Optima in heterogeneous-agent monetary economies

Neil

March 20, 2014

What do we know about good policy or optima in monetary economies? Not much.[1] By *not much* I mean that it is hard to reach conclusions about optima in general monetary economies because optimal policy—even its qualitative aspects—depends on details of the economy that we are unlikely to know much about. This happens because the features that give money a role have at least two consequences: they make it desirable to enhance the return on money (in accord with the so-called Friedman rule) and they inhibit risk-sharing. Moreover, within the class of feasible policies, policies consistent with money having a role, enhancing the return on money and improving risk-sharing are competing goals. Therefore, and not surprisingly, the best way to balance them depends on details of the economy.

1 General monetary economies

A monetary economy is one in which currency-like objects are essential in the sense that their presence allows for the implementation of some good allocations that would otherwise not be implementable. What sort of environments make money essential in this sense? For more than 1,500 years, scholars have pointed to environments with double-coincidence problems. We now know that a double-coincidence problem is only one necessary condition for essentiality of money. (Robinson Crusoe and Friday can encounter double-coincidence problems, but few would say that they need money to overcome those problems.) Many of us now think that a second necessary condition has been identified in the work of Ostroy, Townsend, and Kocherlakota. This necessary condition, which rests on the view that money is nothing but evidence of previous actions, is imperfect monitoring—some lack of common knowledge about previous actions. Therefore, we seem to have two necessary conditions for essentiality—a double-coincidence problem and imperfect monitoring. However, as indicated by recent work on folk theorems, they are far from sufficient. Nevertheless,

[1] If you regard that as an unpleasant message, then too bad. I have a vague recollection of Milton Friedman saying that the problem is not what people don't know; the problem is what they think they know.

those conditions can provide a foundation for constructing examples of monetary economies (see, for example, Araujo 2004). I should also mention that my discussion presumes that the background environment is one in which people in the model cannot commit to future actions—the standard repeated- and dynamic-game assumption.

Now we can turn to what *general* means. A double-coincidence problem requires heterogeneity—as does any model in which there is a motivation for trade. Therefore, let the state of the economy be a joint distribution across the population over types and asset holdings, denoted λ. And, for now, consider the following schematic law of motion in a discrete-time world with a fixed stock of outside (fiat) money and nothing that we would normally identify with policy,

$$\lambda_{t+1} = H(\lambda_t, shocks_t, trades_t). \tag{1}$$

Although I am being vague about the environment, let's treat the order of the arguments of H as depicting the sequence of actions so that the shocks, idiosyncratic and/or aggregate, are realized before the trades occur. The shocks could be about tastes (see SLP section 13.5), endowments (see Levine), or meetings (Kiyotaki-Wright). General in this context means that there are shocks and that λ_t affects the trades at t and that those trades affect λ_{t+1}—the usual situation in heterogeneous-agent models. Put somewhat differently, the trades at t have two roles: they affect date-t period payoffs of the agents and they affect the state at the next date. As a consequence, *good* trades must represent a compromise between what would be best for date-t period payoffs and what would be best for the date-$t+1$ state.

One good way to highlight this meaning of general is to review a list of well-known models that are *not* general in that sense. One is an OLG model of two-period-lived people with one good per date. In that model the distribution of money holdings among the old at a date is not a state variable because they are at a corner; they offer all their money at any price of money. Another is Shi and Trejo-Wright, similar random matching models with money holdings limited to the set $\{0, 1\}$. In those models, the distribution λ is given by the fraction who hold money at the start of a date, a fraction which is determined entirely by the exogenous stock of money. Still another such model is the large-family model in which asset holdings of the members are merged and averaged out across the members (see Shi 1998 for applications to money). And, finally, there are versions of Lagos-Wright (2005) in which periods of centralized trade with quasi-linear preferences or other special assumptions eliminate any inheritied heterogeneity of asset holdings. While these are evidently knife-edge specifications, that might be okay if their consequences for good policies did not depend on that feature. I will try to convince you that that is not the case.

2 Implementability, policy, and the main trade-off

To begin, let's generalize the law of motion to include policy, so that

$$\lambda_{t+1} = H(\lambda_t, shocks_t, trades_t, policy_t). \tag{2}$$

As above, the ordering of arguments depicts the sequence of actions. In accord with that sequencing, I distinguish between the *trade-stage* and the subsequent *policy-stage* at each date. The policy stage has neither production nor consumption; it only has taxes and transfers. Moreover, we want those policies to be consistent with the environment and, in particular, with the assumptions that make money essential.

For the moment, everything I say is about weak implementability. (In order to achieve strong implementability (uniqueness), it is necessary to eliminate non-monetary equilibria and I have nothing new to say about how that might be accomplished.) In order to proceed and to discuss examples, which I will do later, it is helpful to delineate for each stage combinations of the information structure and the notion of defection as depicted in the following table,

Specifications for each stage		
	symmetric	asymmetric
individual defection	I	II
group defection	III	IV

.

The columns here do not refer to histories. Instead, they refer to tastes, endowments, and portfolios. For example, in a model in which people trade in pairs, there is asymmetric information at the trade-stage if the seller has private information about his endowment or his disutility of production. And at either stage, there is asymmetric information if agents can hide assets. As regards rows, under individual defection, the agents get to choose only from the set $\{yes, no\}$ in response to a planner suggestion; under group defection they are free to *cooperatively* defect to any current action that is feasible for the group. Cooperatively is in italics because it is a straightforward notion under symmetric information, but is far from straightforward under asymmetric information.

I favor specifications which permit group defection at the trade stage. Why, after all, should a group be limited to choosing only from $\{yes, no\}$? Such a strategy set is so restrictive that it trivializes the requirement that money be essential and puzzles like coexistence of money and higher-return assets. If the planner suggests cash-in-advance trade and agents can only respond from $\{yes, no\}$, then cash-in-advance trade is almost always implementable. However, group defection and asymmetric information at either stage call for the adoption of some notion of the core under private information, or, perhaps, some notion of renegotiation-proofness. So far as I know, existing work on money has not dealt with such specifications.

In models with a large number of agents, symmetric information, and centralized trade (everyone together), group defection implies price-taking trade.

In one sense, that is convenient because then the planner has nothing to choose at the trade stage. If trade occurs in small groups under symmetric information, then there is a nondegenrate core and the planner has some options at the trade stage. At the policy stage, group defection limits nonneutral transfers to lump-sum transfers. And, if people can hide assets at that stage, then transfers have to be weakly increasing in asset holdings.

Before turning to illustrative environments and examples, I can say a bit about circumstances under which there is a role for policy and the circumstances under which optima depend on details of the economy. First, let's define a welfare criterion and a notion of first-best. My argument that optima depend on details is only strengthened if I use the simplest welfare criterion; namely, ex ante representative-agent welfare. Therefore, I assume a fixed population of infinitely-lived people and treat them as identical prior to the realizations determined by λ_0, the initial state. (We may want to treat λ_0 as arbitrary and given, or we may want to treat the initial asset-holding distribution as something chosen by the planner.) I also assume that the planner can commit to future actions. By first-best, I mean the optimum in the given environment, but with imperfect monitoring replaced by perfect monitoring—but subject to no-commitment by agents. Now, assume that the environment is a general monetary economy. For such an economy, the conjecture is that the best outcome has a role for policy and falls short of the first-best. Therefore, it involves a trade-off between raising the return on money and risk-sharing. Finally, how that trade-off is best resolved depends on all the details of the environment.

It is worth noting that these conclusions are not true in the knife-edge models noted above. In Shi and Trejo-Wright, the models with money holdings limited to the set $\{0, 1\}$, there is no scope for policy. In the large-family model, risk-sharing is accomplished within the family; while in Lagos-Wright (2005), the quasi-linear preferences or other special assumptions eliminate any role for sharing. Thus, in those models, policy does not confront the trade-off between enhancing the return on money and improving risk-sharing. Hence, in those models the first-best can sometimes be attained (see HKW).

3 Pure quid-pro-quo economies

A pure quid-pro-quo economy usually means one in which only spot trade is possible. Such economies have been studied for a long time (examples are Lucas, Bewley, and Kiyotaki and Wright). I will mean by a pure quid-pro-quo economy one with no monitoring so that each person's previous actions are private to that person. Such complete absence of monitoring rules out credit. When combined with group defection at the trade stage and at least individual defection at the transfer stage, it also eliminates taxation—except via inflation.

In Wallace (QJE), I described a conjecture about optimal policy in such economies under specification II for the transfer stage (individual defection and hiding of money) under the assumption that the planner was restricted to a single uniform asset labeled money.

In order to fix ideas I focused on the following two-parameter class of transfer functions at each date: a function τ_t that maps the end-of-trade money holdings of a person into a transfer; namely

$$\tau_t(x) = \max\{0, a_t + b_t x\} \tag{3}$$

where $b_t \geq 0$. When $a_t = 0$, this is equivalent to no transfer (it is just a change in units). When $a_t > 0$, it is equivalent to lump-sum transfers, a scheme with $b_t = 0$. When $a_t < 0$, the transfer is zero for $x \leq -(a_t/b_t)$, and proportional to $x - (a_t/b_t)$ for $x > -(a_t/b_t)$ (the first use of such a scheme is in Andolfatto). The conjecture was that schemes with $a_t \neq 0$ for most dates are optimal for general monetary economies. The very simple idea was that schemes with $a_t \equiv 0$ are equivalent to no intervention and that $a_t = 0$ is interior in the set of feasible policies. Therefore, because no-intervention is not first-best in a general monetary economy, some intervention will help. Moreover, $a_t > 0$ improves risk-sharing, but lowers the return on money; and $a_t < 0$ raises the return on money over some range of holdings while worsening risk-sharing. Which is better depends on all the details of the economy. (Andolafatto works within the general framework of LW in which there is no risk-sharing motive. Therefore, he is able to show a welfare improvement from $a_t < 0$.)

Here, I want to consider a richer class of policies, one involving more than one kind of asset. In particular, I want to discuss a possible role for a second asset, which we should think of as a bond, an asset with a higher rate of return than money. There are two questions relating to policies of that sort. One involves feasibility; in particular, how do we get people to hold both assets? The second is about optimality. As noted above, feasibility is easy if only individual defection is imposed at the trade stage. Under that assumption, Kocherlakota showed that money and bonds are optimal in some economies.

Here, I want to summarize work by Hoonsik Yang that deals with both questions, while allowing for group defection at the trade stage. Yang starts from a specification in Deviatov 2006, a version of the random matching model of Shi and Trejos-Wright with one asset and individual money holdings in $\{0, 1, 2\}$. This is the smallest set of asset holdings which turns that model into a general monetary economy. Deviatov assumes specification III at the trade stage (symmetric information and group defection) and specification II at the transfer stage (allows the hiding of assets, but only individual defection). He studies steady states and looks for the best transfer policy, while modelling inflation as random disintegration of money. He finds that such schemes are optimal in some specifications with high discounting and low risk aversion. Yang takes one of those specifications and shows that having a zero real return asset and money that suffers from inflation does better. Yang assumes that the trade stage is preceded there by a portfolio stage at which individuals freely choose between the two assets at a one-for-one rate. That is followed by a trade stage and then by the policy stage at which transfers and random disintegration of money occur. (The welfare improvement Yang finds is not the result of a weakening of the upper bound because Yang does not permit a seller with two assets to trade.)

The qualitative result that there is improvement is not surprising because, as Yang shows, the Deviatov one-asset specification is a feasible outcome for the planner under his specification. So, having multiple assets is another instance of enriching the set of policy instruments.

4 Monetary economies with credit and taxation

quid pro quo economies are extreme

indeed, for many years we have heard that the world is headed toward being cashless, whatever that means

we want to give it meaning and also bridge the gap between quid pro quo economies and cashless economies

one of our necessary conditions for essentiality gives us a nice candidate for cashless economies: perfect monitoring; that and our specification of quid pro quo economies, economies with no monitoring, give us an obvious way to bridge the gap: imperfect monitoring.

models with two separate parts: a monitored part and a nonmonitored part. (Rachel Kranton, Duke ???? Antinolfi, azariadis, and bullard. (Size of each part does not really matter: the nonmonitored part serves as a threat

Models in which monitored status is not part of the state of the economy: it is a feature of meetings, rather than of people

why do that? simplify the model

aside from that, does it seem like an attractive specification?

Maybe, but it is not the only way

claim: tractability aside, would we do it that way?

So let's not. CW: the sizes matter and the two parts interact in interesting ways

What do we get. This topic has barely been explored. But we have a few examples.

Crucial feature: how do we punish m people?

$\{0,1\}$ and yet an endogenous state variable: the split of money between m and n people. Much more to be done.

see recent QR and RED paper

5 Concluding remarks

Rather than a paper with conclusions, the above lays out a research agenda. ....