DISCUSSION PAPER SERIES

No. 2693

EQUILIBRIUM WELFARE AND GOVERNMENT POLICY WITH QUASI-GEOMETRIC DISCOUNTING

Per Krusell, Burhanettin Kuruscu
and Anthony A Smith, Jr

INTERNATIONAL MACROECONOMICS AND PUBLIC POLICY

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP2693.asp

ISSN 0265-8003

EQUILIBRIUM WELFARE AND GOVERNMENT POLICY WITH QUASI-GEOMETRIC DISCOUNTING

Per Krusell, Stockholm University, University of Rochester and CEPR
Burhanettin Kuruscu, University of Rochester
Anthony A Smith, Jr, Carnegie Mellon University

Discussion Paper No. 2693
February 2001

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 2693

February 2001

ABSTRACT

Equilibrium Welfare and Government Policy with Quasi-Geometric Discounting

We consider a representative-agent equilibrium model where the consumer has quasi-geometric discounting and cannot commit to future actions. With restricted attention to a parametric class for preferences and technology logarithmic utility, Cobb-Douglas production and full depreciation we solve for time-consistent competitive equilibria globally and explicitly. For this class, we characterize the welfare properties of competitive equilibria and compare them to that of a planning problem. The planner is a consumer representative who, without commitment but in a time-consistent way, maximizes his present-value utility subject to resource constraints. The competitive equilibrium results in strictly higher welfare than does the planning problem whenever the discounting is not geometric.

We also explicitly consider taxation in our environment. With a benevolent government that can tax income and capital but cannot commit its future tax rates, time-consistent taxation leads to positive tax rates on capital. These tax rates reproduce the central planning solution, and thus imply a worse outcome in welfare terms than when there is no government.

JEL Classification: D60, E21, H21
Keywords: quasi-geometric discounting, time inconsistency, welfare

Per Krusell
Department of Economics
University of Rochester
Harkness Hall
Rochester, NY 14627-0156
USA
Tel: (1 716) 273 4903/2
Fax: (1 716) 271 3900
Email: pekr@troi.cc.rochester.edu

Burhanettin Kuruscu
Department of Economics
University of Rochester
Harkness Hall
Rochester, NY 14627
USA
Email: kubu@troi.cc.rochester.edu

Anthony A Smith, Jr
Associate Professor of Economics
Graduate School of Industrial
Administration
Carnegie Mellon University
Pittsburgh, PA 15213
USA
Tel: (1 412) 268 7583
Fax: (1 412) 268 7357
Email: smithaa@andrew.cmu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=153042

Submitted 06 November 2000

NON-TECHNICAL SUMMARY

Experimental psychology has gathered a significant body of evidence that 'preference reversals' are a common occurrence in decision making over time. One expression of these findings is that discounting of future rewards is not geometric – the standard case considered in economics – but rather hyperbolic, or 'quasi-hyperbolic'. A view of these findings commonly expressed by economists is that experiments in general are fraught with problems and should be disregarded, or, as a less extreme position, that these particular experiments suffer from specific problems, thus including the possibility that there are alternative interpretations of the experimental results that are not in contradiction with our standard assumptions about preferences. A dismissal of the psychologists' findings, however, seems hazardous since they quite strongly suggest a 'friction' that may be an important part of economic welfare and, possibly, one where government intervention might be helpful. In this Paper we take the latter view: we admit the possibility that consumption–savings decisions are indeed made by agents whose preferences allow reversals. In particular, we assume quasi-geometric preferences, which is the very simplest kind of departure from our standard assumption on discounting. We assume that the time inconsistency in preferences is accompanied by an inability on the part of consumers to commit to future actions, but we view consumers as fully rational; they are aware of their 'internal friction' and do their best to minimize its effects. Our main goal here is the most basic question to an economist: we perform welfare analysis of the market mechanism. The question is: does the 'invisible hand' work as well as in the standard case, or would government intervention be desirable?

First we show that the standard recursive tools used to analyse the neoclassical growth, general equilibrium framework can be employed also in the case of quasi-geometric preferences. A consumer in our economy plays a game with his future selves, with whom he disagrees about how to save, and a key part of his decisions is about how to manipulate his future savings decisions by saving differently today. This is also a way of thinking about the geometric case, with the difference being that no disagreements occur there and manipulation is superfluous. We restrict attention to Markov (time-consistent) equilibria, which are limits of finite-horizon equilibria, thereby obtaining uniqueness in our specific functional-form examples and allowing straightforward welfare analysis and comparative-static exercises.

For the welfare analysis, we assume that the 'visible hand' in our economy – a planner with the ability to command consumption decisions, or a government with the possibility of using tax rates to distort these decisions – is subject to the same friction as are consumers: it cannot commit its future behaviour. So to the extent that it shares the consumers' preferences, it will also have preference reversals, and will want to manipulate its future actions. It is well known that if the government can commit, then it can help consumers achieve higher *ex ante* welfare, but it is not at all clear to us how the government could provide a general commitment technology for consumers' future consumption decisions. If it could commit to paths of future tax rates, it would be helpful, but if it is benevolent and shares the (combination of the current and future) consumers' preferences, such choices are not time-consistent and are thus not likely to be realized. Although we do think that some government commitment may be feasible, we do not think a full ability to commit is likely. Whatever lack of commitment remains is the subject of this Paper.

Our results are rather striking. We find that, whenever there is a time-inconsistency in preferences, not only does a benevolent social planning economy not deliver the same consumption allocation as a laissez faire world, but it delivers strictly lower welfare. Thus, we find a new argument as to why the market mechanism is a particularly good one. The key insight regards price-taking behaviour, as opposed to taking into account the impact of one's decisions on prices (or aggregate allocations). An alternative interpretation of the result is that a decentralization with many identical Robinson Crusoes each operating their own production technology would do as poorly as the visible hand world: a separation of consumption and production decisions is desirable in this economy.

The mechanism behind our main finding is intuitive and can be understood as follows. Suppose that the social planner's preferences coincide exactly with those of the current consumer. They would then both regard their respective future selves as saving too little (the argument focuses on this particular bias for illustration). As a consequence, they would both want to manipulate their future selves. The manipulation occurs *via* savings (their only decision variable here): when an extra dollar is saved, the extra income next period will influence savings next period. With time-consistent preferences, the current saver agrees with his future self about next period's savings decision, and an envelope argument allows him to ignore this effect when deciding on current savings. Here, in contrast, an extra dollar saved today has an additional benefit next period, since (so long as the marginal savings propensity is not negative or zero) it induces more future savings. The difference between the decentralized allocation and the central planning allocation originates in how these induced future increases in savings are perceived: they are perceived as being larger by a price-taker than by a planner. The reason for this is that the price-taker's future behavioural response is 'more linear': the added income they obtain next period is the higher savings times the return, which is perceived as constant. The planner, on the other hand, sees that for every additional dollar saved, the next period's return falls (assuming decreasing marginal productivity to capital). In summary, the incentives to save are higher for the decentralized agents, so they save more, and more saving is better in this economy: it takes us closer to the full commitment outcome. (The argument is parallel when the time inconsistency takes the form of excessive short-run patience.)

The most closely related literature to this Paper is the set of studies started by Strotz (1956) and Phelps and Pollak (1968) and then further developed recently in a set of Papers by Laibson (e.g. Laibson, 1994). Laibson (1996) discusses some aspects of government taxation, but does not consider the case when the government cannot commit its future tax rates, the main focus of our analysis here. In many of Laibson's set-ups (e.g. Laibson, 1997), there are added market frictions such as credit constraints which make illiquid assets play an important role. As a consequence, Ricardian equivalence does not hold. Another common assumption is the existence of partially uninsurable idiosyncratic shocks (Harris and Laibson, 2000). Here, the set-up is the most basic one possible: the only friction is the internal friction in consumers' preferences (naturally accompanied by a lack of commitment), and there is no uncertainty. In particular, all long- and short-term asset markets are operative; although it would be beneficial to close asset markets in the future, without the commitment power to do so, *ex post* it will always be in the interest of the group of consumers as a whole to keep current markets open. As a result, in our environment, Ricardian equivalence holds. Another example of an added friction is the one assumed in O'Donoghue and Rabin (1999), where consumers are not rational, but are rather constantly surprised at their preference reversals.

The analysis in the Paper builds fundamentally on recursive methods, and moreover uses specific functional forms to allow manageable closed-form solutions. The dynamic game thus played between the current and future selves is perhaps the simplest tractable example of a time-consistent equilibrium available in the literature. These literature plans began in Kydland and Prescott (1977) but have not allowed simple closed-form examples of Markov equilibria. We consider the case of a saver with time-inconsistent preferences as very natural grounds for illustrating the main forces at work. Central to this illustration are the recursive set-up and the derivation and discussion of the 'generalized Euler equation'.

Equilibrium Welfare and Government Policy with Quasi-Geometric Discounting

Per Krusell, Burhanettin Kuruşçu, and Anthony A. Smith, Jr.

September 2000

ABSTRACT

We consider a representative-agent equilibrium model where the consumer has quasi-geometric discounting and cannot commit to future actions. With restricted attention to a parametric class for preferences and technology—logarithmic utility, Cobb-Douglas production, and full depreciation—we solve for time-consistent competitive equilibria globally and explicitly. For this class, we characterize the welfare properties of competitive equilibria and compare them to that of a planning problem. The planner is a consumer representative who, without commitment but in a time-consistent way, maximizes his present-value utility subject to resource constraints. The competitive equilibrium results in strictly higher welfare than does the planning problem whenever the discounting is not geometric.

We also explicitly consider taxation in our environment. With a benevolent government that can tax income and capital, but cannot commit its future tax rates, time-consistent taxation leads to positive tax rates on capital. These tax rates reproduce the central planning solution, and thus imply a worse outcome in welfare terms than when there is no government.

1 Introduction

Experimental psychology has gathered a significant body of evidence that "preference reversals" are a common occurrence in decision making over time. One expression of these findings is that discounting of future rewards is not geometric—the standard case considered in economics—but rather hyperbolic, or "quasi-hyperbolic".[1] A view of these findings commonly expressed by economists is that experiments in general are fraught with problems and should be disregarded, or, as a less extreme position, that these particular experiments suffer from specific problems, thus including the possibility that there are alternative interpretations of the experimental results that are not in contradiction with our standard assumptions about preferences. However, a dismissal of the psychologists' findings seems hazardous, since

[1] The term quasi-hyperbolic commonly used in the recent literature—see, e.g., Laibson (1994)—refers to the use of a quasi-geometric discounting function in order to approximate a (generalized) hyperbolic function.

they quite strongly suggest a "friction" that may be an important part of economic welfare and, possibly, also one where government intervention might be helpful. In this paper we take the latter view: we admit the possibility that consumption-savings decisions are indeed made by agents whose preferences allow reversals. In particular, we assume quasi-geometric preferences, which is the very simplest kind of departure from our standard assumption on discounting.[2] We assume that the time-inconsistency in preferences is accompanied by an inability of consumers to commit to future actions, but view consumers as fully rational: they are aware of their "internal friction" and do their best to mimimize its effects. Our main goal here is the most basic question to an economist: we perform welfare analysis of the market mechanism. The question is: does the "invisible hand" work as well as in the standard case, or would government intervention be desirable?

We show, first, that the standard recursive tools used to analyze the neoclassical-growth, general-equilibrium framework can be employed also in the case of quasi-geometric preferences. A consumer in our economy plays a game with his future selves, with whom he disagrees about how to save, and a key part of his decisions is about how to manipulate his future savings decisions by saving differently today. This is also a way of thinking about the geometric case, with the difference that there, no disagreements occur, and manipulation is superfluous. We restrict attention to Markov (time-consistent) equilibria which are limits of finite-horizon equilibria, thereby obtaining uniqueness in our specific functional-form examples and allowing straightforward welfare analysis and comparative-static exercises.

For the welfare analysis, we assume that the "visible hand" in our economy—a planner with the ability to command consumption decisions, or a government with the possibility of using taxes rates to distort these decisions—is subject to the same friction as are consumers: it cannot commit its future behavior. So to the extent that it shares the consumers' preferences, it also will have preference reversals, and will want to manipulate its future actions. It is well known that if the government can commit, then it can help consumers achieve higher ex-ante welfare, but it is not at all clear to us how the government could provide a general commitment technology for consumers' future consumption decisions. If it could commit to paths of future tax rates, it would be helpful, but if it is benevolent and shares the (combination of the current and future) consumers' preferences, such choices are not time-consistent and are thus not likely to be realized. Although we do think that some

[2]Indeed, we look precisely at the class of discounting functions referred to as quasi-hyperbolic. As indicated in the previous footnote, the correct mathematical term for these functions is quasi-geometric; they are quasi-hyperbolic only in the sense that, for certain parameter values, they resemble a generalized hyperbolic function.

government commitment may be feasible, we do not think a full ability to commit is likely. Whatever lack of commitment remains is the subject of this paper.

Our results are rather striking. We find that, whenever there is a time-inconsistency in preferences, not only does a benevolent-social-planning economy not deliver the same consumption allocation as does a laissez-faire world, but it delivers strictly lower welfare. Thus, we find a new argument for why the market mechanism is a particularly good one. The key insight regards price-taking behavior, as opposed to taking into account the impact of one's decisions on prices (or aggregate allocations). An alternative interpretation of the result is that a decentralization with many identical Robinson Crusoes each operating their own production technology would do as poorly as the visible-hand world: a separation of consumption and production decisions is desirable in this economy.

The mechanism behind our main finding is intuitive and can be understood as follows. Suppose that the social planner's preferences coincide exactly with those of the current consumer. They would then both regard their respective future selves as saving too little (the argument focuses on this particular bias for illustration). As a consequence, they would both want to manipulate their future selves. The manipulation occurs via savings (which is their only decision variable here): when an extra dollar is saved, the extra income next period will influence savings next period. With time-consistent preferences, the current saver agrees with his future self about next period's savings decision, and an envelope argument allows him to ignore this effect when deciding on current savings. Here, in contrast, an extra dollar saved today has an additional benefit next period, since (so long as the marginal savings propensity is not negative or zero) it induces more future savings. The difference between the decentralized allocation and the central planning allocation originates in how these induced future increases in savings are perceived: they are perceived as being larger by a price-taker than by a planner. The reason for this is that the price-taker's future behavioral response is "more linear": the added income he obtains next period is the higher savings times the return, which he perceives as constant. The planner, on the other hand, sees that for every additional dollar saved, the next period's return falls (assuming decreasing marginal productivity to capital). In sum, the incentives to save are higher for the decentralized agents, so they save more, and more saving is better in this economy: it takes us closer to the full commitment outcome.[3]

The most closely related literature to this paper is the set of studies started by Strotz (1956) and Phelps and Pollak (1968) and then further developed recently in a set of papers

[3] The argument is parallel when the time-inconsistency takes the form of excessive short-run patience.

Laibson, e.g., Laibson (1994). Laibson (1996) discusses some aspects of government taxation, but does not consider the case when the government cannot commit its future tax rates, the main focus of our analysis here. In many of Laibson's setups (e.g., Laibson (1997)), there are added market frictions, such as credit constraints, which makes illiquid assets play an important role. As a consequence, Ricardian equivalence does not hold. Another common assumption is the existence of partially uninsurable idiosyncratic shocks (Harris and Laibson (2000)). Here, the setup is the most basic one possible: the only friction is the internal friction in consumers' preferences (naturally accompanied with a lack of commitment), and there is no uncertainty. In particular, all long- and short-term asset markets are operative; although it would be beneficial to close asset markets in the future, without the commitment power to do so, ex post it will always be in the interest of the group of consumers as a whole to keep current markets open. As a result, in our environment, Ricardian equivalence holds. Another example of an added friction is the one assumed in O'Donoghue and Rabin (1999), where consumers are not rational, but are rather constantly surprised at their preference reversals.

The analysis in the paper builds fundamentally on recursive methods, and moreover uses specific functional forms so as to allow manageable closed-form solutions. The dynamic game thus played between the current and future selves is perhaps the simplest tractable example of a time-consistent equilibrium available in the literature. This literature plans began in Kydland and Prescott (1977) but has not allowed simple closed-form examples of Markov equilibria. We consider the case with a saver with time-inconsistent preferences as very natural grounds for illustrating the main forces at work. Central to this illustration are the recursive setup and the derivation and discussion of the "generalized Euler equation".[4] Section 2 describes the basic setup and the principles we follow in our analysis. Section 3 introduces recursive competitive equilibrium (without policy), and Section 4 considers the planning problem, including the comparison with the competitive equilibrium outcome. Section 5 looks at time-consistent policy: we show that the planning outcome is also the outcome of a government policy game, provided the government has a sufficiently large set of instruments. Section 7 discusses some weaknesses in the analysis, suggests extensions, and then concludes.

[4] See Laibson (1997).

2 The Setup

2.1 Primitives

Time is discrete and infinite and begins at time 0; there is no uncertainty. An infinitely-lived consumer derives utility from a stream of consumption at different dates. We assume that the preferences of this individual are time-additive, and that they take the form of a sequence of preference profiles:

$$U_0 = u_0 + \beta\left(\delta u_1 + \delta^2 u_2 + \delta^3 u_3 + \ldots\right)$$
$$U_1 = u_1 + \beta\left(\delta u_2 + \delta^2 u_3 + \ldots\right)$$
$$U_2 = u_2 + \beta\left(\delta u_3 + \ldots\right)$$

When $\beta = 1$, we have standard, time-consistent, geometric preferences. When $\beta \neq 1$, there is a time-inconsistency: at date 0, the trade-off between dates 1 and 2 is perceived differently than at date 1, and so on. When $\beta < 1$, we have "excessive short-run impatience": the individual thinks "I want to save, just not right now"; when $\beta > 1$, excessive short-run patience is expressed as "I want to consume, just not right now". We refer to this class of preferences as quasi-geometric, as they are a one-period deviation from the standard geometric case.[5] Successively more complicated extensions are straightforward to analyze within our framework.[6] Figure 1 below illustrates the different cases.

[5]The term quasi-hyperbolic is used in the literature as referring to the same preference setup even though, mathematically, hyperbolic functions take an entirely different form. The reason for the use of the term quasi-hyperbolic is that, for certain parameter values—see Figure 1 below—the discounting function resembles a (generalized) hyperbola. Here, we are interested in the entire class of quasi-geometric preferences.

[6]One would then, for example, assume that the disagreement between current self and the self k periods later disagree not only on the value of date $t+k$ consumption relative to other goods, but on the relative value of date $t+k+1$ consumption as well. This extension would simply require the introduction of "another β". Successive introductions of more βs would then relax the geometric framework more and more.

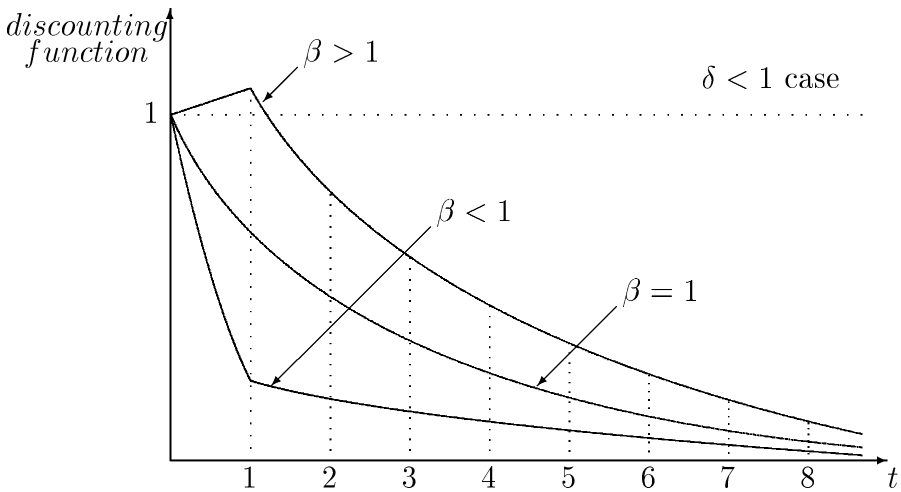

Figure 1

The standard infinite-life setup is often interpreted in terms of an infinitely-lived dynasty. This interpretation is possible here, too. If c_t refers to the entire consumption of generation t, then the assumption of quasi-geometric discounting implies a version of "impure" altruism: although the generation t agent cares about the consumption of all his descendants, he disagrees with his descendants on the weights. For example, he places a higher weight on the consumption of his grandchildren relative to his children than do his children. Thus, with dynasties in mind, the dynamic game we study here can be thought of as a game between different generations.

2.2 Modelling behavior under time-inconsistent preferences

As time progresses, the individual will change his mind about the relative values of consumption at different points in time so long as $\beta \neq 1$. He would, therefore, if he could, commit his future consumption levels. We assume that there is no way for the consumer to do so. This is a rather natural assumption in a framework with time-inconsistent preferences. The reason is that commitment contracts would have to be quite elaborate to avoid renegotiation problems. Suppose two consumers, A and B, agree on a contract whereby consumer B would punish consumer A for any deviation from the planned future consumption. For his services, consumer B would be paid some reward. At the future date, however, consumer A

would be able to convince consumer B not to carry out the punishment: A would just offer slightly more than the original reward to B for tearing the contract instead of adhering to it. Hence, unless the two consumers were playing a infinite game, renegotiation would always make both consumers better off ex post, and any forward-looking consumer would not bother to set up a commitment contract. The fact that a truly infinite horizon is necessary for commitment contracts to work is an argument that makes anonymous markets unlikely suppliers of commitment services in practice, and we believe that it is an important reason why we do not observe such a market.[7] In our model, the horizon is infinite, but we make a general restriction to Markov strategies in our work—one that we discuss below—and one consequence of this restriction is that commitment cannot be achieved.

Further, we assume that the consumer realizes that his preferences will change and makes the current decision taking this into account—this encapsulates our notion of rationality in this framework.[8] This means that we model the decision-making process as a dynamic game, with the agent's current and future selves as players.

For our game, we focus on (first-order) Markov equilibria: at a moment in time, no histories are assumed to matter for outcomes beyond what is summarized in the current stock of wealth held by the agent. The restriction to Markov equilibria is a natural way of restricting the set of equilibria.[9] Further, we restrict attention to those Markov equilibria which are limits of finite-horizon equilibria. This refinement eliminates a large number of equilibria: Krusell and Smith (2000) show that there is a large set of equilibria for this game even when attention is restricted to first-order Markov equilibria. There will typically be a unique remaining equilibrium, one that generalizes the standard preference setup in a continuous manner.

2.3 Assumptions on primitives

In order to obtain closed-form solutions, we restrict preferences and technology to specific functional forms. The period utility function is $u(c) = \log c$. We assume that $\delta < 1$ and that $\beta > 0$. Production is Cobb-Douglas and there is full depreciation so the resource constraint reads

$$c + k' = Ak^{\alpha}.$$

[7]One could argue that commitment mechanisms are more likely to occur within tight social groups/families. Social networks would then potentially play an important role in accumulation decisions.

[8]Others, such as O'Donoghue and Rabin (1999), do not make this assumption.

[9]Laibson (1994) and Bernheim, Ray, and Yeltekin (1999) study trigger-strategy equilibria. We think renegotiation arguments for the Markov assumption are particularly compelling in this "one-agent" game, although we do not have a formal defense for this position.

Primes denote consecutive-period values.

Perfect competition implies marginal-product pricing of the capital and labor inputs:

$$r = \alpha A k^{\alpha-1}$$

$$w = (1-\alpha) A k^{\alpha}.$$

2.4 Recursive formulation of the decision making

Assume that the agent perceives future savings decisions to be given by a function $g(k)$:

$$k_{t+1} = g(k_t).$$

Note that, by the Markov assumption, g is time-independent and only has current capital as an argument.

The agent solves

$$V_0(k) \equiv \max_{k'} u(rk + w - k') + \beta\delta V(k'),$$

where

$$V(k) = u(rk + w - g(k)) + \delta V(g(k)).$$

Notice that successive substitution of V into the agent's objective generates the right objective if the expectations of future behavior are given by the function g.

A solution to the agent's problem is denoted $\tilde{g}(k)$. We have a solution to the agent's problem if the fixed-point condition $\tilde{g}(k) = g(k)$ is satisfied for all k.

Parenthetically, the commitment solution would be obtained if the expression for V instead satisfied the standard dynamic programming functional equation

$$V(k) = \max_{k'} u(rk + w - k') + \delta V(k'),$$

which would yield a g function that differs from $\tilde{g}$ so long as $\beta \neq 1$; $\tilde{g}$ would be used at time 0, and the g forever after.

2.5 Constant prices: an explicit solution

If the prices r and w in the previous section are constant and exogenous, then, with our parametric assumptions, we can solve explicitly for the equilibrium decision rule and corresponding value function. In particular, the value function takes the form $V(k) = a + b\log(k + \frac{w}{r-1})$, where $k + \frac{w}{r-1}$ is proportional to the present value of the agent's lifetime wealth $W \equiv rk + w\sum_{i=0}^{\infty}\frac{1}{r^i}$. The equilibrium decision rule takes the form $g(k) = s\left(rk + \frac{r}{r-1}w\right) - \frac{w}{r-1}$, where $s = \frac{\beta\delta}{1-\delta(1-\beta)}$. It is not too hard to see that this decision rule implies that $W' = srW$, that is, the agent saves a constant fraction of his wealth in each period.

On a stationary point, where the individual's capital stock does not change, we have $g(k) = k$ and $W' = W$. This implies $r = \frac{1}{s}$, a requirement for a steady state in the general equilibrium model, which in turn equals $\frac{1-\delta(1-\beta)}{\beta\delta}$. This means that $\frac{w}{r-1} = s\frac{r}{r-1}w$. The left-hand side of this equality is present value of future labor income; the right-hand side is equal to the savings out of the present value of total labor income. These two need to be equal to each other in equilibrium. As we will show, this kind of condition will also hold off the steady state in the economy with our particular functional-form assumptions.

3 Recursive Competitive Equilibrium

Before we move on to the formal definition of a recursive competitive equilibrium, we need to describe the market structure. Our assumption here is that the consumer rents his capital and labor services to firms, treating prices parametrically. Further, the consumer makes the accumulation decisions for capital. If we assumed that firms made these decisions and that consumers had access to markets for one-period loans, the results would not change. The addition of multiperiod assets would also not change our results: these assets would be priced using arbitrage and their returns would be given by the returns on the relevant one-period assets. If one-period assets did not exist, the results would change. However, with a similar argument as the one used above, it would always be in the interest of consumers ex post to open one-period asset markets, and the spirit we follow here is to treat consumers as fully rational at any moment in time. Thus, the natural benchmark for us is to assume that one-period asset markets are open.

In order to analyze a general equilibrium, on and off its steady state, we need to be explicit about state variables. The agent makes his decision taking as given the prices as functions of the aggregate state $\bar{k}$, $\left(r(\bar{k}), w(\bar{k})\right)$, the law of motion for the aggregate state, $\overline{k}' = G\left(\overline{k}\right)$, and the decision rules of his future selves: $g\left(k, \overline{k}\right)$. The recursive equilibrium requires two state variables for the individual: one for the individual's own capital holdings, k, and one for the average capital holdings in the economy, $\bar{k}$, the latter reflecting prices. The agent's problem can be formulated in a similar way to before:

$$V_0\left(k, \overline{k}\right) = \max_{k'} \log\left(r(\bar{k})k + w(\bar{k}) - k'\right) + \beta\delta V\left(k', \overline{k}'\right)$$

The solution to this problem is given by $\widetilde{g}\left(k, \overline{k}\right)$, where $V\left(k, \overline{k}\right)$ satisfies

$$V\left(k, \overline{k}\right) = \log\left(r(\bar{k})k + w(\bar{k}) - g\left(k, \overline{k}\right)\right) + \delta V\left(g\left(k, \overline{k}\right), \overline{k}'\right).$$

Formally, we have

Definition 1 *A recursive competitive equilibrium for this economy consists of a decision rule,* $g\left(k,\overline{k}\right)$, *a value function,* $V\left(k,\overline{k}\right)$, *pricing functions* $r(\bar{k})$ *and* $w(\bar{k})$, *and a law of motion for aggregate capital,* $\overline{k}' = G\left(\overline{k}\right)$, *such that*

1. *given* $V\left(k,\overline{k}\right)$, $g\left(k,\overline{k}\right)$ *solves the maximization problem above;*
2. *given* $g\left(k,\overline{k}\right)$, $V\left(k,\overline{k}\right)$ *satisfies the functional equation above;*
3. *firms maximize, i.e.,* $r(\bar{k}) = f'\left(\overline{k}\right)$ *and* $w(\bar{k}) = (1-\alpha) f\left(\overline{k}\right)$;
4. *and the law of motion for aggregate capital resulting from the agent's decision is consistent with the law of motion for aggregate capital, i.e.,* $g\left(\overline{k},\overline{k}\right) = G\left(\overline{k}\right)$.

We have

Proposition 1 *For our parametric economy, the recursive competitive equilibrium is given by:*

1. $V\left(k,\overline{k}\right) = a + b\log\overline{k} + c\log\left(k+\varphi\overline{k}\right)$, *where* $c = \frac{1}{1-\delta}$, $b = \frac{\alpha-1}{(1-\delta)(1-\alpha\delta)}$, *and* $\varphi = \frac{(1-\alpha)(1-\delta(1-\beta))}{\alpha(1-\delta)}$;
2. $g\left(k,\overline{k}\right) = \frac{\delta\beta}{1-\delta(1-\beta)} r\left(\overline{k}\right) k$; *and*
3. $G\left(\overline{k}\right) = g\left(\overline{k},\overline{k}\right) = \frac{\alpha\delta\beta}{1-\delta(1-\beta)} A\overline{k}^{\alpha}$

Proof: See the Appendix.

The proof uses the intuition developed in the partial equilibrium case. Notice that the equilibrium savings rates in the partial equilibrium and the general equilibrium solutions are the same. We have also suppressed the discussion of other equilibria; all equilibria calculated in this paper are limits of finite-horizon equilibria, and as such are unique.

We now turn to the equivalence of the aggregate statistics across models with time-consistent and time-inconsistent preferences (this is a result parallel with that found in Barro (1999) for a continuous-time model).

Proposition 2 *The laws of motion for capital are the same for any two models* $\left(\hat{\beta}, \hat{\delta}\right)$ *and* (β, δ) *such that* $\frac{\hat{\delta}\hat{\beta}}{1-\hat{\delta}\left(1-\hat{\beta}\right)} = \frac{\delta\beta}{1-\delta(1-\beta)}$.

The proposition implies that it is not possible to estimate (β, δ) by looking at aggregates (or disaggregated variables, for that matter) for the class of preferences and technology that we concentrate on. As a special case, note that the outcome of the model with (β, δ), is identical to the outcome of the standard growth model with $\widehat{\delta} = \frac{\delta\beta}{1-\delta(1-\beta)}$.

The observational equivalence simplifies presentation here, but does not eliminate the issue we are interested in: the welfare properties of equilibria, and policy analysis. As we will see, welfare properties across models with different underlying preferences but the same equilibrium laws of motion can be very different.

4 Welfare Properties

Welfare properties are usually discussed in terms of the Pareto criterion. Since we are considering an economy with lack of commitment as a central element, it is difficult to formalize this criterion: moving allocations freely around over time violates the assumption of a lack of commitment. We will phrase our discussion in terms of a central planning problem; we will define such a problem, and then compare its solution to the competitive equilibrium. In contrast to a Pareto problem, the central planning problem we formulate does not necessarily have a solution with "good" welfare properties.

There is no obvious best notion of a planner here, since the consumer's different selves disagree. However, although one could think of a "meta-planner" placing positive weights on the lifetime utilities of more than the current self, we do assume here that the planner simply shares the preferences of the current self.[10] That is, by a central planning solution we simply mean a solution which would be chosen by a benevolent representative of the consumers who had the ability to manipulate the current economic choice variables costlessly. Thus, we want to think of the analysis of the benefits of markets as a simple "invisible hand" versus "visible hand" comparison, in the spirit of Adam Smith.

An important reason for adopting this specific planner is found in Section 5 below. There we show that if one formalizes the notion of a government with a sufficiently large set of tax instruments representing the preferences of its current electorate, the resulting allocation coincides with the one we obtain in our fictitious planning economy. Finally, and as motivated

[10]We briefly discuss implications of having meta-planners below.

above, we further assume that the planner cannot directly affect his future choices, thereby having to play a game with his future selves, just like the consumer does in a competitive equilibrium.

4.1 The planning problem

In summary, in this section we assume the following: (i) the planner is a consumer representative: he inherits his (time- inconsistent) preferences; (ii) he faces the same problem as the consumer: he cannot commit to future actions; (iii) we require a time-consistent solution to the planner's problem (future planners' reactions are taken into account in a rational way).

The differences between the consumer's equilibrium problem and the planner's problem are thus as follows: (i) the consumer takes prices as given, whereas the planner has a resource constraint; and (ii) the equilibrium consumer deals with different future players (the consumer's future price-taking selves) than the planner does (the planner's future selves).

The planner's problem can be formulated in the following way:

$$V_0(k) \equiv \max_{k'} u(f(k) - k') + \beta\delta V(k'),$$

where

$$V(k) = u(f(k) - h(k)) + \delta V(h(k)).$$

A solution to the planner's problem is denoted $\widetilde{h}(k)$. We have a solution to the planner's problem if the fixed-point condition $\widetilde{h}(k) = h(k)$ is satisfied for all k.

It is straightforward to show that the following functions are solutions to the planning problem:

$$k' = h(k) = \frac{\beta\delta}{1 - \alpha\delta(1-\beta)} \alpha A k^{\alpha}$$

and

$$V(k) = a + b \log k.$$

The constants a and b are discussed below.

4.2 The planning outcome vs. the competitive outcome

Apparently, whenever $\alpha < 1$, the competitive equilibrium and the planning problem produce different outcomes. When $\beta < 1$, the price-taking agent saves more than the planner does; when $\beta > 1$, the planner saves more. Neither the planner's solution nor the competitive equilibrium outcome coincide with the full commitment solution. The commitment solution for this economy has one savings rate at time 0 and another, higher, one at all future times.

The time-0 commitment rate is equal to the savings rate of the planning problem without commitment; the subsequent savings rate is higher (lower) than both the competitive and the no-commitment planning outcomes when $\beta < (>)\, 1$.

Turning to a welfare comparison between the competitive equilibrium and the planning solution, we have the following.

Proposition 3 *The competitive outcome results in strictly higher welfare for the consumer than the planning outcome does, whenever $\beta \neq 1$ and $\alpha < 1$.*

Proof: See the appendix.

Thus, markets outperform a benevolent social planner in this economy; in particular, competitive behavior results in higher savings (when $\beta < 1$), and since there is undersaving relative to the full commitment case, higher savings moves the economy in the right direction. What explains this finding? As pointed out above, it should not be surprising that the two allocations are different. The planning problem is not a standard planning problem; in particular, the planner faces a different environment than do the competitive consumers: they face different future players, whom they cannot fully control. In order to explain why the planner faces future players that induce worse outcomes, we will make use of the generalized Euler equation in each case.

4.3 The generalized Euler equation

A simple derivation of the generalized Euler equation (GEE), which can be used whenever the value function and the policy function are differentiable, goes as follows. Consider the problem of the planner. His first-order condition reads

$$u'(f(k) - h(k)) = \beta\delta V'(h(k)).$$

To eliminate the unknown function V', take derivatives of the functional equation for V:

$$V'(k) = u'(f(k) - h(k))(f'(k) - h'(k)) + \delta V'(h(k))h'(k).$$

Substitute $V'(h(k))$ from the first of these equations into the second, then update the second and substitute the new expression for $V'(h(k))$ back into the first equation. This gives

$$u'(f(k) - h(k)) =$$

$$\beta\delta u'(f(h(k)) - h(h(k))) \left[f'(h(k)) + (\beta^{-1} - 1)h'(h(k))\right].$$

This is our key behavioral equation: it is a functional equation in the unknown savings function $h(k)$. For readability, consider this equation in sequential form:

$$u'(c_t) = \beta\delta u'(c_{t+1}) \left[f'(k_{t+1}) + (\beta^{-1} - 1)h'(k_{t+1})\right].$$

Notice the $h'(k_{t+1})$ term on the right-hand side: you do not agree with your future self about savings propensities, and therefore value giving your future self more wealth (if $\beta < 1$). When β is equal to 1, this term does not appear: the envelope theorem (which allows the second equation above to be written simply as $V'(k) = u'(f(k) - h(k))f'(k))$ dictates that, since you agree with your future self, you in essence make the future decision yourself, so the indirect effects on savings next period, as captured by h', are second order. This additional "return" to saving a unit today is crucial in what follows.

Similarly, for the competitive equilibrium we obtain

$$u'(f(k) - G(k)) = \delta u'(f(G(k)) - G(G(k))) \left(\beta f'(G(k)) + (1 - \beta)g_1(G(k), G(k))\right),$$

or, in sequential form,

$$u'(c_t) = \beta\delta u'(c_{t+1}) \left[f'(k_{t+1}) + (\beta^{-1} - 1)g_1(k_{t+1}, k_{t+1})\right].$$

The two GEE's look identical except for the derivative term (h' and g_1, respectively). Here is the key difference: h has decreasing returns to its argument,

$$h(k) = \hat{s}\alpha A k^{\alpha}$$

if $\alpha < 1$, whereas g is linear in its first argument,

$$g(k, \bar{k}) = sk\alpha A\bar{k}^{\alpha-1}.$$

So, (if $\beta < 1$) the competitive equilibrium consumer sees a higher benefit from extra saving today than does the planner: everything else equal, the planner sees another unit of savings as yielding a smaller increase in future savings than does the competitive-equilibrium consumer. This makes the competitive agent save more than the planner, because the future selves are undersaving and extra future saving is now a good thing. The argument works for $\beta > 1$ as well. Then, the planner saves more than the competitive equilibrium, which saves too much. An extra saved unit increases future savings more as perceived by the consumer than as perceived by the planner. As a result, the consumer saves less than the planner since extra future savings is now a bad thing.

Behind these arguments is the main difference between a planner and a competitive individual: the planner understands that he affects the "prices", that is, the return to savings, whereas a price-taker does not. This implies that the marginal propensity to save for the planner is decreasing, whereas it is constant for the competitive agent. This is the key difference behind the planner's and the competitive agent's decision rules.

4.4 Utility comparisons for other agents

One might take the view that the consumer's future selves, and their different preferences, ought to be respected and taken separately into account in the comparison between the two allocations above. Suppose, therefore, that we evaluate the utility of the self next in line as he perceives it. Would he prefer the competitive or the central planning allocation?

It is clear from the arguments in the preceding section that, if the next self were given the same amount of capital to start with in both situations, then he would prefer the competitive allocation; it provides higher savings at all times, which is perceived as better (if $\beta < 1$). However, capital is not constant across the two allocations. Depending on the value of β, it is either higher or lower. If $\beta < 1$, it is higher, and the competitive equilibrium dominates, not only in the next period, but in all future periods as well. This is the case most commonly emphasized in the literature on time-inconsistent preferences. In this case, therefore, the competitive equilibrium allocation Pareto dominates the central planning allocation. If, on the other hand, $\beta > 1$, there is an effect in favor of the central planning allocation for all future selves, and the net result is not clear.

The point of view that future selves should be taken into account separately could be expressed also by defining the central planner as someone who shares this view. We return to such a planner in Section 5.3.

5 Policy Analysis

The planning problem in the previous section can perhaps be viewed as artificial. In this section we consider a government with taxation abilities. There are no government expenditures and we impose budget balance; moreover, we assume that the government is benevolent in that it is the representative of the consumer. As the planner, the government cannot commit its future self: future tax rates are set by future governments, and future governments have different preferences, as they are representatives for the futures selves of the consumer.

Assume that the government can tax income and investment at proportional rates τ_y and

τ_i, respectively. The first of these is essentially a lump-sum tax in this model, since income and substitution effects cancel with logarithmic utility. The investment tax, on the other hand, is distortionary.

We first consider the full commitment case for illustration. We then turn to time-consistent taxation. Time-consistent taxation in this model leads to taxes which are constant over time (as are savings rates). We finally consider a third alternative policy experiment, one where there is a possibility of committing to a constant tax rate forever, a "constant-tax constitution", and show how this tax rate would be chosen.

5.1 Full commitment to future taxes

Since the consumer's future selves are undersaving from the perspective of the current self, the government will want to subsidize investment in all future periods. Given the stationarity of the problem and the log/Cobb-Douglas assumptions, these rates will be constant and given by $\tau_i' = \frac{(1-\delta)(\beta-1)}{1-\delta+\delta\beta-\alpha\delta\beta}$ and $\tau_y' = \frac{\alpha\delta(1-\delta)(1-\beta)}{1-\delta+\delta\beta-\alpha\delta\beta}$. These tax rates for future periods will give a savings rate of δ for these periods. For the current period, the government sets the tax rates $\tau_i = \frac{\delta(1-\alpha)(1-\beta)}{1-\alpha\delta-\delta(1-\alpha)(1-\beta)}$ and $\tau_y' = \frac{\alpha\delta^2\beta(1-\alpha)(1-\beta)}{\alpha\delta^2\beta(1-\alpha)(1-\beta)-(1-\alpha\delta)(1-\delta(1-\beta))}$. This gives a savings rate of $\frac{\delta\beta}{1-\alpha\delta(1-\beta)}$ for the current period. This tax sequence generated by fully committed governments is not time-consistent. This creates incentives for future governments to deviate. We now move on to time-consistent taxation.

5.2 No commitment: time-consistent policy

We define a subgame-perfect Markov equilibrium for the government problem. This definition parallels the definition of subgame-perfect Markov equilibria: a tax function describes the tax outcome as a function of the aggregate state, and to support this tax function it is necessary to consider one-period deviations from the equilibrium path.[11] The competitive equilibrium is defined as above for given taxes. Before we go through the definition of the time-consistent policy equilibrium, let us point out that taxes in this equilibrium will be constant, due to our special parametric assumptions. To support them as being constant, however, it is necessary to formally verify all the conditions of a subgame-perfect equilibrium.

Definition 2 *A time-consistent policy equilibrium is defined in several parts; the elements are listed below.*

First, the behavior on the equilibrium path ("outcomes") are as follows:

[11] For a definition in the context of a typical growth model, see Krusell and Ríos-Rull (1999).

- *Tax outcomes are given by* $\tau\left(\overline{k}\right)=\left(\tau_y\left(\overline{k}\right),\tau_i\left(\overline{k}\right)\right)$.
- *Given this tax function, the law of motion for aggregate capital is given by* $G\left(\overline{k}\right)$.
- *Given the tax function and the law of motion for aggregate capital, the individual's decision rule is given by* $g\left(k,\overline{k}\right)$

Second, the one-period deviations to tax rates $\tilde{\tau}=(\tilde{\tau}_y,\tilde{\tau}_i)$ *for the current period, with future taxes given by the tax outcome functions evaluated at the capital stocks implied by the current tax rates and the implied capital accumulation, are given by the following:*

- $\widetilde{G}\left(\overline{k},\tilde{\tau}\right)$ *describes the law of motion for aggregate capital for the one-period deviation.*
- $\tilde{g}\left(k,\overline{k},\tilde{\tau}\right)$ *describes the individual's decision rule for the one-period deviation.*

Third, we have competitive pricing functions $r(\bar{k})$ *and* $w(\bar{k})$ *(equal to the marginal products off the aggregate production function).*

These equilibrium elements have to satisfy:

- *Individual optimization:* $\tilde{g}\left(k,\bar{k},\tilde{\tau}\right)$ *solves*

$$V_0(k,\bar{k},\tilde{\tau})\equiv$$

$$\max_{k'}\left[\log\left(\left(r(\bar{k})k+w(\bar{k})\right)(1-\tilde{\tau}_y)-k'\left(1+\tilde{\tau}_i\right)\right)+\beta\delta V\left(k',\tilde{G}(\bar{k},\tilde{\tau})\right)\right]$$

  and $V\left(k,\bar{k}\right)$ *satisfies*

$$V\left(k,\bar{k}\right)=\log\left(\left(r(\bar{k})k+w(\bar{k})\right)\left(1-\tau_y\left(\bar{k}\right)\right)-g\left(k,\bar{k}\right)\left(1+\tau_i\left(\bar{k}\right)\right)\right)+$$
$$\delta V\left(g\left(k,\bar{k}\right),G(\bar{k})\right).$$

  Note that these requirements imply, as a special case, that $\tilde{g}\left(k,\bar{k},\tau(\bar{k})\right)=g\left(k,\bar{k}\right)$.

- *Consistency between individual and aggregate actions:* $\tilde{g}\left(\bar{k},\bar{k},\tilde{\tau}\right)=\tilde{G}\left(\bar{k},\tilde{\tau}\right)$, *which implies as a special case that* $g\left(\bar{k},\bar{k}\right)=G\left(\bar{k}\right)$.
- *The government maximizes:* $\tau\left(\bar{k}\right)=\left(\tau_y\left(\bar{k}\right),\tau_i\left(\bar{k}\right)\right)$ *solves the following problem:*

$$\max_{(\tilde{\tau}_y,\tilde{\tau}_i)} V_0(\bar{k},\bar{k},\tilde{\tau})$$

  subject to:

$$-\tilde{G}\left(\bar{k},\tilde{\tau}\right)\tilde{\tau}_i=A\bar{k}^{\alpha}\tilde{\tau}_y.$$

Solutions to the problems above can be obtained in the same manner as we derived competitive equilibria above. We conjecture that $\tau_y\left(\bar{k}\right) = \tau_y$, $\tau_i\left(\bar{k}\right) = \tau_i$, i.e., that the tax functions will be constant. This conjecture is straightforward to verify. The one-period deviation decision rule does not depend on future tax rates. If future tax rates were a nontrivial function of aggregate capital then today's tax policy would affect the future tax rates. With the conjecture that the tax outcome function is constant, a number of derivatives become zero. The constancy of the tax function in particular implies that the current capital stock has no importance for how taxes are set; this would not be true in a calibrated growth model. There, the tax function would depend on capital, taxes would change along the transition path, and a one-period deviation in tax rates would alter the tax rates forever after. Needless to say, the parametric assumptions here simplify the analysis tremendously.

Proposition 4 *The time-consistent tax rates are given by*

$$\widetilde{\tau}_i = \frac{\delta(1-\alpha)(1-\beta)}{1-\alpha\delta-\delta(1-\alpha)(1-\beta)} > (<)0 \text{ if } \beta < (>)1$$

and

$$\widetilde{\tau}_y = \frac{\alpha\delta^2\beta(1-\alpha)(1-\beta)}{\alpha\delta^2\beta(1-\alpha)(1-\beta)-(1-\alpha\delta)(1-\delta(1-\beta))} < (>)0 \text{ if } \beta < (>)1.$$

Proof: See the appendix.

It is straightforward to verify that the time-consistent tax rates, perhaps not surprisingly, reproduce the allocations that solve the planner's problem: the government has enough instruments to manipulate current decisions freely, and so chooses the same outcome as if it were a central planner.

As a positive theory of taxation, the model implies

$$\tau_i > (<)\ 0 \text{ if } \beta < (>)\ 1$$

and

$$\tau_y < (>)\ 0 \text{ if } \beta < (>)\ 1.$$

That is, we have positive tax rates on investment when $\beta < 1$. The social planner saves much too little (less than the laissez-faire equilibrium) and so wants to move the equilibrium in the "wrong direction".

If one took the view that the future selves were distinct individuals, then one could also argue that the planner should maximize some other objective, say, some function that puts positive weights on each of the different perceptions of the utility of the chosen consumption path. We consider such a case explicitly in the next section.

5.3 Other government preferences

The government in the time-consistent policy equilibrium that we just studied cares about the utility of the current self. In terms of positive policy analysis, it is interesting to study how the policy implications would change if the government had different preferences. In particular, one can analyze the assumption that the government gives some separate weight to utility as perceived by future generations. Moreover, as a normative issue, it is interesting to look at how governments with different preferences compare in terms of the utility they deliver. If the government always puts some separate weight on future selves, might the utility of the current self be higher in equilibrium than if the government only cares about the current self?

It turns out, as we show in the appendix, that if the government represented both the current and future selves of the consumer—in particular, if it had weights on the utilities of the sequence of future selves—then we would obtain a time-consistent equilibrium with a constant tax rate.[12] As these weights are varied, different constant-tax equilibria are obtained. As the weight on the utility of the future selves increases, tax rates decrease. This is a rather straightforward result: any concern for a future self over the current one simply means a wish to move resources forward in time, thus decreasing current consumption and increasing current investment.

Given this relation between government preferences and tax outcomes, we can simply study the set of constant-tax equilibria. How do these equilibria compare from a normative perspective? Suppose we take the perspective of the current self. The utility of the current self depends on the tax rate. An illustration can be found in Figure 2, where we plot utility against the equilibrium savings rate. The left-hand side of the figure is the $\beta < 1$ case; it plots the utility of the representative (current) consumer as a function of the savings rate. Each such savings rate can thus be thought of as produced by a specific set of constant tax rates (one tax only, in practice, since the other follows from the government budget). This graph does not depend on the current value of capital, again because we are in a special parametric class. The right-hand side, similarly, illustrates the $\beta > 1$ case.

[12] The only restriction on the weights is that they imply a well-defined government objective.

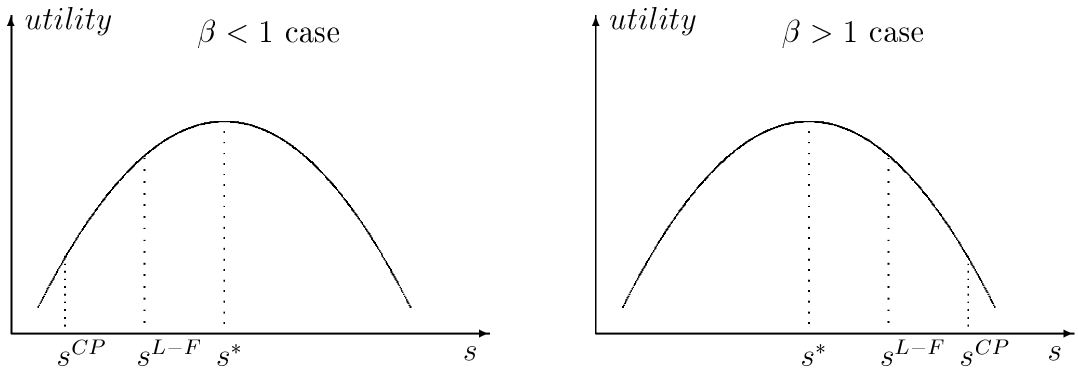

Figure 2

The savings rate which produces the highest utility is labeled $s^\star$. This savings rate comes about with a subsidy on investment. Assuming that $\beta < 1$, the laissez-faire competitive savings rate, i.e., the one resulting from zero taxes, s^{L-F}, is lower. Lower still is the central planning savings rate, s^{CP}, which levies positive taxes on investment.

The best tax rates here—the one that would select the highest utility for the current self, thus leading to $s^\star$ in Figure 2—are given by

$$\tau_i = \frac{-\alpha\delta\left(1-\beta\right)\left(1-\delta\right)}{1-\delta+\delta\beta-\alpha\delta\beta}$$

and

$$\tau_y = \frac{-\alpha\delta\beta\tau_i}{\left(1-\delta\left(1-\beta\right)\right)\left(1+\tau_i\right)-\alpha\delta\beta\tau_i};$$

the decision rule for capital becomes

$$G\left(\overline{k}\right) = \frac{\alpha\delta\beta}{\left(1-\alpha\delta\right)\left(1-\delta\left(1-\beta\right)\right)+\alpha\delta\beta}A\overline{k}^{\alpha}.$$

One could think of the exercise in this section as follows. Suppose that we could choose, now and for the entire future, a government which has taxing rights, and suppose that the government can be selected among a set of altruistic citizens with different relative weights on future selves, weights that are stationary and therefore are time-inconsistent as for the consumer. Starting from zero weight on the future selves, welfare as perceived by all selves would increase in this weight. It would take a citizen with a significant weight on the future selves in order to maximize the utility of the current self—to reach $s^\star$. Such a savings rate would make all selves better off. Further increasing the weight on the future selves would then decrease the utility of the current self, but increase the utility of (at least some of) the future selves. Summarizing, if an "ecologist"—someone with a significant weight on the utility of future generations—could be given permanent tenure as a central planner/government with taxation abilities, we would all be better off.

6 Summary and some concluding remarks

We have studied the performance of the market mechanism relative to a mechanism with a benevolent and potentially active government in what we believe is an interesting new case for economists to study: an economy where consumers have time-inconsistent preferences. Our analysis comes out surprisingly strongly in favor of laissez-faire; an initial reader of the literature on time-inconsistent preferences may get the impression that the consumer needs help and that the government can provide the help. We argue here that this is really only true if the government, or social planner, can help alleviate the commitment problem in a direct way, or indirectly by being able to commit to future tax rates. Further, under the assumption that the government cannot do anything about the commitment problem, we should strictly prefer laissez-faire. It would be important not to have a government that can tax. Although are goal here is not to send a libertarian message, the results here may perhaps be interpreted as a caution against tax policy activism: taxes should not be used to try to correct problems whose underlying cause is a lack of commitment to which the government is also subject.

Could the government provide commitment mechanisms to aid consumers with time-inconsistent preferences? To the extent that it could close short-term credit markets in the future, in effect "creating illiquidity" in the future, then it should, *ceteris paribus*. However, we argued that ex post—when the future arrives—it will be in everybody's interest not to close the markets, and it is unclear how a commitment mechanism for the closing of markets in the future could be created. Moreover, providing illiquidity by offering assets such as the 401(k)—with penalty for early withdrawals—might be helpful, but only if other restrictions on trading/borrowing are simultaneously and credibly put in place. In general, commitment is needed for consumption, and it is not enough to provide assets which are illiquid. Therefore, we take the present analysis as an entirely relevant case.

We found that our tax policy analysis bears resemblance to the "rules vs. discretion" literature. Our government here is benevolent—a samaritan—but it cannot help but adopt policies that are not good. What is the solution to our samaritan's dilemma? In the monetary literature, Rogoff's suggestion of electing a "conservative" does the trick. Here, if one could elect someone to lead government, it would be an “ecologist”: a consumer representative who puts significant independent weight on utility as perceived by consumers in the future or, rather, by the consumer’s future selves.

One weakness of the present approach to representing the preference reversals documented in the experimental psychology literature is that it makes a rather drastic conceptual departure from standard Arrow-Debreu analysis. In particular, it does not build up from axioms of choice for the individual. What would seem to be appropriately captured by an optimization problem—the consumer’s observed behavior—is instead modeled here as the solution to a dynamic game. As such, all the usual problems of (dynamic) game theory appear; for example, indeterminacy of equilibrium is the rule rather than the exception. We “resolve” this problem by focusing on what seems reasonable to us: the limit of the equilibria of finite-horizon economies. An alternative approach, which may ultimately turn out to be more fruitful, is undertaken in a set of papers by Gul and Pesendorfer (1999, 2000). These authors use decision theory, based on axioms over sets of consumption bundles, and arrive at recursive (time-consistent) preference representations of consumer behavior when “temptation” and “self-control” are represented axiomatically. In a related paper, which is in progress, we are considering the effects of policy on equilibrium allocations and welfare when there is “temptation” and “self-control”, as in Gul and Pesendorfer’s framework.

REFERENCES

Barro, Robert (1999), "Ramsey Meets Laibson in the Neoclassical Growth Model", *Quarterly Journal of Economics* 114, 1125–1152.

Bernheim, Douglas B., Debraj Ray, and Sevin Yeltekin (1999), "Self-Control, Saving, and the Low Asset Trap", manuscript.

Gul and Pesendorfer (1999), "Temptation and Self-Control", forthcoming in *Econometrica.*

Gul and Pesendorfer (2000), "Self-Control and the Theory of Consumption", manuscript.

Laibson, David (1994), "Self-Control and Saving", manuscript.

Laibson, David (1997), "Hyperbolic Discount Functions and Time Preference Heterogeneity", manuscript.

Laibson, David and Christopher Harris (2000), "Dynamic Choices and Hyperbolic Consumers", forthcoming in *Econometrica.*

Krusell, Per and José-Víctor Ríos-Rull (1999), "On the Size of U.S. Government: Political Economy in the Neoclassical Growth Model", *American Economic Review* 89, 1156–1181.

Krusell, Per and Anthony A. Smith, Jr. (2000), "Consumption-Savings Decisions with Quasi-Geometric Discounting", manuscript.

Kydland, Finn and Edward C. Prescott (1997), "Rules Rather than Discretion: The Inconsistency of Optimal Plans", *Journal of Political Economy* 85, 473–491.

Laibson, David (1994), "Self-Control and Saving", manuscript.

Laibson, David (1996), "Hyperbolic Discount Functions, Undersaving, and Savings Policy", NBER Working Paper 5635.

Laibson, David (1997), "Hyperbolic Discount Functions and Time Preference Heterogeneity", manuscript.

Laibson, David and Christopher Harris (2000), "Dynamic Choices and Hyperbolic Consumers", forthcoming in *Econometrica.*

O'Donoghue, Edward D. and Matthew Rabin (1999), "Doing It Now or Doing It Later", *American Economic Review* 89, 103–124.

Phelps, Edmund S. and Robert A. Pollak (1968), "On Second-best National Saving and Game-equilibrium Growth", *Review of Economic Studies* 35, 185–199.

Strotz, Robert H. (l1956), "Myopia and Inconsistency in Dynamic Utility Maximization", *Review of Economic Studies* 23, 165–180.

7 APPENDIX

Proof of Proposition 1: The proof follows by using the guess for the value function, $V\left(k,\overline{k}\right) = a + b\log\overline{k} + c\log\left(k+\varphi\overline{k}\right)$, and the guess for the law of motion for aggregate capital, $G\left(\overline{k}\right) = s\alpha A\overline{k}^{\alpha}$. [13] Given these guesses we can solve the agent's problem to obtain $g\left(k,\overline{k}\right) = \frac{\beta\delta c}{1+\beta\delta c}\left(Rk+w\right) - \frac{\varphi\overline{k}'}{1+\beta\delta c}$. Using this decision rule we can verify the guess for the value function and obtain $\varphi = \frac{1-\alpha}{\alpha(1-s)}$, $b = \frac{\alpha-1}{(1-\delta)(1-\alpha\delta)}$, and $c = \frac{1}{1-\delta}$. Inserting $\varphi = \frac{1-\alpha}{\alpha(1-s)}$ into the individual decision rules and setting $g\left(\overline{k},\overline{k}\right) = G\left(\overline{k}\right)$ (which has to hold in competitive equilibrium), we obtain $s = \frac{\delta\beta}{1-\delta(1-\beta)}$. This gives $\varphi = \frac{(1-\alpha)(1-\delta(1-\beta))}{\alpha(1-\delta)}$. Substituting these constants into the agent's decision rule, we obtain $g\left(k,\overline{k}\right)$ and $G\left(\overline{k}\right)$.

Proof of Proposition 3: The proof proceeds as follows: first we derive the value function of the current self, $V_0\left(k\right)$, given a general law of motion of type $k' = sAk^{\alpha}$. We thus obtain a function of s. We then evaluate this value function at $s_1 = \frac{\alpha\delta\beta}{1-\delta(1-\beta)}$ and at $s_2 = \frac{\alpha\delta\beta}{1-\alpha\delta(1-\beta)}$.

Let us first derive $V\left(k\right)$. For the given law of motion for capital, consumption will be given by $c = \left(1-s\right)Ak^{\alpha}$. Then we can write $V\left(k\right)$ as

$$V\left(k\right) = \log\left(\left(1-s\right)Ak^{\alpha}\right) + \delta\log\left(\left(1-s\right)Ak'^{\alpha}\right) + \delta^2\log\left(\left(1-s\right)Ak''^{\alpha}\right) + \ldots.$$

Inserting the law of motion for the capital we obtain

$$\begin{aligned} V\left(k\right) &= \frac{1-\alpha\delta}{\left(1-\alpha\delta\right)\left(1-\delta\right)}\log\left(1-s\right) + \frac{\alpha\delta}{\left(1-\alpha\delta\right)\left(1-\delta\right)}\log s \\ &\quad + \frac{\alpha}{\left(1-\alpha\delta\right)}\log k + \frac{1}{\left(1-\alpha\delta\right)\left(1-\delta\right)}\log A. \end{aligned}$$

$V_0\left(k\right)$ is now given by

$$\begin{aligned} V_0\left(k\right) &= \log\left(\left(1-s\right)Ak^{\alpha}\right) + \beta\delta V\left(sAk^{\alpha}\right) \\ &= \log\left(\left(1-s\right)Ak^{\alpha}\right) + \\ &\quad \beta\delta\left[\begin{array}{c} \frac{1-\alpha\delta}{(1-\alpha\delta)(1-\delta)}\log\left(1-s\right) + \frac{\alpha\delta}{(1-\alpha\delta)(1-\delta)}\log s \\ +\frac{\alpha}{(1-\alpha\delta)}\log\left(sAk^{\alpha}\right) + \frac{1}{(1-\alpha\delta)(1-\delta)}\log A \end{array}\right] \\ &= \frac{1-\delta\left(1-\beta\right)}{1-\delta}\log\left(1-s\right) + \frac{\alpha\delta\beta}{\left(1-\alpha\delta\right)\left(1-\delta\right)}\log s + \ldots. \end{aligned}$$

To evaluate $V_0\left(k\right)$ at s_1 and s_2, we proceed as follows. We first show that there is a unique s^* that maximizes $V_0\left(k\right)$, by showing that $V_0\left(k\right)$ is monotone increasing in s for $s < s^*$ and

[13] Where does this guess come from? $\varphi\overline{k}$ is actually equal to discounted value of lifetime wages (discounted to time -1 in this formulation): $\frac{w}{R} + \frac{w'}{RR'} + \frac{w''}{RR'R''} + \ldots..$ If we use $G\left(\overline{k}\right) = s\alpha A\overline{k}^{\alpha}$ and plug it into the discounted sum, we obtain $\varphi = \frac{1-\alpha}{\alpha(1-s)}$.

monotone decreasing for $s > s^*$. We then complete the proof by showing that $s_2 < s_1 < s^*$ for $\beta < 1$, and $s_2 > s_1 > s^*$ for $\beta > 1$. To do this, first let us look at the function $F(s) = (1-\delta(1-\beta))\log(1-s) + \frac{\alpha\delta\beta}{1-\alpha\delta}\log s$.

$$\begin{aligned} F'(s) &= -\frac{(1-\delta(1-\beta))}{1-s} + \frac{\alpha\delta\beta}{(1-\alpha\delta)s} \\ &= -\frac{[(1-\alpha\delta)(1-\delta(1-\beta))+\alpha\delta\beta]s-\alpha\delta\beta}{s(1-s)(1-\alpha\delta)} \end{aligned}$$

Note that $s^* = \frac{\alpha\delta\beta}{(1-\alpha\delta)(1-\delta(1-\beta))+\alpha\delta\beta}$, $F'(s) < 0$ ($V_0(k)$ is monotone decreasing) for $s > s^*$, and $F'(s) > 0$ ($V_0(k)$ is monotone increasing) for $s < s^*$. We can also easily see that $s_2 < s_1 < s^*$ for $\beta < 1$ and that $s_2 > s_1 > s^*$ for $\beta > 1$.

Proof of Proposition 4: We will demonstrate how to derive the unique solution to the government's problem. The government's problem is

$$\begin{aligned} \max_{(\tilde{\tau}_y,\tilde{\tau}_i)} V_0\left(\overline{k},\overline{k},\tilde{\tau}\right) &= \log\left(\left[A\overline{k}^{\alpha} - \frac{\alpha\delta\beta}{1-\alpha\delta(1-\beta)}A\overline{k}^{\alpha}\right](1-\tilde{\tau}_y)\right) + \beta\delta a \\ &\quad +\beta\delta\frac{\alpha-1}{(1-\delta)(1-\alpha\delta)}\log\left(\frac{\alpha\beta\delta}{1-\delta+\beta\delta}\frac{1-\tilde{\tau}_y}{1+\tilde{\tau}_i}A\overline{k}^{\alpha}\right) \\ &\quad +\frac{\beta\delta}{1-\delta}\log\left(\frac{\alpha\beta\delta}{1-\delta+\beta\delta}\frac{1-\tilde{\tau}_y}{1+\tilde{\tau}_i}A\overline{k}^{\alpha}(1+\varphi)\right). \end{aligned}$$

This problem is equivalent to

$$\max_{(\tilde{\tau}_y,\tilde{\tau}_i)} \frac{1-\alpha\delta+\alpha\delta\beta}{1-\alpha\delta}\log(1-\tilde{\tau}_y) - \frac{\alpha\delta\beta}{1-\alpha\delta}\log(1+\tilde{\tau}_i) + \ldots$$

The government budget reduces to

$$(\tilde{\tau}_y-1)\tilde{\tau}_i\frac{\alpha\delta\beta}{1-\delta(1-\beta)} = \tilde{\tau}_y(1+\tilde{\tau}_i).$$

Solving for $\tilde{\tau}_y$ in terms of $\tilde{\tau}_i$ from the government budget, the government's problem can be written as

$$\max_{\tilde{\tau}_i} Q(\tilde{\tau}_i) = \log(1+\tilde{\tau}_i) - \frac{1-\alpha\delta+\alpha\delta\beta}{1-\alpha\delta}\log((1-\delta(1-\beta))(1+\tilde{\tau}_i) - \alpha\delta\beta\tilde{\tau}_i).$$

$Q(\tilde{\tau}_i)$ is the government's objective, expressed as a function of $\tilde{\tau}_i$. Taking derivatives with respect to $\tilde{\tau}_i$, we can show that $Q'(\tilde{\tau}_i) > 0$ for $\tilde{\tau}_i < \frac{\delta(1-\alpha)(1-\beta)}{1-\alpha\delta-\delta(1-\alpha)(1-\beta)}$ and that $Q'(\tilde{\tau}_i) < 0$ for $\tilde{\tau}_i > \frac{\delta(1-\alpha)(1-\beta)}{1-\alpha\delta-\delta(1-\alpha)(1-\beta)}$. So there is a unique maximum to the government's problem. It is given by the constructed solution to the first-order condition.

A government representing future selves: Let $c_t > 0$ be the weight given by the government to the utility as perceived by self t; c_0 is therefore the weight given to current self. In this case, time-consistent government policy is given by the solution to the following problem:

$$\max_{(\widetilde{\tau}_y, \widetilde{\tau}_i)} c_0 V_0 \left(\overline{k}, \overline{k}, \widetilde{\tau}\right) + \sum_{t=1}^{\infty} c_t V_0 \left(\overline{k}\left(t\right), \overline{k}\left(t\right), \tau\right),$$

where τ is the tax vector selected by the future governments (as before, a time-consistent equilibrium will have a constant tax rate).

We can rewrite this problem as

$$\max_{(\widetilde{\tau}_y, \widetilde{\tau}_i)} c_0 \left[\log\left(1 - \widetilde{\tau}_y\right) + \frac{\alpha\delta\beta}{1 - \alpha\delta} \log\left(\frac{1 - \widetilde{\tau}_y}{1 + \widetilde{\tau}_i}\right)\right] + \frac{1 - \alpha\delta + \alpha\delta\beta}{1 - \alpha\delta} \left(\sum_{t=1}^{\infty} c_t \alpha^t\right) \log\left(\frac{1 - \widetilde{\tau}_y}{1 + \widetilde{\tau}_i}\right).$$

Now let

$$\lambda \equiv \frac{c_0}{c_0 + \frac{1-\alpha\delta+\alpha\delta\beta}{1-\alpha\delta}\left(\sum_{t=1}^{\infty} c_t \alpha^t\right)}.$$

Then the problem can be rewritten as

$$\max_{(\widetilde{\tau}_y, \widetilde{\tau}_i)} \lambda \left[\log\left(1 - \widetilde{\tau}_y\right) + \frac{\alpha\delta\beta}{1 - \alpha\delta} \log\left(\frac{1 - \widetilde{\tau}_y}{1 + \widetilde{\tau}_i}\right)\right] + (1 - \lambda) \log\left(\frac{1 - \widetilde{\tau}_y}{1 + \widetilde{\tau}_i}\right).$$

Time-consistent tax policy for this problem is given by

$$\widetilde{\tau}_i = -\frac{\left(1 - \delta\left(1 - \beta\right)\right)\left(\left(1 - \lambda\right)\left(1 - \alpha\delta\right) + \lambda\alpha\delta\beta\right) - \alpha\delta\beta\left(1 - \alpha\delta + \lambda\alpha\delta\beta\right)}{\left(1 - \delta + \delta\beta - \alpha\delta\beta\right)\left(\left(1 - \lambda\right)\left(1 - \alpha\delta\right) + \lambda\alpha\delta\beta\right)}.$$

Given this, we obtain

$$\frac{1 - \widetilde{\tau}_y}{1 + \widetilde{\tau}_i} = \frac{\left(1 - \delta\left(1 - \beta\right)\right)\left(\left(1 - \lambda\right)\left(1 - \alpha\delta\right) + \lambda\alpha\delta\beta\right)}{\alpha\delta\beta\left(1 - \alpha\delta + \lambda\alpha\delta\beta\right)}.$$

This implies a law of motion

$$G\left(\overline{k}\right) = \frac{\left(\left(1 - \lambda\right)\left(1 - \alpha\delta\right) + \lambda\alpha\delta\beta\right)}{\left(1 - \alpha\delta + \lambda\alpha\delta\beta\right)} A\overline{k}^{\alpha}.$$

It is clear that by varying λ, i.e., the weight on the current self's perceptions, we can obtain a range of tax rates, and therefore savings rates. As special cases, we can obtain the competitive allocation—which places some weight on the future selves' perceptions—and the best constant-tax allocation, which places even more weight on the future. It is also possible to put "too much" emphasis on the future in the sense that the current self may become worse off (but the future selves perhaps better off; they obtain more capital, but are less happy for any given level of capital).