Journal of Monetary Economics 21 (1988) 411–450. North-Holland

INFORMATION CONSTRAINED INSURANCE
The Revelation Principle Extended*

Robert M. TOWNSEND

University of Chicago, Chicago, IL 60637, USA

Received September 1987, final version received December 1987

Private information induced incentive constraints can cause allocations to diverge from full information optimal allocations, often in the direction of limited insurance, but can bring unanticipated anomalies. Related, plausible contract exclusion restrictions can be damaging in welfare terms. More generally, private information optimal allocations can be sensitive to the information structure, suggesting that the latter be specified with an eye toward realism as well as tractability. In an effort to make tractability less of a constraint the paper shows how two apparently difficult information structures – a costly state verification environment and a multi-period multilateral private information environment – can be handled theoretically, by revelation principle methods. The paper also shows how solutions can be generated numerically.

1. Introduction

Authors of real business cycle models have subjected themselves to a constraining but productive discipline. The discipline is that observations are required to be explained by models specified at the level of the primitives of preferences, endowments, and technology. Thus in practice, as in the seminal work of Kydland and Prescott (1982), candidate allocations are required to be Pareto optimal and so to solve the problem of maximizing a weighted sum of the utilities of the agents of the model subject to constraints implied by finite resources.

The fit of the Kydland–Prescott model to U.S. data is not close on some dimensions, and there are phenomena which that model was not designed to explain. This has led Kydland and Prescott (1982) and others to search over

*This paper was presented at the Lisbon conference on Real Business Cycles, the NBER working group on Financial Structure and Economic Activity and workshops at the University of Pennsylvania and the University of Chicago. Helpful comments from the participants of these sessions, especially Robert King and Richard Rogerson, and from a referee, and conversations with Lars Hansen and Roger Myerson are gratefully acknowledged. The idea of computing solutions as solutions to linear programs emerged in earlier work with Edward C. Prescott. The programming talent and effort of Christopher Phelan and Kevin Geraghty are reflected herein. Financial support from the National Science Foundation made this paper possible.

alternative specifications of preferences and technology. Altug (1985) experiments with multiple capital goods, finding that separable preferences may provide a reasonably good fit to the data, unlike Kydland and Prescott. Hansen (1985) introduces non-convexities in labor supply, to deliver higher variability of hours, and Hansen and Sargent (this issue) use non-convexities to deliver a distinction between straight time and overtime employment. It is this interaction between observations and the primitives of the model which is productive, indeed exciting.

It is not obvious where the next round of iterations will take us, but one can take some guesses. The models of Hansen and of Hansen and Sargent treat identical agents differently in *ex post* labor supplies but identically in consumptions, something which follows naturally at an optimum from separable preferences and full insurance. Because full consumption insurance seems inconsistent with actual observations, an obvious next step is to limit insurance in some way. This would seem to have implications for both the dispersion of consumption in a population and for aggregate time series. Indeed, following a long tradition, Bernanke and Gertler (1986), Greenwald and Stiglitz (1986) and Smith (1985) each limit insurance and asset markets in their efforts at modeling business cycle phenomena.

One way to limit insurance is to preclude it exogenously. That is, the modeler can specify what contracts are feasible and what contracts are not. A second way to limit insurance is to find a reason why insurance markets are limited and to deliver insurance arrangements as part of the solution to the optimum problem. Implicitly or explicitly, authors seem to prefer the second approach, for pragmatic and scientific reasons. On the pragmatic side we need some guidance in deciding what contracts or markets to limit or shut down. On the scientific side it is more satisfying to explain the observations of limited insurance from the primitives of preferences and technology than to mimic the observations.

Private information is the prime candidate to use to limit insurance arrangements endogenously. In fact Bernanke and Gertler (1986), Greenwald and Stiglitz (1986), and Smith (1985) all use private information to motivate or derive the insurance contracts and asset markets of their business cycle models.

The purpose of this paper is to aid in these efforts by expositing and extending a method which places private information constraints on resource allocations on the same footing with standard resource constraints. With this accomplished one can solve for optimal arrangements in the obvious way, by maximizing weighted sums of the utilities of the agents of the model subject to these private information and resource constraints. The power of this approach is evident in various ways. First, one does not have to tell private information stories to deliver contract restrictions before laying out the model. Rather, contracts emerge from a global optimum problem. Second, and related, the joint restrictions of private information on contracts and on other phenomena are incorporated. Third, the private information constraints are as weak as possible, so only preferences, endowments, technology, and the information structure dictate the solution. Fourth, and related, these private information optimal solutions, though consistent with limited insurance, sometimes display anomalies, such as enhanced dispersion, the volatilities of artificial lotteries, deliberate scrambling of information, and increased information acquisition in the face of increased cost. Some, if not all, of these anomalies are *ex post* rationalizable, in the sense that once discovered, their existence can be traced back to information contrained insurance motives. But *a priori* logic or 'out of model reasoning' may have missed their rationale. Fifth, and related, plausible but exogenous contract exclusion restrictions can be shown in various contexts to be too strong. Exogenous restrictions are sometimes as damning in welfare terms as private information itself, ironically because the anomolous possibilities just mentioned are ignored. Exogenous restrictions can thus interfere with attempts to explain observations from supposed primitives and to compare alternative policy regimes. Sixth, to the extent that exogenous exclusion restrictions seem more consistent observations, we are led to puzzle over why private information was not enough and to search for a source of further restrictions. In this way the productive interaction between observations and the primitives of economic models can continue.

The private information methods of this paper are the so-called revelation principle methods developed in Myerson (1979) and in Harris and Townsend (1981) and extended here to two apparently difficult private information environments. The first, in section 2, is a costly state verification type environment of Townsend (1979), an interesting but challenging environment in that information which can be made available at a cost is neither entirely private nor entirely public. In fact, Townsend (1979) did not use revelation principle methods in deriving debt contracts as optimal and displayed stochastic audits as a Pareto dominating possibility rather than as an integral part of the analysis. More recent related literature includes Baiman and Demski (1980), Evans (1980), Gale and Hellwig (1985), Mookherjee and Png (1987), and Reinganum and Wilde (1985). But much of this literature stops short of a full, private information optimum, imposing instead various restrictions on contracts, as in Townsend (1979), in an effort to characterize solutions.

The second environment to which revelation principle methods are extended is a multiperiod environment with private information on the part of multiple agents, in section 4. It is sometimes said that the revelation principle does not apply in such contexts, because truth telling is inconsistent with incentives. Afterall, in certain economic environments, agents have a direct need to conceal their information. In such cases one might expect it to be impossible to induce truthtelling. But it can be established that optimal use can be made of concealment devices which scramble information, as in the recent work of Myerson (1986). Here, in fact, an explicit and detailed proof of the revelation principle is offered.

One purpose of the paper, then, is to show how revelation principle arguments can be employed in situations where they might have been thought to be inapplicable. It is hoped that the extensions of the revelation principle, in sections 2 and 4, confirm the power of the approach.

One *caveat* is in order, however. The revelation principle has its limits. In particular, nothing here precludes situations of multiple equilibria, situations in which a direct revelation mechanism has additional, non-truthtelling equilibria. Examples of such situations are contained in Bhattacharya (1984), Demski and Sappington (1984), Postlewaite and Schmeidler (1986), and Repullo (1983), among others. Of course truthtelling *per se* is not the desideratum. The issue is one of implementation, ensuring that the solution to a programming problem with incentive constraints can be achieved as a unique equilibrium outcome for some game. Indeed, Palfrey and Srivastava (1987) have shown how a suitably enriched game or mechanism would have as its unique outcome the desired 'truthtelling' equilibrium, at least up to a limited notion of refinement, for a wide variety of environments. But a further exploration of this idea takes us beyond the scope of the present paper, as would a discussion of limited commitment.

The revelation principle method is also extended here numerically in sections 3 and 5. That is, maximum problems for the determination of private information Pareto optimal allocations are converted to standard linear programs, despite sequential auditing in section 3 or sequential incentive constraints in section 5, and these programs are solved numerically for various parameterized environments. It is from these numerical solutions that the apparent anomalies of private information optimal allocations are discovered and the welfare costs of private information and additional exclusion restrictions are measured. The limitations of existing numerical methods are also discussed.

It might be complained that the two private information environments described above and considered in the paper are special and that implications for time series dynamics are unclear. There is some merit to this complaint: little is known generally. On the other hand, relatively static private information environments are often used as a building blocks in more elaborate dynamic setups. A prime example is the use of the costly state verification environment by Bernanke and Gertler (1986) in their work on financial fluctuations. So it seems that relatively static private information environments should be understood well. In addition, there are growing indications that fully integrated private information dynamic environments can be analyzed. A prime example is the work of Edward Green (1984) who uses a private information pure exchange environment to address the permanent income hypothesis and the adequacy of debt markets. The point is that private information allows some, but not all, interhousehold reallocations, even if aggregates remain unaffected. And, it is relatively easy to generate examples with aggregate effects as well. Indeed, it is shown in section 6 of this paper for a third, two-period, capital accumulation environment that 1) insurance limited by private information can mitigate the welfare gain from intertemporal storage, causing aggregate storage to decrease, and 2) in the context of private information, storage can help mitigate the damning effect of the incentive constraints, causing aggregate storage to increase. Thus aggregate effects emerge rather quickly in numerical examples, but the direction of the effect is dependent on the specific parameters of the environment. The point, again, is not that private information is a wild card producing any outcome whatever, but that the logic of optimal, but information constrained insurance should not be taken for granted.

2. Pareto optimal audits

Imagine an economy with two agents (a and b), two dates, a planning and a consumption period, and a single underlying consumption good. Agent a has preferences over consumption bundle $c^a \geq 0$, as represented by a utility function $U^a(c^a, \theta)$. Here θ is a shock to preferences, or better put, to a household production function, with $U^a(c^a, \theta)$ as an indirect utility function over 'market' goods. Shock θ is observed by agent a alone at the beginning of the consumption period and takes on one of a finite number of values in some set Θ. From the point of view of the planning period, shock θ occurs with probability $p(\theta)$, and this is common knowledge. Agent a has an endowment of e^a units of the consumption good. Preferences of agent b are described by a utility function $U^b(c^b)$ over consumption bundles $c^b \geq 0$. Agent b has an endowment e^b. Consumption vectors $c = (c^a, c^b)$ are bounded by the social endowment $e = e^a + e^b$, and further restricted to a finite number of values, as if consumption took place in discrete units. Finally, if for some reason or other an audit takes place in the consumption period, then K units of the consumption good disappear and parameter value θ is made known to agent b.

The environment here is close to Townsend (1979) in which agent a suffers from a random and privately observed endowment, and agent b in the planning period is in a position to offer insurance to agent a, to smooth his consumption. For the model here the uncertainty and private information concern shocks to preferences, but the insurance motive is the same, with agent a benefiting from a receipt of the consumption good when he is 'urgent' to consume, with high marginal utility, and willing to surrender consumption when he is 'patient'. Shocks to preferences facilitate revelation principle arguments, as one avoids the possibility of pretransfer quantity displays, as in Townsend (1987). Some comments on the original costly state verification environment are given at the end of section 3.

Revelation principle arguments will now be extended to the costly state verification environment of this section. One does this by starting with a rather arbitrary, general game and by reducing it to a simpler, direct revelation game. *Readers who want to see the result, the programming problem for the determination of information constrained optima, can turn directly to program 2.1 or 2.2, with the requisite notation defined therein.*

A general game or resource allocation mechanism for this economy consists of two objects. The first is a message space for agent a, and the second is an outcome function or allocation rule. These are now discussed in turn.

A message space for agent a is a set M. The idea is that after a realization of shock θ agent a may send a message $m \in M$ to agent b, or, better put, to some central computer. For simplicity, the set M may be taken to be a subset of some finite dimensional Euclidean space.

An allocation rule is a set of functions specifying the probability of a completely revealing but costly audit of the shock θ and the probabilities of consumption conditional on whether or not there is an audit and, if there is an audit, conditional also on the revealed value of θ. More formally, let d be a dummy variable taking on values one or zero, depending on whether there is or is not an audit. If there is no audit, so that $d=0$, then the space of possible consumptions is $C^{d=0} \equiv \{(c^a, c^b): c^a \geq 0,\ c^b \geq 0,\ c^a + c^b \leq e\}$, where again this contains a finite number of elements, specified in advance. If there is an audit, so that $d=1$, then K units of the consumption good disappear from the model, and the space of possible consumptions is $C^{d=1} = \{(c^a, c^b): c^a \geq 0,\ c^b \geq 0,\ c^a + c^b \leq e - K\}$. The probability of an audit conditional on a message m is denoted $\Delta^{d=1}(m)$, that is, $\Delta^{d=1}(m)$ is the probability that $d=1$ given message m, and $\Delta^{d=0}(m)$ is the probability that $d=0$ given message m, where of course $\Delta^{d=0}(m) + \Delta^{d=1}(m) = 1$. The probability of consumption bundle c conditional on there not being an audit when the message was m is denoted $\Delta(c|m, d=0)$. Thus $\Delta(m, d=0)$ is a probability measure on the fixed, prespecified space $C^{d=0}$. The probability of consumption bundle c conditioned on there being an audit and the audit-revealed value of the preference shock is θ when the message is m is denoted $\Delta(c|m, \theta, d=1)$. Thus $\Delta(m, \theta, d=1)$ is a probability measure on the space $C^{d=1}$.

Given a game, that is, given a specification of message space M, the probability of audits $\Delta^{d=1}(m)$, $m \in M$, and the probability of consumptions $\Delta(c|m, d=0)$ and $\Delta(c|m, \theta, d=1)$ and given a value of θ, agent a is to choose a message $m \in M$ so as to maximize expected utility. That is, agent a solves a maximization problem: given $\theta \in \Theta$, maximize by choice of $m \in M$

the objective function

$$\Delta^{d=1}(m) \sum_{c\in C^{d=1}} U^a(c^a,\theta)\Delta(c|m,\theta,d=1)$$
$$+\Delta^{d=0}(m) \sum_{c\in C^{d=0}} U^a(c^a,\theta)\Delta(c|m,d=0).$$

For simplicity of notation we shall take the maximizer to this problem, denoted $m^*(\theta)$, to be deterministic though in fact it is easy to accommodate random strategies. Thus, if there is more than one maximizer, we shall suppose some selection rule is in effect.

By virtue of this maximization, we know that $m^*(\theta)\in M$ is the selected message when the preference shock is θ, that $m^*(\tilde{\theta})\in M$ is the selected message when the preference shock is $\tilde{\theta}$, and that therefore $m^*(\theta)$ is weakly preferred to the feasible message $m^*(\tilde{\theta})$ when the preference shock actually is θ. That is, for all $\theta, \tilde{\theta}\in\Theta$,

$$\Delta^{d=1}[m^*(\theta)]\sum_c U^a(c^a,\theta)\Delta[c|m^*(\theta),\theta,d=1]$$
$$+\Delta^{d=0}[m^*(\theta)]\sum_c U^a(c^a,\theta)\Delta[c|m^*(\theta),d=0]$$
$$\geq\Delta^{d=1}[m^*(\tilde{\theta})]\sum_c U^a(c^a,\theta)\Delta[c|m^*(\tilde{\theta}),\theta,d=1]$$
$$+\Delta^{d=0}[m^*(\tilde{\theta})]\sum_c U^a(c^a,\theta)\Delta[c|m^*(\tilde{\theta}),d=0], \qquad (1)$$

where here and below the summation is over either $C^{d=0}$ or $C^{d=1}$ and is understood from the context. Now, as a matter of notation, let

$$\pi^{d=1}(\theta)=\Delta^{d=1}[m^*(\theta)],$$
$$\pi^{d=0}(\theta)=\Delta^{d=0}[m^*(\theta)],$$
$$\pi[c|\theta,d=0]=\Delta[c|m^*(\theta),d=0],$$
$$\pi[c|\tilde{\theta},\theta,d=1]=\Delta[c|m^*(\tilde{\theta}),\theta,d=1],$$

and consider a *new* game with message space Θ and allocation rules π. In the new game $\pi^{d=1}(\tilde{\theta})$ is the probability of an audit given an announced preference shock $\tilde{\theta}$; $\pi[c|\tilde{\theta},d=0]$ is the probability of consumption bundle c given that the announced preference shock value is $\tilde{\theta}$ and there is no audit; and $\pi[c|\tilde{\theta},\theta,d=1]$ is the probability of consumption bundle c given that the announced preference shock value is $\tilde{\theta}$, there is an audit, and the actual preference shock value is θ.

It can now be established that truthtelling is a maximizing strategy under this new game and that the equilibrium allocation of the old game is achieved. For suppose the actual parameter draw is θ and consider inequality (1) under the new notation: for all $\theta,\tilde{\theta}\in\Theta$,

$$\begin{aligned}&\pi^{d=1}(\theta)\sum_c U^a(c^a,\theta)\pi[c|\theta,\theta,d=1]\\&\quad+\pi^{d=0}(\theta)\sum_c U^a(c^a,\theta)\pi[c|\theta,d=0]\\&\geq\pi^{d=1}(\tilde{\theta})\sum_c U^a(c^a,\theta)\pi[c|\tilde{\theta},\theta,d=1]\\&\quad+\pi^{d=0}(\tilde{\theta})\sum_c U^a(c^a,\theta)\pi[c|\tilde{\theta},d=0].\end{aligned}\tag{2}$$

Inequality (2) states simply that given that the actual parameter draw is θ agent a weakly prefers to announce the value θ rather than a counterfactual value $\tilde{\theta}$. Thus truthtelling is maximizing, and one can refer to announced preference shock values and actual preference shock values synonymously. By the construction of π, then, the original equilibrium allocation of resources is achieved.

We have established thus far that any mechanism or game in the class under consideration can be reduced without loss of generality to a so-called *direct revelation mechanism* satisfying inequalities (2). Moreover, any direct revelation mechanism satisfying inequalities (2) is a *bonafide* mechanism. Thus the search for an *ex ante* Pareto optimal mechanism in the class under consideration, with arbitrary message spaces and outcome functions, can be done in the space of direct revelation mechanisms satisfying inequalities (2). That is, for Pareto weights λ^a and λ^b with $\lambda^a+\lambda^b=1$, a Pareto optimal mechanism and a Pareto optimal allocation of resources can be found as a solution to:

Programming problem 2.1

Maximize by choice of the lotteries $\pi^{d=0}(\tilde{\theta})$, expressing the probability of an audit, $d=0$, given parameter $\tilde{\theta}$ is announced; lotteries $\pi[c|\tilde{\theta},\theta,d=1]$, expressing the probability of consumption $c=(c^a,c^b)$ given $\tilde{\theta}$ is announced, θ

is realized, and there is an audit, $d=1$; and lotteries $\pi[c|\tilde{\theta}, d=0]$, expressing the probability of consumption c given $\tilde{\theta}$ is announced and there is no audit, the objective function

$$\lambda^a\left\{\sum_{\theta} p(\theta)\left\{\pi^{d=1}(\theta)\sum_{c} U^a(c^a,\theta)\pi[c|\theta,\theta,d=1]\right.\right.$$

$$\left.\left.+\pi^{d=0}(\theta)\sum_{c} U^a(c^a,\theta)\pi[c|\theta,d=0]\right\}\right\}$$

$$+\lambda^b\left\{\sum_{\theta} p(\theta)\left\{\pi^{d=1}(\theta)\sum_{c} U^b(c^b)\pi[c|\theta,\theta,d=1]\right.\right.$$

$$\left.\left.+\pi^{d=0}(\theta)\sum_{c} U^b(c^b)\pi[c|\theta,d=0]\right\}\right\}, \tag{3}$$

subject to constraint (2) for every $\theta, \tilde{\theta} \in \Theta$.

Revealing of how we solved this problem is the double θ notation in the objective function. Announced and actual θ values coincide. But they do so because constraints (2) are imposed, and these deal with counterfactual claims explicitly.

Obviously, the objective function in program 2.1 is continuous in the choice variables π. Further the set feasible solutions *unrestricted* by (2) is closed and bounded since any probability is a number between zero and one inclusive. Finally, constraints such as (2) are also continuous in π. Thus the constraint set is compact, and so a solution to the program is guaranteed to exist [note that $(c^a, c^b) = (e^a, e^b)$ is presumed to be feasible]. In fact, in eqs. (2) and (3) above the $\pi^{d=1}(\theta)$ and $\pi^{d=0}(\theta)$ always enter multiplicatively. Thus let

$$\pi[c, d=0|\theta] = \pi[c|\theta, d=0]\pi^{d=0}(\theta), \tag{4}$$

$$\pi\left[c, d=1|\tilde{\theta},\theta\right] = \pi\left[c|\tilde{\theta},\theta,d=1\right]\pi^{d=1}(\tilde{\theta}), \tag{5}$$

and program 2.1 can be transformed into a linear program:

Programming problem 2.2

Maximize the objective function

$$\lambda^a\left\{\sum_\theta p(\theta)\left\{\sum_c U^a(c^a,\theta)\pi[c,d=1|\theta,\theta]\right.\right.$$

$$\left.\left.+\sum_c U^a(c^a,\theta)\pi[c,d=0|\theta]\right\}\right\}$$

$$+\lambda^b\left\{\sum_\theta p(\theta)\left\{\sum_c U^b(c^b)\pi[c,d=1|\theta,\theta]\right.\right.$$

$$\left.\left.+\sum_c U^b(c^b)\pi[c,d=0|\theta]\right\}\right\}, \tag{6}$$

subject to

$$\sum_c U^a(c^a,\theta)\pi[c,d=1|\theta,\theta]+\sum_c U^a(c^a,\theta)\pi[c,d=0|\theta]$$

$$\geq \sum_c U^a(c^a,\theta)\pi\left[c,d=1|\tilde{\theta},\theta\right]+\sum_c U^a(c^a,\theta)\pi\left[c,d=0|\tilde{\theta}\right], \tag{7}$$

$$\forall\theta,\tilde{\theta}\in\Theta,$$

where

$$\pi\left[c,d=1|\tilde{\theta},\theta\right]\geq 0, \qquad \forall c\in C^{d=1}, \qquad \forall\tilde{\theta},\theta\in\Theta, \tag{8}$$

$$\pi\left[c,d=0|\tilde{\theta}\right]\geq 0, \qquad \forall c\in C^{d=0}, \qquad \forall\tilde{\theta}\in\Theta, \tag{9}$$

$$\sum_c\pi\left[c,d=1|\tilde{\theta},\theta\right]+\sum_c\pi\left[c,d=0|\tilde{\theta}\right]=1, \qquad \forall\tilde{\theta},\theta\in\Theta. \tag{10}$$

To ensure that solutions to this program weakly dominate autarky one can impose contraints

$$\left\{\sum_\theta p(\theta)\left\{\sum_c U^a(c^a,\theta)\pi[c,d=1|\theta,\theta]\right.\right.$$

$$\left.\left.+\sum_c U^a(c^a,\theta)\pi[c,d=0|\theta]\right\}\right\}$$

$$\geq \sum_\theta p(\theta)U^a(e^a,\theta), \tag{11}$$

and

$$\left\{\sum_{\theta} p(\theta)\left\{\sum_{c} U^b(c^b)\pi[c, d=1|\theta,\theta] + \sum_{c} U^b(c^b)\pi[c, d=0|\theta\right\}\right\}$$

$$\geq U^b(e^b). \tag{12}$$

3. Numerical solutions for the costly state verification environment

Solutions to program 2.2 have been computed for two specifications of preferences. For case 1, preferences of agent a are of the form $U^a(c,\theta) = [c+\delta]^\theta/\theta$ with $\theta \in \{0.2, 0.3, 0.8, 0.9\}$, each θ drawn with probability 0.25. For case 2, $\theta \in \{-3.0, -1.0, +0.10, +0.90\}$. In both cases preferences of agent b are of the form $U^b(c) = [c+\delta]^\psi/\psi$ with $\psi = 0.95$ for sure. Parameter $\delta = 0.05$ is designed to make preferences well defined even for negative θ. Note that for $c^a \geq 1-\delta$, high θ's make agent a more urgent to consume in the sense of having higher marginal utility for a given consumption allocation. In all the examples which follow the aggregate endowment $e = 10$ is split equally *a priori*, with $e^a = e^b = 5$. The cost of an audit is varied from $K = 0.5$ to $K = 10$.

Attention is restricted to a finite number of consumption bundles on the outer edge of the triangle of feasible resource allocations, though this is intended as an approximation to the solution with a continuum of possible consumption values. The first step in the numerical algorithm was to impose a relatively course grid, finding preliminary optimal solutions. Then further refinements in *neighborhoods* of the initial solutions were offered. The guess is that solutions with a relatively fine, *uniform* grid, with increments between consumptions of 0.01, were computed. Also, no case of multiple solutions was uncovered, though this possibility has not been ruled out.

Table 1 displays solutions to program 2.2 for case 1 preferences. These are the unrestricted, private information optimal allocations. Please note the odd usage here: unrestricted means not restricted by anything more than private information. Various additional restrictions are contemplated below. Please note also that here and below private information optimal solutions are rarely time consistent. Here, by the time audits take place, both agents may have full information. Thus one must suppose the existence of a commitment mechanism, perfect costless enforcement.

The full information allocation is also displayed at top of the table, as a base for comparison. Unless otherwise indicated below, the programming weights are $\lambda^a = 1$ and $\lambda^b = 0$, so that the constraint (12), that the utility of agent b be no less than at autarky, is binding. That is, all the gains to trade accrue to agent a, and in effect one maximizes the expected utility of agent a subject to the constraint that the utility of agent b be held at autarky. Also, full 'penalties' can be automatically imposed without loss of utility; that is, $c^a = 0$ if there is an audit, θ is revealed as the true parameter value, and counterfactual value of $\tilde{\theta}$ had been announced. This can always be assumed to be part of an optimal solution, though it matters only when an incentive constraint is binding. In any event the penalty branch is never displayed in the tables.

To be noted in table 1 is the tendency of private information to cause the full information solution to be pressed inward in the direction of less insurance, with higher values of K reflecting greater private information problems. That is, private information implies less indexation. This is 'intuitive' at least in the space of deterministic no audit and audit consumptions. Specifically, the full information solution has agent a eating more when he is relatively urgent to consume, at high θ values, and eating less at low θ values, as promised. Under private information, the no-audit consumption payoffs at $\theta = 0.2$ and $\theta = 0.3$ increase with values of the audit cost K, and the no audit consumption payoffs at $\theta = 0.8$ and $\theta = 0.9$ decrease with K, except as one moves to $K = 5$ at $\theta = 0.8$. Fig. 1 is suggestive, with the arrows indicating movement of no-audit consumptions with increases in K. A similar picture prevails for audit consumptions.

Audit probabilities in table 1 are monotone increasing in θ for each K, with urgency to consume as it were, and monotone decreasing at each θ value as K increases. Again, audit probabilities would seem a natural summary device, and so these results seem intuitively plausible.

The most striking feature of table 1 is the emergence of consumption lotteries when there is no audit. For example, at $K = 5$, there is a genuine lottery at $\theta = 0.8$ and at $\theta = 0.9$. In part consumption lotteries at $K = 5$ replace the random consumption outcomes of probabilitic audits, as at $K = 2$. Also note that these consumption lotteries are separating in that the variance in outcomes is less painful to relatively risk-neutral agent types, $\theta = 0.8, 0.9$, than to the risk-averse agent types, $\theta = 0.2, 0.3$. One gets a sense from the example that consumption lotteries are brought in as a last resort, at high audit costs. But this cannot be true generally. An example in Prescott and Townsend (1984) shows that distinguishing lotteries can sometimes achieve a full information optimal outcome without the need for any audits.

Interestingly, at $K \geq 5$, $\theta = 0.2$ and $\theta = 0.3$ agents are not distinguished from one another, and neither are $\theta = 0.8$ and $\theta = 0.9$ agents. Apparently, $\theta = 0.2$ and $\theta = 0.3$ agents are close to one another, as are $\theta = 0.8$ and $\theta = 0.9$ agents, and the benefit of distinguishing treatment within these groups is outweighed by the costs. That is, within group risk aversions are similar, so distinguishing lotteries would have to be relatively large if they were used.

Table 1

Case 1 preferences.[a]

Full information solution

θ	c^a	$\pi(c^a)$
0.2	1.63	1.0
0.3	1.76	1.0
0.8	6.44	1.0
0.9	9.85	1.0*

Private information solution and audit cost $K = 0.5$

θ	prob (audit)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
0.2	0.0	1.86	1.0	—	—
0.3	0.007964	1.9	0.992036	2.16	0.007964
0.8	0.365536	5.72	0.634464	6.3	0.365536
0.9	0.483195	9.77	0.516805*	9.37	0.483195

Private information solution and audit cost $K = 1$

θ	prob (audit)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
0.2	0.0	2.09	1.0	—	—
0.3	0.001763	2.1	0.998237	2.6	0.001763
0.8	0.304029	5.05	0.695970*	6.28	0.304029
0.9	0.458684	9.69	0.541316	8.91	0.458684

Private information solution and audit cost $K = 2$

θ	prob (audit)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
0.2	0.0	2.48	1.0	—	—
0.3	0.0	2.48	1.0	—	—
0.8	0.207649	4.36	0.792351*	6.7	0.207649
0.9	0.407632	9.21	0.592368*	7.97	0.407632

Private information solution and audit cost $K = 5$

θ	prob (audit)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
0.2	0.0	3.6	1.0*	—	—
0.3	0.0	3.6	1.0*	—	—
0.8	0.0	{0.0, 9.0}	{0.302120, 0.697880}	—	—
0.9	0.0	{0.0, 9.0}	{0.302120, 0.697880}	—	—

[a]The asterisks (*) indicate that the computed solution actually displayed a lottery with support at an adjacent grid point, no doubt an artifact of the grid itself. The indicated solution is at the grid point with highest mass, with the appropriate sum of probabilities. The asterisks are used in this fashion in all tables which follow.

Table 2

Case 2 preferences.

Full information solution

θ	c^a	$\pi(c^a)$
−3.0	1.46	1.0
−1.0	2.23	1.0
0.10	5.97	1.0
0.9	10.0	1.0

Private information solution and audit cost $K = 0.5$

θ	prob($d=1$)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
−3.0	0.00	1.44	1.0	—	—
−1.0	0.000026	2.17	0.999974	2.16	0.000026
0.10	0.014435	6.04	0.985565*	6.01	0.014435
0.9	0.118609	10.0	0.881391	9.5	0.118609

Private information solution and audit cost $K = 1$

θ	prob($d=1$)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
−3.0	0.00	1.42	1.00	—	—
−1.0	0.000027	2.11	0.999973	2.1	0.000027
0.10	0.015158	6.11	0.984841*	6.03	0.015158
0.9	0.115998	10.0	0.884002	9.0	0.115998

Private information solution and audit cost $K = 2$

θ	prob($d=1$)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
−3.0	0.00	1.38	1.0	—	—
−1.0	0.000028	2.01	0.999972	2.00	0.000028
0.10	0.016443	6.23	0.983557*	6.03	0.016443
0.9	0.111653	10.0	0.888347	8.0	0.111653

Private information solution and audit cost $K = 5$

θ	prob($d=1$)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
−3.0	0.0	1.29	1.0	—	—
−1.0	0.000033	1.81	0.999967	1.76	0.000033
0.10	0.019360	6.42	0.980641*	5.0	0.019360
0.9	0.104081	10.0	0.895919	5.0	0.104081

Private information solution and audit cost $K = 9$

θ	prob($d=1$)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
−3.0	0.0	1.28	1.0	—	—
−1.0	0.000037	1.94	0.999963	1.0	0.000037
0.10	0.018372	7.22	0.981628*	1.0	0.018372
0.9	0.0	{ 0.00 { 10.0	0.085140 } 0.914860 }	—	—

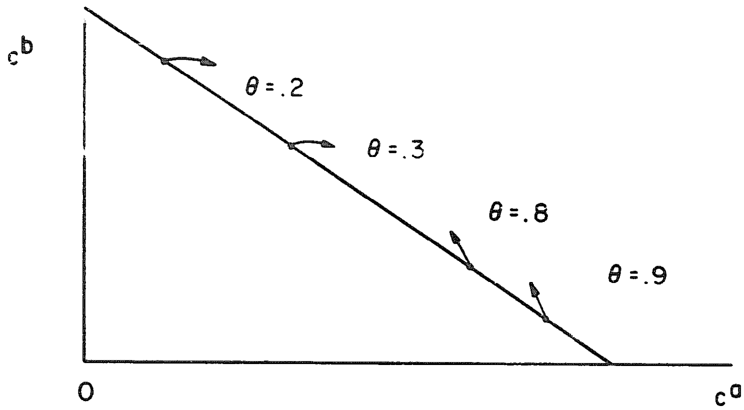

Fig. 1. Case 1 preferences: Less indexation as audit cost K increases.

Solutions are displayed in table 2 for case 2 preferences. Here, however, efforts to provide useful 'summaries' of the 'logic' of the solution proved more difficult. The $\theta = +0.10$ no-audit consumption solution increases with K, while the $\theta = +0.9$ solution stays put at the full information solution, dropping in expected value only at $K=9$. The $\theta = -3.0$ and $\theta = -1.0$ no-audit consumption solutions spread out as K increases (except for very high values of K, e.g., at $\theta = -1.0$ and going to $K=9$). Ignoring the exceptions, fig. 2 displays a type of 'overinsurance'. One guess, unconfirmed in some sense, as to why this happens is that the $\theta = -3.0$ and $\theta = -1.0$ agents are 'relatively easy' to distinguish from one another, while the $\theta = 0.10$ and $\theta = 0.9$ agents are not. Thus $\theta = 0.10$ consumption moves toward $\theta = 0.9$ consumption, and to free up the requisite resources, $\theta = -3$ and $\theta = -1$ consumptions are diminished. In other words, cross-sectional or within population movements dominate the less indexation 'intuition'.

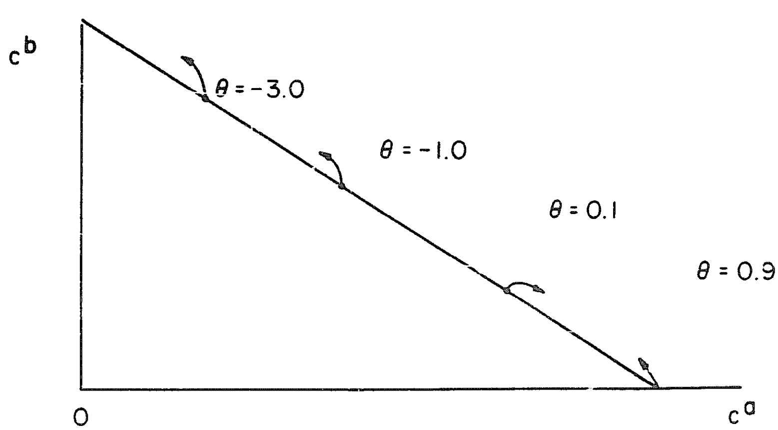

Fig. 2. Case 2 preferences: Counterintuitive indexation as audit cost K increases.

Striking also is the movement of the audit probabilities. From $K=0.5$ up to $K=5$, the probabilities of audits at $\theta=-1.0$ and $\theta=+0.10$ *increase* with increases in K. Apparently, this audit statistic is not a revealing summary of what is going on over the entire solution. Also, at the $K=9$ solution, the audit probability is *non-monotone* in θ, positive at $\theta=-1.0$ and $\theta=+0.10$ but zero at $\theta=0.9$. It is evident from this that the consumption lottery at $\theta=0.9$ in effect replaces a probabilistic audit, as foreshadowed above. This may help to explain the non-monotone audit probability result above. Again, audits are one way but not the only way of inducing random consumption allocations.

Striking also is that audits occur at even $K=9.9$, using 99% of available output (this is not in the table). Apparently, this allows $\theta=-3.0$ and $\theta-1.0$ agents to be distinguished from one another. Afterall, the spread between these types is large, and so distinguishing consumptions is important. And for $K=9.9$ significant and distinct consumption lotteries emerge at $\theta=0.10$ and $\theta=0.9$, so these agent types are also distinguished from one another. Distinguishing lotteries at $\theta=-1$, $\theta=0.10$ and $\theta=0.9$ prevail in the $K=10$ solution where of course there are no audits. Thus, consumption lotteries again appear as a last resort. A large spread among θ's makes distinct consumptions worthwhile, by one device or another. But with relatively high risk aversion, consumption lotteries are costly, and they are resisted until the costs of audits is exorbitant.

Exogenous exclusion restrictions and *a priori* limited contracts matter for consumptions. Table 3 returns to the case 1 preferences of table 1 but restricts audits to be deterministic when they occur, as in most of Townsend (1979). These are computed by a mixed integer linear program. Direct comparison of audit and non-audit branches reveals substantial differences in consumptions almost everywhere. One might also note that in table 3, under the restricted

Table 3

Restriction to deterministic audits, $K=0.5$ (upper part) and $K=1$ (lower part), case 1 preferences.

θ	Audit[a]	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
0.2	N	1.67	1.0*	—	—
0.3	N	1.67	1.0*	—	—
0.8	Y	—	—	6.03	1.0
0.9	Y	—	—	9.32	1.0
0.2	N	2.69	1.0*		
0.3	N	2.69	1.0*		
0.8	N	{0, 3.92}	{0.273724, 0.726271}		
0.9	Y	—	—	9.0	1.0

[a]Y = yes, N = no.

Table 4
Constant probability of auditing, $K = 0.5$, case 1 preferences.

θ^a	prob (audit)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
0.2	0.428382	1.73	0.571618	1.73	0.428382
0.3	0.428382	2.23	0.571618	2.21	0.428382*
0.8	0.428382	5.94	0.571618	5.80	0.428382
0.9	0.428382	8.94	0.571618*	9.27	0.428382

policy, low θ values fail to be distinguished from one another. Also, consumption lotteries emerge earlier, at $\theta = 0.8$ and $K = 1$ in table 3 rather than $K = 5$ as in table 1.

For completeness table 4 shows how a restriction to constant audit probabilities, that is, probabilities independent of θ, can also alter the nature of consumptions. (This is sometimes considered in the literature when revelation-type, announcement possibilities are ignored.) Table 4 should be compared to table 1 at $K = 0.5$. Perhaps surprising is the relatively high audit probability.

Private information optimal solutions were also computed for weights $\lambda^a = 0$ and $\lambda^b = 1$, so that all the gains to trade accrue instead to agent b. Table 5 presents the solution for case 2 preferences and $K = 5$. This should be compared to the $K = 5$ entry in table 2.

Evidently solutions do vary with the weights. In table 5 there tends to be less probabilistic auditing, with a significant consumption lottery at $\theta = 0.9$. This variation of optimal solutions with programming weights is troublesome from the standpoint of developing operational theory. In a large economy, with more households, one might suppose that the outcome be competitive, or in the core, removing the indeterminancy.

Various welfare experiments were also performed. The first experiment delivers a measure of the welfare loss of private information, the percentage of the aggregate endowment which would have to be destroyed in the full

Table 5
Effect of different λ's, $\lambda^a = 0$, $\lambda^b = 1$, $K = 5$, case 2 preferences.

θ^a	prob (audit)	c^a if $d=0$	$\pi(c^a, d=0, \theta)$	c^a if $d=1$	$\pi(c^a, d=1, \theta, \theta)$
−3.0	0.0	1.09	1.00	—	—
−1.0	0.000039	1.35	0.999961	1.32	0.000039
0.1	0.017640	2.67	0.982360*	2.60	0.017640
0.9	0.0	{0.00, 9.93}	{0.296753, 0.703247}		

Table 6

Welfare loss to private information (columns 1) and to restricted, constant audit probability policy (columns 2).

	Case 1 preferences		Case 2 preferences	
	(1)	(2)	(1)	(2)
$K=0.5$	1.2	2.3	0.3	0.6
$K=1$	2.3	4.0	0.7	1.3
$K=2$	4.1	5.5	1.4	2.7
$K=5$	5.9	5.9	3.4	4.5
$K=9$	5.9	5.9	4.7	4.9

information solution in order to lower the objective function to its value in the private information optimal solution. Also computed is the same measure of welfare loss when the private information solution is restricted to audit probabilities which are constant over θ. The results are presented in table 6, for both case 1 and case 2 preferences, as K is varied from 0.5 to 9.0.

For low values of K the welfare loss under the restricted policy is close to twice that of the loss to private information alone. This difference decreases as K increases. At relative high values of K the loss can be identical, as these values imply the constant policy of no auditing is optimal.

A similar welfare comparison is done with the deterministic audit restriction and is displayed in table 7, again for case 1 and case 2 preferences and various values of K. Again, the welfare loss from a restriction to deterministic audits is at least comparable to the welfare loss of private information alone, that is, the loss if often doubled. For case 2 preferences the results are more striking, as audits are driven out prematurely at relatively low values of K, causing dramatic welfare losses.

One policy which is not considered directly is a restriction which precludes consumption lotteries. But table 7 and an additional calculation tell the

Table 7

Welfare loss to private information (columns 1) and to restricted deterministic audit policy (columns 2).

	Case 1 preferences		Case 2 preferences	
	(1)	(2)	(1)	(2)
$K=0.5$	1.2	2.6	0.3	3.1
$K=1$	2.3	4.7	0.7	5.0
$K=2$	4.1	5.9	1.4	5.0
$K=5$	5.9	5.9	3.4	5.0
$K=9$	5.9	5.9	4.7	5.0

Table 8

Welfare loss to private information for different weights λ, $\lambda^a = 1$, $\lambda^b = 0$ (columns 1) and $\lambda^a = 0$, $\lambda^b = 1$ (columns 2).

	Case 1 preferences		Case 2 preferences	
	(1)	(2)	(1)	(2)
$K = 0.5$	1.2	1.2	0.3	0.6
$K = 1$	2.3	2.1	0.7	1.1
$K = 2$	4.1	3.9	1.4	2.0
$K = 5$	5.9	5.6	3.4	2.6
$K = 9$	5.9	5.6	4.7	2.6

story. For case 2 preferences, $K = 9$, and the restricted policy there are no audits at all, and so all beneficial exchange is supported by consumption lotteries. The welfare loss of private information under the restricted policy, relative to full information, is 5%. If random consumptions were also precluded, and audits remained at zero, as seems likely since an audit consumes 90% of output, the solution would be autarkic, necessarily. Then the welfare loss relative to full information, the percentage amount of the endowment which would have to be destroyed in the full information environment in order to lower the objective to its autarkic value, would be a dramatic 21%. This is by far largest welfare loss displayed in any of the examples. For case 1 preferences a similar lottery exclusion increases the welfare loss from 5.9% to 7.0%.

As a check on these welfare measures, the $\lambda^a = 0$ and $\lambda^b = 1$ specification is also adopted (in effect maximizing the *ex ante* utility of agent b subject to the utility of agent a at autarky) and the welfare loss is compared to the $\lambda^a = 1$, $\lambda^b = 0$ specification. Table 8 makes the comparison for the losses to private information for both preference cases. The numbers do differ, though only radically so for the high K, case 2 environment. Still, as we noted earlier, this indeterminancy in solutions is troublesome. More theory is needed to pin down the welfare weights.

We might conclude this discussion of restricted policies and welfare losses by noting the positive and normative implications in a somewhat different way. First, on the positive side, one tends to underpredict the use of audits if restricted policies are considered. Beneficial probabilistic audits occur in the unrestricted solutions even when they can consume 99% of social output, and their exclusion can cause a considerable loss in welfare. Second, on the normative side, suppose private information could be eliminated entirely *ex ante*, if part of social endowment is utilized (thrown away) *ex ante*. Then one could compare the two obvious regimes, *ex ante* elimination of private information versus optimal *ex post* audits with (possible) consumption lotteries. Indeed, the former regime will dominate the latter regime if the percentage of the endowment which is required to eliminate private information in the former is less than the welfare loss to private information displayed in the tables. It is clear from the tables that under the *restricted* policies, policies which enhance the welfare loss of private information, one might favor the use of *ex ante* elimination even though that would not be socially efficient. Alternatively, back to the positive side, if one were to limit attention to restricted policies, one might not make good predictions about choices among regimes.

The numerical examples displayed above make the point that private information optimal solutions can display surprising features, that restrictions which seem reasonable *a priori* can be damning in welfare terms. Of course not every environment will display these features! Indeed, in order to present a more balanced picture, it seems useful to return to the original costly state verification environment of Townsend (1979) and investigate there the impact of probabilistic audits. It may be recalled, again, that Townsend (1979) characterized optimal contracts and showed they looked like standard debt contracts under a restriction to deterministic audits, though probabilistic audits were shown in Townsend (1979) to be a Pareto dominating possibility.

The privately observed random endowment environment of Townsend (1979) can be handled as in program 2.2 here with the exception that an incentive constraint need not be imposed for a counterfactual endowment value greater than the actual endowment value. As an example, suppose $U^a(c^a)=(c^a+0.05)^{0.5}$, $U^b(c^b)=(c^b+0.05)$, $e^a \in \{4,6,8,10\}$ each with probability $\frac{1}{4}$ and $e^b=10$. Then for an audit cost of $K=2$, $\lambda^a=1$, $\lambda^b=0$, and a grid of increments of 0.1 on consumptions the solution is reported in table 9.

As suspected, audits are random. Here, on the other hand, and in contrast with the θ preference shock case, there are no additional (real) consumption lotteries apart from those induced by the grid, and related, no doubt, the solution is driven to autarky even for relatively low values of audit cost K, namely K between 2.95 and 3.0. Also, it is difficult to get high audit probabilities, even as $K \to 0$, though the threat of audit (and penalties) help

Table 9
Standard costly state verification.

e^a	prob (audits)	c^a if $d=0$	$\pi(c^a, d=0, e^a)$	c^a if $d=1$	$\pi(c^a, d=1, e^a, e^a)$
4	0.10520579	5.1	0.89479421	6.2	0.10520579
6	0.02495512	5.8	0.97504488*	7.0	0.02495512
8	0.00506625	7.4	0.11493375*	8.7	0.00506625
10	0.0	9.3	1.0*	—	—

considerably, with the solution approaching the with full information solution as $K \to 0$.

Are there any regularities? Mookherjee and Png (1987) are able to make some progress for a fairly general case, excluding consumption lotteries. They show that *only* the highest endowment will *not* be audited in any scheme which provides the agent with positive consumption, that is, audits will occur with positive probability everywhere else! Further, consumption for agent a in audit branches will always exceed consumption in non-audit branches. Finally, audit probabilities must be monotone decreasing with the magnitude of transfers. Otherwise, anomalies may prevail. In fact, Mookherjee and Png provide an example in which consumption of agent a is non-monotone with values of the endowment in the audit branch, and they leave open the possibility that audit probabilities may be non-monotone with values of the endowment. The exclusion of consumption lotteries may not be warranted, however. If the $U^a(\cdot)$ displays increasing absolute risk aversion, as for the quadratic function, then agent a may be less risk-averse at endowments where he is supposed to receive consumption. That creates the possibility of beneficial separating lotteries.

4. Pareto optimal learning

Imagine an economy with two agents (a and b), two dates (t and $t+1$), and a single underlying consumption good at each date. Given a path for consumptions c_t^i and c_{t+1}^i by agent i at date t and $t+1$, respectively, agent i has preferences as represented by the objective function $U^i(c_t^i, \theta_t^i) + \beta V^i(c_{t+1}^i, \theta_{t+1}^i)$, $i = a, b$. For simplicity here quantities of consumption bundles at each date are presumed to take on at most a finite number of values, as if consumption took place in discrete units. Here also there are supposed to be a finite number of possible shock specifications $(\theta_t^a, \theta_t^b, \theta_{t+1}^a, \theta_{t+1}^b)$ in some finite set Θ, each drawn with *a priori* probabilities $p(\theta_t^a, \theta_t^b, \theta_{t+1}^a, \theta_{t+1}^b)$, which are common knowledge. Shocks θ may be supposed to hit an underlying household production function, so that the $U^i(\cdot)$ and $V^i(\cdot)$ are indirect functions for 'market acquired' consumption. Household specific shocks are naturally private but possibly related.

Shock θ_t^i is seen by agent i *alone* at the beginning of date t, $i = a, b$, and similarly for θ_{t+1}^i at date $t+1$. For simplicity it is supposed that shocks θ_t^i and θ_{t+1}^i still leave agent i in doubt as to the value of shock θ_t^j, $j \neq i$. That is, the conditional probabilities $p(\theta_t^j | \theta_t^i, \theta_{t+1}^i)$ are strictly positive. It is supposed also that the $\theta_{t+1}^a, \theta_{t+1}^b$ are not inferable from contemporary values. That is, the $p(\theta_{t+1}^a, \theta_{t+1}^b | \theta_t^a, \theta_t^b)$ are also strictly positive. For simplicity we may suppose that the aggregate social endowments are deterministic, represented by e_t at date t, with $e_t = e_t^a + e_t^b$ as the sum of endowments of agents a and b, and similarly at date $t+1$.

The environment here may be seen as a special but fully specified economic environment, an analogue to the game theoretic setup studied by Myerson (1986). The environment is special in that the agents take no actions *per se* and special in that events of probability zero cannot occur. The latter is guaranteed by the no-action possibilities specification and the positive probabilities specification. Yet the environment is dynamic – the consumption good must be allocated period by period. These period by period consumption allocations create information revelation possibilities which the mechanism design must accommodate. In this sense the allocation is like an action in Myerson (1986).

Myerson (1986) emphasizes the importance of communication among agents in the context of his dynamic, sequential games. This takes the form, among other things, of allowing agents to commit to certain actions or inactions and communicating this to other agents, either directly, or, indirectly, through a mediator. (See Myerson's example 2.) It also takes the form of allowing agents to coordinate their actions by keying off events, possibly random events, observed by them. This is facilitated by preplay communication. Alternatively, the agents could follow the recommendations of a mediator who passes on such recommendations according to such random events or similar events of his own making. (Again, see Myerson's example 2.) Myerson argues, in fact, that there is no loss of generality in requiring that all communication go through a mediator, because agents can do through a mediator anything they can do on their own. In fact, there may be an advantage to going through a mediator in that the latter can control information flows, preventing agents from acting *ex post* against their own *ex ante* interest. Related, Myerson allows the mediator to scramble his recommendations to the agents. (See Myerson's example 5.)

For the economic environment of this section we accept fully the idea that all communication can go through a mediator without loss of generality. No proof of this is given. Then we shall derive an analogue to the scrambling result allowed by Myerson, showing here that the mediator might well scramble the date 1 consumption allocation in an effort to conceal information. In fact a case of active, non-trivial scrambling is computed. Finally, it proved for the dynamic sequential information environment here that the revelation principle applies. This should help to convince remaining skeptics. More to the point, the argument makes heavy use of dynamic programming considerations and delivers a programming problem and incentive constraint different from the forms presented by Myerson. Though there is an equivalence, the differences in form are instructive. *Again, readers interested in seeing the result first can turn to program 4.1 or 4.2 below, with notation defined therein.*

To continue, then, a game or resource allocation scheme for the economy of this section includes in its specification a message space M_t^i at date t for each agent i ($i = a, b$) and a message space $M^i_{t+1}(m^i_t, c_t)$ at date $t+1$ for each agent i ($i = a, b$) where $c_t = (c^a_t, c^b_t)$. The idea here is that messages sent by agent j at date t may *not* become known to agent i at date $t+1$, so the set of *a priori* feasible messages at date $t+1$ for agent i may *not* be restricted by m^j_t, $j \neq i$. On the other hand past own messages m^i_t and past own consumptions c^i_t are known at date $t+1$, and with deterministic endowments and the assumption that resources cannot be thrown away, c^j_t is known as well. For specificity and simplicity here message spaces are taken to be finite dimensional subsets of Euclidian spaces.

The second class of objects included in the specification of a game is the set of allocation rules. At date t the set of feasible consumptions c_t is restricted to lie in $C_t = \{(c^a_t, c^b_t): c^a_t \geq 0, c^b_t \geq 0,\ c^a_t + c^b_t = e_t\}$, where again space C_t is finite dimensional and similarly at date $t+1$. Let Λ_t denote the set probability measures over C_t. Then an allocation rule at date t is a measure $\Delta_t(m^a_t, m^b_t)$ in Λ_t, a lottery specifying the probability $\Delta_t(c_t | m^a_t, m^b_t)$ of consumption bundle c_t effected by actual messages m^a_t and m^b_t. That is, a computer is preprogrammed to map messages into random outcomes even if messages are privately sent. Similarly, let Λ_{t+1} denote the set of probability measures over consumption set C_{t+1}. Then an allocation rule at date $t+1$ is a measure $\Delta_{t+1}(m^a_t, m^b_t, m^a_{t+1}, m^b_{t+1}, c_t)$ in Λ_{t+1} with typical element $\Delta_{t+1}(c_{t+1} | m^a_t, m^b_t, m^a_{t+1}, m^b_{t+1}, c_t)$.

A strategy at date t for agent i is a possibly random choice of a message in M^i_t given his observed θ^i_t. That is, a strategy is a measure $\sigma^i_t(\theta^i_t)$ conditioned on θ^i_t, a lottery specifying the probability $\sigma^i_t(m^i_t | \theta^i_t)$ of message m^i_t. A strategy at date $t+1$ for agent i is a possibly random choice of a message in $M^i_{t+1}(m^i_t, c_t)$ given his observed shocks θ^i_t and θ^i_{t+1} and conditioned on his past message m^i_t and past observed outcome c_t. That is, a strategy is a measure $\sigma^i_{t+1}(\theta^i_t, \theta^i_{t+1}, m^i_t, c_t)$ over $M^i_{t+1}(m^i_t, c_t)$ with typical element $\sigma^i_{t+1}(m^i_{t+1} | \cdot)$.

Now imagine what happens as the game is played out over time, so that here and below sequential aspects are emphasized. Given some specification of date t strategies $\sigma^b_t(\theta^b_t)$ over M^b_t values and given an (arbitrary) allocation rule $\Delta_t(m^a_t, m^b_t)$ over c_t, agent a observes a particular c_t at the end of date t, knows as well particular values $\tilde{\theta}^a_t$ and $\tilde{m}^a_t$, and at the beginning of date $t+1$ knows a particular $\tilde{\theta}^a_{t+1}$. Inferences about particular values $\tilde{\theta}^b_t$ are then made optimally under Bayes' rule, if possible. That is, in the obvious suggestive notation,

$$\operatorname{Prob}\left(\tilde{\theta}^b_t | \tilde{\theta}^a_t, \tilde{\theta}^a_{t+1}, \tilde{m}^a_t, c_t\right)$$

$$= \frac{\operatorname{Prob}\left(c_t | \tilde{\theta}^a_t, \tilde{\theta}^b_t\right) \operatorname{Prob}\left(\tilde{\theta}^b_t | \tilde{\theta}^a_t, \tilde{\theta}^a_{t+1}\right)}{\sum_{\theta^b_t} \operatorname{Prob}\left(c_t | \tilde{\theta}^a_t, \theta^b_t\right) \operatorname{Prob}\left(\theta^b_t | \tilde{\theta}^a_t, \tilde{\theta}^a_{t+1}\right)}, \tag{13}$$

or, more precisely,

$$p_{t+1}^{a*}(\tilde{\theta}_t^b | \tilde{\theta}_t^a, \tilde{\theta}_{t+1}^a, \tilde{m}_t^a, c_t) = \frac{\sum_{m_t^b} \Delta_t(c_t | \tilde{m}_t^a, m_t^b)\sigma_t^b(m_t^b | \tilde{\theta}_t^b) p(\tilde{\theta}_t^b | \tilde{\theta}_t^a, \tilde{\theta}_{t+1}^a)}{\sum_{\theta^b} \sum_{m_t^b} \Delta_t(c_t | \tilde{m}_t^a, m_t^b)\sigma_t^b(m_t^b | \theta_t^b) p(\theta_t^b | \tilde{\theta}_t^a, \tilde{\theta}_{t+1}^a)}. \tag{14}$$

Here then the left-hand side of (14) denotes the posterior probability of agent a for a particular value $\tilde{\theta}_t^b$ conditioned on $\tilde{\theta}_t^a$, $\tilde{\theta}_{t+1}^a$, $\tilde{m}_t^a$, and c_t. Of course $p(\tilde{\theta}_t^b | \tilde{\theta}_t^a, \tilde{\theta}_{t+1}^a)$ on the right-hand side of (14) is the obvious conditional probability for $\tilde{\theta}_t^b$ based on $\tilde{\theta}_t^a$ and $\tilde{\theta}_{t+1}^a$ alone. The posterior probability of θ_t^b *and* θ_{t+1}^b is then obviously

$$p_{t+1}^{a*}(\theta_t^b, \theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a, m_t^a, c_t)$$
$$\equiv p(\theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a, \theta_t^b) p_{t+1}^{a*}(\theta_t^b | \theta_t^a, \theta_{t+1}^a, m_t^a, c_t).$$

Posteriors for agent b are formed similarly.

Formula (14) and its analogue for agent b do not work for all conditioning values, $\tilde{\theta}_t^a$, $\tilde{\theta}_{t+1}^a$, $\tilde{m}_t^a$, c_t. In particular formula (14) requires that c_t values be generated in conjunction with strategy $\sigma_t^b(m_t^b | \theta_t^b)$ of agent b at date t, for 'on path' values, as it were. Yet the definition of equilibrium, of maximization of agent a at date $t+1$ in particular, requires that agent a have a posterior for *all* conditioning values. In this way, in choosing his maximizing strategy $\sigma_t^b(m_t^b | \theta_t^b)$ at date t, agent b can contemplate *arbitrary* messages m_t^b, knowing there is a mapping for agent a at date $t+1$ from the outcomes generated by such messages to a date $t+1$ decision for agent a. Fortunately, or unfortunately, formula (14) and its analogue for agent b can be filled in any arbitrary way for these so-called off-path values, completing the definition of equilibrium. (This can be unfortunate or troublesome in that these arbitrary specifications matter for the equilibrium outcome!) The intent here, however, is only to deliver any particular equilibrium outcome as a truthtelling equilibrium of a direct mechanism.

A *Bayesian sequential Nash equilibrium* of a particular game is a specification of strategies $\sigma_t^{a*}(\theta_t^a)$, $\sigma_t^{b*}(\theta_t^b)$, $\sigma_{t+1}^{a*}(\theta_t^a, \theta_{t+1}^a, m_t^a, c_t)$ and $\sigma_{t+1}^{b*}(\theta_t^b, \theta_{t+1}^b, m_t^b, c_t)$ and a specification of posterior probabilities $p_{t+1}^{a*}(\theta_t^a, \theta_{t+1}^a, m_t^a, c_t)$ and $p_{t+1}^{b*}(\theta_t^b, \theta_{t+1}^b, m_t^b, c_t)$ such that the following conditions hold:

(i) At date $t+1$, given any θ_t^a, θ_{t+1}^a, m_t^a and c_t, given posterior $p_{t+1}^{c*}(\theta_t^a, \theta_{t+1}^a, m_t^a, c_t)$ and given strategies $\sigma_{t+1}^{b*}(\theta_t^b, \theta_{t+1}^b, m_t^b, c_t)$ and $\sigma_t^{b*}(\theta_t^b)$ for agent b, the strategy $\sigma^{a*}_{t+1}(\theta^a_t, \theta^a_{t+1}, m^a_t, c_t)$ for agent a solves the problem of maximizing the objective function

$$\sum_{\theta^b_t} \sum_{\theta^b_{t+1}} \sum_{m^b_t} \sum_{m^b_{t+1}} \sum_{m^a_{t+1}} \sum_{c_{t+1}} V^a[c^a_{t+1}, \theta^a_{t+1}]$$

$$\times \Delta_{t+1}\left(c_{t+1} | m^a_t, m^b_t, m^a_{t+1}, m^b_{t+1}, c_t\right)$$

$$\times \sigma^a_{t+1}\left(m^a_{t+1} | \theta^a_t, \theta^a_{t+1}, m^a_t, c_t\right) \sigma^{b*}_{t+1}\left(m^b_{t+1} | \theta^b_t, \theta^b_{t+1}, m^b_t, c_t\right)$$

$$\times \sigma^{b*}_t\left(m^b_t | \theta^b_t\right) p^{a*}_{t+1}\left(\theta^b_t, \theta^b_{t+1} | \theta^a_t, \theta^a_{t+1}, m^a_t, c_t\right). \tag{15}$$

Denote the maximized value of the objective function to this problem $V^{a*}[\theta^a_t, \theta^a_{t+1}, m^a_t, c_t]$.

(ii) At date t, given θ^a_t and given strategy $\sigma^{b*}_t(\theta^b_t)$, the strategy $\sigma^{a*}_t(\theta^a_t)$ solves the problem of maximizing the objective function

$$\sum_{\theta^a_{t+1}} \sum_{\theta^b_t} \sum_{m^b_t} \sum_{m^a_t} \sum_{c_t} \{U[c^a_t, \theta^a_t] + \beta V^{a*}[\theta^a_t, \theta^a_{t+1}, m^a_t, c_t]\}$$

$$\times \Delta_t\left(c_t | m^a_t, m^b_t\right) \sigma^a_t\left(m^a_t | \theta^a_t\right) \sigma^{b*}_t\left(m^b_t | \theta^b_t\right) p\left(\theta^a_{t+1}, \theta^b_t | \theta^a_t\right). \tag{16}$$

(iii) The posterior $p^{a*}_{t+1}(\theta^a_t, \theta^a_{t+1}, m^a_t, c_t)$ satisfies Bayes' rule (14) given $\sigma^{b*}_t(\theta^b_t)$ and the specified arguments, if possible, and is arbitrary otherwise.

(iv) Similar conditions apply for strategy $\sigma^{b*}_{t+1}(\theta^b_t, \theta^b_{t+1}, m^b_t, c_t)$ of agent b at $t+1$.

(v) Similar conditions apply for strategy $\sigma^{b*}_t(\theta^b_t)$ of agent b at t.

(vi) Similar conditions apply for the posterior $p^{b*}_{t+1}(\theta^b_t, \theta^b_{t+1}, m^b_t, c_t)$ of agent b.

Given an equilibrium allocation for some specified game we can construct now a new allocation mechanism which has as an equilibrium the same allocation. First, let the (new) message space of agent i at date t be the set of possible values for θ^i_t, and similarly at date $t+1$ let the message space be the set of possible values for θ^i_t *and* θ^i_{t+1}. Each agent i is restricted to announcing values for his privately observed parameter values, or history of parameter values, at each date $t = t, t+1$ and in fact in what follows will be given an incentive to do so truthfully. The new mechanism, however, allows these announcements at date t to be scrambled. In particular, under the old mechanism at date t agent i may have been selecting a message m^i_t at random from M^i_t in an effort to conceal his message and to conceal actual parameter value θ^i_t. By the same token agent $j \neq i$ was attempting to infer m^i_t from c_t

and thus to infer also θ_t^i. From this point of view, then, it is apparent that the old particular message space M_t^i in use was of no special revelance, and one could have replaced it with any alternative that allowed the same degree of scrambling desired by agent i. In short the old message could just as easily have been taken to be Θ_t^i, the set of possible values of parameter θ_t^i, since this allows *any* degree of scrambling including the degree desired by agent i. *This replacement is now assumed.* In addition the new mechanism builds in this degree of scrambling automatically by incorporating it as part of its *internal operation*. The scrambler is a lottery on θ_t^i effected by an announcement of θ_t^i by agent i at date t.

More formally, then, let the allocation rule of the new mechanism, denoted $\pi_t(\cdot)$, be constructed as follows. First, let $\pi_t^i(\theta_t^i) \equiv \sigma_t^{i*}(\theta_t^i)$. Here then $\pi_t^i(\theta_t^i)$ denotes a lottery over internal computer messages $m_t^i \in \Theta_t^i$, here as a function of the value of θ_t^i named by agent i, just as if the internal message had been sent by agent i originally. That is, $\pi_t^i(m_t^i|\theta_t^i)$ is the probability that the computer names m_t^i given that agent i announces θ_t^i. The message m_t^i is however known to agent i, just as if it had been sent by agent i. Second let $\pi_t(m_t^a, m_t^b) \equiv \Delta_t(m_t^a, m_t^b)$. Here then $\pi_t(m_t^a, m_t^b)$ is a lottery over consumptions c_t, so that $\pi_t(c_t|m_t^a, m_t^b)$ is the probability the computer picks c_t given internally generated named values m_t^a and m_t^b for θ_t^a and θ_t^b, respectively. Conditional on messages, then, the computer is doing exactly what it did before, under the old game. Third, let

$$\begin{aligned}
&\pi_{t+1}\left(c_{t+1}|m_t^a, m_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t\right) \\
&\equiv \sum_{m_{t+1}^a} \sum_{m_{t+1}^b} \Delta_{t+1}\left(c_{t+1}|m_t^a, m_t^b, m_{t+1}^a, m_{t+1}^b, c_t\right) \\
&\quad \times \sigma_{t+1}^{a*}\left(m_{t+1}^a|m_t^a, \theta_t^a, \theta_{t+1}^a, c_t\right) \\
&\quad \times \sigma_{t+1}^{b*}\left(m_{t+1}^b|m_t^b, \theta_t^b, \theta_{t+1}^b, c_t\right). \qquad (17)
\end{aligned}$$

Here then $\pi_{t+1}(c_{t+1}|m_t^a, m_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t)$ is the probability of allocation c_{t+1} given earlier computer-generated message m_t^a and m_t^b and named values at date $t+1$ of the parameters θ_t^a and θ_{t+1}^a by agent a and θ_t^b and θ_{t+1}^b by agent b. It may be noted here that unlike the rule at date t, reference to all internally generated messages m_{t+1}^a and m_{t+1}^b at date $t+1$ can be suppressed since date $t+1$ is the last date and there is nothing to learn for the future. That is, the computer is programmed to send internal messages to itself and then to generate random consumption allocations according to the old game. But agents need not keep track of these date $t+1$ messages *per se*. In short, the new allocation rule $\pi_{t+1}(c_{t+1}|\cdot)$ over c_{t+1} entirely encompasses the randomness in old allocation rule $\Delta_{t+1}(c_{t+1}|\cdot)$ *and* in the strategies

$\sigma_{t+1}^{i*}((m_{t+1}^i \mid \cdot)$, $i = a, b$. Finally, note that both θ_t^i and θ_{t+1}^i are incorporated into the allocation rule at date $t+1$, $i = a, b$, but this is simply because both enter the strategy of agent i at date $t+1$.

Now let $\gamma_t^i(\theta_t^i)$ denote a (deterministic) strategy of agent i at date t, a named value in Θ_t^i. Similarly let $\gamma_{t+1}^i(\theta_t^i, \theta_{t+1}^i, m_t^i, c_t)$ denote a (deterministic) strategy for agent i at date $t+1$, again named a doubleton in $\Theta_t^i \times \Theta_{t+1}^i$. It is a relatively easy matter to establish that the new mechanism with message space Θ_t^i and $\Theta_t^i \times \Theta_{t+1}^i$ for agent i at date t and $t+1$, respectively, $i = a, b$, and with rules π specified above, has an equilibrium with truthtelling strategies $\gamma_t^i(\theta_t^i) = \theta_t^i$ and $\gamma_{t+1}^i(\theta_t^i, \theta_{t+1}^i, m_t^i, c_t) = (\theta_t^i, \theta_{t+1}^i)$, that is, with named values equal to actual values. Thus the new mechanism will achieve the same probabilistic allocation of resources as the original mechanism.

For suppose, to emphasize sequential aspects, agent a adopts an arbitrary strategy $\gamma_t^a(\theta_t^a)$ at date t, possibly naming a value of $\tilde{\theta}_t^a$ *distinct* from the actual value θ_t^a. Suppose further this generates a message m_t^a under the scrambler $\pi_t^a(\tilde{\theta}_t^a)$, this message being seen by agent a. Agent a assumes however that agent b announces truthfully at date t, generating message m_t^b under the scrambler $\pi_t^b(\theta_t^b)$. Then at date $t+1$, given actual parameter draw θ_{t+1}^a, agent a has posterior $p_{t+1}^{a*}(\theta_t^b \mid \theta_t^a, \theta_{t+1}^a, m_t^a, c_t)$, exactly the same posterior agent a would have had in the old mechanism for the specified path $(\theta_t^a, \theta_{t+1}^a, m_t^a, c_t)$. That is, for every (m_t^a, c_t) outcome and for parameter draws θ_t^a and θ_{t+1}^a, agent a draws exactly the same inferences about θ_t^b as before. This is apparent from Bayes rule (14) and the assumption that agent b is telling the truth so that in effect *his* messages are sent in exactly the same probabilistic way as before. Off this path, the posterior can be arbitrary and can be taken to be what it was before. Now at date $t+1$, agent a can name the actual values of his parameter draws $(\theta_t^a, \theta_{t+1}^a)$, or he can name counterfactual values $(\tilde{\theta}_t^a, \tilde{\theta}_{t+1}^a)$. If he names truthfully and assumes agent b is naming truthfully, he can generate by construction of the new mechanism the same distribution over outcomes he would have faced under the old mechanism under the specified path. Alternatively, if he announces counterfactually, naming $(\tilde{\theta}_t^a, \tilde{\theta}_{t+1}^a)$, that, by virtue of the construction of the new mechanism, would be like employing the strategy $\sigma_{t+1}^{a*}(\tilde{\theta}_t^a, \tilde{\theta}_{t+1}^a, m_t^a, c_t)$, a random strategy that was available to agent a under the old mechanism under path $(\theta_t^a, \theta_{t+1}^a, m_t^a, c_t)$ but *not chosen*. Thus truthtelling weakly dominates for agent a at date $t+1$ for all paths θ_t^a, θ_{t+1}^a, m_t^a, and c_t. Thus the maximized value of objective function of agent a at date t is given by $V^{a*}[\theta_t^a, \theta_{t+1}^a, m_t^a, c_t]$ as before.

Working backwards in this way to date t, again suppose agent a continues to assume that agent b employs the truthtelling strategy under actual parameter draw θ_t^b, activating the scrambler $\pi_t^b(\theta_t^b)$. Then under actual parameter draw θ_t^a agent a faces essentially the same decision problem as in the old mechanism, because the second period value function is the same. Again he could effect the same distribution of consumption outcomes as before, by announcing truthfully, since his old random strategy and the old probabilistic generation of c_t are incorporated into π_t. Alternatively, he could announce counterfactually, saying $\tilde{\theta}_t^a$, activating scrambler $\pi_t^a(\tilde{\theta}_t^a)$. But this would be like employing the random strategy $\sigma_t^{a*}(\tilde{\theta}_t^a)$ under the old mechanism, a choice available at date t and draw θ_t^a but *not chosen*. Thus truthtelling weakly dominates at date t.

Of course similar arguments apply for agent b at date t and at date $t+1$. Thus both agents adopt truthtelling strategies in the new mechanism, and as is obvious by the above arguments the same distribution of consumption allocation of resources is achieved as under the equilibrium of the old mechanism.

To summarize the discussion thus far, we can restrict ourselves without loss of generality to the above described class of new mechanisms, with truthtelling as equilibrium strategies. That is, the equilibrium outcome of any mechanism in the class given initially, with the specified message spaces and outcome functions, satisfies the constraints just described, namely, for every path $\theta_t^a, \theta_{t+1}^a, m_t^a, c_t$, and for every counterfactual announcement at date $t+1$, $\tilde{\theta}_t^a, \tilde{\theta}_{t+1}^a$,

$$\begin{aligned}
&V^{a*}\left[\theta_t^a, \theta_{t+1}^a, m_t^a, c_t\right] \\
&\equiv \sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{m_t^b} \sum_{c_{t+1}} V^a(c_{t+1}^a, \theta_{t+1}^a) \\
&\quad \times \pi_{t+1}\left(c_{t+1} \mid m_t^a, m_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t\right) \\
&\quad \times \pi_t^b\left(m_t^b \mid \theta_t^b\right) p_{t+1}^{a*}\left(\theta_t^b, \theta_{t+1}^b \mid \theta_t^a, \theta_{t+1}^a, m_t^a, c_t\right) \\
&\geq \sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{m_t^b} \sum_{c_{t+1}} V^a(c_{t+1}^a, \theta_{t+1}^a) \pi_{t+1}\left(c_{t+1} \mid m_t^a, m_t^b, \tilde{\theta}_t^a, \tilde{\theta}_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t\right) \\
&\quad \times \pi_t^b\left(m_t^b \mid \theta_t^b\right) p_{t+1}^{a*}\left(\theta_t^b, \theta_{t+1}^b \mid \theta_t^a, \theta_{t+1}^a, m_t^a, c_t\right),
\end{aligned} \tag{18}$$

where

$$p_{t+1}^{a*}\left(\theta_t^b \mid \theta_t^a, \theta_{t+1}^a, m_t^a, c_t\right) = \frac{\sum_{m_t^b} \pi_t\left(c_t \mid m_t^a, m_t^b\right) \pi_t^b\left(m_t^b \mid \theta_t^b\right) p\left(\theta_t^b \mid \theta_t^a, \theta_{t+1}^a\right)}{\sum_{\hat{\theta}_t^b} \sum_{m_t^b} \pi_t\left(c_t \mid m_t^a, m_t^b\right) \pi_t^b\left(m_t^b \mid \hat{\theta}_t^b\right) p\left(\hat{\theta}_t^b \mid \theta_t^a, \theta_{t+1}^a\right)}, \tag{19}$$

and

$$\begin{aligned}
&p_{t+1}^{a*}\left(\theta_t^b, \theta_{t+1}^b \mid \theta_t^a, \theta_{t+1}^a, m_t^a, c_t\right) \\
&= p\left(\theta_{t+1}^b \mid \theta_t^a, \theta_{t+1}^a, \theta_t^b\right) p_{t+1}^{a*}\left(\theta_t^b \mid \theta_t^a, \theta_{t+1}^a, m_t^a, c_t\right).
\end{aligned}$$

Here, as before, the posterior $p_{t+1}^{a*}(\theta_t^b, \theta_{t+1}^b \mid \cdot)$ should be filled in for *all* $\theta_t^a, \theta_{t+1}^a, m_t^a, c_t$ combinations, creating the possibility that the posterior must be specified, if only in an arbitrary way, for off-path values. Now, however, agents a and b are essentially restricted to deviations which were maximizing actions for some parameter under the old game and for which formulas (19) and (14) apply. Thus there are no off-path values.

Similarly, at date t for every parameter draw θ_t^a and every counterfactual announcement at date t, $\hat{\theta}_t^a$,

$$\sum_{\theta_{t+1}^a} \sum_{\theta_t^b} \sum_{m_t^b} \sum_{m_t^a} \sum_{c_t} \{U^a(c_t, \theta_t^a) + \beta V^{a*}[\theta_t^a, \theta_{t+1}^a, m_t^a, c_t]\}$$

$$\times \pi_t(c_t | m_t^a, m_t^b) \pi_t^a(m_t^a | \theta_t^a) \pi_t^b(m_t^b | \theta_t^b) p(\theta_t^b, \theta_{t+1}^a | \theta_t^a)$$

$$\geq \sum_{\theta_{t+1}^a} \sum_{\theta_t^b} \sum_{m_t^b} \sum_{m_t^a} \sum_{c_t} \{U(c_t, \theta_t^a) + \beta V^{a*}[\theta_t^a, \theta_{t+1}^a, m_t^a, c_t]\}$$

$$\times \pi_t(c_t | m_t^a, m_t^b) \pi_t^a(m_t^a | \hat{\theta}_t^a) \pi_t^b(m_t^b | \theta_t^b) p(\theta_t^b, \theta_{t+1}^a | \theta_t^a). \tag{20}$$

These constraints can be simplified further. In particular under the mechanism $\pi_t(\cdot)$ and $\pi_{t+1}(\cdot)$ the truthtelling strategy is maximizing for agent a for all possible messages m_t^a, given actual parameters θ_t^a and θ_{t+1}^a and previous consumptions c_t, as (18) makes clear. Thus even in the absence of information on m_t^a, truthtelling would be maximizing for agent a. More formally, suppose agent a announced $\hat{\theta}_t^a$ at date t, possibly different from the actual θ_t^a, and though $\hat{\theta}_t^a$ is scrambled via $\pi_t^a(\hat{\theta}_t^a)$, delivering some message m_t^a, this is *not* seen by agent a. For a particular $m_t^a, \hat{\theta}_t^a, \theta_t^a, \theta_{t+1}^a, c_t$ combination, multiply both sides of (18) by $\pi_t^a(m_t^a | \hat{\theta}_t^a)$, delivering in the objective function the new but obvious posterior $p_{t+1}^{a*}(\theta_t^b, \theta_{t+1}^b, m_t^a | \theta_t^a, \theta_{t+1}^a, \hat{\theta}_t^a, c_t)$. Summing over m_t^a values thus yields truthtelling as the maximizing strategy in the absence of information on m_t^a. This makes clear that explicit dependence on m_t^a in the posterior agent a, m_t^b in the posterior of agent b, and m_t^a, m_t^b in the rules π_t and π_{t+1} can be suppressed. That is, let

$$\pi_t^*(c_t | \hat{\theta}_t^a, \hat{\theta}_t^b) \equiv \sum_{m_t^a, m_t^b} \pi_t(c_t | m_t^a, m_t^b) \pi_t^a(m_t^a | \hat{\theta}_t^a) \pi_t^b(m_t^b | \hat{\theta}_t^b),$$

$$\pi_{t+1}^*(c_{t+1} | \hat{\theta}_t^a, \hat{\theta}_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t)$$

$$\equiv \sum_{m_t^a, m_t^b} \pi_{t+1}(c_{t+1} | m_t^a, m_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t)$$

$$\times \pi_t^a(m_t^a | \hat{\theta}_t^a) \pi_t^b(m_t^b | \hat{\theta}_t^b).$$

Then constraints (18) may be written for all actual paths $(\theta_t^a, \theta_{t+1}^a, c_t)$, announcements $\hat{\theta}_t^a$ at date t, and all counterfactual announcements, $\hat{\hat{\theta}}_t^a, \hat{\hat{\theta}}_{t+1}^a$ at date $t+1$,

$$\sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} V^a(c_{t+1}^a, \theta_{t+1}^a) \pi_{t+1}^*(c_{t+1} | \hat{\theta}_t^a, \theta_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t)$$
$$\times p_{t+1}^{a*}(\theta_t^b, \theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a, \hat{\theta}_t^a, c_t)$$
$$\geq \sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} V^a(c_{t+1}^a, \theta_{t+1}^a) \pi_{t+1}^*\left(c_{t+1} | \hat{\theta}_t^a, \theta_t^b, \hat{\hat{\theta}}_t^a, \hat{\hat{\theta}}_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t\right)$$
$$\times p_{t+1}^{a*}(\theta_t^b, \theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a, \hat{\theta}_{t+1}^a, c_t). \tag{21}$$

Even further simplications are now seen to be possible. For recall from the earlier discussion that

$$p_{t+1}^{a*}(\theta_t^b, \theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a, \hat{\theta}_t^a = \delta_t^a(\theta_t^a), c_t)$$
$$= p_{t+1}^{a*}(\theta_t^b | \theta_t^a, \theta_{t+1}^a, c_t, \hat{\theta}_t^a = \delta_t^a(\theta_t^a)) p(\theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a, \theta_t^b), \tag{22}$$

where, given θ_t^a, agent a employs the announcement rule $\hat{\theta}_t^a = \delta_t^a(\theta_t^a)$, with $\delta_t^a(\cdot)$ defined over all θ_t^a. Also, as in previous formula,

$$p_{t+1}^{a*}(\theta_t^b | \theta_t^a, \theta_{t+1}^a, \hat{\theta}_t^a = \delta(\theta_t^a), c_t)$$
$$= \frac{\pi_t^*\left[c_t | \hat{\theta}_t^a = \delta_t^a(\theta_t^a), \theta_t^b\right] p(\theta_t^b | \theta_t^a, \theta_{t+1}^a)}{\sum_{\tilde{\theta}_t^b} \pi_t^*\left[c_t | \hat{\theta}_t^a = \delta_t^a(\theta_t^a), \tilde{\theta}_t^b\right] p(\tilde{\theta}_t^b | \theta_t^a, \theta_{t+1}^a)}. \tag{23}$$

Thus multiplying $p_{t+1}^{a*}(\theta_t^b | \cdot)$ in right-hand side of (22) by the term in the denominator of the right-hand side of (23) yields the numerator of the latter term. In short, multiplying both sides of (21) by the denominator factor yields

$$\sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} V^a(c_{t+1}^a, \theta_{t+1}^a)$$
$$\times \pi_{t+1}^*(c_{t+1} | \hat{\theta}_t^a = \delta_t^a(\theta_t^a), \theta_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t)$$
$$\times \pi_t^*\left[c_t | \hat{\theta}_t^a = \delta_t^a(\theta_t^a), \theta_t^b\right] p(\theta_t^b | \theta_t^a, \theta_{t+1}^a) p(\theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a, \theta_t^b)$$
$$\geq \sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} V^a(c_{t+1}^a, \theta_{t+1}^a)$$
$$\times \pi_{t+1}^*\left(c_{t+1} | \hat{\theta}_t^a = \delta_t^a(\theta_t^a), \theta_t^b, \hat{\hat{\theta}}_t^a, \hat{\hat{\theta}}_{t+1}^a, \theta_t^b, \theta_{t+1}^b, c_t\right)$$
$$\times \pi_t^*\left[c_t | \hat{\theta}_t^a = \delta_t^a(\theta_t^a), \theta_t^b\right] p(\theta_t^b | \theta_t^a, \theta_{t+1}^a) p(\theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a, \theta_t^b), \tag{24}$$

for all announcements $(\hat{\theta}_t^a, \hat{\theta}_{t+1}^a)$. Inequality (24) may be rewritten in the obvious notation as an incentive constraint at date $t+1$,

$$\sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} V^a(c_{t+1}^a, \theta_{t+1}^a) \pi\left[c_t, c_{t+1} | \hat{\theta}_t^a, \theta_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b\right]$$
$$\times p(\theta_t^b, \theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a)$$
$$\geq \sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} V^a(c_{t+1}^a, \theta_{t+1}^a) \pi\left[c_t, c_{t+1} | \hat{\theta}_t^a, \theta_t^b, \hat{\hat{\theta}}_t^a, \hat{\hat{\theta}}_{t+1}^a \theta_t^b, \theta_{t+1}^b\right]$$
$$\times p(\theta_t^b, \theta_{t+1}^b | \theta_t^a, \theta_{t+1}^a), \tag{25}$$

for all paths $\theta_t^a, \theta_{t+1}^a, c_t$, and for all named values $\hat{\theta}_t^a$ and $(\hat{\hat{\theta}}_t^a, \hat{\hat{\theta}}_{t+1}^a)$. Here of course

$$\sum_{c_{t+1}} \pi(c_t, c_{t+1} | \hat{\theta}_t^a, \hat{\theta}_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b)$$
$$= \sum_{c_{t+1}} \pi\left(c_t, c_{t+1} | \hat{\theta}_t^a, \hat{\theta}_t^b, \hat{\hat{\theta}}_t^a, \hat{\hat{\theta}}_{t+1}^a, \hat{\hat{\theta}}_t^b, \hat{\hat{\theta}}_{t+1}^b\right], \tag{26}$$

for all $\hat{\hat{\theta}}_t^a, \hat{\hat{\theta}}_{t+1}^a, \hat{\hat{\theta}}_t^b, \hat{\hat{\theta}}_{t+1}^b$ so that the marginal distribution of c_t is independent of future announcements. Again, inequality (25) ensures agent a would tell the truth in the second period no matter what $\hat{\theta}_t^a$ was announced in the first period, at least on the presumption that agent b follows a truthtelling strategy throughout.

Similarly, with truthtelling established for the second period, so that future parameter announcements are coincident with actual histories and contemporary realizations, eq. (20) reduces for each actual θ_t^a and all named values $\hat{\theta}_t^a$ to an incentive constraint at date t, namely,

$$\sum_{c_t} \sum_{\theta_{t+1}^a} \sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} \left[U^a(c_t^a, \theta_t^a) + \beta V^a(c_{t+1}^a, \theta_{t+1}^a)\right]$$
$$\times \pi\left[c_t, c_{t+1} | \theta_t^a, \theta_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b\right] p(\theta_t^b, \theta_{t+1}^b, \theta_{t+1}^a | \theta_t^a)$$
$$\geq \sum_{c_t} \sum_{\theta_{t+1}^a} \sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} \left[U^a(c_t^a, \theta_t^a) + \beta V^a(c_{t+1}^a, \theta_{t+1}^a)\right]$$
$$\times \pi\left[c_t, c_{t+1} | \hat{\theta}_t^a, \theta_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b\right] p(\theta_t^b, \theta_{t+1}^b, \theta_{t+1}^a | \theta_t^a). \tag{27}$$

In short we are reduced to:

Programming problem 4.1

Find the lotteries $\pi(c_t, c_{t+1} | \hat{\theta}_t^a, \hat{\theta}_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b)$ over consumptions $c_t = (c_t^a, c_t^b)$ and $c_{t+1} = (c_{t+1}^a, c_{t+1}^b)$, given date t parameter announcements

$\hat{\theta}_t^a$ and $\hat{\theta}_t^b$ and date $t+1$ announcements of histories θ_t^a and θ_t^b and contemporary values θ_{t+1}^a and θ_{t+1}^b, which maximize the objective function

$$\lambda^a\Big\{\sum_{\theta_t^a}\sum_{\theta_t^b}\sum_{\theta_{t+1}^a}\sum_{\theta_{t+1}^b}\sum_{c_t}\sum_{c_{t+1}}\left[U^a(c_t^a,\theta_t^a)+\beta V^a(c_{t+1}^a,\theta_{t+1}^a)\right]$$
$$\times\pi\left[c_t,c_{t+1}|\theta_t^a,\theta_t^b,\theta_t^a,\theta_{t+1}^a,\theta_t^b,\theta_{t+1}^b\right]p\left(\theta_t^a,\theta_{t+1}^a,\theta_t^b,\theta_{t+1}^b\right)$$
$$+\lambda^b\Big\{\sum_{\theta_t^a}\sum_{\theta_t^b}\sum_{\theta_{t+1}^a}\sum_{\theta_{t+1}^b}\sum_{c_t}\sum_{c_{t+1}}\left[U^b(c_t^b,\theta_t^b)+\beta V^b(c_{t+1}^b,\theta_{t+1}^b)\right]$$
$$\times\pi\left[c_t,c_{t+1}|\theta_t^a,\theta_t^b,\theta_t^a,\theta_{t+1}^a,\theta_t^b,\theta_{t+1}^b\right]p\left(\theta_t^a,\theta_{t+1}^a,\theta_t^b,\theta_{t+1}^b\right), \tag{28}$$

subject to constraints (25) and (27), to analogue incentive constraints for agent b, and to constraint (26) with

$$\pi\left(c_t,c_{t+1}|\hat{\theta}_t^a,\hat{\theta}_t^b,\theta_t^a,\theta_{t+1}^a,\theta_t^b,\theta_{t+1}^b\right)\geq 0, \tag{29}$$

$$\sum_{c_t}\sum_{c_{t+1}}\pi\left(c_t,c_{t+1}|\hat{\theta}_t^a,\hat{\theta}_t^b,\theta_t^a,\theta_{t+1}^a,\theta_t^b,\theta_{t+1}^b\right)=1. \tag{30}$$

Individual rationality constraints can also be imposed.

Thus far sequential, dynamic programming considerations have been emphasized, with truthtelling at date $t+1$ as a maximizing strategy for agent a for any possible history, and with truthtelling at date t as a maximizing strategy at date t on the premise that the truth will be told in the future. But this is equivalent with saying that if at the beginning of date t agent a contemplated all possible strategies $\delta_t^a(\theta_t^a)$ mapping date t parameter value θ_t^a into date t announcement $\hat{\theta}_t^a$ and at the same time jointly contemplated all possible strategies $\delta_{t+1}^a(\theta_t^a,\theta_{t+1}^a,c_t)$ mapping the specified history into announcement $\hat{\hat{\theta}}_t^a,\hat{\hat{\theta}}_{t+1}^a$, then truthtelling would be a maximizing choice for these strategies. Afterall, the dynamic programming algorithm is just a way to derive optimal strategies. Thus we may write for each actual θ_t^a

$$\sum_{c_t}\sum_{\theta_{t+1}^a}\sum_{\theta_t^b}\sum_{\theta_{t+1}^b}\sum_{c_{t+1}}\left[U^a(c_t^a,\theta_t^a)+\beta V^a(c_{t+1}^a,\theta_{t+1}^a)\right]$$
$$\times\pi\left(c_t,c_{t+1}|\theta_t^a,\theta_t^b,\theta_t^a,\theta_{t+1}^a,\theta_t^b,\theta_{t+1}^b\right)p\left(\theta_t^b,\theta_{t+1}^b,\theta_{t+1}^a|\theta_t^a\right)$$
$$\geq\sum_{c_t}\sum_{\theta_{t+1}^a}\sum_{\theta_t^b}\sum_{\theta_{t+1}^b}\sum_{c_{t+1}}\left[U^a(c_t^a,\theta_t^a)+\beta V^a(c_{t+1}^a,\theta_{t+1}^a)\right]$$
$$\times\pi\left(c_t,c_{t+1}|\hat{\theta}_t^a=\delta_t^a(\theta_t^a),\theta_t^b,\left(\hat{\hat{\theta}}_t^a,\hat{\hat{\theta}}_{t+1}^a\right)=\delta_{t+1}^a(\theta_t^a,\theta_{t+1}^a,c_t),\right.$$
$$\left.\left(\theta_t^b,\theta_{t+1}^b\right)p\left(\theta_t^b,\theta_{t+1}^b,\theta_{t+1}^a|\theta_t^a\right),\right. \tag{31}$$

for all possible strategies $\delta_t^a(\theta_t^a)$ and $\delta_{t+1}^a(\theta_t^a,\theta_{t+1}^a,c_t)$.

In particular, then, truthtelling strategies weakly dominate anything in the more restricted class of strategies under which $\hat{\theta}_t^a$ announcements at date t must necessarily be followed by $(\hat{\theta}_t^a, \hat{\theta}_{t+1}^a)$ announcements at date $t+1$, that is, reannouncements of $\hat{\theta}_t^a$ at $t+1$ coincident with earlier announcement $\hat{\theta}_t^a$ at date t. Under this restriction the parameter announcement $\hat{\theta}_t^a$ merely appears twice, which θ_t^b does also in similar expression. This suggests that one can get by with allocation rule

$$\hat{\pi}(c_t, c_{t+1} | \theta_t^a, \theta_t^b, \theta_{t+1}^a, \theta_{t+1}^b) \equiv \pi(c_t, c_{t+1} | \theta_t^a, \theta_t^b, \theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b),$$

and with announcements of contemporary values *only* each date. In fact, agent a would tell the truth under such an allocation rule, by the restricted version (31), as would agent b by a similar expression, and so the same distribution of consumption allocations would be generated. This leads to the alternative programming problem, now in the form suggested by Myerson (1986),

Programming problem 4.2

Find the lotteries $\hat{\pi}(c_t, c_{t+1} | \theta_t^a, \theta_t^b, \theta_{t+1}^a, \theta_{t+1}^b)$ over consumptions c_t and c_{t+1}, effected by announcements of parameters θ_t^a and θ_t^b at date t and θ_{t+1}^a and θ_{t+1}^b at date $t+1$, which maximize the objective function

$$\begin{aligned}
&\lambda^a \Big\{ \sum_{\theta_t^a} \sum_{\theta_t^b} \sum_{\theta_{t+1}^a} \sum_{\theta_{t+1}^b} \sum_{c_t} \sum_{c_{t+1}} [U^a(c_t^a, \theta_t^a) + \beta V^a(c_{t+1}^a, \theta_{t+1}^a)] \\
&\times \hat{\pi}(c_t, c_{t+1} | \theta_t^a, \theta_t^b, \theta_{t+1}^a, \theta_{t+1}^b) p(\theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b) \Big\} \\
&+ \lambda^b \Big\{ \sum_{\theta_t^a} \sum_{\theta_t^b} \sum_{\theta_{t+1}^a} \sum_{\theta_{t+1}^b} \sum_{c_t} \sum_{c_{t+1}} [U^b(c_t^b, \theta_t^b) + \beta V^b(c_{t+1}^b, \theta_{t+1}^b)] \\
&\times \hat{\pi}(c_t, c_{t+1} | \theta_t^a, \theta_t^b, \theta_{t+1}^a, \theta_{t+1}^b) p(\theta_t^a, \theta_{t+1}^a, \theta_t^b, \theta_{t+1}^b) \Big\},
\end{aligned} \tag{32}$$

subject to constraints, for each actual θ_t^a,

$$\begin{aligned}
&\sum_{c_t} \sum_{\theta_{t+1}^a} \sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} [U^a(c_t^a, \theta_t^a) + \beta V^a(c_{t+1}^a, \theta_{t+1}^a)] \\
&\quad \times \hat{\pi}(c_t, c_{t+1} | \theta_t^a, \theta_t^b, \theta_{t+1}^a, \theta_{t+1}^b) p(\theta_t^b, \theta_{t+1}^b, \theta_{t+1}^a | \theta_t^a) \\
&\geq \sum_{c_t} \sum_{\theta_{t+1}^a} \sum_{\theta_t^b} \sum_{\theta_{t+1}^b} \sum_{c_{t+1}} [U^a(c_t^a, \theta_t^a) + \beta V(c_{t+1}^a, \theta_{t+1}^a)] \\
&\quad \times \hat{\pi}(c_t, c_{t+1} | \gamma_t^a(\theta_t^a), \theta_t^b, \gamma_{t+1}^a(\theta_t^a, \theta_{t+1}^a, c_t), \theta_{t+1}^b) \\
&\quad \times p(\theta_t^b, \theta_{t+1}^b, \theta_{t+1}^a | \theta_t^a),
\end{aligned} \tag{33}$$

for all strategies $\gamma_t^a(\theta_t^a)$ and $\gamma_{t+1}^a(\theta_t^a, \theta_{t+1}^a, c_t)$, and similarly for incentive constraints for agent b.

5. Numerical examples: Optimal learning as optimal scrambling

Numerical computation of solutions to program 4.1 proved tedious due to dimensionality problems. Particularly demanding are the period $t+1$ incentive constraints as these must be conditioned among other things on the date t allocation, c_t. To simplify the problem both θ_t^b and θ_{t+1}^a were set at degenerate values, and this was assumed to be common knowledge. Thus only agent a announces at t, and only agent b announces at $t+1$. Thus reannouncement at date $t+1$ of date t values was no longer necessary. The essential scrambling possibility remained in tact, however, as in announcing θ_{t+1}^b agent b must integrate over θ_t^a values unless θ_t^a is perfectly inferred by the outcome of the lottery over c_t values at date t.

Numerical examples displayed active scrambling, at least for some parameter values. For the example here preferences are defined by

$$U^a(c_t^a, \theta_t^a) = (c_t^a)^{\theta_t^a}, \qquad \theta_t^a \in \{0.2, 0.9\},$$

$$U^b(c_{t+1}^b, \theta_{t+1}^b) = (c_{t+1}^b)^{\theta_{t+1}^b}, \quad \theta_{t+1}^b \in \{0.2, 0.9\},$$

with utility functions of the same form for agent a at $t+1$ and $\theta_{t+1}^a = 0.30$ and for agent b at t with $\theta_t^b = 0.90$. Parameters θ_t^a and θ_{t+1}^b are generated with known prior distribution

$$(\theta_t^a, \theta_{t+1}^b) = \begin{cases} (0.2, 0.9) & \text{with prob. } 0.3, \\ (0.9, 0.9) & \text{with prob. } 0.2, \\ (0.2, 0.2) & \text{with prob. } 0.1, \\ (0.9, 0.2) & \text{with prob. } 0.4. \end{cases}$$

Also let $e_t^a = e_{t+1}^a = e_t^b = e_{t+1}^b = 5$, $\beta = 0.95$, $\lambda^a = 0.5$ and $\lambda^b = 0.5$. The grid on the triangle of feasible consumptions is set at increments of 0.25, computed by successive refinements. To mitigate remaining dimensionality problems, no individual rationality constraints are imposed. The solution is displayed in the obvious notation in table 10. To be noted here is that the allocation $c_t = (1.75, 8.25)$ occurs both for $\theta_t^a = 0.2$ (with prob. 1) and $\theta_t^a = 0.9$ (with prob. 0.0339681) so that at $\theta_t^a = 0.2$, for example, agent b remains uncertain of a's true parameter value.

Revealing of the role being played by this scrambling is the Pareto optimal solution when announcements of agent a at date t are necessarily public. In

Table 10

Optimal scrambling.

θ_t^a	(c_t^a, c_t^b)	$\pi(c_t, \theta_t^a)$	$\theta_t^a, \theta_{t+1}^b$	(c_{t+1}^a, c_{t+1}^b)	$\pi(c_{t+1}, \theta_t^a, \theta_{t+1}^b)$
0.2	(1.75, 8.25)	1.0	(0.2, 0.2)	(4.75, 5.25)	1.0
			(0.2, 0.9)	(2.0, 8.0)	1.0
0.9	(0.0, 10.0)	0.1159346	(0.9, 0.2)	(3.75, 6.25)	1.0
	(1.75, 8.25)	0.0339681			
	(3.25, 6.75)	0.85443844			
			(0.9, 0.9)	(1.0, 9.0)	0.8610602006
				(10.0, 0.0)	0.138397942

that case there is an incentive constraint for agent b at $t+1$ for every θ_t^a announcement. The solution is reported in table 11.

Clearly with θ_t^a announcements public, there is no longer a gain to additional lotteries over c_t values, and the $c_t = (1.75, 8.25)$ possibility at $\theta_t^a = 0.9$ of table 10 is suppressed in table 11. Striking also now, in table 11, is the absence of insurance for agent b at date $t+1$ in the $\theta_t^a = 0.2$ branch, with the shock invariant $c_{t+1} = (2.5, 7.5)$ of table 11 replacing (4.75, 5.25) and (2.0, 8.0) for $\theta_{t+1}^b = 0.2$ and $\theta_{t+1}^b = 0.9$, respectively, in table 10. Evidently, the lottery in table 11 which does distinguish $\theta_{t+1}^b = 0.2$ and $\theta_{t+1}^b = 0.9$ for the $\theta_t^a = 0.9$ branch is not worthwhile in table 11 at the $\theta_t^a = 0.2$ branch, a region in which agent b has relatively low consumption. Without this lottery, as under $\theta_t^a = 0.2$ in table 11, consumption at date $t+1$ cannot depend on agent b's announcement. But since agent b is kept somewhat uninformed about what θ_t^a branch he is in table 10, agent b is never put in a position of choosing between the 5.25 and 8.0 entries of table 10, when $\theta_t^a = 0.2$ and he is to report θ_{t+1}^b, but rather chooses between those points the outcomes at $\theta_t^a = 0.9$ branch, where there is a lottery. Evidently, with somewhat higher expected consumption, lotteries can distinguish effectively.

Table 11

Scrambling suppressed with public announcements.

θ_t^a	(c_t^a, c_t^b)	$\pi(c_t, \theta_t^a)$	$\theta_t^a, \theta_{t+1}^b$	(c_{t+1}^a, c_{t+1}^b)	$\pi(c_{t+1}, \theta_t^a, \theta_{t+1}^b)$
0.2	(1.75, 8.25)	1.0	(0.2, 0.2)	(2.5, 7.5)	1.0
			(0.2, 0.9)	(2.5, 7.5)	1.0
0.9	(0.0, 10.0)	0.1042598	(0.9, 0.2)	(3.75, 6.25)	1.0
	(3.25, 6.75)	0.8957402			
			(0.9, 0.9)	(1.0, 9.0)	0.89189814
				(10.0, 0.0)	0.10810186

6. A capital accumulation environment with endogenously limited insurance

To move away from the pure risk sharing environments considered above consider an economy with one underlying consumption good, two consumption periods, and a continuum of households with names on the unit interval. Each household is subjected to preferences shocks θ_t and θ_{t+1} in the first and second consumption periods, respectively, with prob(θ_t) denoting both the probability of shock θ_t and the fraction of households who experience shock θ_t and similarly at date $t+1$. Shocks θ_t and θ_{t+1} are drawn independently with beginning of period realizations known only to the individual household. Thus preference shocks are the source of private information as above. Each household has endowment $y_t(\varepsilon_t)$ of the consumption good at date t, where ε_t is a common, publicly observed shock. For simplicity $y_{t+1}(\varepsilon_{t+1})$ is supposed to be zero. Finally, the consumption good can be stored in the first period, after the realization of shock ε_t, say in amount $K_{t+1}(\varepsilon_t)$ per capita, yielding $(1-\delta)K_{t+1}(\varepsilon_t)$ units of consumption available per capita in the second period.

To set up the programming problem for the determination of Pareto optimal allocations let $\pi_t(c_t, \theta_t, \varepsilon_t)$ denote the probability at date t of consumption bundle c_t at date t for the representative household conditioned on announced (and actual) preference shock θ_t and on the common endowment shock ε_t. Similarly, let $\pi_{t+1}(c_{t+1}, \theta_t, \theta_{t+1}, \varepsilon_t)$ denote the probability at date $t+1$ of consumption bundle c_{t+1} at date $t+1$ for the representative household conditioned on its announcements of θ_t and θ_{t+1} and on the shock ε_t at date t. An implicit restriction is thus imposed, that prob$(c_{t+1}|c_t, \theta_t, \theta_{t+1}, \varepsilon_t)$ cannot depend on c_t. This simplification allowed computations to be done as before on a personal computer. Of course possible time dependence of the lottery at date $t+1$ on shock θ_t is still incorporated.

With β as the common discount rate and $U[c, \theta]$ as the period-by-period utility function of the representative household, the relevant program from the standpoint of a $t=0$ planning period is:

Programming problem 6.1

Maximize the objective function

$$\sum_{\varepsilon_t} \text{prob}(\varepsilon_t)\Big\{\sum_{\theta_t} \text{prob}(\theta_t) \sum_{c_t} \pi_t(c_t, \theta_t, \varepsilon_t) U[c_t, \theta_t]$$

$$+\beta \sum_{\theta_t} \sum_{\theta_{t+1}} \text{prob}(\theta_t, \theta_{t+1}) \sum_{c_{t+1}} \pi_t(c_{t+1}, \theta_t, \theta_{t+1}, \varepsilon_t) U[c_{t+1}, \theta_{t+1}]\Big\},$$

subject to the date t resource constraint, $\forall \varepsilon_t$,

$$\sum_{\theta_t} \text{prob}(\theta_t) \sum_{c_t} c_t \pi(c_t, \theta_t, \varepsilon_t) = y_t(\varepsilon_t) - K_{t+1}(\varepsilon_t).$$

the date $t+1$ resource constraint, $\forall \varepsilon_t$,

$$\sum_{\theta_t} \text{prob}(\theta_t) \sum_{\theta_{t+1}} \text{prob}(\theta_{t+1}) \sum_{c_{t+1}} c_{t+1}\pi(c_{t+1}, \theta_t, \theta_{t+1}, \varepsilon_t)$$
$$= (1-\delta)K_{t+1}(\varepsilon_t),$$

subject to the date $t+1$ incentive constraints, $\forall \theta_t, \theta_{t+1}, \theta'_{t+1}, \varepsilon_t$,

$$\sum_{c_{t+1}} U(c_{t+1}, \theta_{t+1})\pi(c_{t+1}, \theta_t, \theta_{t+1}, \varepsilon_t)$$
$$\geq \sum_{c_{t+1}} U(c_{t+1}, \theta_{t+1})\pi(c_{t+1}, \theta_t, \theta'_{t+1}, \varepsilon_t),$$

subject to the date t incentive constraints, $\forall \theta_t, \theta'_t, \varepsilon_t$,

$$\sum_{c_t} U(c_t, \theta_t)\pi(c_t, \theta_t, \varepsilon_t) + \sum_{\theta_{t+1}} \beta \, \text{prob}(\theta_{t+1})$$
$$\times \sum_{c_{t+1}} U(c_{t+1}, \theta_{t+1})\pi(c_{t+1}, \theta_t, \theta_{t+1}, \varepsilon_t)$$
$$\geq \sum_{c_t} U(c_t, \theta_t)\pi(c_t, \theta'_t, \varepsilon_t) + \sum_{\theta_{t+1}} \beta \, \text{prob}(\theta_{t+1})$$
$$\times \sum_{c_{t+1}} U(c_{t+1}, \theta_{t+1})\pi(c_{t+1}, \theta'_t, \theta_{t+1}, \varepsilon_t),$$

where of course

$$0 \leq \pi(c_t, \theta_t, \varepsilon_t), \qquad \forall \theta_t, \varepsilon_t, c_t,$$

$$\sum_{c_t} \pi(c_t, \theta_t, \varepsilon_t) = 1, \qquad \forall \theta_t, \varepsilon_t,$$

$$0 \leq \pi(c_{t+1}, \theta_t, \theta_{t+1}, \varepsilon_t), \qquad \forall c_{t+1}, \theta_t, \theta_{t+1}, \varepsilon_t,$$

$$\sum_{c_{t+1}} \pi(c_{t+1}, \theta_t, \theta_{t+1}, \varepsilon_t) = 1, \quad \forall \theta_t, \theta_{t+1}, \varepsilon_t.$$

For numerical examples let $U(c,\theta) = [c^{\theta} - 1]/\theta$ with $\theta \in \{0.5, 0.9\}$ drawn each period with probability $\frac{1}{2}$ for each value. All consumptions were restricted to lie between 0 and 22.5 and on a grid with increments of 0.05. For $\beta = 1$ and $\delta = 0$ and for the $y_t(\varepsilon_t) = 8$ branch (only) the solution is depicted in table 12. In contrast, the full information storage solution for the $y_t(\varepsilon_t) = 8$ branch (only) is depicted in table 13.

Table 12
Private information, optimal solution – under storage.

θ_t	c_t	$\pi(c_t, \theta_t, \varepsilon_t)$	θ_t, θ_{t+1}	c_{t+1}	$\theta_{t+1}(c_{t+1}, \theta_t, \theta_{t+1})$
0.5	2.05	1.0	(0.5,0.5)	3.5	1.0
			(0.5,0.9)	{ 0.0	0.4553 }*
				11.8 }	0.5447
0.9	7.45	1.0*	(0.9,0.5)	1.55	1.0
			(0.9,0.9)	1.55	1.0

The full information solution offers full insurance against θ_t shocks period by period (with no intertemporal dependence). The private information solution offers less insurance against θ_t and θ_{t+1} shocks, especially in the second period. Apparently, as a result of this, the expected utility for a given second period endowment is less under private information than under full information, and the storage solutions reflect this, with storage half of the first period endowment for the full information solution, 4 and 5 for $y_t(\varepsilon_t) = 8$ and $y_t(\varepsilon_t) = 10$, respectively, and at 3.26 and 3.90 respectively for the private information solution. (The coefficient of variation of storage is greater for the full information solution than for the private information solution.)

However, setting $(1-\delta) = 0.8$ and $\beta = 0.9$ yields a different qualitative solution. With storage less productive and the future discounted more, the full information solution for storage is now 0.825 and 0.9375, for $y_t(\varepsilon_t) = 8$ and $y_t(\varepsilon_t) = 10$, respectively. Under private information the storage solution is 2.0 and 2.5, respectively, uniformly *more*. Intuition is aided by table 14, for $y_t(\varepsilon_t) = 8$ branch only, which again reveals a strong intertemporal dependence: $\theta_t = 0.5$ households receive less (expected) consumption at date 1 and more consumption at date 2 than do $\theta_t = 0.9$ households. It is this tie-in of date $t+1$ consumption to date t claims that allows the incentive constraint to be satisfied, and apparently storage is increased in the private information solution to allow date t claims to have more bite. (The coefficient of variation

Table 13
Full information, optimal storage solution.

θ_t	c_t	$\pi(c_t, \theta_t, \varepsilon_t)$	θ_t, θ_{t+1}	c_{t+1}	$\pi_{t+1}(c_{t+1}, \theta_t, \theta_{t+1})$
0.5	1.45	1.0	(0.5,0.5)	1.45	1.0
			(0.5,0.9)	6.55	1.0
0.9	6.55	1.0	(0.9,0.5)	1.45	1.0
			(0.9,0.9)	6.55	1.0

Table 14

Private information, optimal solution – over storage

θ_t	c_t	$\pi(c_t, \theta_t, \varepsilon_t)$	θ_t, θ_{t+1}	c_{t+1}	$\pi_{t+1}(c_{t+1}, \theta_t, \theta_{t+1})$
0.5	3.25	1.0	(0.5,0.5)	2.4	1.0
			(0.5,0.9)	{0.0 4.25	0.2484 0.7516} *
0.9	{0.0 9.8	0.1067 0.8933}	(0.9,0.5) (0.9,0.9)	0.4 0.4	1.0 1.0

of storage is less for full information than for private information, again the opposite of before.)

7. Concluding remarks

It seems we are only at the beginning of efforts to understand the implications of private information for the operation of actual economies. For prototype economies for which revelation principle methods seem most secure, we need to build up experience and intuition, aided by numerical experimentation. And, we should seek to extend existing theoretical and numerical methods not only to more complicated private information environments but also to environments with limited commitment problems and other impediments to trade. The intent should be the development of a widely operational mapping from the supposed primitives of environments to allocations and institutions. In doing so we may seek to explain allocations and institutions and have on hand a reliable base for policy analysis.

References

Altug, Sumru J., 1984, Gestation lags and the business cycle: An empirical analysis, Ph.D. dissertation (Carnegie-Mellon University, Pittsburgh, PA).

Baiman, Stanley and Joel Demski, 1980, Economically optimal performance evaluation and control systems, Journal of Accounting Research 18, suppl., 184–220.

Bernanke, Ben and Mark Gertler, 1986, Financial efficiency, collateral and business fluctuations, Manuscript.

Bhattacharaya, Sudipto, 1984, Tournaments and incentives: Heterogeneity and essentiality, Manuscript.

Demski, Joel S. and David Sappington, 1984, Optimal incentive contracts with multiple agents, Journal of Economic Theory 33, 152–177.

Evans III, J., 1980, Optimal contracts with costly conditional auditing, Journal of Accounting Research 18, suppl., 108–128.

Gale, Douglas and Martin Hellwig, 1985, Incentive-compatible debt contracts: The one period problem, Review of Economic Studies 52, 627–645.

Green, Edward, 1984, Lending and the smoothing of unobservable income, Manuscript.

Greenwald, Bruce and Joseph Stiglitz, 1986, Information, finance constraints, and business fluctuations, NBER manuscript.

Hansen, Gary, 1985, Indivisible labor and the business cycle, Journal of Monetary Economics 6, 309–327.

Harris, Milton and Robert Townsend, 1981, Resource allocation under asymmetric information, Econometrica 49, 33–69.

Kydland, F.E. and E.C. Prescott, 1982, Time to build and aggregate fluctuations, Econometrica 50, 1345–1370.

Mookherjee, Dilip and Ivan Png, 1987, Optimal auditing, insurance, and redistribution, Graduate School of Business research paper no. 887, April (Stanford University, Stanford, CA).

Myerson, Roger, 1979, Incentive compatibility and the bargaining problem, Econometrica 47, 61–74.

Myerson, Roger, 1986, Multistate games with communication, Econometrica 54, 323–358.

Palfrey, Thomas and Sanjay Srivastava, 1987, Mechanism design with incomplete information: A solution to the implementation problem. Manuscript.

Postlewaite, Andrew and David Schmeidler, 1986, Implementation in differential information economies, Journal of Economic Theory 39, 14–33.

Prescott, Edward and Robert Townsend, 1984, General competitive analysis in an economy with private information, International Economic Review 25, 1–20.

Reinganum, J. and L. Wilde, 1985, Income tax compliance in a principal–agent framework, Journal of Public Economics 26, 1–18.

Repullo, Rafael, 1983, Implementation by direct mechanisms under incomplete information, Working paper (London School of Economics, London).

Smith, Bruce, 1985, A business cycle model with private information, Revised manuscript.

Townsend, Robert M., 1979, Optimal contracts and competitive markets with costly state verification, Journal of Economic Theory 21, 265–293.

Townsend, Robert M., 1987, Economic organization with limited communication, American Economic Review 77, 954–971.