Available online at www.sciencedirect.com
ScienceDirect
Journal of Monetary Economics 55 (2008) 21–42
Journal of MONETARY ECONOMICS
www.elsevier.com/locate/jme

Optimal savings distortions with recursive preferences ☆

Emmanuel Farhi[a,*], Iván Werning[b]

[a] *Harvard, USA*
[b] *MIT, USA*

Received 28 May 2007; received in revised form 3 December 2007; accepted 4 December 2007
Available online 23 December 2007

Abstract

This paper derives an intertemporal optimality condition for economies with private information, focusing on a class of recursive preferences. By comparing it to the situation where agents can freely save in a risk-free asset market, we derive the optimal savings distortions necessary for constrained optimality. Our recursive preferences are homogeneous and satisfy a balanced-growth condition, while allowing us to separate the role of risk aversion and intertemporal elasticity of substitution. We perform some quantitative exercises that disentangle the respective roles played by these two parameters in optimal distortions and the implied welfare gains.

Keywords: Optimal savings

1. Introduction

When perfect insurance is unavailable, savings may help individuals smooth the impact on consumption of temporary shocks to income. However, models that derive imperfect insurance from private information suggest banning free access to savings. Constrained efficient allocations in these economies require some distortion in individuals' savings decisions (Diamond and Mirrlees, 1977; Rogerson, 1985; Ligon, 1998; Golosov et al., 2003; Farhi and Werning, 2006). The goal of this paper is to further our understanding of the differences between constrained-efficient allocations and market equilibria. In particular, we investigate the role of preferences.

It is useful to frame the comparison of the market equilibrium and the planning problem in terms of the different variations on consumption plans that are feasible in each case.

With unfettered access to a risk-free asset, agents can perform the following variation to their consumption plans. At any point in time, individuals can lower their current consumption by one unit and increase it in all future periods and contingencies by a constant *absolute* amount, equal to the net rate of return. At a market equilibrium, individuals find themselves at an optimum within this class of variations. The corresponding optimality condition is the familiar intertemporal Euler equation.

☆We thank George-Marios Angeletos for useful comments and insightful discussions. All remaining errors are our own.
*Corresponding author.
E-mail addresses: efarhi@harvard.edu (E. Farhi), iwerning@mit.edu (I. Werning).

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2007.12.002

Instead, a planner must consider the response that any change in the consumption plan may have on work effort, if the latter is not fully under her control due to private information. In general, there exists a class of variations available to the planner on the agent's consumption such that incentives and work effort are preserved. At the constrained-efficient allocation, the planner finds the optimum within this class of variations.

The variations available to the planner do not always, or even typically, coincide with those available to agents in a free-market equilibrium. Differences in these sets of variations lead to optimality conditions that are potentially incompatible. Distortions on savings may then be required to implement the constrained optimum with an asset market.

The first point emphasized by this paper is that the particular form that the set of allowable variations for the planner takes, depends critically on preferences. We begin by showing that there exists a particular class of preferences for which the set of variations available to the planner actually coincides with that available to agents in a free market. As a result, the constrained efficient allocation requires no distortions on agents' savings. The preferences required for this result feature no income effects on work effort. This particular result demonstrates that the form of the discrepancy between the constrained-optimum and the market equilibrium is likely to depend, in general, on preference assumptions.

Next, we propose a class of homogeneous preferences with a balanced-growth condition on work effort that delivers a simple and intuitive class of variations. The allowable variations on consumption for the planner in this case are as follows. At any point in time, the planner can lower the agent's current consumption and increase it in all future periods and contingencies by a constant *proportional* amount. This type of variation is not available to the agent through the asset market, which opens up the possibility for the planner to find Pareto-improvements. The optimal savings distortions are dictated by the difference between the absolute and proportional variations on consumption available to the agent and planner, respectively.

Proportional changes in consumption leave incentives unaltered precisely because preferences are homogeneous and satisfy a balanced-growth condition. We believe that the simplicity and plausibility of these variations is a desirable feature of the preferences we propose. They lead to simple intuitions, transparent theoretical results and a tractable framework for quantitative analysis.

Within this class of variations the resulting optimality condition is extremely simple. It requires that the ratio of current utility to lifetime utility always equal the ratio of current consumption to the expected present discounted value of lifetime consumption. We term this simple optimality condition the *Golden Ratio*. It can also be stated as a *Modified Inverse Euler* equation in a form that resembles the standard Inverse Euler equation that was derived as a necessary condition for optimality for the variations considered in Farhi and Werning (2006).

These preferences have three advantages. First, they are flexible enough to allow us to study the respective impact of two crucial parameters: the coefficient of relative risk aversion and the intertemporal elasticity of substitution. Second, although they feature nonseparability of consumption and work effort, these preferences call for no savings distortions in the absence of recurring uncertainty—just as the separable preferences studied in the literature on the Inverse Euler equation. Third, they lead to a very clean separation result for welfare gains between an idiosyncratic part and an aggregate part.

Towards the end of the paper, we perform some quantitative welfare exercises that compute the gains from optimal savings distortions. We follow Farhi and Werning (2006), where we developed a new approach to analyze the welfare gains from distorting savings and moving away from letting individuals save freely. The method forgoes a complete solution for both consumption and work effort, and focuses, instead, entirely on consumption. We restrict our attention to the case of geometric random walk consumption and constant work effort. Our main goal is to isolate and compare the effects that the intertemporal elasticity of substitution and the coefficient of relative risk aversion have on the size of the intertemporal wedge and the welfare gains from optimal distortions. Thus, although we borrow from Farhi and Werning (2006), the focus in that paper was on the generality in terms of the stochastic process for the baseline allocation of consumption. Instead, our focus here is on a set of stylized baseline allocations that allow us to clearly separate the impact of different preferences assumptions.

Welfare gains depend crucially on four factors: the concavity of the production function, the coefficient of relative risk aversion γ, the intertemporal elasticity of substitution ρ^{-1} and the variance of consumption growth σ_ε^2.

As in Farhi and Werning (2006), we find that gains are decreasing in the concavity of the production function. In partial equilibrium with a linear production function, gains can be extremely large. By contrast, for an endowment economy welfare gains are zero under our hypothesis of a geometric random walk consumption process. For the intermediate case of a neoclassical production function, welfare gains are greatly mitigated.

The steady state of the optimal allocation with savings distortions feature a lower capital stock and a higher interest rate than the corresponding steady state of the market equilibrium, where the precautionary savings motive is at work. The variance of consumption growth and the coefficient of relative risk aversion control the strength of this motive and hence both the interest rate increase and the decrease in capital between the baseline steady state and the optimal steady state. The intertemporal elasticity of substitution on the other hand controls the speed of the transition: the higher ρ^{-1}, the faster the transition, and the higher the welfare gains. The configuration of these three parameters influences greatly the magnitude of the welfare gains.

2. Constrained efficiency vs. free savings

In this section we present a two period economy to introduce the basic concepts and set the stage for the rest of the paper. Against this background, in the next section we turn to an infinite horizon economy with recursive preferences.

Consider a simple economy with two periods $t = 0, 1$. There is no uncertainty at $t = 0$ but at the beginning of period $t = 1$ a state $s_1 \in S$ is realized; we assume S is finite, with $\#S$ values and $p(s)$ is the probability of outcome $s_1 = s$. The agent consumes in the first period and consumes and works in the second. Let c_0 denote consumption in the first period and $(c_1(s), Y_1(s))$ denote consumption and output as a function of the realized state in the second period.

We adopt a general specification of preferences and denote the agent's utility functional over allocations by $U(c_0, c_1(\cdot), Y_1(\cdot))$. Thus, U takes a scalar c_0 and two functions $c_1(\cdot)$ and $Y_1(\cdot)$ as inputs. As a special benchmark case, one can assume the state s_1 determines the worker's productivity and that the worker has an expected-utility function $u(c_0, c_1, e_1)$ over consumption in both periods and work effort $e_1(s) \equiv Y_1(s)/s$. Then $U(c_0, c_1(\cdot), Y_1(\cdot)) = \mathbb{E}[u(c_0, c_1(s), Y_1(s)/s)]$.

Technology is linear

$$c_0 + q\sum_{s\in S} c_1(s)p(s) \leqslant q\sum_{s\in S}{}_1 Y_1(s)p(s) \tag{1}$$

for some $q > 0$. Here, $R = 1/q$ is the rate of return between periods 0 and 1.

2.1. Free savings

2.1.1. First-best

The first-best allocation simply maximizes utility subject only to technology equation (1). At this allocation the first-order conditions for consumption are given by

$$U_{c_0}(c_0, c_1(\cdot), Y_1(\cdot)) = \mu,$$
$$U_{c_1(s)}(c_0, c_1(\cdot), Y_1(\cdot)) = qp(s)\mu,$$

where μ is the multiplier on the resource constraint. The first-order conditions for consumption can be combined into the following generalized Euler equation:

$$1 = \frac{1}{q}\sum_{s\in S}\frac{U_{c_1(s)}(c_0, c_1(\cdot), Y_1(\cdot))}{U_{c_0}(c_0, c_1(\cdot), Y_1(\cdot))}. \tag{2}$$

In the expected-utility case this equation specializes to the familiar Euler equation

$$1 = R\mathbb{E}\left[\frac{u_{c_1}(c_0, c_1(s), e_1(s))}{u_{c_0}(c_0, c_1(s), e_1(s))}\right]. \tag{3}$$

2.1.2. Competitive equilibrium with free savings

The Euler equation (2) also obtains in a free-market economy where individuals have access to saving at rate of return R. For example, suppose that agents live in an incomplete market setting, facing the budget constraints

$$c_0 + k_1 \leqslant 0, \tag{4a}$$

$$c_1(s) \leqslant Y_1(s) + Rk_1 \quad \forall s \in S. \tag{4b}$$

Then the first-order conditions for the agent's utility maximization problem with respect to savings k_1 delivers Eq. (2).[1] Note that the budget constraints (4a)–(4b) imply the resource constraint (1).

2.1.3. A general set-up

More generally, under what conditions does (2) hold? Consider the abstract optimization problem of maximizing utility $U(c_0, c_1(\cdot), Y_1(\cdot))$ subject to

$$(c_0, c_1(\cdot), Y_1(\cdot)) \in \mathscr{F}$$

for some constraint set $\mathscr{F}$. This nests as special cases both the first-best planning problem—with $\mathscr{F} = \mathscr{F}_{\text{fb}}$ defined by the resource constraint (1)—and the agent's optimization in the free-market setting—with $\mathscr{F} = \mathscr{F}_{\text{fm}}$ defined by the budget constraints (4a)–(4b). Suppose that starting from any allocation $(c_0, c_1(\cdot), Y_1(\cdot)) \in \mathscr{F}$ it is possible to define simple variations that maintain the allocation in $\mathscr{F}$:

$$(c_0 - q\Delta, c_1(\cdot) + \Delta, Y_1(\cdot)) \in \mathscr{F} \tag{5}$$

for all Δ in neighborhood of $\Delta = 0$. That is, a feasible allocation can be perturbed by decreasing (increasing) consumption in the first period, while increasing (decreasing) consumption in parallel across all states s in the second period. Note that the same output allocation $Y_1(s)$, $Y_1(s)/s$, is maintained for all states s.

Property (5) holds for both the first-best planning problem and the agent's optimization problem in a free-market setting. More generally, whenever it is satisfied at an optimum, then the generalized Euler equation (2) must be satisfied.

2.1.4. Second-best with private information

Consider next a private-information setting, where the state s is observed only by the agent. By the revelation principle, the best the planner can do is to request a report $r \in S$ from the agent regarding $s \in S$ and assign consumption and output in the second period accordingly. Without loss of generality, one can assume that telling the truth is optimal.

Let $r = \sigma(s)$ denote a reporting strategy for the agent, mapping true states of the world $s \in S$ into reports $r \in S$. Let Σ denote the set of all strategies. The truth-telling strategy is denoted by $\sigma^*(s) = s$ for all $s \in S$. An agent using strategy $\sigma \in \Sigma$ obtains $(c_1^\sigma(s), Y_1^\sigma(s)) = (c_1(\sigma(s)), Y_1(\sigma(s)))$ in state s. Incentive-compatibility can be expressed as

$$U(c_0, c_1(\cdot), Y_1(\cdot)) \geqslant U(c_0, c_1^\sigma(\cdot), Y_1^\sigma(\cdot)) \quad \forall \sigma \in \Sigma. \tag{6}$$

The second-best planning problem corresponds to the case where $\mathscr{F} = \mathscr{F}_{\text{sb}}$ defined by equations (1) and (6). A second-best optimum maximizes utility subject to selecting an allocation in $\mathscr{F}_{\text{sb}}$.

In this general context, typically property (5) with $\mathscr{F}_{\text{sb}}$ fails. The next proposition, however, provides an example where it holds.

Proposition 1. *Let* $U(c_0, c_1(\cdot), Y_1(\cdot)) = \hat{U}(c_0, c_1(\cdot) - v(Y_1(\cdot), \cdot))$ *where* $\hat{U}$ *monotone in its second argument. Then property* (5) *holds for* $\mathscr{F}_{\text{sb}}$ *for all feasible allocations* $(c_0, c_1(\cdot), Y_1(\cdot)) \in \mathscr{F}_{\text{sb}}$.

[1] Indeed, this result holds more generally, even if we assume that there are some taxes and transfers that are a function of output or the state, so that we impose $c_1(s) \leqslant T(Y_1(s), s) + Y_1(s) + Rk_1$ in the second period.

Proof. The result follows by noting that incentive compatibility (6) holds if and only if

$$c(s) - v(Y(s), s) \geqslant c(r) - v(Y(r), s) \quad \forall r, s \in S,$$

which is independent of c_0 and invariant to the operation of exchanging $c(\cdot)$ for $c(\cdot) + \Delta$ for any Δ. □

If property (5) holds for all Δ (not just in a neighborhood around $\Delta = 0$) then it is without loss of generality to allow agents to freely save, in the sense that the planner can allow the agent to select the value for Δ in this variation. It follows that, for the class of preferences identified by the proposition, the planner can allow the agent to save freely, without distortions, at the technological rate of return $R = 1/q$. The economic interpretation of the quasi-linear specification $c - v(Y; s)$ is that there are no income effects on work effort. Savings from the first period do not then affect the choice between work effort and earnings. As a result, they do not disturb incentive compatibility and property (5) holds.

An equivalent way of postulating property (5) is as follows. Any direct mechanism $(c_0 - q\Delta, c_1(r) + \Delta, Y_1(r))$ essentially offers the agent an ex post menu in each state s equal to the loci of points $(c_1(\cdot) + \Delta, Y_1(\cdot))$. In each state s, the agent selects an optimal point on this menu, (c_1^*, Y_1^*). Property (5) then amounts to assuming that this optimum Y_1^* is invariant to Δ. Proposition 1 then identifies the largest class of preferences that guarantee that this is the case for all feasible allocations.

2.2. *Distorted savings*

From the previous subsection, we know that the variations that result from free savings do not generally preserve incentive compatibility. In this situation, what can we say about the desirability of free savings? We approach this question in two complementary ways.

2.2.1. *A Lagrangian approach*

The first is to attach Lagrange multiplier $\mu(\sigma)$ on the incentive constraints (6), leading to an optimality condition that includes the effect that Δ may have on incentive constraints:

$$\begin{aligned}\frac{\partial}{\partial \Delta}\mathscr{L} = {} & \left(1 + \sum_{\sigma \in \Sigma} \mu(\sigma)\right)\left(-qU_{c_0}(c_0, c_1(\cdot), Y_1(\cdot)) + \sum_{s \in S} U_{c_1(s)}(c_0, c_1(\cdot), Y_1(\cdot))\right) \\ & - \sum_{\sigma \in \Sigma} \mu(\sigma)\left(-qU_{c_0}(c_0, c_1(\sigma(\cdot)), Y_1(\sigma(\cdot))) + \sum_{s \in S} U_{c_1(s)}(c_0, c_1(\sigma(\cdot)), Y_1(\sigma(\cdot)))\right) \\ = {} & 0.\end{aligned}$$

Note that if all the incentive constraints are slack, so that $\mu(\sigma) = 0$ for all $\sigma \in \Sigma$, then this expression boils down to the Euler equation (2). Otherwise, the Euler equation (2) will typically not hold. Indeed, if one signs the term $-qU_{c_0}(c_0, c_1(\sigma(\cdot)), Y_1(\sigma(\cdot))) + \sum_{s \in S} U_{c_1(s)}(c_0, c_1(\sigma(\cdot)), Y_1(\sigma(\cdot)))$ for different strategies σ and characterizes which multipliers are nonzero, then one can sign the intertemporal wedge required in the Euler equation.

2.2.2. *Feasible variations*

Another line of attack is to find a different variation, that does preserve incentive compatibility, without changing work effort. This leads to an intertemporal optimality condition that does not involve Lagrange multipliers. One can then compare this optimality condition with the Euler equation (2).

The idea is to find a variation function $\delta(\Delta, s)$ on consumption in the second period that depends on the realized state s so that

$$(c_0 + \Delta, c_1(\cdot) + \delta(\Delta, \cdot), Y_1(\cdot)) \in \mathscr{F} \tag{7}$$

in a neighborhood of $\Delta = 0$. At an optimum we must then have that

$$U_{c_0}(c_0, c_1(\cdot), e_1(\cdot)) + \sum_{s \in S} U_{c_1(s)}(c_0, c_1(\cdot), e_1(\cdot)) \cdot \frac{\partial}{\partial \Delta}\delta(0, s) = 0. \tag{8}$$

For example, with expected utility and $u(c_0, c_1, e_1) = \hat{u}(c_0, c_1) - h(e_1)$ a variation that is feasible is to set $\delta(\Delta, s)$ so that

$$\hat{u}(c_0 + \Delta, c_1(s) + \delta(\Delta, s)) = \hat{u}(c_0, c_1(s)) + A(\Delta) \quad \forall s \in S, \tag{9}$$

where $A(\Delta)$ is such that

$$\sum_{s \in S} (\Delta + \delta(\Delta, s)) p(s) = 0. \tag{10}$$

This variation shifts utility in a parallel way across states $s \in S$. It preserves incentive compatibility because these parallel shifts cancel each other out on both sides of Eq. (6). At an optimum $A'(0) = 0$ so that

$$\frac{\partial}{\partial \Delta} \delta(0, s) = -\frac{\hat{u}_{c_0}(c_0, c_1(s))}{\hat{u}_{c_1}(c_0, c_1(s))}. \tag{11}$$

It then follows that

$$1 = \sum_{s \in S} \frac{\hat{u}_{c_0}(c_0, c_1(s))}{\hat{u}_{c_1}(c_0, c_1(s))} p(s), \tag{12}$$

which is known as the Inverse Euler equation. By Jensen's inequality, this condition is incompatible with the Euler equation (3), except in the special case where there is no uncertainty in the marginal rate of substitution ratio $\hat{u}_{c_0}(c_0, c_1(s))/\hat{u}_{c_1}(c_0, c_1(s))$. Without uncertainty the optimality of no intertemporal distortions follows from Atkinson–Stiglitz's (1976) result on uniform taxation, which requires separability between consumption and effort, as assumed in this case.

2.2.3. *Logarithmic balanced-growth preferences*

Within this class of preferences, an interesting special case with several advantages is the logarithmic balanced-growth specification $u(c_0, c_1) = \log(c_0) + \beta \log(c_1)$. In this case the variations induce parallel multiplicative shifts over second-period consumption:

$$\delta(\Delta, s) = \bar{\delta}(\Delta) c_1(s) \tag{13}$$

for some $\bar{\delta}(\Delta)$. Intuitively, incentives are provided by proportional rewards and punishments. If consumption is scaled up or down by a constant it does not change the incentives for work effort.

In this case, unlike the preference class described in Proposition 1, income effects for work effort are nonzero. Proportional variations are feasible precisely because of the balanced-growth condition, which implies that income and substitution effects exactly cancel each other.

This logarithmic case seems economically appealing, because of the primitives and the simple proportional variations it permits. One simple generalization of this case is to the expected-utility case where

$$u(c_0, c_1, e_1) = \tilde{u}(c_0) + \beta \tilde{u}(c_1) h(e_1) \tag{14}$$

and where $\tilde{u}(c) = c^{1-\alpha}/(1-\alpha)$. This class of preferences also satisfies a balanced-growth condition. It is easily verified that once again the feasible variations are proportional in consumption, as in (13).

In the next section we extend this class to an infinite horizon economy. Preferences that lead to the feasibility of proportional variations turn out to be very tractable. In particular, they lead to a very simple optimality condition. Within a class of baseline allocations, the optimum is easily identified and its welfare improvements quantified.

3. Recursive preferences

We now turn to an infinite horizon and introduce a class of recursive preferences that are homogeneous in the consumption process and separate risk aversion from the intertemporal elasticity of substitution as in Epstein and Zin (1989). Consumption and work effort are not assumed to be separable, but satisfy a balanced-growth condition.

For this class of preferences, we provide simple variations on consumption that maintain incentive compatibility. The variations involve proportional shifts in consumption that do not affect incentives. Both the homogeneity and the balanced-growth specification on preferences are crucial for this result.

Based on these variations we derive the intertemporal optimality condition at the end of the section. The condition is shown to be incompatible with allowing agents to freely save. In this way, an intertemporal wedge on savings is present at the optimal allocation. Thus, some form of distortion on savings is required in any tax implementation of the optimum. In the next section we explore the welfare gains from adhering to this condition for some simple cases.

Our preferences do not satisfy the separability condition required for Atkinson–Stiglitz's uniform taxation theorem. Despite this, it is optimal in the absence of uncertainty to set the intertemporal distortions to zero. Thus, for these preferences, optimal distortions in savings arise from ongoing idiosyncratic uncertainty, just as in the additively separable expected-utility case that leads to the Inverse Euler condition.

3.1. Moral hazard

We build on the following simple static moral-hazard model. At the beginning of the period, the agent first exerts effort a, which is not observable by the planner. The state of nature s is then realized from the distribution $\mathrm{P}(s|a)$. The planner observes s and gives the agent consumption $c(s)$. The agent's expected utility is given by

$$\mathbb{E}[U(c(s)h(a))|a].$$

We suppose the agent's utility $U(c)$ is a power function. This specification satisfies the standard balanced-growth assumption, for which income and substitution effects cancel out. An equivalent reformulation of the agent's objective is

$$U(Ch(a)),$$

where

$$C \equiv \mathbb{CE}[c(s)|a] = U^{-1}(\mathbb{E}[U(c(s))|a])$$

represents the certainty-equivalent obtained from the random consumption $c(s)$.

For our dynamic setting, we proceed analogously. At the start of period t the worker chooses effort a_{t-1}, then the state s_t is realized and observed and the planner allocates consumption $c(s^t)$. Effort affects the distribution of state s_t and lowers utility by a factor $h(a_t)\leqslant 1$ with $h(0)=1$. Preferences are given by the recursion

$$\hat{v}_a(s^{t-1}) = C(s^t)h(a(s^{t-1})),$$

where

$$C(s^t) \equiv \mathbb{CE}[W(c(s^t), \hat{v}_a(s^t))|a(s^{t-1}), s^{t-1}] \tag{15}$$

represents lifetime-certainty-equivalent consumption, with

$$\mathbb{CE} = R^{-1}\mathbb{E}R \tag{16}$$

is the certainty-equivalent function and

$$W(c, \hat{v}) \equiv u^{-1}((1-\beta)u(c) + \beta u(\hat{v})) \tag{17}$$

is a time aggregator, mapping current consumption and future utility into a constant-consumption equivalent.

With this representation of preferences, one can easily see the analogy with the simple static setting. By a change of variables, however, the same preferences can be represented in the following, more convenient, way. For any given effort plan $a \equiv \{a(s^t)\}$, an allocation $c \equiv \{c(s^t)\}$ implies a process for lifetime utility $\{v(s^t|a)\}$ that solves

$$v_a(s^t) = W(c(s^t), \mathbb{CE}[h(a(s^t))v_a(s^{t+1})|a(s^t), s^t]) \quad \forall t, s^t. \tag{18}$$

Incentive compatibility of c, v and a^* requires a^* to maximize initial lifetime utility

$$v_{a^*}(s_0) \geqslant v_a(s_0) \quad \forall a. \tag{19}$$

Since preferences are recursive, this implies that a^* maximizes continuation utility after any history

$$v_{a^*}(s^t) \geqslant v_a(s^t) \quad \forall a, t, s^t. \tag{20}$$

Otherwise, a plan that follows a^* up to s^t and then switches to the actions prescribed by a at and after s^t would be preferable to a^*. That is, Bellman's Principle of Optimality applies to the agent's dynamic program.

We now consider variations in the consumption process that maintain incentive compatibility. After history s^τ the consumption sequence is just shifted proportionally, and this does not affect incentives. At s^τ we shift consumption to compensate, so that incentives are not affected in period τ and earlier periods. The key property we use is the homogeneity of $W(c, v')$ and of $\mathbb{CE}$.

Proposition 2. *Assume $u(x) = x^{1-\rho}/(1-\rho)$ and $R(x) = x^{1-\gamma}/(1-\gamma)$ with $\rho, \gamma \geqslant 0$. Suppose that c, v and a^* satisfy conditions* (18) *and* (19). *Fix a history s^τ. Consider the variation*:

$$\tilde{c}(s^t) = \begin{cases} \Delta c(s^\tau) & \text{for } s^t = s^\tau, \\ \Delta' c(s^t) & \text{for } t > \tau \text{ and } s^t \succ s^\tau, \\ c(s^t) & \text{otherwise.} \end{cases}$$

Then for any Δ' there exists a Δ such that $\tilde{c}$, $\tilde{v}$ and a^* satisfy conditions (18) *and* (19).

Proof. Let $\tilde{v}$ be such that

$$\tilde{v}_a(s^t) = \Delta' v_a(s^t) \quad \text{for } t > \tau \text{ and } s^t \succ s^\tau,$$
$$\tilde{v}_a(s^t) = v_a(s^t) \quad \text{for } t \geqslant \tau \text{ and } s^t \not\succeq s^\tau,$$

so that condition (18) with $\tilde{c}$ is met for all s^t with $t \geqslant \tau$ with $s^t \neq s^\tau$. Now set Δ so that

$$v_{a^*}(s^\tau) = W(\Delta c(s^t), \Delta' \mathbb{CE}[h(a^*(s^\tau)) v_{a^*}(s^{\tau+1}) | a^*(s^\tau), s^\tau]),$$

so that $\tilde{v}_{a^*}(s^\tau) = v_{a^*}(s^\tau)$. Using recursion (18), the inequality (20) evaluated at s^τ implies

$$\mathbb{CE}[h(a^*(s^\tau)) v_{a^*}(s^{\tau+1}) | a^*(s^\tau), s^\tau] \geqslant \mathbb{CE}[h(a(s^\tau)) v_a(s^{\tau+1}) | a(s^\tau), s^\tau],$$

so that

$$\begin{aligned} \tilde{v}_{a^*}(s^\tau) &= W(\Delta c(s^t), \Delta' \mathbb{CE}[h(a^*(s^\tau)) v_{a^*}(s^{\tau+1}) | a^*(s^\tau), s^\tau]) \\ &\geqslant W(\Delta c(s^t), \Delta' \mathbb{CE}[h(a(s^\tau)) v_a(s^{\tau+1}) | a(s^\tau), s^\tau]) = \tilde{v}_a(s^\tau). \end{aligned}$$

Hence, we have that

$$\tilde{v}_a(s^\tau) \leqslant \tilde{v}_{a^*}(s^\tau) = v_{a^*}(s^\tau) \quad \text{for all } a,$$

a^* is optimal from period τ onward and delivers the same continuation utility as previously.

For any plan a define an alternative plan $\hat{a}$ that switches to a^* from period τ onward: $\hat{a}(s^t) = a(s^t)$ for $t < \tau$ and $\hat{a}(s^t) = a^*(s^t)$ for $t \geqslant \tau$. The result above implies that

$$\tilde{v}_a(s_0) \leqslant \tilde{v}_{\hat{a}}(s_0) = v_{\hat{a}}(s_0) \leqslant v_{a^*}(s_0) = \tilde{v}_{a^*}(s_0). \tag{21}$$

That is, $\hat{a}$ dominates a and yields the same utility as without the variation, which in turn is dominated by the recommended action a^* which also yields the same utility as after the variation. This establishes that a^* remains incentive compatible. □

3.2. Private information: a dynamic Mirrleesian economy

Here we build on Mirrlees' static private information model. At the beginning of the period, the agent privately observes productivity θ. The agent then makes a report r and the planner gives the agent consumption $c(r)$ as function of the report. The agent's expected utility is

$$\mathbb{E}[U(c(r)h(r,\theta))|\sigma],$$

where $r = \sigma(\theta)$ is the agent's reporting strategy. We suppose the agent's utility $U(c)$ is a power function. This specification satisfies the standard balanced-growth assumption, for which income and substitution effects cancel out.

For our dynamic setting, we assume the following structure of uncertainty. At the beginning of the period a state s_t is realized and publicly observed by the agent and planner. Then θ_t is realized and observed only by the agent. To simplify we assume that s_t and θ_t take on a finite number of values. After observing the shock θ_t the agent makes a report r_t regarding it to the planner. We collect the variables observed by the planner by $z_t = (s_t, r_t)$ and their histories by $z^t = (s^t, r^t)$.

For any reporting strategy σ

$$v_\sigma(z^t,\theta^t) = W(c(z^t), \mathbb{C}\mathbb{E}[h(z^{t+1},\theta_{t+1})v_\sigma(z^{t+1},\theta^{t+1})|\sigma_{t+1}, z^t, \theta^t]), \tag{22}$$

where $z_{t+1} = (s_{t+1}, \sigma_{t+1}(z^t, \theta^{t+1}))$.

We let σ^* denote the truth-telling strategy $\sigma^*_t(z^t,\theta^t) = \theta_t$. Incentive compatibility requires

$$v_{\sigma^*}(z_0,\theta_0) \geqslant v_\sigma(z_0,\theta_0) \quad \forall \sigma. \tag{23}$$

The proof of the next result is in the Appendix.

Proposition 3. *Assume* $u(x) = x^{1-\rho}/(1-\rho)$ *and* $R(x) = x^{1-\gamma}/(1-\gamma)$ with $\rho, \gamma \geqslant 0$. *For any allocation* (c, h, v) *satisfying* (22) *and* (23), *fix a history* $\hat{z}^\tau$ *and consider the following variation*:

$$\tilde{c}(z^t) = \begin{cases} \Delta c(z^\tau) & \text{for } z^t = \hat{z}^\tau, \\ \Delta' c(z^t) & \text{for } t > \tau \text{ and } z^t \succ \hat{z}^\tau, \\ c(z^t) & \text{otherwise.} \end{cases}$$

Then for any Δ' *there exists a* Δ *such that* $(\tilde{c}, h, \tilde{v})$ *satisfy* (22) *and* (23) *if*: (a) *Conditional on* s^t, *the realization of* θ_t *is independent and identically distributed*; *or* (b) $\rho = 1$ *so that* $u(x) = \log x$.

We do not impose restrictions on the stochastic process for the observable state s_t. Regarding the unobservable shock, the requirement in part (a) does not restrict the process for productivity, and can, in particular, accommodate any degree of persistence. What this requirement does ensure is that the states that affect the evolution of shocks are observable, that there are no hidden states. Although this implies that the observable state s^t is a sufficient statistic for (s^t, θ^t), in the sense that $\Pr(s^{t+n}, \theta^{t+n}|s^t, \theta^t) = \Pr(s^{t+n}, \theta^{t+n}|s^t)$, optimal allocations typically depend on the history θ^t. In this way, the history of reports r^t is relevant. False past reports may then affect the allocation the agent receives, but do not affect the planner's capacity to predict the agent's future productivity. This tractability allows us to find variations that maintain incentive compatibility.

In the logarithmic case, $\rho = 1$, the crucial property is that

$$W(\Delta c, \Delta' v') = \Delta^{1-\beta}(\Delta')^\beta W(c, v').$$

Hence, setting $\Delta^{1-\beta}(\Delta')^\beta = 1$ in the variations does not affect the utility delivered by any reporting strategy. As a result, no assumption on the structure of uncertainty is required.

3.3. The intertemporal optimality condition: the Golden Ratio or the modified Inverse Euler equation

Let us say that an allocation is efficient if it minimizes the present value of consumption $\mathbb{E}\sum_{t=0}^{\infty} q^t c_t$ and delivers a given lifetime utility level in an incentive compatible way. Then any efficient allocation cannot be improved by the variations above. That is, these variations cannot reduce the discounted value of consumption.

Fix a node $\hat{s}^\tau$. Increase consumption at $\hat{s}^\tau$ proportionally by Δ, and increase consumption at all nodes that follow it, $s^t \succ \hat{s}^\tau$, proportionally by Δ'. This variation is permitted by the propositions above. Indexing the variation by Δ' and solving for $\Delta = \delta(\Delta')$ that keeps utility constant, we consider the minimization

$$\min_{\Delta'} \left(\delta(\Delta') c(\hat{s}^\tau) + \Delta' \sum_{t > \tau, s^t} q^t c(s^t) \Pr[s^t | a^*, \hat{s}^\tau] \right). \tag{24}$$

The first-order necessary and sufficient condition for optimality is simply

$$\frac{c_t}{\sum_{s=0}^{\infty} q^s \mathbb{E}_t[c_{t+s}]} = \frac{(1-\beta) u(c_t)}{u(v_t)}. \tag{25}$$

Thus, optimality requires the ratio of current to lifetime utility $(1-\beta)u(c_t)/u(v_t)$ to be equated to the ratio of current consumption with its expected present value $c_t / \sum_{s=0}^{\infty} q^s \mathbb{E}_t[c_{t+s}]$. Rearranging, the ratio of current consumption and utility must be equated to the ratio of the present value of consumption with lifetime utility:

$$\frac{c_t}{(1-\beta) u(c_t)} = \frac{\sum_{s=0}^{\infty} q^s \mathbb{E}_t[c_{t+s}]}{u(v_t)}. \tag{26}$$

Both conditions formalize the optimality of a form of consumption smoothing. We call them the Golden Ratio conditions.

The next result re-expresses the optimality condition above in a way that is more suitable for comparison with the optimality condition—the Euler equation—that results when agents can save freely at the interest rate q^{-1}. We call this condition the Modified Inverse Euler equation.

Proposition 4. *Define*

$$x_{t+1} \equiv \frac{h_{t+1} v_{t+1}}{\mathbb{C}\mathbb{E}_t[h_{t+1} v_{t+1}]}. \tag{27}$$

(a) *At the optimum in* (24) *the following condition holds*:

$$1 = \frac{q}{\beta} \mathbb{E}_t \left[x_{t+1}^{1-\rho} \frac{u'(c_t)}{u'(c_{t+1})} \right]. \tag{28}$$

(b) *If agents can borrow and save freely at the interest rate* q^{-1}, *then the allocation must satisfy the following Euler equation*:

$$1 = \frac{\beta}{q} \mathbb{E}_t \left[x_{t+1}^{\rho-\gamma} \frac{u'(c_{t+1})}{u'(c_t)} \right]. \tag{29}$$

Savings will generally be distorted at the optimal allocation, since the Modified Inverse Euler equation and the Euler equation are incompatible. Thus, in any implementation of the planner's optimum, agents cannot be allowed to borrow and save freely at the interest rate $1/q$.

Suppose that the optimality condition (28) holds. Define the intertemporal wedge τ by solving for the factor $(1-\tau)$ required so that the Euler equation (29) holds when $1/q$ is replaced with $(1-\tau)/q$:

$$1 - \tau = \mathbb{E}_t \left[x_{t+1}^{\rho-\gamma} \frac{u'(c_{t+1})}{u'(c_t)} \right] \mathbb{E}_t \left[x_{t+1}^{1-\rho} \frac{u'(c_t)}{u'(c_{t+1})} \right] \tag{30}$$

so that

$$\tau = -\mathrm{Cov}\left(\frac{u'(c_{t+1})}{u'(c_t)} \frac{x_{t+1}^{1-\gamma}}{x_{t+1}^{1-\rho}}, \frac{u'(c_t)}{u'(c_{t+1})} x_{t+1}^{1-\rho} \right). \tag{31}$$

Importantly, the intertemporal wedge τ is zero whenever there is no uncertainty. For the case of certainty, Atkinson–Stiglitz's uniform-taxation result requires preferences to be separable between consumption and leisure. However, in our recursive specification preferences are not separable. Interestingly, despite this, the absence of resolution of uncertainty between two periods implies that there should be no intertemporal distortion on savings there. In other words, although the separability conditions required by Atkinson–Stiglitz are violated, their uniform commodity taxation result holds under certainty with our preferences. Thus, optimal distortions can be entirely attributed to ongoing idiosyncratic uncertainty, just as in the additively separable expected-utility case that leads to the Inverse Euler equation (Golosov et al., 2003).

Note that if $\gamma = 1$ one gets that $\tau > 0$, guaranteeing that the intertemporal distortion on savings is positive. Another interesting case is when c_t is a geometric random walk at the baseline allocation, so that $c_{t+1} = \varepsilon_{t+1} c_t$. It then follows that v_t is proportional to c_t, and $\tau > 0$. We shall study this case in more detail in the next section.

3.4. Constant absolute risk aversion preferences

In this subsection, we show that for a particular class of preferences with constant absolute risk aversion the optimal distortion on savings is zero. In a static moral-hazard setting, a convenient specification of preferences is

$$\mathbb{E}[U(c - h(a))|a], \tag{32}$$

where $U(x) = -\mathrm{e}^{-\alpha x}$ is exponential. Equivalently, one can express ex ante utility as

$$\mathbb{CE}[c - h(a)|a]. \tag{33}$$

In our dynamic setting, we generalize this specification as follows. Let $u(x) = -\mathrm{e}^{-\rho x}$ and $R(x) = -\mathrm{e}^{-\gamma x}$ and consider the recursion

$$v_a(s^t) = W(c(s^t), \mathbb{CE}[v_a(s^{t+1}) - h(a(s^t))|a(s^t), s^t]), \tag{34}$$

where $W(c, v') = u^{-1}((1-\beta)u(c) + \beta u(v'))$ and $\mathbb{CE} = R^{-1}\mathbb{E}R$. Incentive compatibility requires inequalities (19) as before. The next proposition is proved in the Appendix.

Proposition 5. *Assume* $u(x) = -\mathrm{e}^{-\rho x}$ *and* $R(x) = -\mathrm{e}^{-\gamma x}$ with $\rho, \gamma \geqslant 0$. *Suppose we have* c, v *and* a^* *satisfying conditions* (18) *and* (19). *Fix a history* s^τ. *Consider the variation*:

$$\tilde{c}(s^t) = \begin{cases} c(s^\tau) + \Delta & \text{for } s^t = s^\tau, \\ c(s^t) + \Delta' & \text{for } t > \tau \text{ and } s^t \succ s^\tau, \\ c(s^t) & \text{otherwise.} \end{cases}$$

Then for any Δ' *there exists a* Δ *such that* $\tilde{c}$, $\tilde{v}$ *and* a^* *satisfy conditions* (18) *and* (19).

As above, we say that an allocation is efficient if it minimizes the present value of consumption

$$\sum_{t,s^t} q^t c(s^t) \Pr[s^t|a^*] \tag{35}$$

required to deliver a given lifetime utility level in an incentive compatible way. Then any efficient allocation cannot be improved by the variations above. That is, these variations cannot reduce the net present value of consumption.

Indexing the variation at any node by Δ' and solving for Δ that keeps utility constant we can write the minimization subproblem as in (24). In this case, the first-order necessary and sufficient condition coincides with the condition obtained if the worker could save and borrow freely at a market interest rate q^{-1}.

Proposition 6. *The optimum in* (24) *corresponds to the economy where agents can borrow and save freely at the interest rate* q^{-1}. *The following Euler equation holds*:

$$u'(c_t) = \frac{\beta}{q} u'(\mathbb{CE}(c_{t+1} - h_t)). \tag{36}$$

Hence, for the CARA preferences under consideration, the constrained-optimality condition and the Euler equation coincide. This section focused on a moral hazard setting, but a similar result should hold in a Mirrleesian environment.

4. Welfare gains: quantitative explorations

In this section, we investigate the welfare gains from the optimal savings distortions derived in Section 3. The analysis proceeds along the lines of Farhi and Werning (2006). We focus on the case where the baseline allocation features a geometric random walk consumption process while work effort is constant. The analysis in this section covers both to the private-information and moral-hazard settings.

Assumption 1. The baseline allocation $\{c_t, h_t\}$ is such that $h_t = \bar{h}$ is constant and c_t is a geometric random walk $c_{t+1} = c_t \varepsilon_{t+1}$ with ε_{t+1} identically and independently distributed over time.

4.1. Partial equilibrium

Let us first assume that there is a linear technology to transfer resources from period to period with a gross rate of return $R = q^{-1}$.

The following proposition shows that if the baseline allocation is a pure geometric random walk and h_t is constant, then the cost minimizing allocation attainable through our variations is also a pure geometric random walk.

Proposition 7. *Suppose that Assumption* 1 *holds. Then the cost minimizing allocation* $\{\tilde{c}_t\}$ *is obtained by multiplying* $\{c_t\}$ *by a deterministic drift* g^{-1}:

$$\tilde{c}_t = \alpha g^{-t} c_t$$

with

$$g \equiv (q\hat{\beta}^{-1}\mathbb{E}[\varepsilon](\mathbb{E}[\varepsilon^{1-\gamma}])^{-(1-\rho)/(1-\gamma)})^{1/\rho} \quad \textit{and} \quad \alpha \equiv \left(\frac{1 - qg^{-1}\mathbb{E}[\varepsilon]}{1 - qg^{-\rho}\mathbb{E}[\varepsilon]}\right)^{1/(1-\rho)},$$

where $\hat{\beta} = \beta\bar{h}^{1-\rho}$.

Hence the optimal allocation $\tilde{c}_t$ attainable from the baseline allocation through our variations is such that $\tilde{c}_t$ also follows a geometric random walk, but with a different drift $g^{-1}\mathbb{E}[\varepsilon]$ instead of $\mathbb{E}[\varepsilon]$ for the baseline allocation. This new drift ensures that the constrained-optimality condition—a necessary and sufficient condition for optimality within our class of variations—holds at the optimal allocation $\tilde{c}_t$. Note that β and $\bar{h}^{1-\rho}$ play exactly similar roles in this formula: when $h_t = \bar{h}$ is constant, it acts as a discount factor. This effect is compounded with β to produce an effective discount factor $\hat{\beta} = \beta\bar{h}^{1-\rho}$. It is also useful to note that if $g > 1$, then $\alpha > 1$ and vice versa.

Increasing g while maintaining the value of $q\mathbb{E}[\varepsilon]$ is exactly equivalent to decreasing the effective discount factor $\hat{\beta}$. In other words, the higher g, the lower the effective discount factor $\hat{\beta}$ that makes the constrained-optimality condition hold.

Note also that given $q\mathbb{E}[\varepsilon]$ and g, the intercept α depends only on the intertemporal elasticity of substitution parameter ρ. The risk aversion parameter γ only shifts the effective discount factor $\hat{\beta}$ required for the constrained-optimality condition to hold.

Economists are used to thinking of the discount factor as a primitive of the model, and as the equilibrium interest rate as an outcome. However, contrary to interest rates, discount factors are not directly observable. In fact, most of the evidence concerning discount factors comes from equilibrium values of interest rates. Therefore, in the formula for the intercept α, we prefer to think of the equilibrium interest rate q as the primitive and to solve for the effective discount factor $\hat{\beta}$ that makes the constrained-optimality condition hold given g and $q\mathbb{E}[\varepsilon]$.

Intertemporal wedge: We can compute the optimal wedge in closed form

$$\tau = \frac{-\mathrm{Cov}(\varepsilon, \varepsilon^{-\gamma})}{\mathbb{E}[\varepsilon]\mathbb{E}[\varepsilon^{-\gamma}]}.$$

Note that the wedge is always positive. Its magnitude in this example is independent of ρ and is entirely determined by γ, that is by the agent's attitude toward risk. This highlights that the origin of the wedge is the combination of two factors: the riskiness of tomorrow's consumption from today's perspective and the agent's risk aversion. Absent shocks, there would be no reason to distort savings and the Euler equation would hold. Similarly, if the agent were risk neutral, there would be no reason to distort savings and the wedge would also be zero.

We can re-express the wedge using the formalism of cumulants: let m be the moment generating function of $\log(\varepsilon)$:

$$m(\theta) \equiv \log E[\exp(\theta \log(\varepsilon))] = \log E\left[\varepsilon^{\theta}\right].$$

The nth cumulant of $\log(\varepsilon)$ is given by $\kappa_n \equiv \frac{d^n m}{d\theta^n}(0)$. Cumulants are closely related to moments, as we see from the first four: $\kappa_1 = \mu_1$, $\kappa_2 = \mu_2$, $\kappa_3 = \mu_3$, $\kappa_4 = \mu_4 - 3(\mu_2)^2$. The notation is standard, with μ_1 denoting the conditional mean of $\log(\varepsilon)$ and μ_n, for $n \geqslant 1$, denoting the nth central conditional moment.

Using this notation, we derive a formula that ties the wedge to the higher order moments or cumulants of $\log(\varepsilon)$:

$$-\log(1-\tau) = m(1) + m(-\gamma) - m(1-\gamma) = \sum_{n=2}^{\infty} \kappa_n / n!(1 + (-\gamma)^n - (1-\gamma)^n).$$

In the lognormal case, which we explore below, the higher cumulants κ_n of $\log(\varepsilon)$ are zero for $n \geqslant 3$ and we obtain a closed form for the wedge which depends only on the variance σ_ε^2 of $\log(\varepsilon)$: $\log(1-\tau) = \gamma\sigma_\varepsilon^2$.

Outside of the lognormal case, higher cumulants κ_n of $\log(\varepsilon)$ are non-zero and higher moments of the distribution of consumption growth rates affect the wedge. For example, we can analyze the impact of skewness κ_3. The contribution of this term to the wedge is given by $\kappa_3 \frac{\gamma(1-\gamma)}{2}$. Hence, negative skewness – $\kappa_3 < 0$ – decreases the wedge if $\gamma < 1$ and increases the wedge if $\gamma > 1$.

Welfare gains: The costs $\tilde{k}$ and k of the baseline and the optimal allocations are easily computed to be

$$\tilde{k} = \frac{\alpha c}{1 - qg^{-1}\mathbb{E}[\varepsilon]}$$

and

$$k = \frac{c}{1 - q\mathbb{E}[\varepsilon]}.$$

Combining these two expressions, we can derive the relative reduction in expected discounted cost allowed by our variations.

Proposition 8. *Suppose that Assumption* 1 *holds. Then the relative expected discounted cost reduction achieved by going from the baseline allocation to the optimal allocation is*

$$\frac{k}{\tilde{k}} = \left(\frac{1 - qg^{-\rho}\mathbb{E}[\varepsilon]}{1 - q\mathbb{E}[\varepsilon]}\right)^{1/(1-\rho)} \left(\frac{1 - qg^{-1}\mathbb{E}[\varepsilon]}{1 - q\mathbb{E}[\varepsilon]}\right)^{1-1/(1-\rho)}. \tag{37}$$

By homogeneity, the ratio of the cost of the optimal allocation to the cost of the baseline does not depend on the current level of consumption c. Given the cost of the baseline allocation, or in other words, given $q\mathbb{E}[\varepsilon]$, g is a sufficient statistic for the welfare gains attainable through the variations. It is therefore instructive to perform some comparative statics with respect to g.

Given $q\mathbb{E}[\varepsilon]$ and g, the relative expected cost reduction depends only on the intertemporal elasticity of substitution parameter ρ^{-1}. This is a direct consequence of the fact noted above that given g and $q\mathbb{E}[\varepsilon]$, the intercept α does not depend on the risk aversion parameter.

At $g = 1$, the reduction in cost is 0. This is because in this case, the constrained-optimality condition holds at the baseline allocation. Moreover, a Taylor expansion around $g = 1$ reveals that the cost reduction is zero at the first order in g and increasing in g:

$$\frac{k}{\tilde{k}} \simeq 1 + \frac{1}{2}\frac{q\mathbb{E}[\varepsilon]}{(1 - q\mathbb{E}[\varepsilon])^2}\rho(g-1)^2.$$

When g goes to infinity on the other hand, the cost reduction goes to $1/(1 - q\mathbb{E}[\varepsilon])$. Taking g to infinity is like taking the effective discount factor to 0. In that case, the optimal allocation for $\Delta_{-1} = 1$ is

$$\tilde{c}_t = 0 \text{ for } t \geqslant 1 \quad \text{and} \quad \tilde{c}_0 = c_0.$$

In the limit where ρ goes to 1, we get

$$g=\frac{q}{\beta}\mathbb{E}[\varepsilon] \quad \text{and} \quad \frac{k}{\tilde{k}}=\frac{\beta^{-1}-1}{\beta^{-1}-g}g^{-\beta/(1-\beta)},$$

which is exactly the expression derived in Farhi and Werning (2006).

Euler at the baseline: Given the importance of g, we now investigate its main determinants in the interesting case where the Euler equation holds at the baseline allocation. That the Euler equation holds at the baseline means that

$$c_t^{-\rho}=\beta q^{-1}\mathbb{E}[c_{t+1}^{-\rho}\bar{h}^{1-\rho}v_{t+1}^{\rho-\gamma}](\mathbb{E}[v_{t+1}^{1-\gamma}])^{(\gamma-\rho)/(1-\gamma)},$$

which can be re-expressed as

$$1=\beta q^{-1}\bar{h}^{1-\rho}\mathbb{E}[\varepsilon^{-\gamma}](\mathbb{E}[\varepsilon^{1-\gamma}])^{(\gamma-\rho)/(1-\gamma)}. \tag{38}$$

The effective discount factor $\hat{\beta}=\beta\bar{h}^{1-\rho}$ can then be determined:

$$\hat{\beta}=q(\mathbb{E}[\varepsilon^{\gamma}])^{-1}(\mathbb{E}[\varepsilon^{1-\gamma}])^{(\rho-\gamma)/(1-\gamma)}.$$

Knowing $\hat{\beta}$, the sufficient statistic g for the welfare gains in formula (37) can be derived using the formula in Proposition 7.

Proposition 9. *If Assumption* 1 *holds and the Euler equation holds at the baseline allocation, then*

$$g=(\mathbb{E}[\varepsilon]\mathbb{E}[\varepsilon^{-\gamma}](\mathbb{E}[\varepsilon^{1-\gamma}])^{-1})^{1/\rho}.$$

When ε is lognormally distributed $\log\varepsilon\sim N(\mu,\sigma_\varepsilon^2)$, then the wedge τ, the change in drift from the baseline allocation g and the welfare gains can be computed in terms of the mean μ and the variance σ_ε^2 of consumption growth:

Corollary 1. *Suppose that ε is lognormally distributed* $\log\varepsilon\sim N(\mu,\sigma_\varepsilon^2)$, *then τ and g are given by*

$$\tau=1-\frac{\mathbb{E}[\varepsilon^{1-\gamma}]}{\mathbb{E}[\varepsilon]\mathbb{E}[\varepsilon^{-\gamma}]}=1-\exp[-\gamma\sigma_\varepsilon^2]\simeq\gamma\sigma_\varepsilon^2$$

and

$$g=\exp\left(\frac{\gamma}{\rho}\sigma_\varepsilon^2\right)\simeq 1+\frac{\gamma}{\rho}\sigma_\varepsilon^2.$$

As we already discussed, the wedge is increasing in the degree of risk aversion γ and in the magnitude of the shocks σ_ε^2. Moreover, γ and σ_ε^2 affect the wedge in a complementary way. When shocks are lognormal, the formula takes the remarkably simple form $\tau=1-\exp[-\gamma\sigma_\varepsilon^2]$.

The crucial parameter g is associated with $(\gamma/\rho)\sigma_\varepsilon^2$. The higher the variance of the shocks, and the higher risk aversion, the higher the required change in drift g between the baseline and the optimum. Similarly, the higher the intertemporal elasticity of substitution ρ^{-1}, the higher g.

Intuitively, this can be seen by taking the limit as ρ goes to 0, so that consumption at different dates become perfect substitutes. The Euler equation and the optimality condition are incompatible in the limit where ρ goes to 0, since the required change in drift g goes to infinity. Note, however, that in this case, the intercept α converges to $1-q(\mathbb{E}[\varepsilon^{-\gamma}])^{-1}\mathbb{E}[\varepsilon^{1-\gamma}]$. Intuitively, when ρ goes to 0, it is optimal to front-load consumption more and more. In the limit, it is best to deliver all consumption in the first period so that agents are entirely shielded from consumption risk. The cost reduction is nontrivial. Indeed, we have

$$\lim_{\rho\to 0}\frac{k}{\tilde{k}}=\frac{1-q\exp\left[\mu+(1-2\gamma)\frac{\sigma_\varepsilon^2}{2}\right]}{1-q\exp\left[\mu+\frac{\sigma_\varepsilon^2}{2}\right]}\simeq 1+\frac{qe^{\mu}}{1-qe^{\mu}}\gamma\sigma_\varepsilon^2.$$

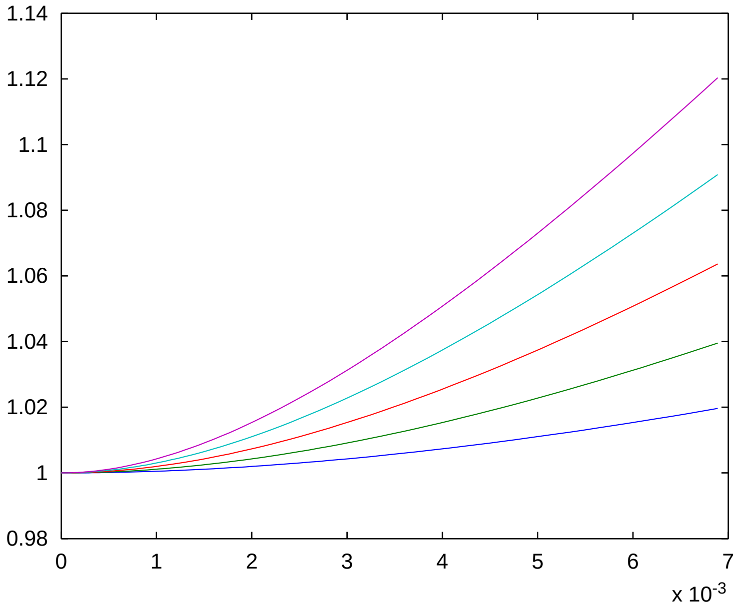

Fig. 1. Welfare gains as a function of σ_ε^2. Baseline consumption is a geometric random walk and h_t is constant. The Euler equation holds. The different curves correspond to different values of $\hat{\sigma}$ ranging from 1 to 3.

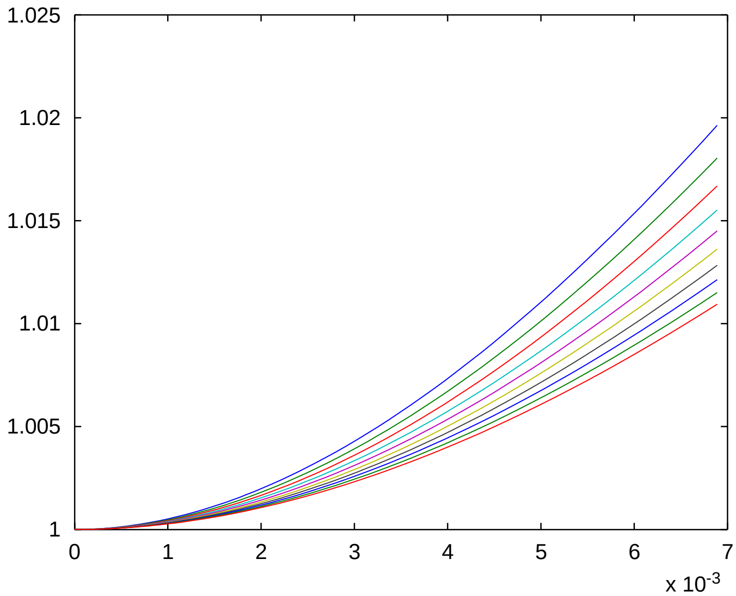

Fig. 2. Welfare gains as a function of σ_ε^2 when baseline consumption is a geometric random walk and h_t is constant. The different curves correspond to different values of ρ^{-1} ranging from 0.5 to 0.9.

The gains are increasing in the intertemporal elasticity of substitution ρ^{-1}: intuitively, as consumption at different dates become more substitutable, it becomes easier to compensate the agent for a decrease in the drift in consumption in order to lower his exposure to risk. In fact, we can derive a simple formula for small σ_ε:

$$\frac{k}{\tilde{k}} \simeq 1 + \frac{q\mathrm{e}^{\mu}}{(1 - q\mathrm{e}^{\mu})^2}\frac{\gamma^2}{\rho}\sigma_\varepsilon^4. \tag{39}$$

From this formula it is apparent that at the first relevant order, risk aversion and the intertemporal elasticity of substitution enter the formula for the gains only through γ^2/ρ.

Quantitative exploration: Figs. 1 and 2 plot the reciprocal of the relative cost reduction using Eq. (37) as a measure of the relative welfare gains as a function of σ_ε^2. The figures use an empirically relevant range for σ_ε^2 which is taken to vary between 0 and 0.007. The value of $q\mathbb{E}[\varepsilon]$ is set to 0.97.

In Fig. 1, the intertemporal elasticity of substitution ρ^{-1} is set to 1 and the different curves correspond to different values of the relative risk aversion coefficient γ ranging from 1 to 3 in increments of 0.5. The gains are increasing in γ: Increasing γ by 10% is exactly equivalent to increasing σ_ε^2 by 10%.

In Fig. 2, the relative risk aversion coefficient γ is set to 1, and the different curves correspond to different values of the intertemporal elasticity of substitution ρ^{-1} ranging from 0.5 to 1 in increments of 0.1. The gains are increasing in ρ^{-1}. Increasing ρ^{-1} by 10% is roughly equivalent to increasing σ_ε^2 by 5%.

Two lessons emerge from our simple exercise. First, welfare gains range from small to potentially large. Second, they depend a lot on three parameters of the model: γ, ρ and σ_ε^2. The coefficient of relative risk aversion γ and the variance of consumption growth σ_ε^2 play an especially important role over the range consistent with the available empirical evidence concerning these two parameters. The intertemporal elasticity of substitution ρ^{-1} is important, but its influence over the empirically relevant range is somewhat less dramatic. This is both because the range for this parameter is smaller and because ρ^{-1} enters with a smaller power than γ and σ_ε^2 as can be seen from (39).

4.2. General equilibrium

Up to now we have restricted the analysis to partial equilibrium. Alternatively, one can think of the results we have derived so far as applying to an economy facing some given constant rate of return to capital. In Farhi and Werning (2006), we argue that neglecting general equilibrium effects magnifies the welfare gains from reforming the consumption allocation. Here we explore the joint influence of risk aversion and the intertemporal elasticity of substitution on general equilibrium welfare gains.

Planning problem: Consider a baseline allocation $\{c_t, h_t\}$. In order to set-up the planning problem, it is useful to introduce the following notation: let $\Upsilon(\{c_t, h_t\}, \Delta_{-1})$ be the set of allocations $\tilde{c}_t$ attainable through our variations from the baseline allocation $\{\Delta_{-1}c_t, h_t\}$. Note that the shifted allocation $\{\Delta_{-1}c_t, h_t\}$ is incentive compatible and delivers a value lifetime utility increased by a multiplicative factor Δ_{-1} to the agent. In general equilibrium, the planning problem can be set-up as

$$W(K_0) = \max_{\{\tilde{c}_t, \tilde{K}_{t+1}\}} \tilde{v}_0 \tag{40}$$

subject to

$$\tilde{v}_t = h_t((1-\beta)\tilde{c}_t^{1-\rho} + \beta(\mathbb{E}[\tilde{v}_{t+1}^{1-\gamma}])^{1/(1-\gamma)})^{1/(1-\rho)} \quad \text{for } t = 0, 1, \ldots,$$

$$\{\tilde{c}_t\} \in \Upsilon(\{c_t, h_t\}, \Delta_{-1}),$$

$$\tilde{K}_{t+1} + \mathbb{E}[\tilde{c}_t] \leqslant F(\tilde{K}_t, \tilde{N}_t) + (1-\delta)\tilde{K}_t \quad \text{for } t = 0, 1, \ldots,$$

$$\tilde{K}_0 = K_0.$$

Necessary and sufficient conditions for this problem are

$$\tilde{c}_t^\rho = \frac{1}{\beta h_t^{1-\rho}[F_K(\tilde{K}_t, \tilde{N}_t) + (1-\delta)]} \mathbb{E}_t\left[\frac{\tilde{v}_{t+1}}{(\mathbb{E}_t[\tilde{v}_{t+1}^{1-\gamma}])^{1/(1-\gamma)}} \tilde{c}_{t+1}^\rho\right] \quad \text{for } t = 0, 1, \ldots .$$

Of course, we have $W(K_0) = \Delta_{-1}W$ where W is the welfare achieved at the baseline allocation and Δ_{-1} is the maximand in (40).

Note that we can always decompose $\tilde{c}_t = \tilde{c}_t^{\mathrm{i}}\tilde{C}_t$ with the property that $\mathbb{E}[\tilde{c}_t^{\mathrm{i}}] = 1$ and $\tilde{C}_t = \mathbb{E}[\tilde{c}_t]$, where the superscript i stands for idiosyncratic. Since our variations allow for deterministic parallel shifts in consumption, we have that $\{\tilde{c}_t\} \in \Upsilon(\{c_t, h_t\}, \Delta_{-1})$ for some Δ_{-1} if and only if $\{\tilde{c}_t^{\mathrm{i}}\} \in \Upsilon(\{c_t^{\mathrm{i}}, h_t\}, \Delta_{-1})$.

The analysis of this planning problem is tackled in full generality in Farhi and Werning (2006), where we also explore nongeometric random walk baseline allocations: we provide cases where (40) can be separated into two different planning problems, one involving only the idiosyncratic part of the allocation $\tilde{c}_t^{\mathrm{i}}$ and the other only the aggregate part $\tilde{C}_t$. Here instead, we focus on the special case where the baseline allocation features geometric random walk consumption with constant h_t.

Geometric random walk with constant h_t: Suppose that the baseline allocation features geometric random walk consumption with constant h_t and constant aggregate consumption:

$$c_{t+1} = c_t \varepsilon_{t+1} \quad \text{and} \quad h_t = \bar{h},$$

where ε_{t+1} is independently and identically distributed across agents and time and with $\mathbb{E}[\varepsilon_{t+1}] = 1$. In other words, Assumption 1 holds and $\mathbb{E}[\varepsilon] = 1$.

Define

$$\beta_\varepsilon \equiv \beta(\mathbb{E}[\varepsilon^{1-\gamma}])^{(1-\rho)/(1-\gamma)} \quad \text{and} \quad \hat{\beta}_\varepsilon \equiv \bar{h}^{1-\rho} \beta_\varepsilon.$$

Proposition 10. *Suppose that Assumption* 1 *holds and* $\mathbb{E}[\varepsilon] = 1$. *The solution to* (40) *is* $\tilde{c}_t = \tilde{C}_t c_t^{\mathrm{i}}$ *where* $\tilde{C}_t$ *and* Δ_{-1} *are the solutions of the standard neoclassical growth model with CRRA preferences*:

$$\frac{(\Delta_{-1} C_0)^{1-\rho}}{1-\rho} = (1-\hat{\beta}_\varepsilon) \max_{\{\tilde{C}_t, \tilde{K}_{t+1}\}} \sum_{t=0}^{\infty} \hat{\beta}_\varepsilon^t \frac{\tilde{C}_t^{1-\rho}}{1-\rho} \tag{41}$$

subject to

$$\tilde{K}_{t+1} + \tilde{C}_t \leqslant F(\tilde{K}_t, \tilde{N}_t) + (1-\delta)\tilde{K}_t \quad \text{for } t = 0, 1, \ldots,$$

$$\tilde{K}_0 = K_0.$$

The property that the idiosyncratic component of the baseline allocation is already optimal relies crucially on the assumption of geometric random walk with constant h_t. Intuitively, as we saw above, the planner only wants to affect the drift of $\{\tilde{c}_t^{\mathrm{i}}\}$, which is impossible in the case of an endowment economy where $1 = \mathbb{E}[\tilde{c}_t^{\mathrm{i}}]$.

In the case where the baseline allocation is a geometric random walk with constant h_t, we can therefore restrict our attention to the aggregate part of the allocation: all the potential welfare gains come from modifying the aggregate component of the allocation.

Euler equation at the baseline: Suppose that in addition, the baseline allocation represents a steady state where the Euler equation holds (Table 1).

Let $q_{\mathrm{SS}} = (1 - \delta + F_K(K_{\mathrm{SS}}, N_{\mathrm{SS}}))^{-1}$ be the inverse of the steady state interest rate. In that case, we can derive as above an expression for $\hat{\beta}_\varepsilon$:

$$\hat{\beta}_\varepsilon = q_{\mathrm{SS}}(\mathbb{E}[\varepsilon^{-\gamma}])^{-1}\mathbb{E}[\varepsilon^{1-\gamma}](\mathbb{E}[\varepsilon])^{\rho-1} = q_{\mathrm{SS}}(\mathbb{E}[\varepsilon])^{\rho}\frac{\mathbb{E}[\varepsilon^{1-\gamma}]}{\mathbb{E}[\varepsilon]\mathbb{E}[\varepsilon^{-\gamma}]}.$$

Table 1
Welfare gains

γ	$\rho^{-1} = 0.5$			$\rho^{-1} = 0.75$			$\rho^{-1} = 1$		
	δW_{PE} (%)	δW_{GE} (%)	$\tilde{r}_{\mathrm{SS}}$ (%)	δW_{PE} (%)	δW_{GE} (%)	$\tilde{r}_{\mathrm{SS}}$ (%)	δW_{PE} (%)	δW_{GE} (%)	$\tilde{r}_{\mathrm{SS}}$ (%)
1	2.02	0.09	3.82	1.56	0.10	3.82	1.07	0.10	3.82
2	3.62	0.34	4.55	5.15	0.37	4.55	6.53	0.38	5.28
3	7.03	0.69	5.28	9.85	0.75	5.28	12.33	0.79	5.28

That the baseline allocation is a steady state implies in particular that $\mathbb{E}[\varepsilon] = 1$. We can therefore simplify the formula for $\hat{\beta}_\varepsilon$:

$$\hat{\beta}_\varepsilon = q_{SS} \frac{\mathbb{E}[\varepsilon^{1-\gamma}]}{\mathbb{E}[\varepsilon]\mathbb{E}[\varepsilon^{-\gamma}]}.$$

The optimal allocation will eventually reach a steady state where the inverse of the interest rate $\tilde{q}_{SS}$ is given by $\tilde{q}_{SS} = \hat{\beta}_\varepsilon$.

When ε is lognormally distributed $\log \varepsilon \sim \mathrm{N}(\mu, \sigma_\varepsilon^2)$, then we can compute $\hat{\beta}_\varepsilon$ and $\tilde{q}_{SS}$ in terms of μ and σ_ε^2. We get the remarkably simple formula:

$$\tilde{q}_{SS} = \hat{\beta}_\varepsilon = q_{SS} \exp(-\gamma\sigma_\varepsilon^2). \tag{42}$$

Eq. (42) shows that the new interest rate is higher than the initial interest rate (that is, $\tilde{K}_{SS} < K_{SS}$) by a factor given by $\exp(\gamma\sigma_\varepsilon^2)$. The higher risk aversion and the variance of consumption growth, the higher the increase in steady state interest rates, and the higher the reduction in steady state capital stock. Because the baseline allocation has no trend, the intertemporal elasticity of substitution does not affect the level of the new interest rate $\tilde{q}_{SS}^{-1}$. The only thing our variations allow in this case is to correct the externality created by the precautionary savings motive, the intensity of which is controlled only by the relative risk aversion γ and the variance of consumption growth σ_ε^2.

As we just discussed, the coefficient of relative risk aversion γ and the variance of consumption growth σ_ε^2 control the decrease in capital between the baseline steady state and the optimal steady state. The intertemporal elasticity of substitution, on the other hand, controls the speed of the transition: the higher ρ^{-1}, the faster the transition, and the higher the welfare gains.

We now compute the welfare gains in general equilibrium for the neoclassical production function $F(K,N) = K^\alpha N^{1-\alpha} + (1-\delta)K$. We set $\alpha = 0.36$, $\delta = 0.09$. We set the variance of consumption growth at the highest end of the values we used in our partial equilibrium computations: $\sigma_\varepsilon^2 = 0.007$. We take the initial interest rate at the baseline allocation to be $r_{SS} = q_{SS}^{-1} - 1 = 3.07\%$. We perform the computations of welfare gains for three different values of the intertemporal elasticity of substitution ρ^{-1}—0.5, 0.75 and 1—and three different values for the relative risk aversion coefficient γ—1, 2 and 3. For each configuration of these parameters, we report the welfare gains in partial equilibrium δW_{PE} if the interest rate were fixed at r_{SS}, the welfare gains in general equilibrium δW_{GE} and the interest rate $\tilde{r}_{SS}$ at the new steady state for the optimal allocation.

An important general lesson from this exercise, as pointed out in Farhi and Werning (2006), is that taking into account the concavity of the production function—that is, taking into account general equilibrium effects—greatly mitigates the welfare gains. This is because in general equilibrium, reducing the drift of the consumption process—the optimal policy under partial equilibrium—yields lower and lower gains as consumption and capital go down over time and the equilibrium interest rate increases. As a consequence, it is optimal to reduce the drift differential. Eventually, under the optimal allocation, the drift differential goes to 0 and the economy reaches the new steady state with a higher interest rate and a lower capital stock.

Even though the partial equilibrium welfare gains can be as high as 12.33%, the general equilibrium welfare gains never go above 0.79%. The highest gains are reached for the highest value of the intertemporal elasticity of substitution $\rho^{-1} = 1$ and the highest value of the relative risk aversion coefficient $\gamma = 3$. For those parameter values, the new interest rate is substantially higher than the initial interest rate: $\tilde{r}_{SS} = 5.28\%$, whereas $r_{SS} = 3.07\%$. Despite this large difference in interest rates and therefore in steady state capital stocks, the general equilibrium welfare gains are moderate at 0.79%.

5. Conclusion

This paper studied constrained efficient allocations in private information economies. We focused on how the optimal savings distortions featured in those allocations depend on individuals' preferences. We introduced a recursive class of preferences that allowed a separation of risk aversion from intertemporal substitution, and derived general results on the nature of optimal distortions.

We then performed a quantitative investigation for a class of geometric random walk consumption allocations. We showed that savings distortions depend only on risk aversion and the variance of the shocks to consumption. However, the welfare gains from these distortions depend on both parameters, although we found greater sensitivity to risk aversion.

The purpose of the quantitative exercise was to illustrate the role preferences, but it was limited in terms of the consumption allocations it considered. In Farhi and Werning (2006) we undertake a comprehensive exploration of savings distortions and welfare gains for general consumption processes.

Appendix

Proof of Proposition 3. Part (a). The proof parallels the proof of Proposition 2 closely. First note that since preferences are recursive an incentive compatible allocation satisfies

$$v_{\sigma^*}(z^t,\theta^t) \geqslant v_\sigma(z^t,\theta^t) \quad \forall z^t,\theta^t,\sigma, \tag{43}$$

so that truth-telling maximizes continuation utility after any history of reports.

Let

$$\tilde{v}_\sigma(z^t,\theta^t) = \Delta' v_\sigma(z^t,\theta^t) \quad \text{for } t>\tau \text{ and } z^t \succ \hat{z}^\tau,$$
$$\tilde{v}_\sigma(z^t,\theta^t) = v_\sigma(z^t,\theta^t) \quad \text{for } t\geqslant\tau \text{ and } z^t \not\succeq \hat{z}^\tau,$$

so that condition equation (22) with $\tilde{c}$ is met for all θ^t with $t\geqslant\tau$ with $\theta^t \neq \theta^\tau$. Let Δ solve

$$v_{\sigma^*}(\hat{z}^\tau,\theta^\tau) = W(\Delta c(\hat{z}^\tau), \Delta'\mathbb{CE}[h(z^{\tau+1},\theta_{\tau+1})v_{\sigma^*}(z^{\tau+1},\theta^{\tau+1})|\sigma^*_{t+1},\hat{z}^\tau,\theta^\tau]).$$

So that $v_{\sigma^*}(\hat{z}^\tau,\theta^\tau) = \tilde{v}_{\sigma^*}(\hat{z}^\tau,\theta^\tau)$ for all θ^τ. Using the recursion (22), the inequality (43) evaluated at $\hat{z}^\tau$ implies

$$\mathbb{CE}[\Delta'(s^{\tau+1})h(z^{\tau+1},\theta_{\tau+1})v_{\sigma^*}(z^{\tau+1},\theta^{\tau+1})|\sigma^*_{t+1},\hat{z}^\tau,\theta^\tau]$$
$$\geqslant \mathbb{CE}[h(z^{\tau+1},\theta_{\tau+1})v_\sigma(z^{\tau+1},\theta^{\tau+1})|\sigma_{t+1},\hat{z}^\tau,\theta^\tau]$$

for all histories $\theta^{\tau+1}$ and reporting plans σ. Hence,

$$\tilde{v}_{\sigma^*}(z^\tau,\theta^\tau) = W(\Delta c(z^\tau), \mathbb{CE}[\Delta'(s^{\tau+1})h(z^{\tau+1},\theta_{\tau+1})v_{\sigma^*}(z^{\tau+1},\theta^{\tau+1})|\sigma^*_{t+1},\hat{z}^\tau,\theta^\tau])$$
$$\geqslant W(\Delta c(s^\tau), \Delta'(s^{\tau+1})\mathbb{CE}[h(z^{\tau+1},\theta_{\tau+1})v_\sigma(z^{\tau+1},\theta^{\tau+1})|\sigma_{t+1},\hat{z}^\tau,\theta^\tau]) = \tilde{v}_\sigma(z^\tau,\theta^\tau).$$

Collecting the inequalities, we have shown that in period τ

$$\tilde{v}_\sigma(z^\tau,\theta^t) \leqslant \tilde{v}_{\sigma^*}(z^\tau,\theta^t) = v_{\sigma^*}(z^\tau,\theta^t) \quad \text{for all } z^\tau,\theta^t,\sigma.$$

Thus, σ^* is optimal from period τ onward and delivers the same continuation utility as previously.

For any plan σ define an alternative plan $\hat{\sigma}$ that starts at σ and then switches to σ^* from period τ onward: $\hat{\sigma}_t(z^\tau,\theta^t) = \sigma_t(z^\tau,\theta^t)$ for $t<\tau$ and $\hat{\sigma}_t(z^\tau,\theta^t) = \sigma^*_t(z^\tau,\theta^t)$ for $t\geqslant\tau$. The result above implies that

$$\tilde{v}_\sigma(z_0,\theta_0) \leqslant \tilde{v}_{\hat{\sigma}}(z_0,\theta_0) = v_{\hat{\sigma}}(z_0,\theta_0) \leqslant v_{\sigma^*}(z_0,\theta_0) = \tilde{v}_{\sigma^*}(z_0,\theta_0). \tag{44}$$

That is, $\hat{\sigma}$ dominates σ and yields the same utility as without the variation, which in turn is dominated by the recommended action σ^* which also yields the same utility as after the variation. This establishes that σ^* remains incentive compatible.

Part (b). Note that

$$\tilde{v}_\sigma(z^t,\theta^t) = \Delta' v_\sigma(z^t,\theta^t) \quad \text{for } t>\tau \text{ and } z^t \succ \hat{z}^\tau,$$
$$\tilde{v}_\sigma(z^t,\theta^t) = v_\sigma(z^t,\theta^t) \quad \text{for } t\geqslant\tau \text{ and } z^t \not\succeq \hat{z}^\tau.$$

Set $\Delta = (\Delta')^{-\beta}$ so that

$$\tilde{v}_\sigma(\hat{z}^\tau,\theta^\tau) = h(z^\tau,\theta_\tau)W(\Delta c(\hat{z}^\tau), \Delta'\mathbb{CE}[v_{\sigma^*}(z^{\tau+1},\theta^{\tau+1})|\sigma^*_{t+1},\hat{z}^\tau,\theta^\tau])$$
$$= h(z^\tau,\theta_\tau)W(c(\hat{z}^\tau), \mathbb{CE}[v_{\sigma^*}(z^{\tau+1},\theta^{\tau+1})|\sigma^*_{t+1},\hat{z}^\tau,\theta^\tau]) = v_\sigma(\hat{z}^\tau,\theta^\tau).$$

It follows by backward induction that

$$\tilde{v}_\sigma(z^t, \theta^t) = v_\sigma(z^t, \theta^t) \quad \text{for all } s^t,\ t \leqslant \tau.$$

In particular, $\tilde{v}_\sigma(z_0, \theta_0) = v_\sigma(z_0, \theta_0)$, so that the result follows from incentive compatibility of the original allocation. □

Proof of Proposition 5. Let $\tilde{v}$ be such that

$$\tilde{v}_a(s^t) = v_a(s^t) + \Delta' \quad \text{for } t > \tau \text{ and } s^t \succ s^\tau,$$
$$\tilde{v}_a(s^t) = v_a(s^t) \quad \text{for } t \geqslant \tau \text{ and } s^t \nsucc s^\tau,$$

so that condition (18) with $\tilde{c}$ is met for all s^t with $t \geqslant \tau$ with $s^t \neq s^\tau$. Now set Δ so that

$$\begin{aligned} v_{a^*}(s^\tau) &= W(c(s^t) + \Delta, \mathbb{CE}[v_{a^*}(s^{\tau+1}) + \Delta' - h(a^*(s^\tau))|a^*(s^\tau), s^\tau]) \\ &= W(c(s^t) + \Delta, \Delta' + \mathbb{CE}[v_{a^*}(s^{\tau+1}) - h(a^*(s^\tau))|a^*(s^\tau), s^\tau]), \end{aligned}$$

so that $\tilde{v}_{a^*}(s^\tau) = v_{a^*}(s^\tau)$. Using recursion (18), the inequality (20) evaluated at s^τ implies

$$\mathbb{CE}[v_{a^*}(s^{\tau+1}) - h(a^*(s^\tau))|a^*(s^\tau), s^\tau] \geqslant \mathbb{CE}[v_a(s^{\tau+1}) - h(a(s^\tau))|a(s^\tau), s^\tau],$$

so that

$$\begin{aligned} \tilde{v}_{a^*}(s^\tau) &= W(c(s^t) + \Delta, \Delta' + \mathbb{CE}[v_{a^*}(s^{\tau+1}) - h(a^*(s^\tau))|a^*(s^\tau), s^\tau]) \\ &\geqslant W(c(s^t) + \Delta, \Delta' + \mathbb{CE}[v_a(s^{\tau+1}) - h(a(s^\tau))|a(s^\tau), s^\tau]) = \tilde{v}_a(s^\tau). \end{aligned}$$

Hence, we have that

$$\tilde{v}_a(s^\tau) \leqslant \tilde{v}_{a^*}(s^\tau) = v_{a^*}(s^\tau) \quad \text{for all } a,$$

a^* is optimal from period τ onward and delivers the same continuation utility as previously.

For any plan a define an alternative plan $\hat{a}$ that switches to a^* from period τ onward: $\hat{a}(s^t) = a(s^t)$ for $t < \tau$ and $\hat{a}(s^t) = a^*(s^t)$ for $t \geqslant \tau$. The result above implies that

$$\tilde{v}_a(s_0) \leqslant \tilde{v}_{\hat{a}}(s_0) = v_{\hat{a}}(s_0) \leqslant v_{a^*}(s_0) = \tilde{v}_{a^*}(s_0). \tag{45}$$

That is, $\hat{a}$ dominates a and yields the same utility as without the variation, which in turn is dominated by the recommended action a^* which also yields the same utility as after the variation. This establishes that a^* remains incentive compatible. □

Proof of Proposition 6. The equation that defines Δ as a function of Δ' is

$$-(1-\beta)u(c_t)(\mathrm{e}^{-\rho\Delta} - 1) = \beta u(\mathbb{CE}(v_{t+1} - h_t))(\mathrm{e}^{-\rho\Delta'} - 1).$$

From this equation we get that at $\Delta' = 0$:

$$\frac{\mathrm{d}\Delta}{\mathrm{d}\Delta'} = -\frac{\beta}{1-\beta}\frac{u(\mathbb{CE}(v_{t+1} - h_t))}{u(c_t)}.$$

At the optimum, we must have that at $\Delta' = 0$:

$$\frac{\mathrm{d}\Delta}{\mathrm{d}\Delta'} = -\frac{1}{r},$$

where r is defined by $r = q^{-1} - 1$. Therefore, the following optimality condition must hold:

$$(1-\beta)u(c_t) = \beta r u(\mathbb{CE}(v_{t+1} - h_t)).$$

Noting that $u(c) = (-1/\rho)u'(c)$, this condition is equivalent to

$$(1-\beta)u'(c_t) = \beta r u'(\mathbb{CE}(v_{t+1} - h_t)),$$

which is the optimality condition in the problem where the agents can borrow and save freely at the interest rate r. Transforming these two equivalent conditions into the Euler equation in the text is straightforward.

Consider the constrained efficient allocation. We can rewrite the equation that defines Δ as a function of Δ' in the following way:

$$-\Delta' = \frac{1}{\rho}\log(1 + r(1 - \exp(-\rho\Delta))).$$

This defines $-\Delta'$ as a concave function of Δ. Therefore, $\Delta' \geqslant -r\Delta$. Now consider giving the agents the constrained efficient allocation and allowing them to not only choose a reporting strategy but also to borrow and save between history s^τ and subsequent periods. The following variations are then available to the agents:

$$\tilde{c}(s^t) = \begin{cases} c(s^\tau) + \Delta & \text{for } s^t = s^\tau, \\ c(s^t) + r\Delta & \text{for } t > \tau \text{ and } s^t \succ s^\tau, \\ c(s^t) & \text{otherwise.} \end{cases}$$

Since $\Delta' \geqslant -r\Delta$, whatever reporting strategy the agent chooses when these variations are permissible, he will always achieve lower utility than under the same reporting strategy if he were given the variations allowed for the planner. Since the constrained efficient allocation is incentive compatible, he cannot achieve higher utility than under the constrained efficient allocation without any additional saving or borrowing. Generalizing that argument to any history s^τ, this proves the proposition. $\square$

Proof of Proposition 7. When consumption is a geometric random walk and h_t is constant, it is possible to derive lifetime utility in closed form:

Lemma 1. *Suppose that Assumption* 1 *holds. Then* $v_t = A\bar{h}c_t$ *with*

$$A = \left(\frac{1-\beta}{1-\beta\bar{h}^{1-\rho}(\mathbb{E}[\varepsilon^{1-\gamma}])^{(1-\rho)/(1-\gamma)}}\right)^{1/(1-\rho)}.$$

The key feature that delivers this result is the homogeneity of agents' preferences. For a given $\bar{h}$, a proportional shift in consumption today moves consumption in every future period by a proportional factor, thereby shifting lifetime utility in consumption equivalent units by the same multiplicative factor. The constant disutility $\bar{h}$ on the other hand, acts exactly like a discount factor. Hence utility in consumption equivalent units v_t is directly proportional to consumption and to the disutility from effort or work. This is reminiscent of the static settings in Sections 2.1 and 2.2.

It is then easy to guess and verify that the solution proposed in Proposition 7 both preserves the level of utility and satisfies the constrained-optimality condition. $\square$

Proof of Proposition 10. Before proving this proposition, it is useful to establish the following lemma.

Lemma 2. *Consider the allocation in Proposition* 10. *We can write* $\tilde{v}_t = \tilde{V}_t v_t^{\mathrm{i}}$ *where* v_t^{i} *is the lifetime utility derived from* $\{c_t^{\mathrm{i}}, \bar{h}\}$: $v_t^{\mathrm{i}} = A\bar{h}c_t^{\mathrm{i}}$ *with* $A = ((1-\beta)/(1-\hat{\beta}_\varepsilon))^{1/(1-\rho)}$ *and*

$$\frac{\tilde{V}_t^{1-\rho}}{1-\rho} = (1-\hat{\beta}_\varepsilon)\sum_{s=t}^{\infty}\hat{\beta}_\varepsilon^{s-t}\frac{\tilde{C}_t^{1-\rho}}{1-\rho}.$$

Let us now prove Proposition 10. We only need to check that

$$\tilde{c}_t^\rho = \frac{1}{\beta\bar{h}^{1-\rho}[F_K(\tilde{K}_t, \tilde{N}_t) + (1-\delta)]}\mathbb{E}_t\left[\frac{\tilde{v}_{t+1}}{(\mathbb{E}_t[\tilde{v}_{t+1}^{1-\gamma}])^{1/1-\gamma}}\tilde{c}_{t+1}^\rho\right]$$

holds. Decomposing $\tilde{c}_t^{1-\rho}$ into the product $\tilde{C}_t c_t^{\mathrm{i}}$ and using Lemma 2, we can express this condition as

$$1 = \left(\frac{\hat{\beta}_\varepsilon \tilde{C}_{t+1}^{-\rho}[F_K(\tilde{K}_t, \tilde{N}_t) + (1-\delta)]}{\tilde{C}_t^{-\rho}} \right)^{-1}.$$

This is the standard Euler equation that is trivially verified by the solution of the neoclassical growth problem (41). This concludes the proof of Proposition 10. □

References

Diamond, P.A., Mirrlees, J.A., 1977. A model of social insurance with variable retirement. Working Papers 210, Massachusetts Institute of Technology, Department of Economics.

Epstein, L., Zin, S., 1989. Substitution, risk aversion and the temporal behavior of consumption and asset returns: a theoretical framework. Econometrica 57, 937–968.

Farhi, E., Werning, I., 2006. Capital taxation: quantitative explorations of the Inverse Euler equation. Mimeo.

Golosov, M., Kocherlakota, N., Tsyvinski, A., 2003. Optimal indirect and capital taxation. Review of Economic Studies 70 (3), 569–587.

Ligon, E., 1998. Risk sharing and information in village economics. Review of Economic Studies 65 (4), 847–864.

Rogerson, W.P., 1985. Repeated moral hazard. Econometrica 53 (1), 69–76.