

The Review of Economic Studies Ltd.

Computing Multi-Period, Information-Constrained Optima

Author(s): Christopher Phelan and Robert M. Townsend

Source: *The Review of Economic Studies*, Vol. 58, No. 5 (Oct., 1991), pp. 853-881

Published by: [The Review of Economic Studies Ltd.](#)

Stable URL: <http://www.jstor.org/stable/2297941>

Accessed: 08/12/2010 08:28

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=resl>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The Review of Economic Studies Ltd. is collaborating with JSTOR to digitize, preserve and extend access to *The Review of Economic Studies*.

<http://www.jstor.org>

Computing Multi-Period, Information-Constrained Optima

CHRISTOPHER PHELAN and ROBERT M. TOWNSEND
University of Chicago

First version received September 1989; final version accepted October 1990 (Eds.)

This paper presents a detailed theoretical derivation and justification for methods used to compute solutions to a multi-period (including infinite-period), continuum-agent, unobserved-effort economy. Actual solutions are displayed illustrating cross-sectional variability in consumption and labour effort in the population at a point in time and variability for a typical individual over time. The optimal tradeoff between insurance and incentives is explored and the issue of excess variability is addressed by consideration of the analogue full-information economy and various restricted-contracting regimes.

1. INTRODUCTION

A strong but troublesome prediction of stochastic general equilibrium macroeconomic models is the extreme level of risk-sharing. Even in macroeconomic models with endogenously derived population diversity in labour effort (e.g. Rogerson (1988), Hansen (1985)) there is, nevertheless, full consumption insurance. That is, an individual's consumption does not depend on his labour effort or the resulting productivity of that effort. In the absence of either aggregate shocks or individual preference shocks this implies constant consumption for each individual over time. In the absence of preference or initial wealth or Pareto-weight diversity this implies equal consumption levels across individuals at a point in time. Such predictions seem implausible.

Introducing unobserved actions or other "incentive problems" has long held promise of helping to reconcile stochastic general equilibrium macro models with apparently incomplete insurance. The idea, simple enough, is that full consumption insurance has adverse incentive effects on effort. Thus consumption dependent on labour effort or labour productivity may be desirable *ex ante*. Unfortunately, the characterization of solutions to incentive-constrained economies, and especially dynamic incentive-constrained economies, has proved difficult. Analytic treatments have tended to provide disappointingly weak characterizations, even with unintuitive and stringent restrictions on technologies and preferences. Analytic treatments of repeated incentive-constrained economies have been hampered by the history-dependent nature of solutions.¹

Some recent contributions have made important progress. Perhaps the most salient among these is Fudenberg, Holmstrom, and Milgrom (1986) which shows that under some assumptions there need be no history-dependence and that an optimal long-term agreement can be replicated by a series of single-period contracts. These assumptions are (1) exponential utility in consumption for the agent, (2) uncontrolled access to credit on the part of the agent, as if this were unobserved, and (3) a disutility cost for action

1. See Rogerson (1985) and Townsend (1982) for a proof of the necessity of such history-dependence.

which is like a financial cost and hence subtracted linearly from consumption. As they note, the ability to control consumption is crucial, and much of the earlier literature is exploiting this possibility. The exponential utility function removes wealth effects, and this also seems to be crucial. In any event, our paper should be viewed as a complementary effort; we are able to display solutions without any of these restrictions.

Our alternative to further search for analytic theorems is to present and justify a method for computing solutions to a repeated, continuum-agent version of the classic agency problem. We believe the method we present is sufficiently general to readily extend to a less restricted set of incentive constrained frameworks. Again, such methods are important because in the absence of analytic solutions, computing solutions is the only way to discover the characteristics of a model and the only way to allow for the eventual comparison of a model to data.

Any method for computing solutions must be formally justified. Specifically for us, this means proving we have defined a problem for which a solution exists, that our reduction to a dynamic programming framework is valid, that our computation method is guaranteed to converge, and that this converged solution is the solution to the original problem. This paper presents a sequence of theorems to that effect.

We also explore the characteristics of the classic prototype agency model by displaying computed solutions to it and by contrasting these to the solutions of the analogue full-information economy. In full-information solutions, consumption and labour efforts (actions) are constant for individuals over time (full insurance) and population distributions are degenerate if there is no initial diversity. As foreshadowed above, computed solutions to the incentive-constrained model overturn these predictions. Computed solutions are also contrasted with solutions (also computed) to the agency problem restricted to one-period contracts. This restriction induces greater consumption variability and lower average output than the incentive-constrained solution without this restriction. We thus illustrate the (computed) optimal tradeoff between insurance and incentives and the effect of multi-period contracting on such environments.

This work is derived mainly from three somewhat distinct lines of research in the literature. One line is represented by Prescott and Townsend (1984*a, b*), who demonstrate the ability to analyze incentive-constrained economies in the space of measures of economic variables. Relatedly, Townsend (1987, 1988) numerically simulates a wide variety of static or short-horizon incentive-constrained economies using linear programming techniques.

Another line of research is represented by Abreu, Pearce, and Stacchetti (1986), and by Spear and Srivastava (1987) who provide the key insight into the nature of history-dependent solutions needed to yield a computationally-feasible repeated formulation of the problem. Abreu, Pearce, and Stacchetti show for an unobserved-action, cartel problem that a given history can be sufficiently represented by the expected profit of the cartel members from the given time on. This gives a state variable similar to the amount of capital in a growth problem. Spear and Srivastava extend this to the infinite-horizon principal-agent problem where the state variable is the expected utility of the agent.

Lastly, Green (1987) analyzes a closely related continuum-agent endowment economy with private information concerning the individual endowments. In our conclusion, we argue that our methods apply to this economy as well. This paper builds directly on Spear and Srivastava, but extends the principal-agent problem to a continuum of agents, and uses a formulation consistent with the measure-space techniques of Prescott and Townsend (1984*a, b*). These steps allow the numerical computation of solutions. We borrow from Green the idea and justification of maximizing a discounted social surplus

in a context in which society has access to an unlimited linear borrowing-lending technology. Finally, we emphasize here a sequence of theorems which justify the recursive formulation of our framework.

We will focus heavily on the relationships among the full-information (or fully-observed-action) economy, the static unobserved-action economy, and the repeated unobserved-action economy. Section 2 introduces the full-information problem, in our somewhat unorthodox, more general manner of stating the problem, with lotteries. Section 3 alters this environment to the static unobserved-action economy and then discusses and displays the nature of solutions. Section 4 considers the repeated version of the agency problem and displays and discusses computed solutions. Section 5 gives a sequence of theorems justifying the recursive formulation used to solve the repeated problems. Section 6 gives an overview of the computational methods used, and Section 7 offers a conclusion.

II. THE FULL-INFORMATION ECONOMY

In this section, we consider the social planning problem for an economy with a continuum of identical agents each having an identical but independent production technology taking in an agent's own labour and producing the single consumption good as a function of this labour and an independent shock. Our method of finding the Pareto optima for the economy is to maximize social surplus subject to each of the agents receiving exactly a prespecified ex ante expected utility. Social surplus is defined as the total amount of the consumption good that is produced but not eaten. A Pareto optimum for this economy is then a solution to this problem for an initial distribution of required utilities which gives a non-negative surplus, so that it is feasible, and has the property that any other distribution of initial utilities that Pareto-dominates the initial distribution has a negative surplus and thus is not feasible.²

The social contract or planning problem assigns to each agent an action, a , from a finite set of possible actions $A \subset \mathbb{R}_+$.³ We assume the contractual rule for determining this action may be probabilistic. Such an action for a given agent results in the realization of an output quantity q , an element of a finite set $Q \subset \mathbb{R}_+$. For a given action, $a \in A$, output is determined by an exogenous probability $P(q|a)$. Probability $P(q|a)$ is assumed strictly positive so that any output is possible given any action. No other restrictions are put on the technology such as likelihood ratio or convexity conditions common in the literature. The social contract, again allowing for mixed outcomes, then gives each agent some consumption amount $c \in C$ where C is again some finite subset of \mathbb{R}_+ .

Mixed or probabilistic outcomes for the individual are assumed to be possible without having uncertainty in the aggregate due to the continuum of agents. This is justified formally in Judd (1985).⁴ Indeed, all probability measures used in our discussion, here and below, are viewed at the aggregate level as representing population fractions and at the individual level as representing probabilities.

The utility function for each agent is denoted by $U(a, c) : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$. For given $A \times C \subset \mathbb{R}_+ \times \mathbb{R}_+$, $U(a, c)$ is assumed bounded on $A \times C$. Further, for all else equal, the agent prefers lower action and higher consumption. We need not assume separability in

2. Green uses and justifies this method as the dual programming problem to that of maximizing utilities subject to an initial wealth constraint.

3. Sets are denoted by upper-case bold type. Subsets are denoted by lower case bold type.

4. Judd proves the existence of a measure assigning joint-output probabilities over all agents taking a given action which is consistent with independence, the individual measures $P(q|a)$, and where a law of large numbers can be stated and holds.

the a, c components except to allow quicker computation methods. Preferences towards risk are characterized by expected utility.

Lastly, we make additional assumptions when discussing the nature of solutions. First, we assume that $U(a, c)$ is strictly concave over changes in c , holding the action a constant. Second, we assume that higher actions imply greater expected output. These assumptions do not affect the validity of our methods of finding solutions.

The lowest ex ante utility level possible for an agent is that of receiving with certainty the lowest consumption in C and the highest labour assignment in A . We denote this utility by w . The highest utility is that of receiving with certainty the highest consumption amount and the lowest labour. This utility is denoted by \bar{w} . Any utility level w between w and \bar{w} and hence representable as $\alpha w + (1 - \alpha)\bar{w}$, for some number α , $0 \leq \alpha \leq 1$, is possible by giving fraction of α of those required to receive w the plan yielding \hat{w} specified above and fraction $(1 - \alpha)$ the plan yielding \bar{w} specified above. As fractions are viewed as probabilities by the individual agents, expected utility for the individual would be $w = \alpha w + (1 - \alpha)\bar{w}$ as required. Thus the set of possible ex ante expected utilities is an interval, $W = [w, \bar{w}]$. We assume that the initial distribution of required utilities in the population has finite support and define $d_0(w)$ to be the fraction of agents whose initial required utility is w . Object d_0 is then the associated distribution of exact ex ante expected utility levels as opposed to lower bounds on ex ante expected utility.

For each utility such that $d_0(w) > 0$ the choice variables for society in the planning problem can be written as $\Pi^w(a, q, c)$: the probability for an agent required to receive w of taking action a , having output q occur in his own production technology and receiving consumption amount c . Object Π^w thus defines a probability measure. Again, $\Pi^w(a, q, c)$ is also the fraction of those agents promised w who will be assigned action a , get output q , and receive consumption c . For a given $w \in W$, we define a *contract* as such a function Π^w which satisfies the following constraints.

First, that the discounted expected utility for the agents required to get w is actually w , or,

$$C1. \quad w = \sum_{A \times Q \times C} U[a, c] \Pi^w(a, q, c).$$

Second, since a given probability measure Π^w implies conditional probabilities of outputs given an action but is nevertheless a choice object, it must be constrained so that these implied conditional probabilities coincide with the ones imposed by nature, namely $P(q|a)$. That is, for all $(\bar{a}, \bar{q}) \in A \times Q$,

$$C2. \quad \sum_C \Pi^w(\bar{a}, \bar{q}, c) = P(\bar{q}|\bar{a}) \sum_{Q \times C} \Pi^w(\bar{a}, q, c).$$

Lastly, we require that Π^w actually represent a probability measure, or,

$$C3. \quad \sum_{A \times Q \times C} \Pi^w(a, q, c) = 1 \quad \text{and for all } (a, q, c) \in A \times Q \times C, \Pi^w(a, q, c) \geq 0.$$

For a given distribution d_0 over required ex ante expected utility $w \in W$, we define an allocation as a collection of contracts for each w in the support $d_0(w)$. The planning problem is then to maximize total social surplus by separately maximizing the surplus from each utility group. This separation is valid because constraints C1 through C3 must hold separately for each $w \in W$. The choice of one contract Π^w does not affect the constraint set for choosing the contract for another ex ante utility w . Again, since one is just determining fractions in a population, the surplus from each utility group and thus total social surplus is a non-random number—there is no aggregate uncertainty. We formally define the programming problem for finding surplus-maximizing contracts as

Full Information Problem (FIP). Maximize by choice of probability measure Π^w the objective function, $s(w) = \sum_{A \times Q \times C} (q - c) \Pi^w(a, q, c)$ subject to Π^w satisfying C1 through C3.

Call solutions to these problems Π^{w*} with corresponding optimized surplus values $s^*(w)$. Total social surplus for initial distribution d_0 is then $S^*(d_0) = \sum_w s^*(w) d_0(w)$. An initial distribution of ex ante expected utilities d_0 is feasible if $S^*(d_0) \geq 0$.

Lastly, we must specify when solutions to the above programming problem correspond to Pareto optima. An initial distribution of utilities d_0 and its associated surplus maximizing plans $\{\Pi^{w*}\}_{w \in W}$ represent a Pareto optimum if the support of d_0 lies within the non-increasing portion of $s^*(w)$.

The first step in the argument is to show that the function $s^*(w): W \rightarrow \mathbb{R}_+$ is weakly concave. Linear combinations (where the weights add to unity) of Π functions satisfying C2 and C3 themselves satisfy C2 and C3, and satisfy C1, the utility constraint, where w is the interpolated ex ante utility for the agent. Thus this linear combination is a feasible plan for delivering this interpolated w . Since the linear combination yields the interpolated surplus, the optimized surplus function $s^*(w)$ is at least linear.

Given concavity, $s^*(w)$ has at most one maximum. For an allocation to be Pareto optimal, we need only to require that all initial required utilities are on the downward-sloping portion of $s^*(w)$, and that aggregate surplus is zero. That is, it is not possible to raise the utility of a subset of agents without lowering aggregate surplus and thus violating aggregate feasibility. Formally, an initial distribution of utilities $d_0(w)$ and its associated surplus maximizing plans $\{\Pi^{w*}\}_{w \in W}$ represent a Pareto optimum if aggregate surplus is zero and all points in the support of $d_0(w)$ are weakly greater than the argmax of $s^*(w)$.⁵

It is important to note that for a given w , FIP is a *linear* programme. The expression $(q - c)$ in the objective function and expressions $P(q|a)$ and $U[a, c]$ in the constraints are simply coefficients on the choice variables $\Pi^w(a, q, c)$. This allows FIP to be solved numerically for each $w \in W$ using a standard revised simplex algorithm.

Solutions to FIP display the full-insurance characteristics common to full-information programmes. If a given contract specified large variations in leisure or consumption among a set of agents required to receive a given ex ante utility, another contract with slightly less mean leisure or consumption but less variability would keep the ex ante expected utilities of the agents constant while raising the social surplus. Thus variation will exist only where such tightenings are not possible, that is when further tightenings are impeded by the finiteness of sets A and C .

For example, suppose there is only one action (full leisure), the consumption set C consists of the integers $\{0, 1, 4, 5\}$, and $U(a, c) = c^{0.5}$. If all agents are required to receive an ex ante expected utility of 1.5, then some uncertainty in consumption is required to deliver this ex ante utility. Nevertheless, the curvature of the utility function will make it optimal to mix between the adjacent points 1 and 4 (with 0.5 probability on each). This is the smallest possible spread.

If $Q = \{q, \bar{q}\}$ and each of these outputs occurs with 0.5 probability, then one solution to this FIP is to set $\Pi^{1.5}(0, q, 1) = 0.5$, and $\Pi^{1.5}(0, \bar{q}, 4) = 0.5$ (all other $\Pi^{1.5}(0, q, c) = 0$). That is, give high consumption to those with high output. However, it is also a solution to set $\Pi^{1.5}(0, q, 4) = 0.5$, and $\Pi^{1.5}(0, \bar{q}, 1) = 0.5$ (all other $\Pi^{1.5}(0, q, c) = 0$). That is, give high consumption to those with low output. It is also optimal to set $\Pi^{1.5}(0, q, 1) = 0.25$,

5. In the full-information programme, the argmax of $s^*(w)$ is w . This property does not hold when incentive constraints are introduced.

$\Pi^{1.5}(0, \bar{q}, 1) = 0.25$, $\Pi^{1.5}(0, \underline{q}, 4) = 0.25$, and $\Pi^{1.5}(0, \bar{q}, 4) = 0.25$ (all other $\Pi^{1.5}(0, q, c) = 0$). Specifically, there will be at most coincidental dependence of an agent's consumption on his output. Further all variation will diminish as more and more intermediate points are added to C . Recall again that aggregate output for such a group of agents is predetermined by specified actions—there is no aggregate risk and thus no need for individuals to bear any.

III. THE STATIC UNOBSERVED-ACTION ECONOMY.

The static unobserved-action economy is identical to the full information economy except that an agent's action is unobservable by everyone other than the agent himself. The social contract can recommend actions, but since actions are unobserved, each agent can contemplate deviating from the action recommended for him. The contract-design problem for a given w thus is *identical* to the full-information problem *except* for additional constraints requiring that obeying the action recommendation is always weakly preferred.⁶

For a given $w \in W$, these additional constraints require that once the recommended action a is announced, that the expected utility of obeying the recommendation is greater than that of each possible deviation. These constraints take the general form for all assigned and possible alternative action pairs $(a, \hat{a}) \in A \times A$,

$$\sum_{Q \times C} U[a, c] \{ \Pi^w(c|q, a) P(q|a) \} \geq \sum_{Q \times C} U[\hat{a}, c] \{ \Pi^w(c|q, a) P(q|\hat{a}) \},$$

where $\Pi^w(c|q, a)$ is the conditional probability implied by $\Pi^w(a, q, c)$. The expression $\{ \Pi^w(c|q, a) P(q|a) \}$ is the probability of a given (q, c) combination given that action a is recommended and that this action is taken. Likewise, the expression $\{ \Pi^w(c|q, a) P(q|\hat{a}) \}$ is the probability of a given (q, c) combination given that action a is announced and deviation action \hat{a} is taken instead. Given that $\Pi^w(q, c|a) = \Pi^w(c|q, a) P(q|a)$ we can solve for $\Pi^w(c|q, a)$, substitute this on each side of the incentive constraint, and simplify to get

$$\sum_{Q \times C} U[a, c] \Pi^w(q, c|a) \geq \sum_{Q \times C} U[\hat{a}, c] \frac{P(q|\hat{a})}{P(q|a)} \Pi^w(q, c|a).$$

Finally, one can multiply both sides of this by the marginal probability of action a , $\Pi^w(a)$ (where this derived from $\Pi^w(a, q, c)$ by summing over q and c). $\Pi^w(a)$ is either zero, in which case the incentive constraint does not matter because a is never recommended, or positive, in which case the inequality is essentially unaltered. This yields

$$C4. \quad \sum_{Q \times C} U[a, c] \Pi^w(a, q, c) \geq \sum_{Q \times C} U[\hat{a}, c] \frac{P(q|\hat{a})}{P(q|a)} \Pi^w(a, q, c).$$

Ratio $P(q|\hat{a})/P(q|a)$ gives how many more times likely it is that output q will occur given deviation action \hat{a} as opposed to recommended action a , and thus updates the joint probability of observing recommended action a , output q , and consumption c .

6. Another difference is that because of these constraints the set of possible ex ante expected utilities W has a higher lower bound. Since actions must be induced, the lowest possible ex ante expected utility for the unobserved action economy is that of receiving the lowest consumption and the *lowest* labour amount: at least this utility can be achieved unilaterally by any agent under any contract specifying consumptions c as a function of outputs q simply by taking the lowest action. For any utility lower than this, every plan will violate C4 (see below). It is important to note that this is not a participation constraint but simply a consequence of assuming unobserved actions. Again, for reasons given earlier all utilities between this lower bound and the previous upper bound on ex ante utility are possible through randomization schemes.

The most significant change in the character of solutions to the unobserved action problem as opposed to solutions to the full-information problem is that if the full-information solution specifies any action other than the lower bound on actions, then the full-insurance outcome violates an incentive constraint. If the consumption of an individual does not depend on the output of that individual, as in the full-information solution, then it is always preferable for this agent to take a lower action; this raises his utility directly and does not affect his consumption, even though such an action would change the probabilities of his own outputs. Thus, in general, solutions will not display full insurance.

Figures 1 through 4 display solutions to both the full information and unobserved action planning problems. The period-by-period utility function for each agent is $U[a, c] \equiv [(c^{0.5}/0.5) + ((1-a)^{0.5}/0.5)]$. The set of feasible actions, A , equals $\{0, 0.2, 0.4, 0.6\}$, and possible output quantities $Q = \{1, 2\}$. Consumption can take on any of 81 equally spaced values between zero and 2.25. Lastly, the technology relating action to the probability of each output is represented by

a	$P(q = 1)$	$P(q = 2)$
0	0.9	0.1
0.2	0.6	0.4
0.4	0.4	0.6
0.6	0.25	0.75

These parameters, along with those used for later examples, were chosen to display the possible characteristics of the economies we consider, and thus the main choice criteria was that they resulted in readable graphs. The figures presented are “typical” in that their general characteristics did not change over the set of parameters we explored.

The function $s^*(w)$ for each case is displayed in Figure 1. As can be seen in the maximization problems, randomizing between any two feasible schemes yields the interpolated surplus along with the interpolated utility for the agent while satisfying the other constraints. Thus the surplus functions must always be at least weakly concave over W . Further, the surplus for the unobserved-action environment cannot exceed that for the full-information environment since we have only added constraints. Note that since the y -axis is in terms of real consumption units, the loss to imposing these constraints (the vertical distance between the functions) is also in terms of real consumption. As Figure 1 shows, this loss is extremely small at the egalitarian optimum ($s^*(w) = 0$) for this set of parameters.⁷ Nevertheless, the characteristics of optimal plans given unobserved action is markedly different than given full information.

The upwards-sloping portion for the unobserved-action environment exists because the only incentive-compatible way to give an agent the ex ante utility of the lower bound on consumption and the lower bound on action is to actually give the lower bound of consumption and recommend the lower bound on action with certainty. For the left end-point in required utility, a higher required utility can deliver a higher surplus due to a higher expected level of action (Figure 2), and thus higher aggregate output from agents assigned that required utility. This higher output is more than enough to compensate the

7. This result may be somewhat akin to large literature which computes welfare losses and almost always finds small numbers.

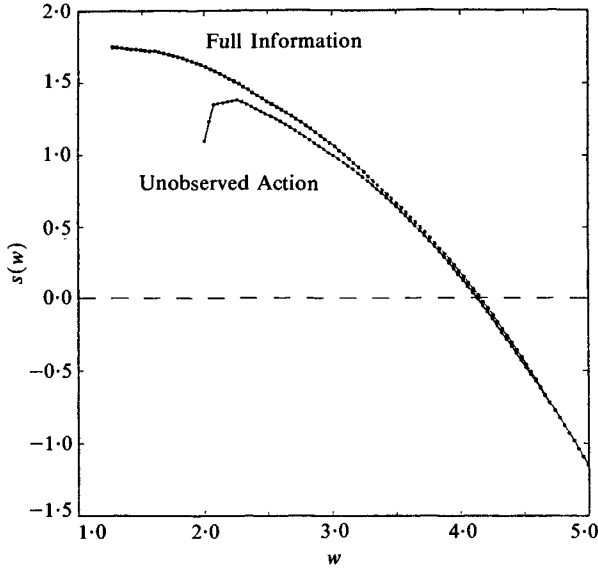


FIGURE 1
Optimized surplus functions

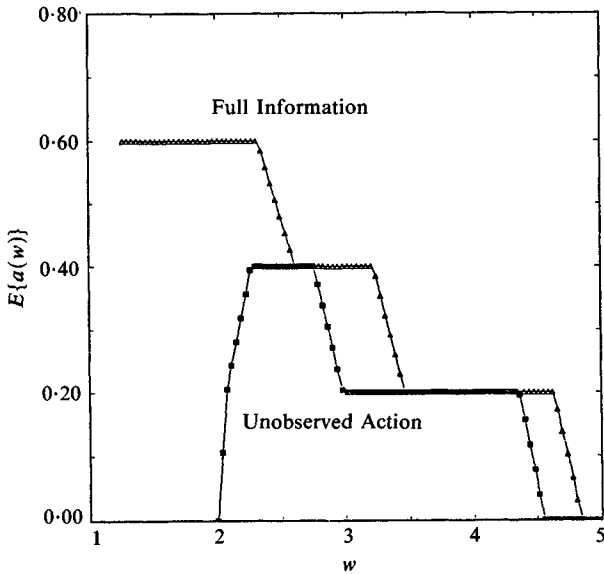


FIGURE 2
Actions

agents for the higher action. Note that the solution smooths over discrete points in A by mixing over adjacent points in A . For yet higher values of w , leisure is a normal good and thus decreases for both problems as w increases.

Figures 3 and 4 respectively display the unobserved-action and full-information consumption values c as functions of w , a , and q . As anticipated in the discussion of




FIGURE 3
Unobserved action consumption




FIGURE 4
Full information consumption

incentive constraints for the unobserved action economy, whenever action is greater than zero, higher consumption is associated with higher outputs. That is, an individual is now bearing risk. Further, for a given w , the difference in consumption over outputs is greater when the assigned action is greater. For the full-information solution, other than coincidental variations due to the finiteness of sets C and A , consumption does not vary

over outputs or assigned action.⁸ Lastly, note that given a distribution of required utilities d_0 , one can derive from these pictures cross-sectional figures in consumption and effort.

IV. REPEATED UNOBSERVED ACTION ECONOMIES.

The ability of incentive-constrained model economies to help explain effort and consumption variability is more pronounced when we examine the dynamic or repeated versions of the previous model. The repeated version allows a given individual's consumption and effort to be variable over time and for the distribution of agent characteristics in the population to change over time. We emphasize here that such dynamic economies can be solved by iteratively solving for value functions, or in our case, surplus functions, much like in optimal capital accumulation problems, but each step in the iteration defines linear programmes much like the static environments considered above.

Here we assume that time is discrete and agents are assumed to discount at a common rate $\beta < 1$. The social problem is to maximize the discounted sum of social surpluses subject to each agent receiving a given *ex ante* expected discounted utility. The possible initial required discounted expected utilities can be derived along the same lines as for the static economies. For the full-information case the lower bound on the possible set of such utilities, w_T , is that of receiving the lower bound on consumption and the upper bound on effort with certainty at every date. For the unobserved-action case, the lower bound on utility is that of receiving the lower bound on consumption and the *lower* bound on effort with certainty at every date, since as before, this situation can be achieved unilaterally by each agent or

$$w_T = \sum_{t=1}^T \beta^{t-1} U(a, \underline{c}), \quad \bar{w}_T = \sum_{t=1}^T \beta^{t-1} U(a, \bar{c}).$$

Again, all utilities in between can be achieved by randomizing between these two extremes. We denote these sets $W_T = [w, \bar{w}_T]$ for the T -period environment where T is a positive integer or infinity.

In using discounted social surplus as the objective function we are implicitly assuming (as in Green) that society can borrow and lend at a constant rate of interest. For convenience, we set this rate equal to $(1/\beta) - 1$ so that the full-information solution will have no aggregate borrowing or lending. A plan with zero (or positive) discounted social surplus is said to be feasible even though the associated contemporary surplus may be negative. In this case future resources are brought forward though they must be repaid with interest at some point through the horizon T . Alternatively, this can be viewed as society having the ability to invest resources in a productive technology, accumulating reserves for subsequent consumption. For the unobserved-action economy, computed solutions have this latter property.⁹

As before, we take as given $d_0(w)$, the fraction of agents at each initial w . As we shall establish carefully in the sections which follow, for a finite T -period economy, the social problem at date 1 can be defined in terms of the surplus function for the $(T-1)$ -period economy. For the infinitely-repeated economy, the social problem is recursively

8. If $U(a, c)$ were not separable then consumption could vary with assigned action for the full-information case. With unobserved action, consumption varies with assigned action even with separability.

9. The assumption of a linear storage technology is crucial in allowing us to derive a feasible method for computing solutions to the repeated unobserved-action economy. Specifically, without this assumption, it is no longer the case that an optimal allocation can be found by separately maximizing the social surplus from each utility group. Instead, one must solve for all contracts simultaneously as a function of the initial distribution of utilities $d_0(w)$. In the recursive formulation, it is the distribution of utilities at any given time which acts as the state variable for the economy. Distributions do not make convenient state variables. This type of linearity was also used in the seminal paper of Green (1987).

defined in terms of its own surplus function. The key, which again will be established carefully in the following sections, is to consider the utility of an agent from tomorrow on as a choice variable today. Thus the choice problem for society is to find for an agent required to have an ex ante expected discounted utility of w today the probability him taking action $a \in A$ today, having output $q \in Q$ today, receiving consumption $c \in C$ today, and being required to receive expected discounted utility $w' \in W_{T-1}$ from tomorrow on. We denote this probability $\Pi_T^w(a, q, c, w')$.¹⁰ As in the static problems, for given $w \in W_T$, these chosen probabilities must satisfy certain conditions:

First, the analogue of C1, that ex ante expected discounted utilities are satisfied, or,

$$C5. \quad w = \sum_{A \times Q \times C \times W_{T-1}} \{U[a, c] + \beta w'\} \Pi_T^w(a, q, c, w').$$

Second, the analogue of C2, that the chosen probabilities of outputs given actions are consistent with nature, or, for all $(\bar{a}, \bar{q}) \in A \times Q$,

$$C6. \quad \sum_{C \times W_{T-1}} \Pi_T^w(\bar{a}, \bar{q}, c, w') = P(\bar{q} | \bar{a}) \sum_{Q \times C \times W_{T-1}} \Pi_T^w(\bar{a}, q, c, w').$$

Third, the analogue of C3 are that these represent valid probability measures, or,

$$C7. \quad \sum_{A \times Q \times C \times W_{T-1}} \Pi_T^w(a, q, c, w') = 1 \quad \text{and} \quad \Pi_T^w(a, q, c, w') \geq 0$$

for all $(a, q, c, w') \in A \times Q \times C \times W_{T-1}$.

Lastly, for the unobserved-action problem, the analogue of C4, that the contract be incentive compatible, or, for all assigned and alternative action pairs $(a, \hat{a}) \in A \times A$,

$$C8. \quad \sum_{Q \times C \times W_{T-1}} \{U[a, c] + \beta w'\} \Pi_T^w(a, q, c, w') \\ \geq \sum_{Q \times C \times W_{T-1}} \{U[\hat{a}, c] + \beta w'\} \frac{P(q | \hat{a})}{P(q | a)} \Pi_T^w(a, q, c, w').$$

Again each of these constraints holds separately for each w type and thus the optimal social contract can be found by separately maximizing the social surplus from each w type, $w \in W_T$. At any given date, then, the plan that maximizes social surplus satisfies,

Repeated Problem (RP). Maximize by choice of probability measures Π_T^w the objective function $s_T(w) = \sum_{A \times Q \times C \times W_{T-1}} \{(q - c) + \beta s_{T-1}^*(w')\} \Pi_T^w(a, q, c, w')$ subject to Π_T^w satisfying C5 through C7, and for the unobserved action economy, C8.

The function $s_{T-1}^*(w') : W_{T-1} \rightarrow \mathbb{R}$ on the right-hand side is assumed to give the solution surplus values for the $(T-1)$ -period economy. For the infinite-period economy, it is assumed to give its own solution values. That is, for here and the rest of the paper, when T is set equal to infinity, the notation $T-1$ is meant to equal infinity as well.

The T -period economy is solved by solving the one-period (or static) problem and using its solution values to solve the two-period problem and so on. The infinite-period economy is solved by finding a surplus function $s_\infty(\cdot)$ that, when put in the objective function of RP, returns solution surplus values equal to $s_\infty(\cdot)$. This can be done by choosing any initial guess for $s_\infty(\cdot)$ and using it in the objective function of RP for every $w \in W_\infty$. The solution values of these programmes are then used as the next guess for $s_\infty(\cdot)$. As with standard dynamic programming for the capital accumulation problem, this sequence of guessed functions will converge to the true $s_\infty(\cdot)$. In the next section we guarantee this from theory.

10. There is some abuse of the notation here. Since w' can potentially take on a continuum of values, it is not proper to choose probabilities defined on all possible (a, q, c, w') points. Probabilities should be defined using more general measure notation. We do this in the later sections. For now, the notation is implicitly assuming a restriction that the number of possible future utilities is finite.

Both for the finite- and the infinite-period problems, we “grid” the relevant sets of possible utilities (W_∞ or W_T, W_{T-1}, \dots) allowing only a finite number of points in these. Then with actual finite or finite approximations to sets A, Q , and C , any stage in the computation process amounts to finding a solution to the finite linear-programming problem RP for a given w and a given function $s(\cdot)$ arrived at either by solving the $(T-1)$ -period economy for a finite-period problem or from the last iteration for the infinite-period problem. These linear programmes have a finite number of variables and a finite number of constraints. The only potential problem is that the number of variables and constraints may be large. Further, for each iteration there are as many programmes to solve as the number of grid points we allow in W_T .

Solutions for the full-information problem FIP, in the static environment, are also solutions to the full-information repeated environments. Specifically, there is full insurance so that action and consumption levels are variable only where tightenings are not possible due to the finiteness of A and C as noted before. Moreover, future utilities w' are chosen with certainty to yield the pre-specified initial utility w and thus guarantee the same action and consumption levels over time. If the future utility assignment of some measure Π^w were variable, then the mean could be assigned with certainty without affecting the initial prespecified utility w of the agent, but at least weakly raising surplus because the surplus function is at least weakly concave, and for non-trivial examples is strongly concave. If this certain w' were not equal to the initial required w then consumption and work probabilities would vary over time. From the concavity of the agents utility, this stream of consumption and leisure levels could be smoothed while holding ex ante utility constant but raising surplus. Thus to repeat, a solution to the static full-information problem is a solution to all the repeated problems whether the horizon T is finite or infinite.

However, and much to the point of this paper, solutions to the repeated unobserved-action economy do not have this property and can display variation in an individual's effort and consumptions over time with positive autocorrelation and variability in population distributions at a point in time. In particular, these distributions can now be non-trivial despite initial uniformity in the population. Figures 5 through 8 display solutions to the infinitely-repeated unobserved-action problem using the same parameters as in the static problems for the point-in-time utility function, the technology $P(q|a)$, and the sets A, Q , and C . We assume a discount rate of $\beta = 0.8$.

The optimized surplus function for the infinitely-repeated unobserved action problem is displayed in Figure 5 along with the full-information surplus function and the surplus function associated with infinitely repeating the solution to the $T=1$ economy. (Note that a static plan giving w to the agent with surplus $s(w)$ gives an expected discounted utility of $(1/(1-\beta)) \cdot w$ to the agent and discounted surplus $(1/(1-\beta)) \cdot s(w)$ if infinitely repeated. This allows us to virtually copy the surplus functions from Figure 1 onto Figure 5). The $T=\infty$ unobserved action surplus lies everywhere below the surplus from the full-information problem due to the added incentive constraints. It lies everywhere above the $T=1$ surplus because disallowing history-dependence in consumption and recommended action is a binding constraint. In fact, Figure 5 shows that much of the surplus lost from the agency problem is recovered when history-dependence is allowed even for the relatively high rate of discount implied by $\beta = 0.8$.

Allowing history-dependence also induces higher actions (Figure 6). As in the static problem, the lower bound on utility w can be achieved only by giving the lower bound on consumption with certainty and thus, for the incentive constraint to hold, having to recommend the lowest action. However, allowing history-dependence does lessen the severity of this for w near the lower bound. (Note again that points not on the action

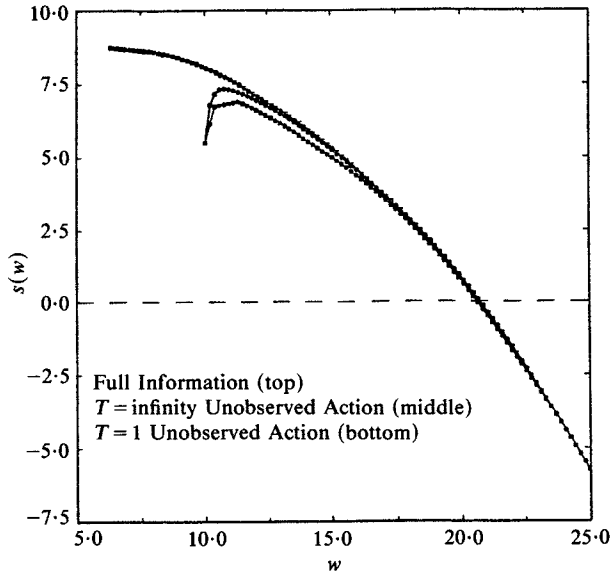


FIGURE 5
Optimized surplus functions

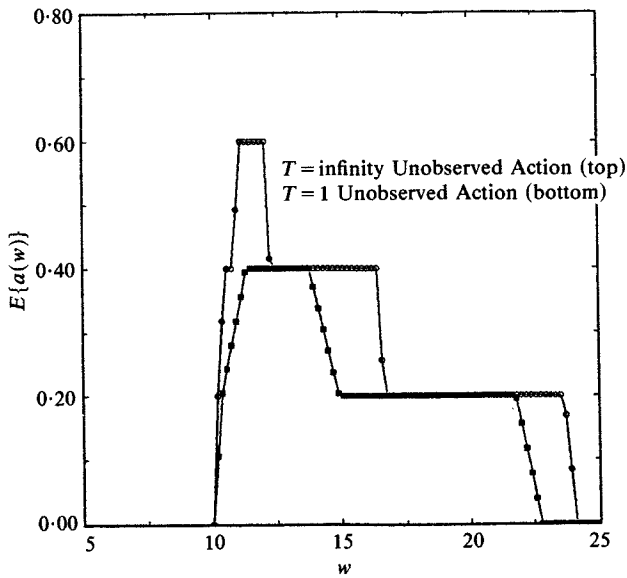


FIGURE 6
Actions

grid are achieved through randomization.) Again, away from these points, action decreases as w increases due to the normality of leisure.

Allowing history-dependence in consumption (Figure 7) greatly reduces the effect of contemporary output on consumption as compared to the static problem (Figure 3), that is, it allows better contemporaneous smoothing.




FIGURE 7

Unobserved action consumption

Figure 8 graphs the transition function in w as a function of w , a , and q . At both corners, this function lies on the black 45° line since the only way to achieve the endpoints of W_∞ is period-by-period repetition of the extreme schemes used to generate W_1 . At all other points where action is greater than zero, high output implies higher w' tomorrow, low output lower w' tomorrow.




FIGURE 8

Future utility

Given an initial w and a random-number generator, one can use the functions $\Pi_T^w(a, q, c, w')$ to generate individual time series of utilities, consumptions, or actions. One simply keeps track of the w' generated each period and uses it as the initial condition for the next. Figures 9 and 10 display four such series for consumption and expected discounted utility w for the same parameters as used earlier except for the discount factors $\beta = 0.95$. (The parameters $\beta = 0.8$ were chosen for Figure 5 through Figure 8 because




FIGURE 9
Individual consumptions




FIGURE 10
Individual utilities




FIGURE 11
Consumptions over time

these graphs are not easily readable at higher β values. Figures 9 through 12 are most readable for higher β values.) Each agent's initial w is set such that $s_{\infty}(w) = 0$. This is the highest ex ante expected discounted utility which can be assigned equally to the agents and thus represents the ex ante "fair" Pareto optimum.

The probability measures Π_T^w also allow one to chart population distributions over time. Figures 11 and 12 (where again β and $\delta = 0.95$) demonstrate that if every agent is given "fair" w initially that over time expected utilities and consumptions spread out giving non-trivial diversity in the population at a point in time. For other chosen parameters, diversity in action levels will exist for the same reasons. Although for our computed examples, the steady-state population distribution is degenerate due to absorbing states at the endpoints, it takes hundreds of periods for this absorption to significantly affect the distribution when starting at the fair w .¹¹ The endogenously generated diversity we display suggests that incentive problems may account for a non-trivial portion of the diversity in consumption and work levels present in the economy.

V. JUSTIFICATION OF METHODS.

In this section we justify the recursive methods used to formulate and solve the repeated problems. Specifically, we formulate the planning problem directly on the set of measures

11. We are not sure if there exist non-degenerate steady states in the environment of this paper. In computed solutions, agents seem to simply walk away from each other until they hit an absorbing state at one or the other endpoint in utility. We have, however, just discovered Thomas and Worrall (1990) who show analytically in the infinite-horizon version of Townsend (1982) that the steady state is always degenerate. The agent's (borrower's) utility becomes arbitrarily negative with probability one.




FIGURE 12
Utilities over time

defined on the possibly infinite-dimensional space of individual histories. We show in Theorem 1 that a solution exists. We then show that at any point in time and any history to that point, the expected discounted utility of any given agent from that time on is sufficient to describe the history of that agent along an optimal path; in a sense then one can start the problem over at any point. (Theorem 2). This allows us to show the optimal surplus or value function satisfies a recursive functional relationship. Further, one can then convert over to a space which replaces lotteries over contemporary outcomes conditioned on histories to lotteries conditioned on expected discounted utilities. This conversion requires a reduction of the incentive constraint from one defined on contingent sequences of deviations over the entire time-horizon to an apparently simpler incentive constraint concerning only one-date-at-time deviations given obedience from that date on. That this reduction is valid is the thrust of Theorem 3. Finally, for the case of an infinite time-horizon, Theorem 4 argues the existence of a unique surplus or value function satisfying the functional equation, with an associated maximal policy for the original problem.

The use of measure notation is necessary because the space of possible ex ante utilities is not finite but we want to define probabilities on it. Given this, there are no additional complications to allowing the sets A , Q , and C to be general compact subsets of the real line, finite or not.

The most general formulation of the problem of finding the optimal social contract for the repeated environments is to have as the social choice variable the joint probability measure over entire lifetime sequences of events for each agent given his initial required ex ante expected discounted utility. This allows the dependence of time t choice variables

on all previous information. If we let H be the set of triplets $H = \{h = (a, q, c) | a \in A, q \in Q, c \in C\}$, then we can define H^T , (T a finite integer or infinity) as the possibly infinite cross-product space $H^T \equiv \times_1^T H$, where a generic element $h^T \in H^T$ is a sequence (or history) of recommended actions, outputs, and consumptions for a given agent starting at date 1 and going through date T , possibly infinity. We use superscripts to denote sequences up to and including the superscripted date, as in $h^T \in H^T$, and subscripts to denote the elements of such sequences. That is, for a given sequence h^T , notation h_t denotes the (a, q, c) triplet associated with date t , and for the cross-product set H^T , H_t denotes the set of possible date t triplets. Lastly, wherever we use notation indicating a series from date t to date T , if T equals infinity we mean the infinite sequence from date t on.

With real-valued compact sets A and Q , we define the exogenous technology relating an agent's effort to his output as a probability measure $P(q|a)$. Measure P is defined on measure space $[Q, \mathcal{B}(Q)]$ where $\mathcal{B}(Q)$ is used to denote the Borel subsets of Q . Thus the object q in $P(q|a)$ is such a Borel subset, $q \in \mathcal{B}(q)$. The number $P(q|a)$ in the interval $[0, 1]$ gives the fraction of agents taking action a whose resulting output q is an element of set q .¹² Fractions $P(q|a)$ are viewed also as probabilities by the individual agents. Further, for any action, measure $P(\cdot|a)$ is associated with strictly positive continuous density over the set of outputs Q . This guarantees that all observed quantity sequences are possible given any action strategies, and thus we avoid any "off-the-equilibrium-path" considerations. Finally, no additional restrictions such as monotone likelihood ratio or convexity conditions on $P(\cdot|a)$ are required; these are made in the literature on principal-agent problems in efforts to secure analytical solutions.

Let D_0 define the initial measure of lifetime expected discounted utilities owed to the agents. For each generic subset w in the Borel subsets of W_T , $\mathcal{B}(W_T)$, number $D_0(w)$ gives the fraction of agents whose initial required utilities w are elements of set w .¹³ Unlike the introductory sections, we no longer assume measure D_0 has finite support, but still assume a countable support for technical reasons. Lastly, for real-valued compact sets A and C , we require that $U[a, c]$ both be bounded on $A \times C$ and be continuous.

A measure $F_T^w(h^T)$ returns for a given initially required discounted expected utility $w \in W_T$ and each generic subset $h^T \in \mathcal{B}(H^T)$ the fraction of agents whose actual possibly-infinite sequences of actions, outputs, and consumptions are elements of h^T . It is this measure for each $w \in W_T$ that is taken as the choice variable for society.

For simplicity of notation, we will denote the marginal measures implied by a given F_T^w on subsequences of events starting at date 1 and going to date $t < T$ as¹⁴

$$f_t^w(h^t) \equiv \int_{\{h^t \times H_{t+1} \times \dots \times H_T\}} dF_T^w(h^T). \tag{1}$$

A social contract is a set of measures F_T^w for all $w \in W_T$ satisfying the constraints defined below, namely C9, C10, and C11, the analogues of C1, C2, and C4 for the static economy. The first constraint is that the required discounted expected utilities for the agents are actually satisfied if agents follow the recommended actions, or for all $w \in W_T$,

C9.
$$w = \int_{H^t} \left\{ \sum_{i=1}^T \beta^{i-1} U[a_i, c_i] \right\} dF_T^w(h^T).$$

12. Note here and below that subsets are denoted as lower case bold type.
 13. We will continue the notation $\mathcal{B}(\cdot)$ to denote the Borel subsets of the indicated sets.
 14. The use of sets under the integral sign denotes integration over the elements of that set.

For the second constraint note that as in the static problem, the probability weights on outputs are choice variables in F_T^w and thus must be constrained to be consistent with those imposed by nature. For every date, and almost every history up to that date, a measure F_T^w implies a conditional probability measure of output given that history. Again, each of these measures must be constrained to be consistent with the exogenous technical measure $P(q|a)$. That is, for all t , $h^{t-1} \in \mathcal{B}(H^{t-1})$, $a_t \in \mathcal{B}(A)$, $q_t \in \mathcal{B}(Q)$,

$$C10. \quad \int_{\{h^{t-1} \times a_t \times q_t \times C\}} df_t^w(h^t) = \int_{\{h^{t-1} \times a_t \times Q \times C\}} P(q_t|a) df_t^w(h_t).$$

The interpretation follows. Integrating measure f_t^w over C on the left-hand side of C10 gives a probability measure defined on $\mathcal{B}(H^{t-1} \times A \times Q)$. Hence the entire term on the left-hand side gives the probability of outcomes in subset $h^{t-1} \times a_t \times q_t$. This measure must agree with the measure arrived at on the right-hand side by integrating f_t^w over both Q and C to get a measure defined on $\mathcal{B}(H^{t-1} \times A)$, and then deriving the joint probability over $h^{t-1} \times a_t \times q_t$ by using the underlying technology P , a measure defined on q and conditioned on contemporary action and trivially on history h_{t-1} .¹⁵ This constraint is essentially analogous to constraint C2 for the static problems.

The third constraint concerns incentive compatibility. For repeated environments each agent can have a deviation strategy depending on all information available to the agent at the time of the possible deviation. An agent's strategy does not need to depend on any societal aggregates since these are predetermined by a given plan $\{F_T^w\}^{w \in W_T}$. It is impossible for a single agent with zero weight to affect these. The aggregates are predetermined because each agent believes the other agents will obey their action recommendations. Incentive compatibility requires that given this it is weakly optimal for each agent to have the strategy of always obeying his own action recommendation.

To formally describe possible deviations, let $\hat{\Delta}_t$ be the set of functions mapping $H^{t-1} \times A$ to A . Define $\hat{\Delta}^T$ to be the set of T length sequences $\{\hat{\delta}_t\}_{t=1, T}$ of such functions. An element $\hat{\delta}^T \in \hat{\Delta}^T$ represents a strategy of a given agent. So again, an agent's strategy has as arguments the actual recommended actions, quantities, and consumptions the agent realizes up to and including $t - 1$, and the action recommended to him at time t . A strategy $\hat{\delta}^T$ thus gives the agent's action at all times under all possible histories.

Again, incentive compatibility requires that the strategy of taking the recommended action at all times under all histories is weakly optimal for all $w \in W_T$, and $\hat{\delta}^T \in \hat{\Delta}^T$, or that¹⁶

$$C11. \quad w = \int_{H^T} \left\{ \sum_{t=1}^T \beta^{t-1} U[a_t, c_t] \right\} dF_T^w(h^T) \\ \geq \int_{H^T} \left\{ \sum_{t=1}^T \beta^{t-1} U[\hat{\delta}_t(h^{t-1}, a_t), c_t] \right\} dF_T^w(h^T | \hat{\delta}^T) \equiv w(\hat{\delta}^T)$$

where $F_T^w(h^T | \hat{\delta}^T)$ is the probability measure facing an agent following deviation strategy $\hat{\delta}^T$. This differs from measure $F_T^w(h^t)$ in that deviation actions alter the probabilities of outputs and thus must be taken into account, as in the incentive constraint for the static

15. Given a probability measure $\mu(a \times b) : \mathcal{B}(A \times B) \rightarrow [0, 1]$ and associated marginal $\omega(a) : \mathcal{B}(A) \rightarrow [0, 1]$ where $\omega(a) = \int_{a \times B} d\mu(a, b)$, conditional measures $\gamma(b|a)$ can be defined (or chosen) for almost all a with respect to ω such that $\mu(a \times b) = \int_a \gamma(b|a) d\omega(a)$. (Billingsley (1986, Section 33)). We require that it is possible to choose such conditional probabilities to equal P .

16. These can be imposed, and $\hat{\Delta}^T$ can be restricted to deterministic strategies, without loss of generality using revelation principle or direct mechanism arguments. See Harris and Townsend (1981), Myerson (1979, 1982), and Townsend (1982).

environment. Measure $F_T^w(h^T | \hat{\delta}^T)$ can be defined such that for all $a_1 \in \mathcal{B}(A)$, $q_1 \in \mathcal{B}(Q)$, $c_1 \in \mathcal{B}(C)$, ..., $a_T \in \mathcal{B}(A)$, $q_T \in \mathcal{B}(Q)$, $c_T \in \mathcal{B}(C)$, and $\hat{\delta}^T \in \hat{\Delta}^T$,

$$F_T^w(h^T | \hat{\delta}^T) \equiv \int_{\{h^T = a_1 \times q_1 \times c_1 \times \dots \times a_T \times q_T \times c_T\}} \frac{P[q_1 | \hat{\delta}_1(a_1)] P[q_2 | \hat{\delta}_2(h_1, a_2)]}{P[q_1 | a_1] P[q_2 | a_2]} \dots \frac{P[q_T | \hat{\delta}_T(h^{T-1}, a_T)]}{P[q_T | a_T]} dF_T^w(h^T). \tag{2}$$

Ratio $P[q_i | \hat{\delta}_i(h^{i-1}, a_i)] / P[q_i | a_i]$ gives how many times more likely it is that $q_i \in q_i$ under strategy $\hat{\delta}^T$ than under the recommended strategy δ^T ; that is, this ratio updates the weights on sequences with $q_i \in q_i$ to account for the use of strategy $\hat{\delta}^T$.

For a given distribution of required ex ante expected utilities D_0 , the social problem is then to maximize by the choice of measures $\{F_T^w\}^{w \in W_T}$ over histories h^T , discounted social surplus, or,

Problem 1 (P1). Maximize by choice of $\{F_T^w\}^{w \in W_T}$ the objective function

$$S_T(D_0) \equiv \int_{W_T} \int_{H^T} \left\{ \sum_{t=1}^T \beta^{t-1} (q_t - c_t) \right\} dF_T^w(h^T) dD_0(w)$$

subject to the F_T^w satisfying C9 through C11 for all $w \in W_T$.

Note that since C9 through C11 must each hold separately for all $w \in W_T$, we can bring the maximization inside the integral over W_T and consider the simpler problem of maximizing the surplus from each w type.¹⁷ Although to guarantee a economy-wide optimum it is only necessary to maximize the surplus from the w elements of a support of D_0 , there is no loss to requiring that surplus be maximized for all $w = W_T$. This gives us

Problem 2 (P2). Maximize by choice of F_T^w the objective function

$$s_T(w) \equiv \int_{H^T} \left\{ \sum_{t=1}^T \beta^{t-1} (q_t - c_t) \right\} dF_T^w(h^T)$$

subject to F_T^w satisfying C9 through C11.

Denote the set of solutions to this problem $\{F_T^{w*}\}^{w \in W_T}$ with solution surplus values $s_T^*(w)$.

Theorem 1. *A solution to Problem 2 exists.*

Proof. (Outline of argument with some intuition). Since A , Q , and C are compact metric spaces, set H is compact metric space, and by a theorem of Tychonoff the infinite cross-product of H , H^T is a compact metric space, even if infinite-dimensional. This implies that the space of measures Π on H^T is also compact (and metrizable) relative to the weak topology.¹⁸

17. The assumption of a countable support ensures that the outer integration is always defined. This could have been avoided if we considered our choice variable one measure, F_T , over $\mathcal{B}(W_T \times H^T)$ instead of a continuum of measures F_T^w for each $w \in W_T$ each over $\mathcal{B}(H^T)$ and constrained F_T to be consistent with D_0 . This was avoided for expositional reasons.

18. Probability measures on compact sets are "tight" (Billingsley (1968)) and then by Prohorov's theorem are weakly compact, or compact relative to the weak topology.

The objective function of P2 can be shown to be continuous at any point \bar{F}_T^w since it is defined by the integration of bounded continuous functions. Similarly, the space Π restricted by C9 through C11 can be shown to be closed because one is dealing with linear equality and inequality constraints. If every element in some sequence were to satisfy some linear inequality constraint but the limit point did not, one can obtain a contradiction. As closed subsets of compact sets are compact, the constraint set for Problem 2 is compact. Continuous real-valued functions on non-empty compact topological spaces achieve maxima. Note that the constraint set is non-empty because for each $w \in W_T$ an F_T^w satisfying C9 through C11 can be created by mixing over the plans described earlier yielding the utility endpoints \underline{w} and \bar{w} . Note for this, and more generally below, that linear combinations of measures satisfying C10 and C11 also satisfy C10 and C11. \parallel

Further simplification is now possible. Specifically, the optimal F_T^{w*} are recursive in the following sense: If an individual starts with expected discounted utility w and subsequently has history h^τ (which under F_T^{w*} gives him an expected discounted utility from $\tau+1$ on of \hat{w}), F_T^{w*} might as well treat him from $\tau+1$ on as if it were the first period of a $(T-\tau)$ -period economy and the individual is required to receive \hat{w} . That is, the solution to a T -period economy can be stated in terms of the solution to $(T-\tau)$ -period economies, and the solution to the infinite-period economy can be stated in terms of itself.

For a given F_T^w and almost every history up to time τ , h^τ , we can define the discounted expected utility of an agent from time $\tau+1$ on, $W_\tau^w(h^\tau)$ as

$$W_\tau^w(h^\tau) \equiv \int_{\{H_{\tau+1} \times \dots \times H_T\}} \left\{ \sum_{i=\tau+1}^T \beta^{i-\tau-1} U(a_i, c_i) \right\} dF_T^w(h_{\tau+1}, \dots, h_T | h^\tau), \tag{3}$$

where conditional probability measure $F_T^w(h_{\tau+1} \times \dots \times h_T | h^\tau)$ is defined (or chosen) such that for all $h^\tau \times h_{\tau+1} \times \dots \times h_T$,

$$F_T^w(h^\tau \times h_{\tau+1} \times \dots \times h_T) = \int_{h^\tau} F_T^w(h_{\tau+1} \times \dots \times h_T | h^\tau) df_\tau^w(h^\tau). \tag{4}$$

Then, formally, we can state the recursive nature of the F_T^{w*} in Theorem 2.

Theorem 2. *For all $w \in W_T$, there exists an optimal plan F_T^{w*} with the following property: for any finite-length history h^τ with associated expected discounted utility $W_\tau^w(h^\tau)$, the probability measure on future events given that history, $F_T^{w*}(\cdot | h^\tau)$ can be chosen to be both consistent with equation 4 and to equal $F_{T-\tau}^{\hat{w}}$, an optimal probability measure for the $(T-\tau)$ -period economy for an individual required to receive $\hat{w} \equiv W_\tau^w(h^\tau)$. (Note that for the infinite-period economy the solution to the $(T-\tau)$ -period economy is the same as the solution to the T -period economy, namely $F_T^{w*} = F_{T-\tau}^{\hat{w}}$ for all $w \in W_\infty$.)*

*Proof.*¹⁹ Fix a positive measure subset of time- τ histories with respect to the optimal plan F_T^{w*} , denoted $\bar{h}^\tau \in \mathcal{B}(H^\tau)$. Suppose for almost all of the $\bar{h}^\tau \in \bar{h}^\tau$ there exists a new plan for determining the continuation fractions, $\tilde{F}_T^w(\cdot | \bar{h}^\tau)$, that generates a higher discounted surplus from $\tau+1$ on than the continuation plan $F_T^{w*}(\cdot | \bar{h}^\tau)$, that is,

$$\begin{aligned} \tilde{\sigma}_\tau(\bar{h}^\tau) &\equiv \int_{\{H_{\tau+1} \times \dots \times H_T\}} \left\{ \sum_{i=\tau+1}^T \beta^{i-\tau-1} (q_i - c_i) \right\} d\tilde{F}_T^w(h_{\tau+1}, \dots, h_T | \bar{h}^\tau) \\ &> \int_{\{H_{\tau+1} \times \dots \times H_T\}} \left\{ \sum_{i=\tau+1}^T \beta^{i-\tau-1} (q_i - c_i) \right\} dF_T^{w*}(h_{\tau+1}, \dots, h_T | \bar{h}^\tau) \equiv \sigma_\tau^*(\bar{h}^\tau), \end{aligned} \tag{5}$$

19. This proof follows a line of argument in Spear and Srivastava (1987).

with history-contingent surpluses $\tilde{\sigma}_\tau(\bar{h}^\tau)$ and $\sigma_\tau^*(\bar{h}^\tau)$ defined on the left and right equalities. Further suppose that this new plan $\tilde{F}_T^w(\cdot|\bar{h}^\tau)$ satisfies conditions C9 through C11 if time τ is taken as time zero and $W_\tau^w(\bar{h}^\tau)$ is taken as the initial required utility.

The collection of measures $\tilde{F}_T^w(\cdot|\bar{h}^\tau)$ for each $\bar{h}^\tau \in \bar{h}^\tau$ can be used to define a time-zero plan \tilde{F}_T^w . First, simply define the collection of conditional measures $\tilde{F}_T^w(\cdot|h^\tau)$ over all of H^τ , not just the preselected subset \bar{h}^τ , by letting $\tilde{F}_T^w(\cdot|h^\tau) \equiv F_T^{w*}(\cdot|h^\tau)$ for $h^\tau \notin \bar{h}^\tau$ so that the old optimal plan is *not* changed *off* the branches in \bar{h}^τ . Also no change is made to the old optimum on or before date τ . This allows us then to define \tilde{F}_T^w for all $h^\tau \times h_{\tau+1} \times \dots \times h_T \in \mathcal{B}(H^T)$ as

$$\tilde{F}_T^w(h^\tau \times h_{\tau+1} \times \dots \times h_T) \equiv \int_{h^\tau} \tilde{F}_T^w(h_{\tau+1} \times \dots \times h_T | h^\tau) df_\tau^{w*}(h^\tau). \tag{6}$$

The fact that for each original w type, $\tilde{F}_T^w(\cdot|h^\tau)$ gives exactly the same continuation utility to the agent after any time τ history implies that the ex ante utility of the agent is the same under \tilde{F}_T^w as under F_T^{w*} , so that \tilde{F}_T^w satisfies C9. Further, that the $\tilde{F}_T^w(\cdot|h^\tau)$ each satisfy C10 insures that \tilde{F}_T^w also does.

Discounted surplus from the perspective of time zero is greater under \tilde{F}_T^w than under F_T^{w*} because surplus at time zero is the discounted sum of outputs over consumptions through date τ and the discounted surplus from $\tau+1$ on. That is,

$$\begin{aligned} \tilde{s}_T(w) &\equiv \int_{H^\tau} \left\{ \sum_{t=1}^\tau \beta^{t-1} (q_t - c_t) + \beta^\tau \tilde{\sigma}_\tau(h^\tau) \right\} d\tilde{f}_\tau^w(h^\tau) \\ &> \int_{H^\tau} \left\{ \sum_{t=1}^\tau \beta^{t-1} (q_t - c_t) + \beta^\tau \sigma_\tau^*(h^\tau) \right\} df_\tau^{w*}(h^\tau) \equiv s_T^*(w), \end{aligned} \tag{7}$$

since by the definition of \tilde{F}_T^w (equation (6)), $\tilde{f}_\tau^w(h^\tau) = f_\tau^{w*}(h^\tau)$ and also that $\tilde{\sigma}_\tau(h^\tau) > \sigma_\tau^*(h^\tau)$ for $h^\tau \in \bar{h}^\tau$ and equal elsewhere.

We need only to show that \tilde{F}_T^w is incentive compatible from the perspective of time zero, and we contradict our original assumptions, since if \tilde{F}_T^w is within the constraint of the time-zero problem and gives a higher discounted surplus than F_T^{w*} , then F_T^{w*} could not have been optimal.

Claim. \tilde{F}_T^w is incentive compatible from time zero.

Proof of Claim. Suppose not. Then there exists a deviation plan $\hat{\delta}^T \in \hat{\Delta}^T$ such that

$$w^* = \tilde{w} < \int_{H^\tau} \left\{ \sum_{t=1}^\tau \beta^{t-1} U[\hat{\delta}_t(h^{t-1}, a_t), c_t] \right\} d\tilde{F}_T^w(h^T | \hat{\delta}^T) \equiv \tilde{w}(\hat{\delta}^T) \tag{8}$$

where w^* is the discounted expected utility of the agent under obedience and plan F_T^{w*} , \tilde{w} is the same but under plan \tilde{F}_T^w and obedience, and $\tilde{F}_T^w(\cdot|\hat{\delta}^T)$ is the measure defined by \tilde{F}_T^w and the deviation strategy $\hat{\delta}^T$ as in equation (2). For a given history at time τ , the agent has a discounted expected utility from that point on of following the deviation strategy that can be denoted $\tilde{W}_\tau^w(h^\tau, \hat{\delta}^T)$. Equation (8) can be rewritten as

$$\tilde{w} < \tilde{w}(\hat{\delta}^T) \equiv \int_{H^\tau} \left\{ \sum_{t=1}^\tau \beta^{t-1} U[\hat{\delta}_t(h^{t-1}, a_t), c_t] + \beta^\tau \tilde{W}_\tau^w(h^\tau, \hat{\delta}^T) \right\} d\tilde{f}_\tau^w(h^\tau | \hat{\delta}^T). \tag{9}$$

From the incentive compatibility of the $\tilde{F}_T^w(\cdot|h^\tau)$ we can replace $\tilde{W}_\tau^w(h^\tau, \hat{\delta}^T)$ in (9) with the discounted expected utility of following a strategy of obedience from $\tau+1$ on, denoted $\tilde{W}_\tau^w(h^\tau)$, without violating the inequality since obedience weakly dominates from $\tau+1$

on. This in turn can be replaced by $W_\tau^w(h^\tau)$, the discounted expected utility of the agent if plan $F_\tau^w(\cdot|h^\tau)$ is followed, due to the fact that plan $\tilde{F}_\tau^w(\cdot|h^\tau)$ satisfies constraint C9 from the perspective of time τ , that is that it delivers the same utility to the agent from the perspective of h^τ having occurred. Lastly, it follows that for any particular deviation plan $\hat{\delta}^T$, $\tilde{f}_\tau^w(h^\tau|\hat{\delta}^T) = f_\tau^{w*}(h^\tau|\hat{\delta}^T)$ (since no change is made in the old optimum through date τ and thus $\tilde{f}_\tau^w(h^\tau) = f_\tau^{w*}(h^\tau)$) and this replacement can be made in (9) also. This allows us to write

$$w^* < \int_{H^\tau} \left\{ \sum_{t=1}^\tau \beta^{t-1} U[\delta_t(h^{t-1}, a_t), c_t] + \beta^\tau W_\tau^{w*}(h^\tau) \right\} df_\tau^{w*}(h^\tau|\hat{\delta}^T). \tag{10}$$

The right-hand side of equation (10) is the utility of following $\hat{\delta}^T$ up to time τ and obedience afterward under the old plan F_τ^{w*} . If $\hat{\delta}^T$ specifies obedience up to and including time τ then equation (10) would be an equality and we have a contradiction. Assuming $\hat{\delta}^T$ does imply deviation on or before time τ , if we alter strategy $\hat{\delta}^T$ to obedience after date τ , we have derived another deviation strategy $\bar{\delta}^T$ such that $w^* < w^*(\bar{\delta}^T)$, which is a contradiction due to the assumed incentive compatibility of the time zero optimal plan.

This proves the claim of the incentive compatibility of \tilde{F}_τ^w , implying \tilde{F}_τ^w is within the constraint set of the time zero problem, contradicting the optimality of F_τ^{w*} . Thus, for all τ , and almost all histories h_τ , the surplus from $\tau+1$ on of following the original optimum, $\sigma_\tau^*(h^\tau)$, is the maximal surplus from $\tau+1$ on. Thus there is no loss in surplus from at every date τ pretending it is the first date of a $(T-\tau)$ economy and the initial required utility w is $W_\tau^{w*}(h^\tau)$. All other information in h^τ can be ignored. ||

We are now well on the way to reducing the original maximum problem to finding a function satisfying a recursive functional relationship. In particular, we now know from Theorem 2 that optimized surpluses $s_\tau^*(w)$ satisfy the recursive relationship,

$$s_\tau^*(w) = \int_{H_1} \{(q_1 - c_1) + \beta s_{\tau-1}^*[W_1^*(h_1)]\} df_1^{w*}(h_1). \tag{11}$$

Note from this that the optimized surplus today is determined only by the measure on one-period histories $f_1^{w*}(h_1)$, the function $s_{\tau-1}^*(w)$ and the function $W_1^*(h_1)$.²⁰ For a T -period economy, we can derive from $f_1^{w*}(h_1)$ and $W_1^*(h_1)$ a probability measure $g_T^{w*}(h_1 \times w)$ determining the joint probability weights on an agent having a given history h_1 today and expected discounted utility w from tomorrow. Measure $g_T^{w*}(h_1 \times w)$ gives the fraction of agents whose one-period history paired with *end-of-period* expected discounted utility is an element of $h_1 \times w$, and can be defined for all $h_1 \times w \in \mathcal{B}(H_1 \times W_{T-1})$ by

$$g_T^{w*}(h_1 \times w) \equiv \int_{h_1} \int_w I(h_1, w') dw' df_1^{w*}(h_1) \tag{12}$$

where for indicator function $I: H_1 \times W_{T-1} \rightarrow \{0, 1\}$, $I(h_1, w') = 1$ if $W_1^*(h_1) = w'$ and 0 otherwise, and the F_τ^w implying f_1^{w*} is now assumed to satisfy the recursive property in Theorem 2. Notation w' is used to denote expected utility as of the end of the first period from tomorrow on and w is for initial required utility. Note (12) implies that for almost all h_1 with respect to f_1^{w*} the implied conditional of w' on h_1 is degenerate; a particular number w' is always assigned for each history.

(12) allows us to express (11) as

$$s_\tau^*(w) = \int_{\{H_1 \times W_{T-1}\}} \{(q_1 - c_1) + \beta s_{\tau-1}^*(w')\} dg_T^{w*}(h_1, w'). \tag{13}$$

20. This is analogous to what Abreu, Pearce, and Stacchetti call factorization.

This allows us to find solutions to finite T -period problems by finding the functions $g_t^{w*}(h_1, w')$ for all $t \leq T$ and all $w \in W_t$. For the infinitely-repeated problem we need only to find the function $g_\infty^{w*}(h_1, w')$ for all $w \in W_\infty$. To do either we need to search over the class of possible measures $g_T^w(h_1, w')$ on $H_1 \times W_{T-1}$, where again T may be infinity, (and thus $T - 1$ equals infinity as well.)

A collection of arbitrary measures g_t^w for all $t \leq T$ and all $w \in W_t$, each degenerate on some end-of-period utility w' for almost all h_1 , defines a collection of measures F_T^w over entire histories h^T (for all $w \in W_T$) by successively applying the appropriate g_t^w . Marginal measure f_1^w , over initial h_1 , is defined by integrating g_T^w over w' . Measure f_2^w , over (h_1, h_2) can then be obtained by integrating g_{T-1}^w over w' and using f_1^w to determine the weights on the first-period histories, while keeping track of the end-of-period utilities associated with those first-period histories. This is repeated for f_3^w and further.²¹

In order to be able to compute optimal plans we need to know what restrictions must be put on an arbitrary collection of g_t^w measures, $\{g_t^w\}_{t \leq T}^{w \in W_t}$, so that all of the F_T^w plans that $\{g_t^w\}_{t \leq T}^{w \in W_t}$ imply satisfy constraints C9 through C11 if and only if these restrictions hold. Condition C9 is satisfied by the implied F_T^w if and only if for all $t \leq T$, $w \in W_t$,

$$C12. \quad w = \int_{\{H_1 \times W_{t-1}\}} \{U[a_1, c_1] + \beta w'\} dg_t^w(h_1, w').$$

Constraint C12 simply requires that at the beginning of all dates, the w assigned in the last period, and used as a state variable now to select from the $\{g_t^w\}$ family, is actually the expected discounted utility of the agent assigned it. Constraint C10 is satisfied by the implied F_T^w if and only if all the g_t^w generate the correct conditional probabilities of outputs on actions, or for all $a \in \mathcal{B}(A)$, $q \in \mathcal{B}(Q)$, $t \leq T$, and $w \in W_t$,

$$C13. \quad \int_{\{a \times q \times C \times W_{t-1}\}} dg_t^w(a, q, c, w') = \int_{\{a \times Q \times C \times W_{t-1}\}} P(q|a) dg_t^w(a, q, c, w').$$

The interpretation is analogous to C2.

Lastly, the most subtle equivalence concerns the incentive constraints. The result is summarized in Theorem 3.

Theorem 3. *For any collection of measures $\{g_t^w\}_{t \leq T}^{w \in W_t}$ satisfying conditions C12 and C13, the corresponding explicit history measures $\{F_T^w\}_{w \in W_T}$ satisfy incentive compatibility (constraint C11) if and only if for all $t \leq T$, $w \in W_t$ and all functions $\hat{\delta}: A \rightarrow A$,*

$$(C14) \quad \begin{aligned} w &\equiv \int_{\{H_1 \times W_{t-1}\}} \{U[a_1, c_1] + \beta w'\} dg_t^w(h_1, w') \\ &\equiv \int_{\{H_1 \times W_{t-1}\}} \{U[\hat{\delta}(a_1), c_1] + \beta w'\} dg_t^w(h_1, w' | \hat{\delta}), \end{aligned}$$

where for $a \times q \times c \times w \in \mathcal{B}(A \times Q \times C \times W_{t-1})$, $g_t^w(a \times q \times c \times w | \hat{\delta})$ is defined by

$$g_t^w(a \times q \times c \times w | \hat{\delta}) \equiv \int_{\{a \times q \times c \times w\}} \frac{P(q | \hat{\delta}(a))}{P(q | a)} dg_t^w(a, q, c, w') \tag{14}$$

the obvious analogue to (2).

21. Note for the case $T = \infty$ that the $\mathcal{B}(H^t)$ on which the f_t^w are defined can be considered partitions or sub-sigma-algebras of $\mathcal{B}(H^\infty)$. Further, this sequence of partitions converges to $\mathcal{B}(H^\infty)$ and thus the sequence of f_t^w functions converges to F_∞^w as $t \rightarrow \infty$.

Proof. Object w' is by condition C12 the expected utility from tomorrow on of a non-deviation strategy from tomorrow on and thus condition C14 requires that there be no gain to deviation at any history given that the agent will not deviate in the future. Since condition C11 requires that all contingent strategies be weakly worse than non-deviation, this includes the specific strategy of deviating at a given history and then following non-deviation from then on. Thus we already have the necessity of C14 or that every F_T^w satisfies C11 only if every $g_t^w, t \leq T$ and $w \in W_T$ satisfies C14.

To get the sufficiency of C14 suppose there exists a collection $\{g_t^w\}_{t \leq T}^{w \in W_t}$ that all satisfy C14 but a corresponding $\{F_T^w\}^{w \in W_t}$ implied by this collection does not satisfy C11 for some initial utility w_0 . This implies there exists an $\hat{\delta}^T \in \hat{\Delta}^T$ such that $w(\hat{\delta}^T) = w_0 + \varepsilon$ ($\varepsilon > 0$), where $w(\hat{\delta}^T)$ denotes the expected discounted utility of the agent following strategy $\hat{\delta}^T$. Consider an alternative action strategy $\tilde{\delta}^T$ which follows $\hat{\delta}^T$ up through some finite date $\tau \leq T$ but assumes obedience from $\tau + 1$ on. Condition C14 which assumes obedience from tomorrow on thus implies for all $h^{\tau-1}$,

$$\begin{aligned} W_{\tau-1}^w(h^{\tau-1}) &\equiv \int_{\{H_\tau \times \dots \times H_T\}} \{U[a_\tau, c_\tau] + \beta W_\tau^w(h^{\tau-1}, h_\tau)\} dF_T^w(h_\tau, \dots, h_T | h^{\tau-1}) \\ &\cong \int_{\{H_\tau \times \dots \times H_T\}} \{U[\tilde{\delta}_\tau, c_\tau] \\ &\quad + \beta W_\tau^w(h^{\tau-1}, h_\tau)\} dF_T^w(h_\tau, \dots, h_T | h^{\tau-1}, \tilde{\delta}^T) \\ &\equiv W_{\tau-1}^w(h^{\tau-1} | \tilde{\delta}^T). \end{aligned} \tag{15}$$

Stepping back one date, condition C14 again implies for all $h_{\tau-2}$,

$$\begin{aligned} W_{\tau-2}^w(h^{\tau-2}) &\equiv \int_{\{H_{\tau-1} \times \dots \times H_T\}} \{U[a_{\tau-1}, c_{\tau-1}] + \beta W_{\tau-1}^w(h^{\tau-2}, h_{\tau-1})\} dF_T^w(h_{\tau-1}, \dots, h_T | h^{\tau-2}) \\ &\cong \int_{\{H_{\tau-1} \times \dots \times H_T\}} \{U[\tilde{\delta}_{\tau-1}, c_{\tau-1}] \\ &\quad + \beta W_{\tau-1}^w(h^{\tau-2}, h_{\tau-1})\} dF_T^w(h_{\tau-1}, \dots, h_T | h^{\tau-2}, \tilde{\delta}^T). \end{aligned} \tag{16}$$

The expression $W_{\tau-1}^w(h^{\tau-2}, h_{\tau-1})$ on the right-hand side of (16) represents the utility corresponding to obedience from date τ on, even though $\tilde{\delta}^T$ allows disobedience at date τ . If we replace $W_{\tau-1}^w(h^{\tau-2}, h_{\tau-1})$ on the right-hand side with $W_{\tau-1}^w(h^{\tau-2}, h_{\tau-1} | \tilde{\delta}^T)$, that is, the utility associated with disobedience at date τ , (15) ensures the inequality is maintained. This gives

$$\begin{aligned} W_{\tau-2}^w(h^{\tau-2}) &\cong \int_{\{H_{\tau-1} \times \dots \times H_T\}} \{U[\tilde{\delta}_{\tau-1}, c_{\tau-1}] \\ &\quad + \beta W_{\tau-1}^w(h^{\tau-2}, h_{\tau-1} | \tilde{\delta}^T)\} dF_T^w(h_{\tau-1}, \dots, h_T | h^{\tau-2}, \tilde{\delta}^T) \\ &\equiv W_{\tau-2}^w(h^{\tau-2} | \tilde{\delta}^T) \end{aligned} \tag{17}$$

Repeating this stepping back to $t = 1$ gives

$$w_0 \cong w(\tilde{\delta}^T). \tag{18}$$

For a finite T , if we choose $\tau \approx T$ we have generated a contradiction since $\tilde{\delta}^T = \hat{\delta}^T$ but the utility associated with strategy $\tilde{\delta}^T$ is weakly worse than obedience (18) and the utility of strategy $\hat{\delta}^T$ is assumed strictly greater than obedience. For the infinitely-repeated case, denote \bar{w} as the greatest element of W_∞ and \underline{w} as its least element. If we choose

date τ sufficiently great $\beta^\tau \cdot (\bar{w} - w) < \varepsilon$, we have $|w(\tilde{\delta}^\tau) - w(\hat{\delta}^\tau)| < \varepsilon$, or $w(\tilde{\delta}^\tau) > w(\hat{\delta}^\tau) - \varepsilon$. Substitution from (18) gives $w_0 > w(\hat{\delta}^\tau) - \varepsilon$, which is a contradiction due to the initially assumed equality $w(\hat{\delta}^\tau) = w_0 + \varepsilon$. \parallel

Now again let $\{F_T^{w*}\}^{w \in W_T}$ denote a solution to P2 satisfying the recursive property of Theorem 2, and $s_T^*(w)$ denote the solution values. We know from above that the associated $\{g_t^{w*}\}_{t \leq T}^{w \in W_t}$ satisfying C12 through C14 and any other set $\{g_t^w\}_{t \leq T}^{w \in W_t}$ satisfying C12 through C14 must induce a set $\{F_T^w\}^{w \in W_T}$ with a lower surplus from the optimality of set $\{F_T^{w*}\}^{w \in W_T}$. This implies that $s_T^*(w)$ must be the maximal value of the objective function in

Problem 3 (P3). Maximize by choice of g_T^w the objective function

$$s_T(w) \equiv \int_{\{H_1 \times W_{T-1}\}} \{(q - c) + \beta s_{T-1}^*(w')\} dg_T^w(h_1, w')$$

subject to constraints C12 through C14.

For finite-period problems, one can start with the one-period problem (where $s_0^* = 0$ everywhere and $W_0 = \{0\}$) to find s_1^* and the set $\{g_1^{w*}\}^{w \in W_1}$ and then solve P3 given s_1 to solve the two-period problem and repeat until period T . For infinite-period problem, we need to derive a method of finding the maximal surplus function $s_\infty^*(\cdot)$. We can then characterize a solution to Problem P1 as the probability measure on infinite sequences h^∞ implied by the set $\{g_\infty^w\}^{w \in W_\infty}$ which solves P3 given $s_\infty^*(\cdot)$. To find $s_\infty^*(\cdot)$, we can define an operator T mapping functions $s_\infty(\cdot)$ into functions $Ts_\infty(\cdot)$ by letting, for every $w \in W_\infty$,

$$Ts_\infty(w) \equiv \max_{g_\infty^w} \int_{\{H_1 \times W_\infty\}} \{(q_1 - c_1) + \beta s_\infty(w')\} dg_\infty^w(h_1, w') \tag{19}$$

subject to conditions C12 through C14.

We know from the statement of Problem P3 that $Ts_\infty^*(\cdot) = s_\infty^*(\cdot)$. Theorem 4 states that no other function $s_\infty(\cdot)$ has this property and delivers a method of finding $s_\infty^*(\cdot)$.

Theorem 4. *Mapping T as defined in (19) maps bounded continuous functions on W_∞ into bounded continuous functions on W_∞ . Moreover, T is a contraction mapping, thus there exists a unique function s_∞^* such that $Ts_\infty^*(\cdot) = s_\infty^*(\cdot)$.*

Proof. A maximum exists over g_∞^w for the right-hand side of (19) because the objective function is an integral over bounded continuous functions defined on $H_1 \times W_\infty$ and hence is continuous in the weak topology, and the constraint set $\varphi(w)$ is compact. Set $\varphi(w)$ is compact because the constraints are linear equalities or inequalities and hence $\varphi(w)$ is closed and the set of unrestricted measures on $H_1 \times W_\infty$ is compact. That is, the constraint set is a closed subset of a compact set and hence is itself compact. Continuous functions on compact sets achieve their maxima.

That $Ts_\infty(\cdot)$ is bounded is trivial given the compactness of H_1 and the supposed boundedness of $s_\infty(\cdot)$. The continuity of $Ts_\infty(\cdot)$ in w follows from familiar arguments. The objective function of (19) is trivially continuous in w as well as in the choice objects; utility level w doesn't enter as an argument. The constraint set, $\varphi(w)$, is lower semicontinuous. This follows because any point g_∞^w in set $\varphi(w)$ can be attained as the limit of a sequence of $g_\infty^{w^n} \in \varphi(w^n)$ as w^n goes to w . If w^n were to approach w from above, for example, then one can construct the $g_\infty^{w^n}$ as the appropriate weighted combinations of g_∞^w and the scheme yielding the upper end point \bar{w} . This uses the fact that linear combinations of measures satisfying constraints C13 and C14 also satisfy C13 and C14

and thus one simply chooses the appropriate weights such that the linear combination gives the appropriate expected utility to satisfy constraint C12. The argument generalizes to all $w^n \rightarrow w$.

Also, the constraint set $\varphi(w)$ is upper semicontinuous. That is, if $w^n \rightarrow w$, and $g_\infty^{w^n} \rightarrow g_\infty^w$ with $g_\infty^{w^n} \in \varphi(w^n)$ then $g_\infty^w \in \varphi(w)$. This follows as with the establishment of the closedness of the constraint set above. Then the analogue of the maximum theorem given in Debreu (1959), here for a metric space, gives the desired result.

To establish T is a contraction mapping on this space of bounded continuous functions, it suffices to note that the Blackwell sufficient conditions are satisfied:

- (i) If $s' \geq s''$ then $Ts' \geq Ts''$, and
- (ii) For any constant K , $T(s + K) = Ts + \beta K$, $0 < \beta < 1$.

Existence and uniqueness of a function s_∞^* satisfying $s_\infty^* = Ts_\infty^*$ follows from the contraction property. Further iterations from s to Ts to T^2s and so on are ensured to converge to s_∞^* for any starting value s_∞ . ||

VI. NOTES ON COMPUTATION.

This section outlines the nature of the numerical algorithm used to compute solutions for a grid of utilities on W_t , $t \leq T$, and grids on the sets A , Q , and C . We also make clear the sense in which our solutions approximate the solution where these sets can take on a continuum of values.

For a given function $s_{t-1}^*(\cdot)$ or guess $s_\infty(\cdot)$ and initial w , finding $s_t(w)$ or $Ts_\infty(w)$ is an infinite linear programme literally having a continuum of choice variables, the mass to put on every point in the support of g_t^w , and if set A or Q is continuous, a continuum of constraints since C14 must hold for all possible strategies, and C5 for all $q \in \mathcal{B}(Q)$. If one imposes a grid on A and Q , that is allows (a, q) to take on only a finite number of points, then the number of constraints becomes finite. If one imposes a grid on W_{t-1} and C as well, then the number of choice variables becomes finite, and function $s_{t-1}(\cdot)$ defined on W_{t-1} is a finite-length vector. This makes $s_t^*(w)$ or $Ts_\infty(w)$ the solution to a finite linear programme for a given function $s_{t-1}^*(\cdot)$ or guess $s_\infty(\cdot)$. Such programmes can be computed using standard revised simplex algorithms. The entire function $s_t^*(\cdot)$ or $Ts_\infty(\cdot)$ is obtained by finding $s_t^*(w)$ or $Ts(w)$ for all $w \in \bar{W}_t$, where \bar{W}_t is the grid on the appropriate W_t . Further, the contraction theorem (Theorem 4) still applies even when W_∞ is restricted to a grid, so iteration over these functions for a given grid \bar{W}_∞ is ensured to converge. The converged solutions and associated policies are the ones we report.

Now return to the unrestricted set $H_1 \times W_{t-1}$. By imposing finer and finer grids on H_1 and W_{t-1} , for instance by uniformly distributing the elements and successively doubling the number of values they can take on (call these sets $(\bar{H}_1^n \times \bar{W}_{t-1}^n)$: $n = 1, \infty$). We get a sequence of solutions $\{g_t^w\}^{w \in \bar{W}_t^n}$, $n = 1, \infty$ which are valid measures on the unrestricted space $H_1 \times W_{t-1}$. Moreover, as a sequence on a compact metric space, at least some subsequence must converge, say to $\{g_t^w\}^{w \in W_t}$. Yet this must yield the same surplus as $\{g_t^{w*}\}^{w \in W_t}$ since any true maximizer for an unrestricted H_1 and W_{t-1} can be approximated arbitrarily closely by measures in sufficiently refined grids given that the actual technology P , (a measure on the unrestricted $Q \times A$) can be closely approximated on grided sets \bar{Q} and \bar{A} .

Finally, large gains are realized in the size and speed of computation of these linear programmes if one can separate each period into sub-periods as is easily the case with a separable utility function. Here it is the case that the expected utility of the agent also

sufficiently describes history *in the middle of a period*, after the output is realized but before the consumption is handed out, for the same reasons as in the proof of Theorem 2. This allows us to separately choose the joint probability of (a, q, w^m) triplets given w , where w^m is the utility of the agent in the middle of the period, and the joint probability of (c, w') pairs given w^m where, as before, w' is the expected discounted utility of the agent at the beginning of the next period.

This breaks each linear programme into two smaller ones. The programme choosing the probabilities over (a, q, w^m) triplets has the same number of constraints since the number of consumption points did not affect the number of constraints but has far fewer variables. For a gridded set \bar{C} with n_c elements, the number of variables in the smaller programme will be $1/n_c$ the number of variables in the original problem. The second linear programme choosing the probabilities of (c, w') pairs has only two constraints (that required utility is satisfied and that probabilities add to one) and thus runs extremely quickly. The adoption of this method allowed us to move from solving these problems on a CRAY supercomputer to a fast PC.

VII. CONCLUDING REMARKS.

The methods we present are more general than they may appear. For instance, the source of the incentive problem does not appear important. This allows similar arguments to those given to apply to repeated private-preference shock economies, or repeated private-endowment shock economies such as Green (1987). Since Green's model does not have actions, simply let the set A have one element and thus remove the incentive constraints on actions. Because Green has unobserved endowments (or unobserved q) we need to add constraints to require that those with high outputs do not claim to have low outputs, or, for all $(q, \hat{q}) \in Q \times Q$ (where $\hat{q} < q$),

$$\sum_{c \times w_{i-1}} \{U[a, c] + \beta w'\} \Pi_i^*(a, q, c, w') \geq \sum_{c \times w_{i-1}} \{U[a, c + q - \hat{q}] + \beta w'\} \Pi_i^*(a, \hat{q}, c, w') \quad (20)$$

This replacement of constraints does not upset the logic of the earlier arguments.

We have also shown in an earlier working paper version (Phelan and Townsend (1988)) that optimal renegotiation-proof contracts can be found by suitably restricting the utility sets W such that they induce non-increasing surplus functions over w , but otherwise solving the same programming problems. But the restriction to renegotiation-proof contracts does make a difference. Without it, as in the body of this paper, the optimal solution enters the upwards surplus regions, in which case, *ex post*, all agents can be made better off by starting over. In any event, this method of restricting utility sets also appears promising for computing other limited commitment environments such as Atkeson (1988) where international borrowers (agents) can withdraw from the credit system at any time with the utility associated with autarky given their present amount of capital.

Acknowledgement. Research support for Townsend from the National Science Foundation and the National Center For Supercomputing Applications (Champaign-Urbana, IL) and for Phelan from the Earhart Foundation and Bradley Foundation is gratefully acknowledged, as are helpful comments from In-Koo Cho, John Cochrane, Lars Hansen, John Heaton, Robert Lucas, Ed Prescott, Tom Sargent and participants in workshops at the University of Chicago, Carnegie-Mellon University, the IMSS Institute at Stanford and the Cornell Conference on Supercomputing and Economic Theory, and the Santander Workshop on New Developments in Macroeconomic Theory.

REFERENCES

- ABREU, D., PEARCE, D. and STACCHETTI, E. (1986), "Optimal Cartell Equilibria with Imperfect Monitoring", *Journal of Economic Theory*, **39**, 251-269.
- ATKESON, A. (1988), "International Lending with Moral Hazard and Risk of Repudiation" (Manuscript).
- BILLINGSLEY, P. (1968) *Convergence of Probability Measures*. (New York: John Wiley & Sons).
- BILLINGSLEY, P. (1986) *Probability and Measure*. (New York: John Wiley & Sons).
- DEBREU, G. (1959) *The Theory of Value* (New Haven: Yale University Press).
- FELDMAN, M. and GILLES, C. (1985), "An Expository Note on Individual Risk without Aggregate Uncertainty", *Journal of Economic Theory*, **35**, 26-32.
- FUDENBERG, D., HOLMSTROM, B. and MILGROM, P. (1986), "Repeated Moral Hazard with Borrowing and Saving" (Manuscript).
- GREEN, E. (1987), "Lending and the Smoothing of Uninsurable Income", in Prescott, E. and Wallace, N. (eds.), *Contractual Arrangements for Intertemporal Trade*. (Minneapolis: University of Minnesota Press), 3-25.
- HANSEN, G. (1985), "Indivisible Labor and the Business Cycle", *Journal of Monetary Economics*, **16**, 309-327.
- HARRIS, M. and TOWNSEND, R. (1981), "Resource Allocation Under Asymmetric Information", *Econometrica*, **49**, 33-64.
- JUDD, K. (1985), "The Law of Large Numbers With A Continuum Of IID Random Variables", *Journal of Economic Theory*, **35**, 19-25.
- MYERSON, R. (1979), "Incentive Compatibility and the Bargaining Problem", *Econometrica*, **47**, 61-74.
- MYERSON, R. (1982), "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems", *Journal of Mathematical Economics*, **10**, 67-82.
- PEARCE, D. (1987), "Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation" (Manuscript).
- PHELAN, C. and TOWNSEND, R. (1988), "Computing Solutions To Infinite-Period, Incentive-Constrained Economies" (Manuscript).
- PRESCOTT, E. and TOWNSEND, R. (1984a), "General Competitive Analysis in an Economy with Private Information", *International Economic Review*, **25**, 1-20.
- PRESCOTT, E. and TOWNSEND, R. (1984b), "Pareto Optima and Competitive Equilibria with Adverse Selection Moral Hazard", *Econometrica*, **52**, 21-46.
- RADNER, R. (1981), "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship", *Econometrica*, **49**, 1127-1148.
- RADNER, R. (1985), "Repeated Principal-Agent Games with Discounting", *Econometrica*, **53**, 1173-1198.
- ROGERSON, R. (1988), "Indivisible Labor, Lotteries and Equilibrium", *Journal of Monetary Economics*, **21**, 3-17.
- ROGERSON, W. (1985), "Repeated Moral Hazard", *Econometrica*, **53**, 69-76.
- RUBENSTEIN, A. and YAARI, M. (1983), "Repeated Insurance Contracts and Moral Hazard", *Journal of Economic Theory*, **30**, 74-97.
- SPEAR, S. and STRIVASTAVA, S. (1987), "On Repeated Moral Hazard with Discounting", *Review of Economic Studies*, **54**, 599-618.
- THOMAS, J. and WORRALL, T. (1990), "Income Fluctuations and Asymmetric Information: An Example of a Repeated Principal-Agent Problem", *Journal of Economic Theory*, **51**, 367-390.
- TOWNSEND, R. (1982), "Optimal Multiperiod Contracts and the Gain to Enduring Relationships Under Private Information", *Journal of Political Economy*, **90**, 1116-1186.
- TOWNSEND, R. (1987), "Economic Organization with Limited Communication", *American Economic Review*, **77**, 954-971.
- TOWNSEND, R. (1988). "Information Constrained Insurance: The Revelation Principle Extended", *Journal of Monetary Economics*, **21**, 411-450.