the relaxation of borrowing constraints, not to the introduction of savings distortions.

To get an idea of the efficiency gains that are obtained from the relaxation of borrowing constraints, we performed the following exercise. The basic idea is to use the perturbations to construct a new allocation where the Euler equation always holds with equality. That is, the new allocation stops short of satisfying the Inverse Euler equation, so it does not introduce positive intertemporal wedges. Instead, it removes the negative intertemporal wedges that were present due to borrowing constraints. We compute the resource savings from this perturbations as a simple measure of the efficiency gains due to the relaxation of borrowing constraints.²⁴

²⁴Theoretically, it isn’t evident that these perturbations actually produce positive efficiency gains, as op-

$\sigma = 5$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .2$				
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains	
			Total	Borrowing
0	4.01%	5.34%	0.0%	0.0%
0.3	3.89%	5.39%	0.0%	0.0%
0.6	3.61%	5.48%	0.2%	0.2%
0.9	2.66%	5.72%	1.0%	1.0%

$\sigma = 5$ and $\text{Std}(\log(n_t) - \log(n_{t-1})) = .4$				
ρ	r_{ss}	\tilde{r}_{ss}	Efficiency Gains	
			Total	Borrowing
0	3.43%	5.54%	0.2%	0.2%
0.3	2.90%	5.55%	0.7%	0.7%
0.6	1.95%	5.58%	2.1%	2.0%
0.9	-0.16%	5.52%	8.4%	6.3%

Table 4: Efficiency Gains for replication of [Aiyagari \[1994\]](#) when $\sigma = 5$.

More precisely, we seek to determine the unique allocation $\{\tilde{u}_t^{i,B}, n_t^i, \tilde{K}_t^B, \tilde{\lambda}^B C_{ss}\}$, where B stands for borrowing, that satisfies the following constraints. First, we impose that the allocation be achievable through parallel perturbations of the baseline allocation, and deliver the same utility

$$\{\tilde{u}_t^{i,B}\} \in Y(\{u_t^i\}, 0).$$

Second, we impose that the resource constraints hold

$$\tilde{K}_{t+1}^B + \int \mathbb{E}^i [c(\tilde{u}_t^{i,B}) \Pr(\theta^{i,t})] d\psi + \tilde{\lambda}^B C_{ss} \leq (1 - \delta)\tilde{K}_t^B + F(\tilde{K}_t^B, N_{ss}) \quad t = 0, 1, \dots$$

and require that the initial capital be equal the initial capital $\tilde{K}_0^B = K_0$ of the baseline allocation. Finally, we impose that the Euler equation hold for every agent in every period

$$U'(c(\tilde{u}_t^{i,B})) = \beta[1 - \delta + F_K(\tilde{K}_{t+1}^B, N_{ss})]\mathbb{E}_t^i[U'(c(u_{t+1}^{i,B}))].$$

The allocation $\{\tilde{u}_t^{i,B}, n_t^i, \tilde{K}_t^B, \tilde{\lambda}^B C_t\}$ can improve on the baseline allocation by insuring that the Euler equation holds for every agent in every period. Note that this allocation does not satisfy the Inverse Euler equation and is therefore not Δ -efficient. We adopt $\tilde{\lambda}^B$ as our measure of the efficiency gains deriving from the relaxation of borrowing constraints.

The rightmost column in each of table reports the efficiency gains $\tilde{\lambda}^B$ that can be attributed to negative ones. In our cases, the measure always came out to be positive.

tributed to the relaxation of borrowing constraints. Perhaps the most important conclusion of our numerical exercise can be drawn by comparing the efficiency gains $\tilde{\lambda}$ with the efficiency gains $\tilde{\lambda}^B$ deriving from the relaxation of borrowing constraints. We find that for all size and persistence of shocks, and all utility specifications, most of the efficiency gains come from the relaxation of borrowing constraints. Indeed, the difference between $\tilde{\lambda}$ and $\tilde{\lambda}^B$ is less than 0.1% except for the largest and most persistent shocks—a standard deviation of labor income growth of 40%, and a mean-reversion coefficient ρ greater than 0.9. In other words, most of the efficiency come from allowing agents to better smooth their consumption over time by alleviating borrowing constraints, rather than by optimally distorting savings as prescribed by the substitution the Inverse Euler equation to the Euler equation.

6 Conclusions

This paper provided a method for evaluating the auxiliary role that savings distortions play in social insurance arrangements. We put it to use to evaluate the welfare importance of recent arguments for capital taxation based on the Inverse Euler equation. Our main finding is that, once general equilibrium effects are taken into account, these gains are small.

The method developed here is flexible enough to accommodate several extensions and it may be of interest to investigate how these may affect the quantitative conclusions found here for our benchmark calibration. In a separate paper [Farhi and Werning, 2008] we pursued such an extension using Epstein-Zin preferences that separate risk aversion from the intertemporal elasticity of substitution. In ongoing work, we study an environment with overlapping-generations instead of the infinite horizon dynastic setup used here. The basic methodology and decompositions introduced here are useful for such extensions.

Appendix

A Proofs for Section 3.2

To simplify the arguments, we first assume that the horizon is finite with terminal period $T > 0$. We also assume that $V(n; \theta)$ is continuously differentiable with respect to n and θ . We also assume throughout that shocks are continuous.

Consider an assignment for utility and labor $\{u_t^i, n_t^i\}$. We define the continuation value $U(\theta^{i,t})$ conditional on history $\theta^{i,t}$ as follows

$$U^i(\theta^{i,t}) = \mathbb{E} \left[\sum_{s=0}^T \beta^s \left[u_{t+s}^i(\theta^{i,t+s}) - V(n_{t+s}^i(\theta^{i,t+s}); \theta_{t+s}^i) \right] \mid \theta^{i,t} \right].$$

Fix a history $\theta^{i,t}$ and consider a report $\hat{\theta}_t^i \in \Theta$. Define the strategy $\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}$ as follows: $\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}(\theta^{i,s}) = \theta^{i,s}$ except if $\theta^{i,s} \succeq \theta^{i,t}$, in which case

$$\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}(\theta^{i,t-1}, \theta_t^i, \theta_{t+1}^i, \theta_{t+2}^i, \dots) = (\theta^{i,t-1}, \hat{\theta}_t^i, \theta_{t+1}^i, \theta_{t+2}^i, \dots).$$

This strategy coincides with truth-telling except in period t after history $\theta^{i,t}$ where the report $\hat{\theta}_t^i$ can be different from the true shock θ_t^i .

For clarity and brevity, we use the notation $\hat{\theta}^{i,s}$ for $\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}(\theta^{i,s})$. We denote the continuation utility after history $\theta^{i,s} \succeq \theta^{i,t}$ under the strategy $\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}$, by $U^i(\sigma_{\hat{\theta}_t^i}^{\theta^{i,t}}(\theta^{i,s}))$ or $U^i(\hat{\theta}^{i,s}; \theta^{i,t})$ for short. Similarly, we denote by $\mathbb{E}[U^i(\hat{\theta}^{i,s}; \theta^{i,t})]$ or $\mathbb{E}[U^i(\hat{\theta}^{i,s}) \mid \theta^{i,t}]$ the expectation of this continuation utility, conditional on the realized history $\theta^{i,t}$ at date t .

We say that the assignment for utility and labor $\{u_t^i, n_t^i\}$ is *regular* if for all $t \geq 0$ and $\theta^{i,t-1} \in \Theta^t$, the continuation utility $U^i(\hat{\theta}^{i,t}; \theta_{i,t}^i)$ is absolutely continuous in the true shock θ_t^i , differentiable with respect to the true shock θ_t^i , and the derivative, which we denote by

$$-\frac{\partial}{\partial \theta_t^i} V(n(\hat{\theta}^{i,t}); \theta_t^i) + \beta \frac{\partial}{\partial \theta_t^i} \mathbb{E}[U^i(\hat{\theta}^{i,t+1}) \mid \theta^{i,t}]$$

is bounded by a function $b(\theta_t^i; \theta^{i,t-1})$ which is integrable with respect to θ_t^i .

Consider two regular, incentive compatible, assignments for utility and labor $\{u_t^i, n_t^i\}$ and $\{\tilde{u}_t^i, \tilde{n}_t^i\}$, that share the same assignment for labor $\{n_t^i\}$. Denote by $\{U^i\}$ and $\{\tilde{U}^i\}$ the corresponding continuation utilities. We now show that there exists $\Delta_{-1} \in \mathbb{R}$ such that $\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, \Delta_{-1})$.

Note we can restate the result as follows: for all t with $0 \leq t \leq T-1$ and for all

$\theta^{i,t-1} \in \Theta^t$, there exists $\Delta(\theta^{i,t-1}) \in \mathbb{R}$ such that

$$\tilde{U}^i(\theta^{i,t}) = U^i(\theta^{i,t}) + \Delta(\theta^{i,t-1}). \quad (17)$$

Fix a history $\theta^{i,t}$ and consider the strategies $\sigma_{\tilde{\theta}_t^i}^{\theta^{i,t}}$. Using Theorem 4 in [Milgrom and Segal \[2002\]](#), Incentive compatibility and regularity implies the following Envelope condition

$$U^i(\theta^{i,t-1}, \theta_t^i) = U^i(\theta^{i,t-1}, \theta_t^{i'}) + \int_{\theta_t^{i'}}^{\theta_t^i} \left[-\frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} V(n(\tilde{\theta}^{i,t}); \tilde{\theta}_t^i) + \beta \frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} \mathbb{E} \left[U^i(\tilde{\theta}^{i,t+1}) \mid (\tilde{\theta}_{t-1}^i, \tilde{\theta}_t^i) \right] d\tilde{\theta}_t^i \right]. \quad (18)$$

where $\tilde{\theta}^{i,t+s} = (\theta^{i,t-1}, \tilde{\theta}_t^i, \theta_{t+1}^i, \dots, \theta_{t+s}^i)$ for all $s \geq 0$ and $\tilde{\theta}^{i,t-s} = \theta^{i,t-s}$ for all $s > 0$.

The same expression holds for $\{\tilde{u}_t^i, n_t^i\}$. Using the fact that shocks are continuous and that $\{u_t^i, n_t^i\}$ and $\{\tilde{u}_t^i, n_t^i\}$ share the same labor assignment, this implies that there exists $\Delta(\theta^{i,t-1}) \in \mathbb{R}$ such that

$$\begin{aligned} \tilde{U}^i(\theta^{i,t}) &= U^i(\theta^{i,t}) + \Delta(\theta^{i,t-1}) \\ &\quad + \beta \int_{\underline{\theta}}^{\theta_t^i} \frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} \mathbb{E} \left[\tilde{U}^i(\tilde{\theta}^{i,t+1}) - U^i(\tilde{\theta}^{i,t+1}) \mid (\tilde{\theta}_{t-1}^i, \tilde{\theta}_t^i) \right] d\tilde{\theta}_t^i. \end{aligned}$$

For $t = T$, the last term on the right hand side vanishes and we have that for all $\theta^{i,T} \in \Theta^{T+1}$, there exists $\Delta(\theta^{i,T-1}) \in \mathbb{R}$ such that (17) holds. We proceed by induction on t . Suppose that $t \geq 0$ and that for all $\theta^{i,t} \in \Theta^{t+1}$, there exists $\Delta(\theta^{i,t}) \in \mathbb{R}$ such that (17) holds. Then applying (18) and using the fact that

$$\frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \tilde{\theta}_t^i} \mathbb{E} \left[\Delta(\tilde{\theta}^{i,t}) \mid (\tilde{\theta}^{i,t-1}, \tilde{\theta}_t^i) \right] = 0,$$

we immediately find that for all $\theta^{i,t-1} \in \Theta^t$ there exists $\Delta(\theta^{i,t-1}) \in \mathbb{R}$ such that

$$\tilde{U}^i(\theta^{i,t}) = U^i(\theta^{i,t}) + \Delta(\theta^{i,t-1}).$$

The proof follows by induction.

When the horizon is infinite, the regularity conditions must be complemented with a limit condition. For example, expanding our Envelope condition over two periods, we find

$$\begin{aligned} \tilde{U}^i(\theta^{i,t}) &= U^i(\theta^{i,t}) + \Delta(\theta^{i,t-1}) + \beta^2 \int_{\underline{\theta}}^{\theta^i} \frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \bar{\theta}_t^i} \\ \mathbb{E} \left[\int_{\underline{\theta}}^{\bar{\theta}_{t+1}^i} \frac{\partial}{\partial \tilde{\theta}_{t+1}^i} \Big|_{\tilde{\theta}_{t+1}^i = \overleftarrow{\theta}_{t+1}^i} \mathbb{E} \left[(\tilde{U}^i - U^i) \left(\overleftarrow{\theta}^{i,t+2} \right) \Big| (\overleftarrow{\theta}^{i,t}, \tilde{\theta}^{i,t+1}) \right] d\overleftarrow{\theta}_t^i \Big| (\bar{\theta}^{i,t-1}, \tilde{\theta}^{i,t}) \right] d\bar{\theta}_t^i. \end{aligned}$$

Consider the sequence which is constructed by iterating our Envelope condition for the utility and labor assignment $\{u_t^i, n_t^i\}$. The first element of the sequence is

$$\beta \int_{\underline{\theta}}^{\theta^i} \frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \bar{\theta}_t^i} \mathbb{E} \left[U^i(\bar{\theta}^{i,t+1}) \Big| (\bar{\theta}^{i,t-1}, \tilde{\theta}^{i,t}) \right] d\bar{\theta}_t^i.$$

Similarly, the second element of the sequence is

$$\begin{aligned} \beta^2 \int_{\underline{\theta}}^{\theta^i} \frac{\partial}{\partial \tilde{\theta}_t^i} \Big|_{\tilde{\theta}_t^i = \bar{\theta}_t^i} \\ \mathbb{E} \left[\int_{\underline{\theta}}^{\bar{\theta}_{t+1}^i} \frac{\partial}{\partial \tilde{\theta}_{t+1}^i} \Big|_{\tilde{\theta}_{t+1}^i = \overleftarrow{\theta}_{t+1}^i} \mathbb{E} \left[U^i \left(\overleftarrow{\theta}^{i,t+2} \right) \Big| (\overleftarrow{\theta}^{i,t}, \tilde{\theta}^{i,t+1}) \right] d\overleftarrow{\theta}_t^i \Big| (\bar{\theta}^{i,t-1}, \tilde{\theta}^{i,t}) \right] d\bar{\theta}_t^i. \end{aligned}$$

Our proof then carries over if the condition that the limit of these terms when the number of iterations goes to infinity is equal to zero is added to the requirements for regularity.

B Solving the Planning Problem: Relaxed Planning Problem and Component Planning Problems

B.1 Relaxed Problem

We now construct a relaxed planning problem that replaces the resource constraints with a single present value condition. This problem is indexed by a sequence of intertemporal prices or interest rates, which encodes the scarcity of aggregate resources in every period. The relaxed planning problem can be further decomposed in a series of component planning problems corresponding to the different types $i \in I$.

Given some positive intertemporal prices $\{\tilde{Q}_t\}$ with the normalization that $\sum_{t=0} \tilde{Q}_t C_t = 1$ we replace the sequence of resource constraints (16) with the single present value condition

$$\tilde{\lambda} = \sum_{t=0}^{\infty} \tilde{Q}_t \left(F(\tilde{K}_t, N_t) + (1 - \delta)\tilde{K}_t - \tilde{K}_{t+1} - \int \mathbb{E}^i [c(\tilde{u}_t^i)] d\psi \right) \quad (19)$$

which is obtained by multiplying equation (16) by \tilde{Q}_t and summing over $t = 0, 1, \dots$

Formally, the *relaxed planning problem* seeks the allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ where $\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, 0)$ that maximizes $\tilde{\lambda}$ given by [equation \(19\)](#).

The connection with the original planning problem is the following. Suppose that, given some $\{\tilde{Q}_t\}$, the optimal allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ for the relaxed problem satisfies the resource constraints [\(16\)](#). Then, this allocation solves the original planning problem. This relaxed problem approach is adapted from [Farhi and Werning \[2007\]](#) and is related to the first welfare theorem proved in [Atkeson and Lucas \[1992\]](#).

The converse is also true. Indeed, the prices $\{\tilde{Q}_t\}$ are Lagrange multipliers and Lagrangian necessity theorems guarantee the existence of prices $\{\tilde{Q}_t\}$ for the relaxed problem. The following lemma, which follows immediately from Theorem 1, Section 8.3 in [Luenberger \[1969\]](#), provides one such result.

Lemma 1. *Suppose $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda} C_t\}$ solves the planning problem and the resource constraints [\(16\)](#) hold with equality. Then there exists a sequence of prices $\{\tilde{Q}_t\}$ such that this same allocation solves the relaxed planning problem.*

The relaxed planning problem can be decomposed into a subproblem for capital $\{\tilde{K}_t\}$ and a series of component planning problems for the utility assignment $\{\tilde{u}_t^i\}$. The subproblem for capital maximizes the right hand side of [equation \(19\)](#) with respect to $\{\tilde{K}_{t+1}\}$. The first-order conditions, which are necessary and sufficient for an interior optimum, are

$$1 = \tilde{q}_t (F_K(\tilde{K}_{t+1}, N_{t+1}) + 1 - \delta) \quad t = 0, 1, \dots$$

where $\tilde{q}_t \equiv \tilde{Q}_{t+1}/\tilde{Q}_t$. This, together with the normalization that $\sum_{t=0}^{\infty} \tilde{Q}_t C_t = 1$, implies a one-to-one relationship between $\{\tilde{K}_{t+1}\}$ and $\{\tilde{Q}_t\}$.

The component planning problem for type $i \in I$ maximizes the right hand side of [equation \(19\)](#) with respect to the utility assignment $\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, 0)$. The objective reduces to minimizing the present value of consumption:

$$\sum_{t=0}^{\infty} \frac{\tilde{Q}_t}{\tilde{Q}_0} \mathbb{E}^i [c(\tilde{u}_t^i)]. \quad (20)$$

Each of these component planning problems can be solved independently. Since both the objective and the constraints are convex it follows immediately that, given $\{\tilde{q}_t\}$, the first order conditions for optimality at an interior solution in the component planning problem [\(20\)](#) coincide exactly with the Inverse Euler equation [\(14\)](#):

$$c'(\tilde{u}_t^i) = \frac{\tilde{q}_t}{\beta} \mathbb{E}_t^i [c'(\tilde{u}_{t+1}^i)] \quad t = 0, 1, \dots$$

B.2 A Bellman Equation

In most situations of interest, the baseline allocation admits a recursive representation for some endogenous state variable. This is the case whenever $\{\theta_t^i\}$ is a Markov process and the baseline allocation depends on the history of shocks $\theta^{i,t-1}$ in a way that can be summarized by an endogenous state x_t^i , with law of motion $x_t^i = M(x_{t-1}^i, \theta_t^i)$ and given initial condition x_0^i . Note that the transition matrix M is assumed to be independent of i . Types then correspond to different initial value x_0^i . This is the only thing that distinguishes them. The endogenous state x_t^i is a function of the history of exogenous shocks $\theta^{i,t}$. Defining the state vector $s_t^i = (x_t^i, \theta_t^i)$ there must exist a function \bar{u} such that $u_t^i(\theta^{i,t}) = \bar{u}(s_t^i)$ for all $\theta^{i,t}$. In what follows we drop the hat notation and we stop indexing the allocation by the type i of the agents. We denote a baseline allocation by $u(s_t)$ and use the notation $c(s_t)$ for $c(u(s_t))$.

For such baseline allocations we can reformulate the component planning problems recursively as follows.

The idea is to take s_t as an exogenous process and keep track of the additional lifetime utility Δ_{t-1} previously promised as an endogeneous state variable.

For any date τ define the continuation plans $\{u_t^\tau\}_{t=0}^\infty$ with $u_t^\tau(s) = u_{t+\tau}(s)$ and the value function corresponding to the component planning problem starting at date τ with state s

$$K(\Delta_-, s; \tau) \equiv \inf \sum_{t=0}^{\infty} \frac{\tilde{Q}_{\tau+t}}{\tilde{Q}_\tau} \mathbb{E} [c(\tilde{u}_t^\tau) \mid s_\tau = s] \quad \text{s.t. } \{\tilde{u}_t^\tau\} \in \mathcal{Y}(\{u_t^\tau\}, \Delta_-).$$

This value function satisfies the Bellman equation

$$K(\Delta_-, s; \tau) = \min_{\Delta} [c(u(s) + \Delta_- - \beta\Delta) + \tilde{q}_\tau \mathbb{E}[K(\Delta, s'; \tau + 1) \mid s]]. \quad (21)$$

The optimization over Δ is one-dimensional and convex and delivers a policy function $g(\Delta_-, s; \tau)$. Combining the necessary and sufficient first-order condition for Δ with the envelope condition yields the Inverse Euler equation in recursive form

$$c'(u(s) + \Delta_- - \beta g(\Delta_-, s; \tau)) = \frac{\tilde{q}_\tau}{\beta} \mathbb{E} [c'(u(s') + g(\Delta_-, s; \tau) - \beta g(g(\Delta_-, s; \tau), s'; \tau + 1)) \mid s]$$

This condition can be used to compute $g(\cdot, \cdot; \tau)$ for given $g(\cdot, \cdot, \tau + 1)$.

An optimal plan for $\{\tilde{u}_t\}$ can then be generated from the sequence of policy functions $\{g(\cdot, \cdot; \tau)\}$ by setting $\tilde{u}(s^t) = u(s_t) + \Delta(s^{t-1}) - \beta\Delta(s^t)$ and using the recursion $\Delta(s^t) = g(\Delta(s^{t-1}), s_t; t)$ with initial condition $\Delta_{-1} = 0$.

With a fixed discount factor q , the value function is independent of time, $K(\Delta_-, s)$, and solves the stationary Bellman equation

$$K(\Delta_-, s) = \min_{\Delta} [c(u(s) + \Delta_- - \beta\Delta) + q\mathbb{E}[K(\Delta, s') \mid s]]. \quad (22)$$

This dynamic program admits an analogy with a consumer's income fluctuation problem that is both convenient and enlightening. We transform variables by changing signs and switch the minimization to a maximization. Let $\tilde{\Delta}_- \equiv -\Delta_-$, $\tilde{K}(\tilde{\Delta}_-; s) \equiv -K(-\tilde{\Delta}_-; s)$ and $\tilde{U}(x) \equiv -c(-x)$. Note that the pseudo utility function \tilde{U} is increasing, concave and satisfies Inada conditions at the extremes of its domain.²⁵ Reexpressing the Bellman equation (22) using these transformation yields:

$$\tilde{K}(\tilde{\Delta}_-; s) = \max_{\tilde{\Delta}} [\tilde{U}(-u(s) + \tilde{\Delta}_- - \beta\tilde{\Delta}) + q\mathbb{E}[\tilde{K}(\tilde{\Delta}; s') \mid s]].$$

This reformulation can be read as the problem of a consumer with a constant discount factor q facing a constant gross interest rate $1 + r = \beta^{-1}$, entering the period with pseudo financial wealth $\tilde{\Delta}_-$, receives a pseudo labor income shock $-u(s)$. The fictitious consumer must decide how much to save $\beta\tilde{\Delta}$; pseudo consumption is then $x = -u(s) + \tilde{\Delta}_- - \beta\tilde{\Delta}$.

The benefit of this analogy is that the income fluctuations problem has been extensively studied and used; it is at the heart of most general equilibrium incomplete market models [Aiyagari, 1994].

With logarithmic utility, the Bellman equation can be simplified considerably. The idea is best seen through the analogy, noting that the pseudo utility function is exponential $-e^{-x}$ in this case. It is well known that for a consumer with CARA preferences a one unit increase in financial wealth, $\tilde{\Delta}$, results in an increase in pseudo-consumption, x , of $r/(1+r) = 1 - \beta$ in parallel across all periods and states of nature. It is not hard to see that this implies that the value function takes the form $\tilde{K}(\tilde{\Delta}_-; s) = e^{-(1-\beta)\tilde{\Delta}_-} \tilde{k}(s)$. These ideas are behind the following result.

Proposition 4. *With logarithmic utility and constant discount q , the value function in equation (22) is given by*

$$K(\Delta_-; s) = e^{(1-\beta)\Delta_-} k(s),$$

²⁵ An important case is when the original utility function is CRRA $U(c) = c^{1-\sigma}/(1-\sigma)$ for $\sigma > 0$ and $c \geq 0$. Then for $\sigma > 1$ the function $\tilde{U}(x)$ is proportional to a CRRA with coefficient of relative risk aversion $\tilde{\sigma} = \sigma/(\sigma - 1)$ and $x \in (0, \infty)$. For $\sigma < 1$ the pseudo utility \tilde{U} is "quadratic-like", in that it is proportional to $-(-x)^\rho$ for some $\rho > 1$, and $x \in (-\infty, 0]$.

where function $k(s)$ solves the Bellman equation

$$k(s) = Ac(s)^{1-\beta} (\mathbb{E}[k(s') | s])^\beta, \quad (23)$$

where $A \equiv (q/\beta)^\beta / (1-\beta)^{1-\beta}$.

The optimal policy for Δ can be obtained from $k(s)$ using

$$\Delta = \Delta_- - \frac{1}{\beta} \log \left(\frac{(1-\beta)k(s)}{c(s)} \right). \quad (24)$$

Proof. See [Appendix C](#). ■

This solution is nearly closed form: one needs only compute $k(s)$ using the recursion in [equation \(23\)](#), which requires no optimization. No simplifications on the stochastic process for skills are required.

B.3 Idiosyncratic Gains

The idiosyncratic planning problem lends itself to a similar analysis. In order to avoid repetitions, we only sketch the corresponding analysis. We can define a corresponding relaxed problem, given a sequence of prices $\{\hat{Q}_t\}$. We can also prove an analogue of [Lemma \(1\)](#). This relaxed planning problem can then be decomposed into a series of component planning problems. When the baseline allocation is recursive, and given a sequence of prices $\{\hat{Q}_t\}$, we can study the component planning problems using a Bellman equation as in [equation \(21\)](#). The corresponding first-order conditions also take the form of an Inverse Euler equation

$$c'(\hat{u}_t^i) = \frac{\hat{q}_t}{\beta} \mathbb{E}_t^i [c'(\hat{u}_{t+1}^i)] \quad t = 0, 1, \dots$$

where $\hat{q}_t \equiv \hat{Q}_{t+1} / \hat{Q}_t$.

B.4 A Complete Solution When the Baseline is a Recursive Steady State

Suppose that the baseline allocation is recursive with state s and features constant aggregate consumption $C_t = C_{ss}$. It is not necessary for the argument that aggregate capital K_t and labor N_t be constant at the baseline allocation. The idiosyncratic efficiency gains can then be easily computed. [Propositions ??](#) implies that we can use $\hat{q}_t = \beta$ to compute the idiosyncratic allocation. [Proposition 4](#) then allows us to compute the solution in closed

form. This leads to

$$\tilde{\lambda}^I = 1 - \frac{(1 - \beta) \int k(s) d\psi}{C_{ss}}$$

where $k(s)$ solves [equation \(22\)](#) with $q = \beta$. We can then apply [Proposition 2](#) and compute the efficiency gains $\tilde{\lambda}$ and the corresponding Δ -efficient allocation by solving the Aggregate planning problem, a version of the neoclassical growth model. This provides a complete solution to the planning problem when the baseline allocation is recursive and features constant aggregate consumption.

C Proof of [Proposition 4](#)

With logarithmic utility the Bellman equation is

$$\begin{aligned} K(s, \Delta_-) &= \min_{\Delta} [s \exp(\Delta_- - \beta\Delta) + q\mathbb{E}[K(s', \Delta) | s]] \\ &= \min_{\Delta} [s \exp((1 - \beta)\Delta_- + \beta(\Delta_- - \Delta)) + q\mathbb{E}[K(s', \Delta) | s]] \end{aligned}$$

Substituting that $K(\Delta_-, s) = k(s) \exp((1 - \beta)\Delta_-)$ gives

$$k(s) \exp((1 - \beta)\Delta_-) = \min_{\Delta} [s \exp((1 - \beta)\Delta_- + \beta(\Delta_- - \Delta)) + q\mathbb{E}[k(s') \exp((1 - \beta)\Delta) | s]],$$

and cancelling terms:

$$\begin{aligned} k(s) &= \min_{\Delta} [s \exp(\beta(\Delta_- - \Delta)) + q\mathbb{E}[k(s') \exp((1 - \beta)(\Delta - \Delta_-)) | s]] \\ &= \min_d [s \exp(-\beta d) + q\mathbb{E}[k(s') \exp((1 - \beta)d) | s]] \\ &= \min_d [s \exp(-\beta d) + q\mathbb{E}[k(s') | s] \exp((1 - \beta)d)] \end{aligned}$$

where $d \equiv \Delta - \Delta_-$. We can simplify this one dimensional Bellman equation further. Define $\hat{q}(s) \equiv q\mathbb{E}[k(s') | s]/s$ and

$$M(\hat{q}) \equiv \min[\exp(-\beta d) + \hat{q} \exp((1 - \beta)d)].$$

The first-order conditions gives

$$\beta \exp(-\beta d) = \hat{q}(1 - \beta) \exp((1 - \beta)d) \quad \Rightarrow \quad d = \log \frac{\beta}{(1 - \beta)\hat{q}}. \quad (25)$$

Substituting back into the objective we find that

$$\begin{aligned} M(\hat{q}) &= \frac{1}{1-\beta} \exp(-\beta d) = \frac{1}{1-\beta} \exp\left(-\beta \log \frac{\beta}{(1-\beta)\hat{q}}\right) \\ &= \frac{1}{(1-\beta)^{1-\beta} \beta^\beta} \hat{q}^\beta = B\hat{q}^\beta, \end{aligned}$$

where B is a constant defined in the obvious way in terms of β .

The operator associated with the Bellman equation is then

$$T[k](s) = sM\left(q \frac{\mathbb{E}[k(s') | s]}{s}\right) = As^{1-\beta} (\mathbb{E}[k(s') | s])^\beta,$$

where $A \equiv Bq^\beta = (q/\beta)^\beta / (1-\beta)^{1-\beta}$.

Combining the Bellman $k(s)/s = M(\hat{q}) = A\hat{q}^\beta$ with [equation \(25\)](#) yields the policy function as a function of $K(s)$. This completes the proof.

D Proof of Proposition 2

Consider the aggregate allocation $\{\tilde{C}_t, N_t, \tilde{K}_t, \tilde{\lambda}C_t\}$ that solves the Aggregate planning problem and the utility assignment $\{\hat{u}_t^i \in Y(\{u_t^i\}, 0)$ that solve the Idiosyncratic planning problem as well as the corresponding idiosyncratic efficiency gains $\tilde{\lambda}^I$.

Since the resource constraints hold with equality in the Idiosyncratic planning problem, we know from [Lemma 1](#) that there exists a sequence of prices $\{\hat{Q}_t\}$, such that

$$c'(\hat{u}_t^i) = \frac{\hat{q}_t}{\beta} \mathbb{E}_t^i[c'(\hat{u}_{t+1}^i)] \quad t = 0, 1, \dots \quad (26)$$

where $\hat{q}_t \equiv \hat{Q}_{t+1}/\hat{Q}_t$. Moreover, the sequence $\{\hat{q}_t\}$ is given by $\hat{q}_t = \beta C_t/C_{t+1}$.

The aggregate allocation $\{\tilde{C}_t, N_t, \tilde{K}_t, \tilde{\lambda}C_t\}$ satisfies the necessary and sufficient first-order conditions

$$U'(\tilde{C}_t) = \beta (1 - \delta + F_K(\tilde{K}_{t+1}, N_{t+1})) U'(\tilde{C}_{t+1}) \quad t = 0, 1, \dots \quad (27)$$

Define the following sequence $\{\delta_t\}$

$$\delta_t = -U((1 - \tilde{\lambda}^I)C_t) + U(\tilde{C}_t) \quad t = 0, 1, \dots$$

We have $\sum_{t=0}^{\infty} \beta^t \delta_t = 0$. With this choice of $\{\delta_t\}$, we then define a utility assignment

$\{\tilde{u}_t^i\} \in Y(\{u_t^i\}, 0)$ as follows $\tilde{u}_t^i = \hat{u}_t^i + \delta_t$. The allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda}C_t\}$ then satisfies all the constraints of the original planning problem. Moreover using $c'(\tilde{u}_t^i) = c'(\delta_t)c'(\hat{u}_t^i)$, [equation \(26\)](#) and [equation \(27\)](#), we find that

$$c'(\tilde{u}_t^i) = \frac{\tilde{q}_t}{\beta} \mathbb{E}_t^i[c'(\tilde{u}_{t+1}^i)] \quad t = 0, 1, \dots$$

where

$$1 = \tilde{q}_t (F_K(\tilde{K}_{t+1}, N_{t+1}) + 1 - \delta) \quad t = 0, 1, \dots$$

Hence the allocation $\{\tilde{u}_t^i, n_t^i, \tilde{K}_t, \tilde{\lambda}C_t\}$ satisfies the sufficient first order conditions in the planning problem. It therefore represents the optimum.

E Numerical Method

To solve the planning problem we use the result developed in [Appendix B.1](#).

To apply this result, we seek the appropriate sequence of discount factors $\{\tilde{q}_t\}$ as follows. For any given $\{\tilde{q}_t\}$, we solve the non-stationary Bellman [equation \(21\)](#) using a policy iteration method with Δ_{t-1} as the endogenous state variable and s_t as the exogenous state. Using the underlying policy function for consumption in [equation \(21\)](#), and integrating in every period over ψ , we compute an aggregate sequence of consumption $\{\tilde{C}_t\}$. Using the resource constraints, we can solve for $\tilde{\lambda}$ and a sequence for capital $\{\tilde{K}_t\}$ that has $\tilde{K}_0 = K_{ss}$.

From [Appendix B.1](#), if the condition $1/\tilde{q}_t = 1 + F_K(\tilde{K}_t, N_{ss}) - \delta$ is met this constitutes a solution. Otherwise, we take a new sequence of discount factors given by $\tilde{q}'_t = (1 + F_K(\tilde{K}_t, N_{ss}) - \delta)^{-1}$ and iterate until convergence to a fixed point.

References

- John M Abowd and David Card. Intertemporal labor supply and long-term employment contracts. *American Economic Review*, 77(1):50–68, March 1987. 32
- John M Abowd and David Card. On the covariance structure of earnings and hours changes. *Econometrica*, 57(2):411–45, March 1989. 32
- Mark Aguiar and Erik Hurst. Consumption versus expenditure. *Journal of Political Economy*, 113(5):919–948, 2005. 28
- S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics*, 109(3):659–684, 1994. 22, 27, 31, 32, 34, 35, 36, 37, 44
- Andrew Atkeson and Robert E., Jr. Lucas. On efficient distribution with private information. *Review of Economic Studies*, v59(3):427–453, 1992. 42
- A. B. Atkinson and J. E. Stiglitz. The design of tax structure: Direct versus indirect taxation. *Journal of Public Economics*, 6(1-2):55–75, 1976. 2
- Richard Blundell, Luigi Pistaferri, and Ian Preston. Consumption inequality and partial insurance. IFS Working Papers W04/28, Institute for Fiscal Studies, 2004. 28
- Christophe Chamley. Optimal taxation of capital in general equilibrium. *Econometrica*, 54:607–22, 1986. 2
- Juan Carlos Conesa and Dirk Krueger. On the optimal progressivity of the income tax code. *Journal of Monetary Economics*, forthcoming, 2005. 6
- Julio Davila, Jay H. Hong, Per Krusell, and Jose-Victor Rios-Rull. Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks. PIER Working Paper Archive 05-023, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania, July 2005. available at <http://ideas.repec.org/p/pen/papers/05-023.html>. 6
- Angus Deaton and Christina Paxson. Intertemporal choice and inequality. *Journal of Political Economy*, 102(3):437–67, 1994. 29, 32
- Peter A. Diamond and James A. Mirrlees. A model of social insurance with variable retirement. Working papers 210, Massachusetts Institute of Technology, Department of Economics, 1977. 2, 3
- Emmanuel Farhi and Iván Werning. Inequality and social discounting. *Journal of Political Economy*, 115(3):365–402, 2007. 42
- Emmanuel Farhi and Iván Werning. Optimal savings distortions with recursive preferences. *Journal of Monetary Economics*, 55:21–42, 2008. 38
- Emmanuel Farhi and Ivan Werning. Insurance and taxation over the life cycle. Working paper 16749, NBER, 2010. 5

- Milton. Friedman. *Capitalism and Freedom*. University of Chicago Press, 1962. 2
- John Geanakoplos and Heracles M. Polemarchakis. Existence, regularity, and constrained suboptimality of competitive allocations when the asset market is incomplete. Cowles Foundation Discussion Papers 764, Cowles Foundation, Yale University, 1985. available at <http://ideas.repec.org/p/cwl/cwldpp/764.html>. 6
- Michael Golosov, Maxim Troshkin, and Aleh Tsyvinski. Optimal dynamic taxes. Working paper, Yale University, 2010. 5
- Mikhail Golosov and Aleh Tsyvinski. Designing optimal disability insurance: A case for asset testing. *Journal of Political Economy*, 114(2):257–279, April 2006. 3, 6
- Mikhail Golosov, Narayana Kocherlakota, and Aleh Tsyvinski. Optimal indirect and capital taxation. *Review of Economic Studies*, 70(3):569–587, 2003. 2, 15, 20
- Mikhail Golosov, Aleh Tsyvinski, and Iván Werning. New dynamic public finance: A user’s guide. *NBER Macroeconomics Annual 2006*, 2006. 5
- Fatih Guvenen. Learning your earning: Are labor income shocks really very persistent? *American Economic Review*, 97(3):687–712, June 2007. 32
- Fatih Guvenen. An empirical investigation of labor income processes. *Review of Economic Dynamics*, 12(1):58–79, January 2009. 32
- John C. Hause. The fine structure of earnings and the on-the-job training hypothesis. *Econometrica*, 48:1013–1029, 1980. 32
- Jonathan Heathcote, Kjetil Storesletten, and Giovanni L. Violante. Two views of inequality over the life-cycle. *Journal of the European Economic Association (Papers and Proceedings)*, 3:543–52, 2004. 29
- John Heaton and Deborah J Lucas. Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy*, 104(3):443–87, June 1996. 32
- Hugo A Hopenhayn and Juan Pablo Nicolini. Optimal unemployment insurance. *Journal of Political Economy*, 105(2):412–38, April 1997. 2
- Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5–6):953–969, 1993. 22, 27
- Kenneth L. Judd. Redistributive taxation in a perfect foresight model. *Journal of Public Economics*, 28:59–83, 1985. 2
- Dirk Krueger and Fabrizio Perri. On the welfare consequences of the increase in inequality in the united states. In Mark Gertler and Kenneth Rogoff, editors, *NBER Macroeconomics Annual 2003*, pages 83–121. MIT Press, Cambridge, MA, 2004. 29
- Finn E. Kydland. A clarification: Using the growth model to account for fluctuations: Reply. *Carnegie-Rochester Conference Series on Public Policy*, 21:225–252, 1984. 32

- Ethan Ligon. Risk sharing and information in village economics. *Review of Economic Studies*, 65(4):847–64, 1998. 2
- Lee A. Lillard and Yoram A. Weiss. Components of variation in panel earnings data: American scientists, 1960-70. *Econometrica*, 47:437–454, 1979. 32
- Jr Lucas, Robert E. Supply-side economics: An analytical review. *Oxford Economic Papers*, 42(2):293–316, April 1990. 2, 24, 28
- David G. Luenberger. *Optimization by Vector Space Methods*. John Wiley & Sons, Inc, 1969. 42
- Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, March 2002. 40
- James A. Mirrlees. An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38(2):175–208, 1971. 7
- William P. Rogerson. Repeated moral hazard. *Econometrica*, 53(1):69–76, 1985. 2
- Eytan Sheshinski. The optimal linear income-tax. *Review of Economic Studies*, 39(3):297–302, July 1972. 2
- Robert Shimer and Ivan Werning. Liquidity and insurance for the unemployed. *American Economic Review*, 98(5):1922–42, December 2008. 3, 6
- Daniel T. Slesnick and Aydogan Ulker. Inequality and the life-cycle: Consumption. Technical report, 2004. 29
- Stephen E. Spear and Sanjay Srivastava. On repeated moral hazard with discounting. *Review of Economic Studies*, 54(4):599–617, 1987. 22
- Nancy L. Stokey and Sergio Rebelo. Growth effects of flat-rate taxes. *Journal of Political Economy*, 103(3):519–50, 1995. 27
- Kjetil Storesletten, Chris I. Telmer, and Amir Yaron. Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717, 2004a. 25, 28, 32
- Kjetil Storesletten, Christopher I. Telmer, and Amir Yaron. Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633, April 2004b. 32