Available online at www.sciencedirect.com
ScienceDirect
Journal of Monetary Economics 55 (2008) 43–47
Journal of MONETARY ECONOMICS
www.elsevier.com/locate/jme

Discussion

Comment on: "Optimal saving distortions with recursive preferences" by Emmanuel Fahri and Iván Werning ☆

Fabrizio Perri*

University of Minnesota, Federal Reserve Bank of Minneapolis, NBER and CEPR, USA

Received 10 December 2007; accepted 10 December 2007

Available online 23 December 2007

1. Introduction

This paper is part of a research agenda that deals with the normative analysis of Bewley/Aiyagari (henceforth BA) economies.[1] The distinguishing feature of these economies is that agents do not have access to a full set of contingent claims and can only use an uncontingent bond to transfer resources across dates. BA economies are nowadays a successful tool for positive analysis, and they are used for understanding a variety of issues such as asset pricing, welfare costs of business cycles or wealth and consumption inequality.

The positive success of BA model economies leads naturally to normative issues and in particular to the question of whether the equilibrium allocations in BA economies can be improved upon. The answer to this question depends crucially on the set of allocations available to a planner/policy maker who wants to improve on the equilibrium.

Three papers consider different sets, and they give a very different answer to this question. Cole and Kocherlakota (2001) assume that agents have unobservable income shocks and also can access a private storage technology. In this case the equilibrium allocation emerging from a BA trading arrangement is also efficient in the sense that there is nothing that a planner (subject to the same informational constraints of private agents) could do to improve on the equilibrium.

Davila et al. (2005) instead consider the case in which the planner can freely distort individual saving (still respecting individual budget constraints and measurability of saving payoffs with respect to the state space) and find that the equilibrium allocation can be improved upon and that the welfare gains of doing so can be very large.

Fahri and Werning (2006) finally consider a world in which the planner can perturb the BA equilibrium allocation but can do so in a restricted fashion, i.e., only leaving equilibrium effort choices unchanged. They show that, in such a world, a planner can improve on the equilibrium allocation of a BA economy and that the

☆The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

*Tel.: +1 612 625 7504; fax: +1 612 624 0209.
E-mail address: fperri@umn.edu

[1]The reference papers for BA economies are Bewley (1986) and Aiyagari (1994).

0304-3932/$ - see front matter
doi:10.1016/j.jmoneco.2007.12.005

size of the welfare improvement that a planner could achieve depends crucially on the technological constraints of transferring resources across dates.

Fahri and Werning (2008), henceforth FW, extends and complements Fahri and Werning (2006). It extends it as it considers a more general environment with recursive and non-separable preferences and characterizes how the optimal perturbation from the BA equilibrium changes as preferences change. It complements it because, by looking at different preferences, it helps understand the logic of how perturbation works in improving efficiency and why a particular perturbation preserves incentive compatibility. In this note I will analyze a special case of the two period setup discussed in Section 2 of FW. The hope is that this example will help the reader better understand the overall message of Fahri and Werning (2006) and FW. I will conclude by discussing some future directions this research agenda could take.

2. A two period BA economy

Consider a small open economy inhabited by a continuum of agents who live for two periods (periods 1 and 2), consume in period 1, consume and work in period 2. In the baseline allocation (the one FW wants to improve upon), all agents have first period consumption $c_1 = 1$, while in the second period each agent consumes $c_2^{\mathrm{h}} = 1 + \varepsilon$ with probability 0.5 and $c_2^{\mathrm{l}} = 1 - \varepsilon$ with probability 0.5, $0 < \varepsilon < 1$. Consumption is financed with output which is produced using effort, but for the purpose of this note we do not need to model explicitly the production process nor the disutility from effort. Utility from consumption is common across agents and is given by

$$u(c_1) + \frac{\beta}{2}(u(c_2^{\mathrm{h}}) + u(c_2^{\mathrm{l}})), \tag{1}$$

where $u(.)$ is a function satisfying $u' > 0, u'' < 0, u''' > 0$ and $\beta > 0$ is the discount factor.[2] Any agent in the economy (including the planner) can transfer resources across dates at the exogenously given rate R. Let us also assume that at the baseline allocation the Euler equation holds with equality, i.e.,

$$u'(1) = \frac{\beta R}{2}(u'(1+\varepsilon) + u'(1-\varepsilon)) \tag{2}$$

so that this allocation is consistent with a BA trading arrangement. It is immediate to see that this allocation is inefficient in the sense that with the same resources currently used for consumption, a higher utility could be achieved or the same utility could be achieved with less resources. The key source of inefficiency is that agents do not equalize consumption across states in period 2. This type of inefficiency is typical in BA economies, due to the lack of state contingent assets.

Before we consider in detail how the Fahri Werning approach would work in this economy, it is helpful to consider what an unconstrained planner can do. Let us focus first on the second period. Suppose that the planner sets consumption in period 2 equal for each agent and equal to

$$c_{\mathrm{CE}} = u^{-1}\left(\tfrac{1}{2}(u(1+\varepsilon) + u(1-\varepsilon))\right) < 1, \tag{3}$$

where the subscript CE stands for certainty equivalent. This would give agents the same utility as in the baseline allocation but would save the planner an amount, in terms of period 1 consumption, $(1/R)(1 - c_{\mathrm{CE}}) > 0$.

This though is not all a planner can do. When individual risk in the second period is eliminated, the marginal value of resources in the second period is lower than the value of resources in the first period. To see this, note that

$$u'(1) = \frac{\beta R}{2}(u'(1+\varepsilon) + u'(1-\varepsilon)) > \beta R u'(1) > \beta R u'(c_{\mathrm{CE}}), \tag{4}$$

[2] Here, for simplicity it is assumed that disutility from effort is separable from utility from consumption. In FW most of the analysis is conducted with non-separable preferences. But the goal of this note is to illustrate some principles that can be applied also in the non-separable case.

where the first equality is the Euler equation of the baseline allocation, the first inequality follows from the strict convexity of marginal utility and the second inequality from the definition of c_{CE}. This implies that the planner would find it optimal to shift consumption from the second to the first period. In particular, it is easy to show that the optimal allocation (i.e., the one that minimizes the resources necessary to provide the same utility to agent as in the baseline allocation) of an unconstrained planner is characterized by two numbers c_1^* and c_2^*, representing consumption in the first and second period, and that $c_1^* > 1 = c_1$ and $c_2^* < c_{\mathrm{CE}}$. In other words, the planner improves efficiency of a BA economy using two key features: *risk reduction* and *consumption front loading*. Risk reduction improves efficiency equalizing consumption, and hence marginal utilities, across states. Front loading is a consequence of risk reduction: the presence of risk raises the marginal value of resources in the future, thus inducing agents in BA economies to save more. Once risk is eliminated, the marginal value of future resources is reduced and it is efficient (both for agents and for the planner) to front load (relative to the BA allocation) consumption.

In FW risk reduction and consumption front loading are also at work but, relative to the previous case, in a very constrained fashion. In particular, FW imposes constraints on planner choices using an incentive compatibility argument. They assume that labor effort of individual agents is not observable by the planner. They also assume that the set of allocations available to the planner is the set in which labor effort is the same as in the baseline economy. Since labor effort is not observable, this assumption imposes that incentive compatibility of the original BA allocation is preserved. This is the case if differences in expected utility that each agent receives at every node are the same as in the baseline allocation. This implies that the only perturbation that a planner can do is to add and subtract, for each date, a constant amount of utility for each state. In the simplified setup analyzed here, this perturbation can be characterized by a single number, namely the amount of additional utility that is subtracted to each agent in period 2. If we denote this number by Δ, then the efficient perturbation, i.e., the perturbation that delivers the same utility as the baseline with the minimum cost solves

$$\min_{\Delta} u^{-1}(u(1) + \beta\Delta) + \frac{1}{2R}(u^{-1}(u(1+\varepsilon) - \Delta) + u^{-1}(u(1-\varepsilon) - \Delta)). \tag{5}$$

It is immediate to show that the optimally perturbed allocation satisfies the so-called inverse Euler equation (Rogerson, 1985)

$$\frac{1}{u'(\tilde{c}_1)} = \frac{1}{2\beta R}\left(\frac{1}{u'(\tilde{c}_2^{\mathrm{h}})} + \frac{1}{u'(\tilde{c}_2^{\mathrm{l}})}\right), \tag{6}$$

where $\tilde{c}_1$, $\tilde{c}_2^{\mathrm{h}}$, $\tilde{c}_2^{\mathrm{l}}$ represent consumption at the optimally perturbed allocation. Evaluating the inverse Euler equation at the baseline allocation and substituting the expression for βR that can be derived from the regular Euler equation, it is easy to show that the right-hand side of the inverse Euler equation (6) is greater than the left-hand side if and only if

$$\tfrac{1}{2}(u'(1+\varepsilon) + u'(1-\varepsilon)) > 1, \tag{7}$$

which is always true if $u''' > 0$. If the right-hand side of (6) is greater than the left-hand side at the baseline, the constrained planner will increase current consumption, just as the unconstrained planner. Why does the constrained planner front load consumption? Consider an agent in a BA economy who wants to borrow to increase its consumption in period 1. Borrowing 1 unit today requires repaying R units in all possible states tomorrow. This is particularly costly in state $1-\varepsilon$ where the marginal value of consumption is high. Consider instead the problem of the planner who wants to provide an extra unit of consumption in period 1. Incentive compatibility requires that the unit must be "repaid" in a state contingent fashion; in particular it will be repaid in such a way that utility loss from repayment is the same across states. Concavity of utility implies that repayment is higher in state $1+\varepsilon$ and lower in state $1-\varepsilon$ so that it will be less costly, relative to the repayment in a BA economy. Since repayment is less costly, a constrained planner will use more borrowing (i.e., front load consumption) than private agents in a BA economy.

Notice though that the constrained planner does not necessarily achieve risk reduction with its optimal perturbation. In this regard, one case which is particularly instructive, as it also relates to many results in FW, is the case in which $u(.) = \log(.)$. In this case the minimization problem (5) reduces to

$$\min_{\Delta} e^{\beta\Delta} + \frac{1}{2R}((1+\varepsilon)e^{-\Delta} + (1-\varepsilon)e^{-\Delta}) = \min_{\Delta} e^{\beta\Delta} + \frac{1}{R}e^{-\Delta}$$

which shows how the optimal perturbation in the log case amounts to (just as in the case of recursive utility) a proportional shift in consumption in each state. Solving for the optimal Δ yields $\Delta = -\log(\beta R)/(1+\beta) = -\log(1-\varepsilon^2)/(1+\beta) > 0$. Note that, as in the case of an unconstrained efficient allocation, consumption is front loaded relative to the baseline, but, unlike in the unconstrained efficient case, there is no risk reduction: indeed it is immediate to see that the ratio between marginal utility in the low state and in the high state is the same in the baseline and in the constrained efficient allocation. In the slightly more general case in which $u(c) = c^{1-\sigma}/(1-\sigma)$, one can show that the ratio of marginal utilities across states in the optimally perturbed allocation is given by

$$\frac{u'(\tilde{c}_2^{\mathrm{l}})}{u'(\tilde{c}_2^{\mathrm{h}})} = \left(\frac{(1-\varepsilon)^{1-\sigma} - \Delta^*(1-\sigma)}{(1+\varepsilon)^{1-\sigma} - \Delta^*(1-\sigma)}\right)^{-\sigma/(1-\sigma)}, \tag{8}$$

where $\Delta^* > 0$ is the optimal perturbation to the baseline economy. Note that

$$\frac{u'(\tilde{c}_2^{\mathrm{l}})}{u'(\tilde{c}_2^{\mathrm{h}})} \lesseqqgtr \frac{u'(1-\varepsilon)}{u'(1+\varepsilon)} = \left(\frac{1+\varepsilon}{1-\varepsilon}\right)^{\sigma} \quad \text{if and only if } \sigma \gtreqqless 1, \tag{9}$$

i.e., whether the optimal perturbation reduces or increases the marginal utility ratio, i.e., risk, relative to the baseline depends on whether the relative risk aversion coefficient σ is bigger or smaller than 1. The fact that in general risk reduction can be moderate (or absent) is another explanation of why at the optimally perturbed allocation free saving cannot be allowed. The perturbed allocation front loads consumption relative to the BA allocation but it does not necessarily reduce risk, so private agents would have an incentive to undo the front loading in order to cope with the risk.

To conclude, the simple example discussed here highlights two key features, *risk reduction* and *consumption front loading*, which are useful to understand how and how much the optimal perturbation proposed in FW can improve welfare in a BA economy. In particular it shows that welfare gains in FW are mostly obtained through consumption front loading, and this clarifies why the size of the welfare gain depends crucially on the ability of the planner of transferring resources across dates.

3. Future research directions

As I mentioned earlier, I view the FW paper as an essentially normative one, and as such I see it very fitting for the Carnegie Rochester Public Policy series. Although the paper does contain some critical and general theoretical advances, such as the characterization of a generalized inverse Euler equation for a class of recursive and non-separable preferences, I view the general intent of the paper as very practical, i.e., how can we improve risk allocations in economies with private information frictions and where risk sharing is severely limited by the lack of contingent assets.

In particular, in such a setting FW focuses on a particular class of perturbations that improve over an initial allocationm, but it does not attempt to characterize the global constrained optimum. This is obviously a limitation of the approach, but their choice has a very good justification, which again is of practical nature. Characterizing the FW optimal perturbation is much easier than computing the global constrained optimum. Fahri and Werning (2006) shows how this type of optimal perturbation can be easily computed and evaluated for a very general class of stochastic processes driving the underlying risk. FW instead shows how the optimal perturbation can be easily computed and evaluated for a general class of preferences. In many of the environments considered in FW, we do not know how to completely characterize the global constrained optimum, so the FW approach is a real contribution; it teaches us how to make a step in the right direction, and it is in some cases the biggest step we know how to take.

In terms of where to go next, two issues come to mind. One is dependence of the optimally perturbed allocation on preferences and the second is its market implementation.

Regarding the first issue, one of the main messages of FW is that preferences are hugely important both for the characterization of the optimally perturbed allocation and for its welfare impact. They present cases ranging from the one discussed in Proposition 1, in which there is no feasible perturbation that improves over the baseline BA economy, to some (plausible) parameterization of recursive preferences in which welfare gain from the perturbation exceeds 10% of lifetime consumption. Also in the case of recursive utility the optimal perturbation takes a particularly simple form; in other cases it can be a fairly complicated function of observables and of preference parameters.

Although this is a commom issue, I think that, if one wants to take the perturbation argument more seriously, one needs to bring more data to bear in order to pin down the structure and the parameters of preferences. And since the key features of preferences that determine the optimal perturbation are attitudes of agents toward risk and intertemporal substitution, it seems to me that asset prices could be just the right data to look at.

Regarding implementation, FW shows that, in order to implement optimal perturbation, restrictions to savings (such as a tax on savings) are necessary, but they are obviously not sufficient. As it is clear from the example discussed in Section 2, a tax on saving can induce agents to front load consumption but cannot get the "right" period 2 consumption. Obviously one could think of many, possibly complicated, systems of tax and transfers that could implement the optimal perturbation (for related work on this see, for example, Albanesi and Sleet, 2006). But I was intrigued by the particular simple form, i.e., proportional shifts to consumption, that the perturbation takes in the case of log utility and in the case of recursive preferences. That form suggests that the optimally perturbed allocation might be decentralized by introducing the right kind of financial borrowing instrument, i.e. an instrument whose repayment is contingent on consumption of an agent. Obviously such decentralization would require consumption to be observable and would still require restrictions on savings (otherwise agents would just go short on the new asset and long on the traditional uncontingent bond). But, if it worked, it could be a way of extracting from the markets information on preference parameters, which are necessary to design the optimal perturbation but on which we know relatively little about.

References

Aiyagari, S.R., 1994. Uninsured idiosyncratic risk and aggregate saving. Quarterly Journal of Economics 109, 659–684.

Albanesi, S., Sleet, C., 2006. Dynamic optimal taxation with private information. Review of Economic Studies 73, 1–30.

Bewley, T., 1986. Stationary monetary equilibrium with a continuum of independently fluctuating consumers. In: Hildenbrand, W., Mas-Colell, A. (Eds.), Contribution to Mathematical Economics in Honor of Gerard Debreu. North-Holland, Amsterdam.

Cole, H., Kocherlakota, N., 2001. Efficient allocations with hidden income and hidden storage. Review of Economic Studies 68, 523–542.

Davila, J., Hong, J., Krusell, P., Ríos-Rull, V., 2005. Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks, manuscript.

Fahri, E., Werning, I., 2006. Capital taxation: quantitative explorations of the inverse Euler equation, manuscript.

Fahri, E., Werning, I., 2008. Optimal savings distortions with recursive preferences. Journal of Monetary Economics, in press. doi:10.1016/j.jmoneco.2007.12.002.

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