Dynamic Mechanism Design with Hidden Income and Hidden Actions *

Matthias Doepke
UCLA and CEPR

Robert M. Townsend
University of Chicago

August 2004

Abstract

We develop general recursive methods to solve for optimal contracts in dynamic principal-agent environments with hidden states and hidden actions. Starting from a general mechanism with arbitrary communication, randomization, full history dependence, and without restrictions on preferences or technology, we show that the optimal contract can be implemented as a recursive direct mechanism. A curse of dimensionality which arises from the interaction of hidden income and hidden actions can be overcome by introducing utility bounds for behavior off the equilibrium path. Environments with multiple actions are implemented using multiple layers of such off-path utility bounds.

Keywords: Mechanism Design, Dynamic Contracts, Recursive Contracts
JEL Classification: C63, C73, D82

*We are grateful to Harold Cole, Hugo Hopenhayn, Ken Judd, Marek Kapicka, Lars Ljungqvist, Roger Myerson, Nicola Pavoni, Ned Prescott, Phil Reny, Tom Sargent, Steve Tadelis, Iván Werning, Ruilin Zhou, and seminar participants at Boston University, Chicago, Cleveland Fed, Gerzensee, IIES, MIT, Richmond Fed, Stanford, UCL, UCLA, UC San Diego, the SED meeting in San José, and the SIEPR conference on "Credit Market Frictions in the Macroeconomy" for helpful comments. We also would like to thank the co-editor and two anonymous referees for suggestions that substantially improved the paper. Matthias Doepke acknowledges support by the NSF (grant SES-0217051), and Robert Townsend was supported by the NSF (grant 318340) and the NICHD (grant R01 HD027638). Doepke: Department of Economics, UCLA, 405 Hilgard Ave, Los Angeles, CA 90095. E-mail address: doepke@econ.ucla.edu. Townsend: Department of Economics, The University of Chicago, 1126 East 59th Street, Chicago, IL 60637. E-mail address: rtownsen@midway.uchicago.edu.

1 Introduction

The aim of this paper is to develop general recursive methods for computing optimal contracts in dynamic principal-agent environments with hidden states and hidden actions. A key feature of the environments that we consider is that the evolution of the state of the economy is *endogenous*, since it depends on hidden actions taken by the agent. This kind of intertemporal link is a natural feature of many private-information problems. Well-known examples include environments with hidden income which may depend on prior actions such as storage or investment, or environments with hidden search where success may depend on search effort in multiple periods. Despite the ubiquity of such interactions, most of the existing literature has avoided situations where hidden states and hidden actions interact, mainly because of the computational difficulties posed by these interactions.

We attack the problem by starting from a general planning problem which allows for history dependence and unrestricted communication and does not impose restrictions on preferences or technology. We show that this problem can be reduced to a recursive version with direct mechanisms and vectors of utility promises as the state variable. In this recursive formulation, however, a "curse of dimensionality" leads to an excessively large number of incentive constraints, which renders numerical computation of optimal allocations infeasible. We solve this difficulty by providing an equivalent alternative formulation of the planning problem in which the planner can specify behavior off the equilibrium path. This leads to a dramatic reduction in the number of constraints that need to be imposed when computing the optimal contract. With the methods developed in this paper, a wide range of dynamic mechanism design problems can be analyzed which were previously considered to be intractable.

As our baseline environment, we concentrate on the case of a risk-averse agent who has unobserved income shocks and can take unobserved actions such as storage or investment which influence future income realizations. There is a risk-neutral planner who wants to provide optimal incentive-compatible insurance against the income shocks. The design of optimal incentive-compatible mechanisms in environments with privately observed income shocks has been studied by a number of previous authors, including Townsend (1979), Green (1987), Thomas and Worrall (1990), Wang (1995), and Phelan (1995). In such environments, the principal extracts information from the agents about their income and uses this information to provide optimal incentive-compatible insurance. Much of the existing literature has concentrated on environments in which the agent is not able to take unobserved actions that influence the probability distribution over future income. The reason for this limitation is mainly technical; with hidden actions, the agent and the planner do not have common knowledge over probabilities of future states, which renders standard methods of computing optimal mechanisms inapplicable.

Theoretical considerations suggest that the presence of hidden actions can have large effects on the constrained-optimal mechanism. To induce truthful reporting of income, the planner needs control over the agent's intertemporal rate of substitution. When hidden actions are introduced, the planner has less control over rates of substitution, so providing insurance becomes more difficult. Indeed, there are some special cases where the presence of hidden actions causes insurance to break down completely. Allen (1985) shows that if an agent with hidden income shocks has unobserved access to a perfect credit market (both borrowing and lending), the principal cannot offer any insurance beyond the borrowing-lending solution. The reason is that any agent will choose the reporting scheme that yields the highest net discounted transfer regardless of what the actual income is.[1] Cole and Kocherlakota (2001b) consider a related closed economy environment in which the agent can only save, but not borrow. The hidden savings technology is assumed to have the same return as a public storage technology to which the planner has access. Remarkably, under additional assumptions, Cole and Kocherlakota obtain the same result as Allen: the optimal solution is equivalent to unrestricted access to credit markets. Ljungqvist and Sargent (2003) extend Cole and Kocherlakota to an open economy and establish the equivalence result with even fewer assumptions.

The methods developed in this paper can be used to compute optimal incentive-constrained mechanisms in more general environments with hidden actions in which self-insurance via investment or storage is an imperfect substitute for saving in an outside financial market, either because the return on storage is lower than the financial market return, or because the return on investment is high but with uncertain and unobserved returns. We show that in such environments, optimal information-constrained insurance can improve over the borrowing-and-lending solution despite severe information prob-

[1]The issue of credit-market access is also analyzed in Fudenberg, Holmstrom, and Milgrom (1990), who show that if principal and agent can access the credit market on equal terms, the optimal dynamic incentive contracts can be implemented as a sequence of short-term contracts. In Abraham and Pavoni (2002), trade is facilitated by giving the agent imperfect control over publicly observed output or over messages to the planner.

lems. As in Townsend (1982), the agent can be given incentives to correctly announce the underlying state, since he is aware of his own intertemporal rate of substitution. Unobserved investment subject to moral hazard with unobserved random returns does not undercut this intuition. The optimal outcome reduces to pure borrowing and lending only for a relatively narrow range of returns where the expected return on uncertain investment is roughly equal to (slightly above) the return in outside financial markets. Outside this range, insurance in addition to borrowing and lending is possible due to variability in the intertemporal rate of substitution.

Much of the existing literature on dynamic mechanism design has been restricted to environments where the curse of dimensionality that arises in our model is avoided from the outset. The paper most closely related to ours is Fernandes and Phelan (2000). Fernandes and Phelan develop recursive methods to deal with dynamic incentive problems with Markov income links between the periods, and as we do, they use a vector of utility promises as the state variable. There are three key differences between our paper and the approach of Fernandes and Phelan. First, the recursive formulation has a different structure. While in our setup utility promises are conditional on the realized income, in Fernandes and Phelan the promises are conditional on the unobserved action in the preceding period. Both methods should lead to the same results if the informational environment is the same, but they differ in computational efficiency. Even if income was observed, our method would be preferable in environments with many possible actions, but relatively few income values. Indeed, Kocherlakota (2004) points out that with a continuum of possible values for savings, the approach of Fernandes and Phelan is not computationally practical. A second more fundamental difference is that Fernandes and Phelan do not consider environments where both actions and states are unobservable to the planner, which is the main source of complexity in our setup (they consider each problem separately). Our methods are therefore applicable to a much wider range of problems, including those where the planner (or larger community) knows virtually nothing, only what has been transferred to the agent. Finally, our analysis is different from Fernandes and Phelan in that we allow for randomization and unrestricted communication, and we show from first principles that various recursive formulations are equivalent to the general formulation of the planning problem.

Another possible approach for the class of problems considered here is the recursive first-order method used by Werning (2001) and Abraham and Pavoni (2002) to analyze a dynamic moral hazard problem with storage. Kocherlakota (2004) casts doubts on the wider applicability of first-order methods, because the agent's decision problem is not necessarily concave when there are complementarities between savings and other decision variables. Werning and Abraham and Pavoni consider a restricted setting in which output is observed, the return on storage does not exceed the credit-market return, and saving is not subject to random shocks. To the extent that first-order methods are justified in a given application, they are computationally more efficient than our approach.[2] However, the range of possible applications is quite different. While the first-order approach (if it is justified) works well with continuous choice variables, we can allow arbitrary utility functions and technologies, discrete actions, and chunky investment decisions. Our methods also generalize to much more complicated environments with multiple actions, additional observable states, history-dependent technologies, and non-separable preferences.

In the following section, we provide an overview of the methods developed in this paper. In Section 3 we introduce the baseline environment that underlies our theoretical analysis and formulate a general planning problem with unrestricted communication and full history dependence. In Section 4, we invoke and prove the revelation principle to reformulate the planning problem using direct message spaces and enforcing truth-telling and obedience. We then provide a recursive formulation of this problem with a vector of utility promises as the state variable (Program 1). In Section 5, we develop an alternative formulation which allows the planner to specify utility bounds off the equilibrium path. This method leads to a dramatic reduction in the number of incentive constraints that need to be imposed when computing the optimal contract. In Section 6, we extend our methods to an environment with multiple unobserved actions and additional observable states, and show that additional actions can be handled by introducing multiple layers of off-path utility bounds. In Section 7 we provide two computed examples that put our methods to work. The first environment is a version of our baseline model with hidden storage with a fixed return, which allows us to contrast our numerical findings to some of the theoretical results pertaining to the same environment. We also compute a two-action environment with hidden effort, training, and productivity. Section 8 concludes, and all proofs are contained in the mathematical appendix.

[2] Abraham and Pavoni (2002) check a computed example to verify whether within numerical limits the agent is following a maximizing strategy under the prescribed contract.

2 Outline of Methods

As our baseline environment, we consider a dynamic mechanism design problem with hidden income and hidden actions. The agent realizes an unobserved income shock at the beginning of the period. Next, the planner may make a transfer to the agent, and at the end of period the agent takes an action which influences the probability distribution over future income realizations. We formulate a general planning problem of providing optimal incentive-constrained insurance to the agent. Our ultimate aim is to derive a formulation of the planning problem which can be solved numerically, using linear programming techniques as in Phelan and Townsend (1991). In order to make the problem tractable, we employ dynamic programming to convert the general planning problem to problems which are recursive and of relatively low dimension. Dynamic programming is applied at two different levels. First, building on Spear and Srivastava (1987), we use utility promises as a state variable to gain a recursive formulation of the planning problem (Section 4). In addition, we apply similar techniques within the period as a method for reducing the dimensionality of the resulting programming problem (Section 5).

A key feature of our approach is that we start from a general setup which allows randomization, full history dependence, and unrestricted communication. We formulate a general planning problem in the unrestricted setup (Section 3.2), and show from first principles that our recursive formulations are equivalent to the original formulation. Rather than imposing truth-telling and obedience from the outset, we prove a version of the revelation principle for our environment (Proposition 2).[3] Truth-telling and obedience are thus derived as endogenous constraints capturing all the information problems inherent in our setup.

After proving the revelation principle, the next step is to work towards a recursive formulation of the planning problem. Given that in our model both income and actions are hidden, so that in fact the planner does not observe anything, standard methods need to be extended to be applicable to our environment. The main problem is that with hidden income and actions a *scalar* utility promise is not a sufficient state variable, since agent and planner do not have common knowledge over probabilities of current and future

[3]Initially, derivations of the revelation principle as in Harris and Townsend (1977), Harris and Townsend (1981), Myerson (1979), and Myerson (1982) were for the most part static formulations. Here we are more explicit about the deviation possibilities in dynamic games, as in Townsend (1988). Unlike Myerson (1986), we do not focus on zero-probability events, but instead concentrate on within-period maximization operators using off-path utility bounds to summarize all possible deviation behavior.

states. However, our underlying model environment is Markov in the sense that the unobserved action only affects the probability of tomorrow's income. Once that income is realized, it completely describes the current state of affairs for the agent, apart from any history dependence generated by the mechanism itself. We thus show that the planning problem can still be formulated recursively by using a vector of income-specific utility promises as the state variable (see Proposition 4 and the ensuing discussion).

It is a crucial if well-understood result that the equilibrium of a mechanism generates utility realizations. That is, along the equilibrium path a utility vector is implicitly being assigned, a scalar number for each possible income (though a function of the realized history). If the planner were to take that vector of utility promises as given and reoptimize so as to potentially increase surplus, the planner could do no better and no worse than under the original mechanism. Thus, equivalently, we can assign utility promises explicitly and allow the planner to reoptimize at the beginning of each date. Using a vector of utility promises as the state variable introduces an additional complication, since the set of feasible utility vectors is not known in advance. To this end, we show that the set of feasible utility vectors can be computed recursively as well by applying a variant the methods developed in Abreu, Pearce, and Stacchetti (1990) (Propositions 7 and 8 in Appendix A.2).[4]

Starting from our recursive formulation, we discretize the state space to formulate a version of the planning problem which can be solved using linear programming and value-function iteration (Program 1 in Section 4.3 below). However, we now face the problem that in this "standard" recursive formulation a "curse of dimensionality" arises, in the sense that the number of constraints that need to be imposed when computing the optimal mechanism becomes very large. The problem is caused by the truth-telling constraints, which describe that reporting the true income at the beginning of the period is optimal for the agent. In these constraints, the utility resulting from truthful reporting has to be compared to all possible deviations. When both income and actions are unobserved, the number of such deviations is large. A deviation consists of lying about income, combined with an "action plan" which specifies which action the agent will take, given any transfer and recommendation he may receive from the planner before taking the action. The number of such action plans is equal to the number of possible actions taken to the power of the product of the number of actions and the number of transfers (recall that

[4]See also Cole and Kocherlakota (2001a) in an application to dynamic games with hidden states and actions.

there are finite grids for all choice variables to allow linear programming). The number of constraints therefore grows exponentially in the number of transfers and the number of actions. Thus even for moderate sizes of these grids, the number of constraints becomes too large to be handled by any computer.

To deal with this problem, we show that the number of constraints can be reduced dramatically by allowing the planner to specify outcomes off the equilibrium path (see Program 2 in Section 5 below). The intuition is that we use a maximization operator which makes it unnecessary to check all possible deviations, which can be done by imposing off-path utility bounds. The advantage of specifying behavior off the equilibrium path is that optimal behavior will be at least partly defined even if the agent misreports, so that not all possible deviations need to be checked. This technique derives from Prescott (2003), who uses the same approach in a static moral-hazard framework. The planner specifies upper bounds to the utility an agent can get by lying and receiving a specific transfer and recommendation afterwards. The truth-telling constraints can then be formulated in a particularly simple way by summing over the utility bounds. Additional constraints ensure that the utility bounds hold, i.e., the actual utility of deviating must not exceed the utility bound regardless of what action the agent takes. The number of such constraints is equal to the product of the number of transfers and the square of the number of actions. The total number of constraints in Program 2 is approximately linear in the number transfers and quadratic in the number of actions. In Proposition 5, we show that Program 1 and Program 2 are equivalent. With Program 2, the planning problem can be solved with fine grids for all choice variables.[5]

Notice that the advantages of using utility bounds are similar to the advantages of using a recursive formulation in the first place. One of the key advantages of a recursive formulation with utility promises as a state variable is that only one-shot deviations need to be considered. The agent knows that his future utility promise will be delivered, and therefore does not need to consider deviations that extend over multiple periods. Using off-path utility bounds applies a similar intuition to the incentive constraints within a period. The agent knows that it will be in his interest to return to the equilibrium path later in the period, which simplifies the incentive constraints in the beginning of the period.

A number of authors have considered environments where hidden savings interact with additional unobserved actions. For example, in the optimal unemployment insurance an-

[5]We are still restricted to small grids for the state variable, however, since otherwise the number of possible utility assignments becomes very large.

alyzed by Werning (2001) both job-search effort and savings are unobservable, and Abraham and Pavoni (2002) consider an environment with hidden work effort and hidden savings. These environments are also characterized by additional observable states, such as the outcome of the job search or production. In Section 6, we outline how our methods can be naturally extended to richer environments with additional actions and observed states. All our results carry over to the more general case. We show that multiple unobserved actions can be handled at a relatively low computational cost by introducing multiple layers of off-path utility bounds. At the beginning of the period, the truth-telling constraints are implemented using a set of utility bounds that are conditional on the first action only. These utility bounds, in turn, are themselves bounded by a second layer of utility bounds, which are conditional on the second action.

3 The Model

In the following sections we develop a number of recursive formulations for a mechanism design problem with hidden states and hidden actions. As the baseline case, we consider an environment with a single unobserved action. When deriving the different recursive formulations, we concentrate on the case of infinitely many periods with unobserved income and actions in every period. With little change in notation, the formulations can be adapted to models with finitely many periods, partially observable income and actions, and multiple actions. In particular, in Section 6, we show how our methods can be generalized to an environment with an additional action and an observable state. Figures 1 and 2 (in Section 6 below) show the time line for the baseline environment and the extended double-action environment. In Section 6, we will see that multiple actions within the period complicate the within-period incentive constraints, and we show how these complications can be handled by introducing multiple layers of utility bounds. The intertemporal aspect of the incentive problem, however, does not depend on the number of actions. We therefore start out with a simpler environment in which there is only a single unobserved action to concentrate on the intertemporal issues.

3.1 The Single Action Environment

In our baseline environment, there is an agent who is subject to income shocks, and can take actions that affect the probability distribution over future income. The plan-

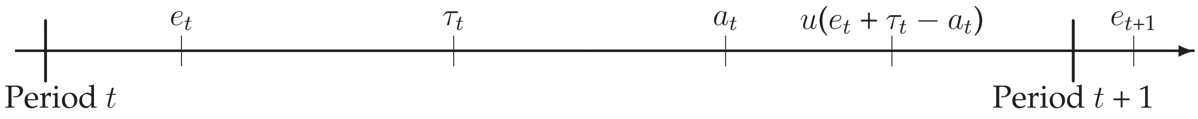

Figure 1: The Sequence of Events in Period t

ner, despite observing neither income nor actions, wants to provide optimal incentive-compatible insurance to the agent. The timing of events is summarized in Figure 1. At the beginning of each period the agent receives an income e from a finite set E. The income cannot be observed by the planner. Then the planner gives a transfer τ from a finite set T to the agent. At the end of the period, the agent takes an action a from a finite set A. Again, this action is unobservable for the planner. In most examples below we will concentrate on positive a and interpret that as storage or investment[6]. We can think of approximating a continuum of possible values of actions A by a finite set, or alternatively we can take indivisibilities seriously as in the investment literature. Similarly, the set of realized incomes E is approximated by a finite set of Markov states.[7] The agent consumes the amount $e + \tau - a$ and enjoys period utility $u(e + \tau - a)$, where the utility function is defined on $E \times T \times A$. Our methods do not require any specific assumptions on the utility function $u(\cdot)$, apart from it being real-valued. Indeed, we could separate consumption $e + \tau$ from the action a and write $U(e + \tau, a)$, though we would lose the credit market or investment interpretation.

The action a influences the probability distribution over income in the next period. Probability $\mu(e|a)$ denotes the probability of income e if the agent took action a in the previous period. In the initial period the probability $\mu(e)$ of income e does not depend on prior actions. For tractability, and consistent with the classic formulation of a moral-hazard problem, we assume that all states occur with positive probability, regardless of the action:

Assumption 1 *The probability distribution over income e satisfies $\mu(e|a) > 0$ for all $e \in E$, all*

[6]Without any changes in the setup we could also allow for a to be negative, which can be interpreted as borrowing.

[7]With non-stochastic saving, next period's beginning-of-period "income" would be random income plus the return on saving. If income has finite support, then for a given saving decision, certain income realizations would be impossible. This complicates but does not invalidate our methods—we could impose zero consumption for detectable off-path behavior and continue as below.

$a \in A$.

Otherwise, we place no restrictions on the investment technology. Note that since yesterday's action affects probabilities over income only in the current period[8], the income realization completely describes the current environment as far as the agent is concerned. Likewise, with time-separable utility, previous actions and incomes do not enter current utility as a state variables (though we shall generalize below). Apart from physical transactions, there is also communication taking place between the agent and the planner. We do not place any prior restrictions on this communication, in order not to limit the attainable outcomes. At a minimum, the agent has to be able send a signal about his beginning-of-period income, and the planner has to be able to send a recommendation for the investment or unobserved action.

The discount factor of the planner is denoted Q and β is the discount factor of the agent. The planner is risk-neutral and minimizes the expected discounted transfer, while the agent maximizes expected discounted utility. The objective function of the planner can be given the usual small-economy interpretation, i.e. the planner has unrestricted access to outside credit markets. In this case, the discount factor Q is given by $Q = \frac{1}{1+r}$, where r is taken to be the outside credit-market interest rate. We assume that both discount factors are less than one so that utility is finite and our problem is well defined.

Assumption 2 *The discount factors Q and β of the planner and the agent satisfy* $0 < Q < 1$ *and* $0 < \beta < 1$.

When there are only finitely many periods, we only require that both discount factors be bigger than zero.

While we formulate the model in terms of a single agent, another powerful interpretation is that there is a continuum of agents with mass equal to unity. In that case, the probability of an event represents the fractions in the population experiencing that event. Here the planner is merely a programming device to compute an optimal allocation: when the discounted surplus of the continuum of agents is zero, we have attained a Pareto optimum.

[8]This assumption could be weakened as in Fernandes and Phelan (2000). If the action had effects over multiple periods, additional state variables would have to be introduced to recover a recursive formulation.

3.2 The Planning Problem

We now want to formulate the Pareto problem of the planner maximizing surplus subject to providing reservation utility to the agent. Since the planner does not have any information on income and actions of the agent, we need to take a stand on what kind of communication is possible between the planner and the agent. In order not to impose any constraints from the outset, we start with a general communication game with arbitrary message spaces and full history-dependence. At the beginning of each period the agent realizes an income e. Then the agent sends a message or report m_1 to the planner, where m_1 is in a finite set M_1. Given the message, the planner assigns a transfer $\tau \in T$, possibly at random. Next, the planner sends a message or recommendation $m_2 \in M_2$ to the agent, and M_2 is finite as well. Finally, the agent takes an action $a \in A$. In the direct mechanism that we will introduce later, m_1 will be a report on income e, while m_2 will be a recommendation for the action a.

We will use h_t to denote the realized income and all choices of planner and agent within period t:

$$h_t \equiv \{e_t, m_{1t}, \tau_t, m_{2t}, a_t\}.$$

We denote the space of all possible h_t by H_t. The history up through time t will be denoted by h^t:

$$h^t \equiv \{h_{-1}, h_0, h_1, \ldots, h_t\}.$$

Here $t = 0$ is the initial date. The set of all possible histories up through time t is denoted by H^t and is thus given by:

$$H^t \equiv H_{-1} \times H_0 \times H_1 \times \ldots \times H_t.$$

At any time t, the agent knows the entire history up through time $t - 1$. On the other hand, the planner never sees the true income or the true action. We will use s_t and s^t to denote the part of the history known to the planner. We therefore have within period t

$$s_t \equiv \{m_{1t}, \tau_t, m_{2t}\},$$

where the planner's history of the game up through time t will be denoted by s^t, and the set S^t of all histories up through time t is defined analogously to the set H^t above. Since the planner sees a subset of what the agent sees, the history of the planner is uniquely determined by the history of the agent. We will therefore write the history of the planner as a function $s^t(h^t)$ of the history h^t of the agent. There is no information present at the beginning of time, and consequently we define $h_{-1} \equiv s_{-1} \equiv \emptyset$.

The choices of the planner are described by a pair of outcome functions $\pi(\tau_t|m_{1t}, s^{t-1})$ and $\pi(m_{2t}|m_{1t}, \tau_t, s^{t-1})$ which map the history up to the last period as known by the planner and events (messages and transfers) that already occurred in the current period into a probability distribution over transfer τ_t and a report m_{2t}. The choices of the agent are described by a strategy. A strategy consists of a function $\sigma(m_{1t}|e_t, h^{t-1})$ which maps the history up to the last period as known by the agent and the current income into a probability distribution over the first report m_{1t}, and a function $\sigma(a_t|e_t, m_{1t}, \tau_t, m_{2t}, h^{t-1})$ which determines the action.

We use $p(h^t|\pi, \sigma)$ to denote the probability of history h^t under a given outcome function π and strategy σ. The probabilities over histories are defined recursively, given history h^{t-1} and action $a_{t-1}(h^{t-1})$, by:

$$\begin{aligned}p(h^t|\pi, \sigma) = p(h^{t-1}|\pi, \sigma)\ \mu(e_t|a_{t-1}(h^{t-1}))\ \sigma(m_{1t}|e_t, h^{t-1})\ \pi(\tau_t|m_{1t}, s^{t-1}(h^{t-1}))\\ \pi(m_{2t}|m_{1t}, \tau_t, s^{t-1}(h^{t-1}))\ \sigma(a_t|e_t, m_{1t}, \tau_t, m_{2t}, h^{t-1}).\end{aligned}$$

Also, $p(h^t|\pi, \sigma, h^k)$ is the conditional probability of history h^t given that history h^k occurred with $k \leq t$, and conditional probabilities are defined analogously. In the initial period, probabilities are given by:

$$p(h^0|\pi, \sigma) = \mu(e_0)\ \sigma(m_{10}|e_0)\ \pi(\tau_0|m_{10})\pi(m_{20}|m_{10}, \tau_0)\ \sigma(a_0|e_0, m_{10}, \tau_0, m_{20}).$$

For a given outcome function π and strategy σ, the expected utility of the agent is:

$$U(\pi, \sigma) \equiv \sum_{t=0}^{\infty} \beta^t \left[\sum_{H^t} p(h^t|\pi, \sigma) u(e_t + \tau_t - a_t)\right]. \tag{1}$$

The expression above represents the utility of the agent as of time zero. We will also require that the agent use a maximizing strategy at all other nodes, even if they occur with probability zero. The utility of the agent given that history h^k has already been realized is given by:

$$U(\pi, \sigma | h^k) \equiv \sum_{t=k+1}^{\infty} \beta^{t-1-k} \left[\sum_{H^t} p(h^t | \pi, \sigma, h^k) u(e_t + \tau_t - a_t) \right]. \tag{2}$$

We now define an *optimal strategy* σ for a given outcome function π as a strategy that maximizes the utility of the agent at all nodes. The requirement that the strategy be utility maximizing can be described by a set of inequality constraints. Specifically, for a given outcome function π, for any alternative strategy $\hat{\sigma}$, and any history h^k, an optimal strategy σ has to satisfy:

$$\forall \hat{\sigma}, h^k : \quad U(\pi, \hat{\sigma} | h^k) \leq U(\pi, \sigma | h^k). \tag{3}$$

Inequality (3) thus imposes or describes optimization from any history h^k onward.

We are now able to provide a formal definition of an optimal strategy:

Definition 1 *Given an outcome function π, an optimal strategy σ is a strategy such that inequality* (3) *is satisfied for all k, all $h^k \in H^k$, and all alternative strategies $\hat{\sigma}$.*

Of course, for $h^k = h^{-1}$ this condition includes the maximization of expected utility (1) at time zero.

We imagine the planner as choosing an outcome function and a corresponding optimal strategy subject to the requirement that the agent realize at least reservation utility, W_0:

$$U(\pi, \sigma) \geq W_0. \tag{4}$$

Definition 2 *An equilibrium $\{\pi, \sigma\}$ is an outcome function π together with a corresponding optimal strategy σ such that* (4) *holds, i.e., the agent realizes at least his reservation utility. A feasible allocation is a probability distribution over income, transfers and actions that is generated by an equilibrium.*

The set of equilibria is characterized by the promise-keeping constraint (4), by the optimality condition (3), and of course a number of adding-up constraints that ensure that both outcome function and strategy consist of probability measures. For brevity these latter constraints are not written explicitly.

The objective function of the planner is:

$$V(\pi, \sigma) \equiv \sum_{t=0}^{\infty} Q^t \left[\sum_{H^t} p(h^t|\pi, \sigma)(-\tau_t) \right] \tag{5}$$

When there is a continuum of agents there is no aggregate uncertainty and (5) is the actual surplus of the planner, or equivalently, the surplus of the community as a whole. The Pareto problem for a continuum of agents can be solved by setting that surplus to zero. In the single-agent interpretation there is uncertainty about the realization of transfers, and (5) is the expected surplus. In either case, the planner's problem is to choose an equilibrium that maximizes (5). By construction, this equilibrium will be Pareto optimal. The Pareto frontier can be traced out by varying reservation utility W_0, and with a continuum of households, by picking the W_0 that generates zero surplus for the planner.

Definition 3 *An optimal equilibrium is an equilibrium that solves the planner's problem.*

Proposition 1 *There are reservation utilities $W_0 \in R$ such that an optimal equilibrium exists.*

4 Deriving a Recursive Formulation

4.1 The Revelation Principle

Our ultimate aim is to find a computable, recursive formulation of the planning problem. We begin by showing that without loss of generality we can restrict attention to a direct mechanism where there is just one message space each for the agent and the planner. The message space of the agent will be equal to the space of incomes E, and the agent will be induced to tell the truth. The message space for the planner will be equal to the space of actions A, and it will be in the interest of the agent to follow the recommended action. Since we fix the message spaces and require that truth-telling and obedience be optimal for the agent, instead of allowing any optimal strategy as before, it has to be the case that the set of feasible allocations in this setup is no larger than in the general setup with arbitrary message spaces. The purpose of this section is to show that the set of feasible allocations is in fact identical. Therefore there is no loss of generality in restricting attention to truth-telling and obedience from the outset.

More formally, we consider the planning problem described above under the restriction that $M_1 = E$ and $M_2 = A$. We can then express the contemporary part of the history of the planner as:

$$s_t \equiv \{e_t, \tau_t, a_t\},$$

with history s^t up through time t defined as above. Notice that since we are considering the history of the planner, e_t is the *reported*, not necessarily actual income, and a_t is the *recommended* action, not necessarily the one actually taken. The planner only knows the transfer and the messages. This will be different once we arrive at the mechanism with truth-telling and obedience, where reported and actual income and recommended and actual action always coincide.

As before, the planner chooses an outcome function consisting of probability distributions over transfers and reports. For notational convenience, we express the outcome function as the joint probability distribution over combinations of transfer and recommendation. This is equivalent to choosing marginal probabilities as above. The planner therefore chooses probabilities $\pi(\tau_t, a_t | e_t, s^{t-1})$ that determine the transfer τ_t and the recommended action a_t as a function of the reported income e_t and the history up to the last period s^{t-1}.

We now impose constraints on the outcome function π that ensure that the outcome function together with a specific strategy of the agent, namely truth-telling and obedience, are an equilibrium. First, the outcome function has to define probability measures. We require that $\pi(\tau_t, a_t | e_t, s^{t-1}) \geq 0$ for all transfers, actions, incomes and histories, and that:

$$\forall e_t, s^{t-1}: \quad \sum_{T,A} \pi(\tau_t, a_t | e_t, s^{t-1}) = 1. \tag{6}$$

Given an outcome function, we define probabilities $p(s^t|\pi)$ over histories in the obvious way, where the notation for σ is suppressed on the premise that the agent is honest and obedient. Given these probabilities, as in (4), the outcome function has to deliver reservation utility W_0 to the agent, provided that the agent reports truthfully and takes the recommended actions:

$$\sum_{t=0}^{\infty} \beta^t \left[\sum_{S^t} p(s^t|\pi) u(e_t + \tau_t - a_t) \right] \geq W_0. \tag{7}$$

Finally, it has to be optimal for the agent to tell the truth and follow the recommended action, so that (3) holds for the outcome function and the maximizing strategy σ of the agent, which is to be truthful and obedient. In particular, the utility of honesty and obedience must weakly dominate the utility derived from any possible deviation strategy mapping any realized history, which may be generated by possible earlier lies and disobedient actions, into possible lies and alternative actions today, with plans for possible deviations in the future. We write a possible deviation strategy δ, which is allowed to be fully history-dependent, as a set of functions $\delta_e(h^{t-1}, e_t)$ that determine the reported income as a function of the actual history h^{t-1} and the true income e_t, and functions $\delta_a(h^{t-1}, e_t, \tau_t, a_t)$ that determine the actual action as a function of the history h^{t-1}, income e_t, transfer τ_t, and recommended action a_t. Since the actual action may be different from the recommendation, this deviation also changes the probability distribution over histories and states. The agent takes this change into account, and the changed probabilities are denoted as $p(h^t|\pi, \delta)$, with the inclusion of other conditioning elements where appropriate. In particular, we require that the actions of the agent be optimal from any history of the planner s^k onward. It will also be useful to write down separate constraints for each possible income e_{k+1} in period $k+1$. Then for every possible deviation (δ_e, δ_a), any history s^k, and any e_{k+1}, the outcome function has to satisfy:

$$\forall \delta, s^k, e_{k+1}: \quad \sum_{t=k+1}^{\infty} \beta^t \left[\sum_{H^t} p(h^t|\pi, \delta, s^k, e_{k+1}) u(e_t + \tau_t - \delta_a(h^{t-1}, e_t, \tau_t, a_t)) \right]$$
$$\leq \sum_{t=k+1}^{\infty} \beta^t \left[\sum_{S^t} p(s^t|\pi, s^k, e_{k+1}) u(e_t + \tau_t - a_t) \right]. \quad (8)$$

Here $p(h^t|\pi, \delta, s^k, e_{k+1})$ on the left-hand side is the probability of actual history h^t implied by outcome function π and deviation δ conditional on the planner's history s^k and realized income e_{k+1}, and the $p(s^t|\pi, s^k, e_{k+1})$ on the right-hand side is the probability under truth-telling and obedience as above, but now conditioned on s^k and e_{k+1}. Condition (8) imposes or describes honesty and obedience on the equilibrium path, similar to (3).

It might seem at first sight that (8) is less restrictive than (3), because only a subset of possible deviations is considered. Specifically, deviations are non-random, and a constraint is imposed only at every s^t node instead of every node h^k of the agent's history. However, none of these limitations are restrictive. Allowing for randomized deviations would lead to constraints which are linear combinations of the constraints already imposed. Imposing (8) is therefore sufficient to ensure that the agent cannot gain from randomized deviations. Also, notice that the conditioning history s^k enters (8) only by affecting probabilities over future states s^t through the history-dependent outcome function π. These probabilities are identical for all h^k that coincide in the s^k part once e_{k+1} is realized. The agent's private information on past incomes e and actions a affects the present only through the probabilities over different future incomes. Imposing a separate constraint for each h^k therefore would not put additional restrictions on π.

Definition 4 *An outcome function is an equilibrium outcome function under truth-telling and obedience if it satisfies the constraints* (6), (7) *and* (8) *above. A feasible allocation in the truth-telling mechanism is a probability distribution over income, transfers and actions that is implied by an equilibrium outcome function.*

Feasible allocations under truth-telling and obedience are a subset of the feasible allocations in the general setup, since (7) implies that (4) holds, and (8) implies that (3) holds. In fact, we can show that the set of feasible allocations in the general and the restricted setup are identical.

Proposition 2 (Revelation Principle) *For any message spaces M_1 and M_2, any allocation that is feasible in the general mechanism is also feasible in the truth-telling-and-obedience mechanism.*

The proof (outlined in the appendix) takes the usual approach of mapping an equilibrium of the general setup into an equilibrium outcome function in the restricted setup. Specifically, given an equilibrium (π, σ) in the general setup, the corresponding outcome function in the restricted setup is gained by prescribing the outcomes on the equilibrium path while integrating out all the message spaces:

$$\pi^\star(\tau_t, a_t | e_t, s^{t-1}) \equiv \sum_{H^{t-1}(s^{t-1}), M_1, M_2} p(h^{t-1}|s^{t-1})\, \sigma(m_{1t}|e_t, h^{t-1})\, \pi(\tau_t|m_{1t}, s^{t-1}(h^{t-1}))$$
$$\pi(m_{2t}|m_{1t}, \tau_t, s^{t-1}(h^{t-1}))\, \sigma(a_t|e_t, m_{1t}, \tau_t, m_{2t}, h^{t-1}).$$

The proof then proceeds by showing that the outcome function $\pi^\star$ on the left-hand side satisfies all the required constraints. The essence of the matter is that lying or deviating under the new outcome function would be equivalent to using the optimizing strategy function under the original outcome function, but evaluated at a counterfactual realization. For example, an agent who has income e but reports $\hat{e}$ will face the same probability distribution over transfers and recommendations as an agent who under the original outcome function behaved "as if" the income were $\hat{e}$. The agent can never gain this way, since σ is an optimal strategy, and it is therefore preferable to receive the transfers and recommendations intended for income e instead of $\hat{e}$.

We are therefore justified in continuing with the restricted setup which imposes truth-telling and obedience. The objective function of the planner is now:

$$V(\pi) \equiv \sum_{t=0}^{\infty} Q^t \left[\sum_{S^t} p(s^t|\pi)(-\tau_t) \right], \tag{9}$$

and the original planning problem can be expressed as maximizing (9) subject to (6), (7), and (8) above.

4.2 Utility Vectors as State Variables

We now have a representation of the planning problem that requires truth-telling and obedience and yet does not constitute any loss of generality. However, we still allow fully history-dependent outcome functions. The next step is to reduce the planning problem to a recursive version with a vector of promised utilities as the state variable.

We wish to work towards a problem in which the planner has to deliver a vector of promised utilities at the beginning of period k, with elements depending on income e_k. It will be useful to consider an auxiliary problem in which the planner has to deliver a vector of reservation utilities $\mathbf{w}_0$ depending on income in the initial period. The original planning problem can then be cast, as we shall see below, as choosing the vector of initial utility assignments $\mathbf{w}_0$ which yields the highest expected surplus for the planner, given the initial exogenous probability distribution over states $e \in E$ at time $t = 0$.

In the auxiliary planning problem, we impose the same probability constraints (6) and incentive constraints (8) as before. However, instead of a single promise-keeping constraint (7) there is now a separate promise-keeping constraint for each possible initial income. For all e_0, we require:

$$\forall e_0 :$$

$$\sum_{T,A} \pi(\tau_0, a_0|\mathbf{w}_0, e_0) \left[u(e_0 + \tau_0 - a_0) + \sum_{t=1}^{\infty} \beta^t \Big[\sum_{S^t} p(s^t|\pi, s_0) u(e_t + \tau_t - a_t) \Big] \right] = w_0(e_0). \tag{10}$$

Here the vector $\mathbf{w}_0$ of income-specific utility promises $w_0(e_0)$ is taken as given. Notice that we write the outcome function π as a function of the vector of initial utility promises $\mathbf{w}_0$. In period 0, there is no prior history, but in a subsequent period t the outcome function also depends on the history up to period $t-1$, so the outcome function would be written as $\pi(\tau_t, a_t|\mathbf{w}_0, e_t, s^{t-1})$.

In principle, specifying a separate utility promise for each income is more restrictive than a requiring that a scalar utility promise be delivered in expected value across incomes. However, the original planning problem can be recovered by introducing an initial stage at which the initial utility vector is chosen by the planner.

Since the vector of promised utilities $\mathbf{w}_0$ will serve as our state variable, it will be important to show that the set of all feasible utility vectors has nice properties.

Definition 5 *The set* $\mathbf{W}$ *is given by all vectors* $\mathbf{w}_0 \in R^{\#E}$ *that satisfy constraints* (6), (8), *and* (10) *for some outcome function* $\pi(\tau_t, a_t|e_t, s^{t-1})$.

Proposition 3 *The set* $\mathbf{W}$ *is nonempty and compact.*

Now we consider the problem of a planner who has promised utility vector $\mathbf{w}_0 \in \mathbf{W}$ and has received report e_0 from the agent. In the auxiliary planning problem, the maximized surplus of the planner is given by:

$$V(\mathbf{w}_0, e_0) = \max_{\pi} \sum_{T,A} \pi(\tau_0, a_0|\mathbf{w}_0, e_0) \left[-\tau_0 + \sum_{t=1}^{\infty} Q^t \Big[\sum_{S^t} p(s^t|\pi, s_0)(-\tau_t) \Big] \right], \tag{11}$$

where the maximization over current and future π is subject to constraints (6), (8), and (10) above, for a given $\mathbf{w}_0 \in \mathbf{W}$ and $e_0 \in E$.

We want to show that this problem has a recursive structure. To do this, we need to define on-path future utilities that result from a given choice of π. For all s^{k-1}, e^k, let:

$$w(e_k, s^{k-1}, \pi) =$$
$$\sum_{T,A} \pi(\tau_k, a_k|\mathbf{w}_0, e_k, s^{k-1}) \left[u(e_k + \tau_k - a_k) + \sum_{t=k+1}^{\infty} \beta^{t-k} \Big[\sum_{S^t} p(s^t|\pi, s^k) u(e_t + \tau_t - a_t) \Big] \right], \tag{12}$$

and let $\mathbf{w}(s^{k-1}, \pi)$ be the vector of these utilities over all e_k. We can now show a version of the principle of optimality for our environment:

Proposition 4 *For all* $\mathbf{w}_0 \in \mathbf{W}$ *and* $e_0 \in E$*, and for any* s^{k-1} *and* e_k*, there is an optimal contract* $\pi^\star$ *such that the remaining contract from* s^{k-1} *and* e_k *onward is an optimal contract for the auxiliary planning problem with* $e_0 = e_k$ *and* $\mathbf{w}_0 = \mathbf{w}(s^{k-1}, \pi^\star)$.

Thus the planner is able to reoptimize the contract at any future node. For Proposition 4 to go through, it is essential that we chose a vector of utility promises as the state variable, as opposed to the usual scalar utility promise which is realized in expected value across states. If the planner reoptimized given a scalar utility promise at a given date, the distribution of expected utilities across states might be different than in the original contract. Such a reallocation of utilities would change the incentives for lying and disobedience in the preceding period, so incentive-compatibility of the complete contract would no longer be guaranteed. This problem is avoided by specifying a separate utility promise for each possible income. Likewise, in implementing the utility promises it does not matter whether the agent lied or was disobedient in the past, since the agent has to report the realized income in either case, and once income is realized past actions have no further effects.[9]

Given Proposition 4, we know that the maximized surplus of the planner can be written as:

$$V(\mathbf{w}_0, e_0) = \sum_{A,T} \pi^\star(\tau_0, a_0|\mathbf{w}_0, e_0)\left[-\tau_0 + Q\sum_E \mu(e_1|a_0(s^0))V(\mathbf{w}(s^0, \pi^\star), e_1)\right]. \tag{13}$$

In light of (13), we can cast the auxiliary planning problem as choosing transfers and actions in the initial period, and choosing continuation utilities from the set $\mathbf{W}$, conditional on history $s^0 = \{e_0, \tau_0, a_0\}$.

We are now close to the recursive formulation of the planning problem that we are looking for. We will drop the time subscripts from here on, and write the choices of the planner as a function of the current state, namely the vector of promised utilities $\mathbf{w}$ that has to be delivered in the current period and the reported income e. The choices are functions $\pi(\tau, a|\mathbf{w}, e)$ and $\mathbf{w}'(\mathbf{w}, e, \tau, a)$, where $\mathbf{w}'$ is the vector of utilities promised from tomorrow onward, which is restricted to lie in $\mathbf{W}$. Assuming that the value function V is known (it needs to be computed in practice), the auxiliary planning problem can be solved by solving a static optimization problem for all vectors in $\mathbf{W}$. An optimal contract for the non-recursive auxiliary planning problem can be found by assembling the appropriate solutions of the static problem.

[9] The state space would have to be extended further if the action affected outcomes for more than one period into the future.

We still need to determine which constraints need to be placed on the static choices $\pi(\tau, a|\mathbf{w}, e)$ and $\mathbf{w}'(\mathbf{w}, e, \tau, a)$ in order to guarantee that the implied contract satisfies the probability measure constraints (6), the maximization constraint (8), and the promise keeping constraint (10) above. In order to reproduce (6), we need to impose:

$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) = 1. \tag{14}$$

The promise-keeping constraint (10) will be satisfied if we impose:

$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) \left[u(e + \tau - a) + \beta \sum_{E} \mu(e'|a) w'(\mathbf{w}, e, \tau, a)(e') \right] = w(e) \tag{15}$$

where along the equilibrium path, honesty and obedience prevails in reports e and actions a. Note that $w'(\mathbf{w}, e, \tau, a)(e')$ is the appropriate scalar utility promise for e'. The incentive constraints are imposed in two parts. We first require that the agent cannot gain by following another action strategy $\delta_a(\tau, a)$, assuming that the reported income e was correct. Note that e enters the utility function as the actual value and as the conditioning element in π as the reported value.

$$\forall \delta_a :$$
$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) \left[u(e + \tau - \delta_a(\tau, a)) + \beta \sum_{E} \mu(e'|\delta_a(\tau, a)) w'(\mathbf{w}, e, \tau, a)(e') \right] \leq w(e). \tag{16}$$

A similar constraint on disobedience is also required if the initial report was e, but the true state was $\hat{e}$, i.e., false reporting. Note that $\hat{e}$ enters the utility function as the actual value but e is the conditioning element in π on the left-hand side, and $w(\hat{e})$ is the on-path realized utility under honesty and obedience at $\hat{e}$.

$$\forall \hat{e} \neq e, \delta_a :$$
$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) \left[u(\hat{e} + \tau - \delta_a(\tau, a)) + \beta \sum_{E} \mu(e'|\delta_a(\tau, a)) w'(\mathbf{w}, e, \tau, a)(e') \right] \leq w(\hat{e}). \tag{17}$$

Conditions (16) and (17) impose a sequence of period-by-period incentive constraints on the implied full contract. The constraints rule out that the agent can gain from disobedience or misreporting in any period, given that he goes back to truth-telling and obedience from the next period on. Equations (16) and (17) therefore imply that (8) holds for one-shot deviations. We still have to show that (16) and (17) are sufficient to prevent deviations in multiple periods, but the argument follows as in Phelan and Townsend (1991). That is, for a finite number of deviations, we can show that the original constraints are satisfied by backward induction. The agent clearly does not gain in the last period when he deviates, since this is just a one-time deviation and by (16) and (17) is not optimal. Going back one period, the agent has merely worsened his future expected utility by lying or being disobedient in the last period. Since one-shot deviations do not improve utility, the agent cannot make up for this. Going on this way, we can show by induction that any finite number of deviations does not improve utility. Lastly, consider an infinite number of deviations. Let us assume that there is a deviation that gives a gain of ϵ. Since $\beta < 1$, there is a period T such that at most $\epsilon/2$ utils can be gained from period T on. This implies that at least $\epsilon/2$ utils have to be gained until period T. But this contradicts our result that there cannot be any gain from deviations with a finite horizon.

Thus we are justified to pose the auxiliary planning problem as solving:

$$V(\mathbf{w}, e) = \max_{\pi \geq 0, \mathbf{w}'} \sum_{A,T} \pi(\tau, a|\mathbf{w}, e) \left[-\tau + Q \sum_{E} \mu(e'|a) V(\mathbf{w}'(\mathbf{w}, e, \tau, a), e') \right] \tag{18}$$

by choice of π and $\mathbf{w}'$, subject to constraints (14) to (17) above. Program 1 below is a version of this problem with a discrete grid for promised utilities as an approximation. We have assumed that the function $V(\mathbf{w}, e)$ is known. In practice, $V(\mathbf{w}, e)$ can be computed with standard dynamic programming techniques. Specifically, the right-hand side of (18) defines an operator T that maps functions $V(\mathbf{w}, e)$ into $TV(\mathbf{w}, e)$. It is easy to show, as in Phelan and Townsend (1991), that T maps bounded continuous functions into bounded continuous functions, and that T is a contraction. It then follows that T has a unique fixed point, and the fixed point can be computed by iterating on the operator T.

The preceding discussion was based on the assumption that the set $\mathbf{W}$ of feasible utility vectors is known in advance. In practice, $\mathbf{W}$ is not known and needs to be computed alongside the value function $V(\mathbf{w}, e)$. $\mathbf{W}$ can be computed with the dynamic-programming methods described in detail in Abreu, Pearce, and Stacchetti (1990). An outline of the method is given in Appendix A.2. The appendix also discusses issues related to the discretization of the set of feasible utility vectors which is necessary for numerical implementation of our algorithm.

Finally, the entire discussion is easily specialized to the case of a finite horizon T. V_T would be the value function for period T, V_{T-1} for period $T-1$, $\mathbf{W_{T-1}}$ the set of feasible promised utilities at time $T-1$, and so on.

4.3 The Discretized Version

For numerical implementation of the recursive formulation of the planning problem, we require finite grids for all choice variables in order to employ linear programming techniques. #E is the number of grid points for income, #T is the number of possible transfers, and #A is the number of actions. The vector of promised utilities is also assumed to be in a finite set $\mathbf{W}$, and the number of possible choices is #$\mathbf{W}$. To stay in the linear programming framework, we let the planner choose a probability distribution over vectors of utility promises, instead of choosing a specific utility vector.[10] That is, τ, a, and $\mathbf{w}'$ are chosen jointly under π. Notice that while the finite grids for income, transfer, and action are features of the physical setup of the model, the finite grid for utility promises is merely a numerical approximation of the continuous set in our theoretical formulation (see Appendix A.2 for a discussion of this issue).

With finite grids, the optimization problem **Program 1** of a planner who has promised vector $\mathbf{w}$ and *has received report* e is:

$$V(\mathbf{w}, e) = \max_{\pi \geq 0} \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[-\tau + Q \sum\nolimits_E \mu(e'|a) V(\mathbf{w}', e')\right] \tag{19}$$

subject to the constraints (20) to (23) below. The first constraint is that the $\pi(\cdot)$ sum to one to form a probability measure, as in (14):

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) = 1. \tag{20}$$

Second, the contract has to deliver the utility that was promised for state e, as in (15):

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(e + \tau - a) + \beta \sum\nolimits_E \mu(e'|a) w'(e')\right] = w(e). \tag{21}$$

[10] This imposes no loss of generality. If the contract puts weight on more than one utility vector, the corresponding mixed contract is a feasible choice for the planner who wants to implement the implied expected utility vector. The planner therefore cannot do better by choosing lotteries. It is also not possible to do worse, since the planner is free to place all weight on just one utility vector.

Third, the agent needs incentives to be obedient. Corresponding to (16), for each transfer τ and recommended action a, the agent has to prefer to take action a over any other action $\hat{a} \neq a$:

$$\forall \tau, a, \hat{a} \neq a: \quad \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(e + \tau - \hat{a}) + \beta \sum\nolimits_E \mu(e'|\hat{a}) w'(e')\right] \leq \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(e + \tau - a) + \beta \sum\nolimits_E \mu(e'|a) w'(e')\right]. \tag{22}$$

Finally, the agent needs incentives to tell the truth, so that no agent with income $\hat{e} \neq e$ would find this branch attractive. Under the promised utility vector $\mathbf{w}$, agents at $\hat{e}$ should get $w(\hat{e})$. Thus, an agent who actually has income $\hat{e}$ but says e nevertheless must not get more utility than was promised for state $\hat{e}$. This has to be the case regardless whether the agent follows the recommendations for the action or not. Thus, for all states $\hat{e} \neq e$ and all functions $\delta : T \times A \to A$ mapping transfer τ and recommended action a into an action $\delta(\tau, a)$ actually taken, we require as in (17):

$$\forall \hat{e} \neq e, \delta: \quad \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \delta(\tau, a)) + \beta \sum\nolimits_E \mu(e'|\delta(\tau, a)) w'(e')\right] \leq w(\hat{e}). \tag{23}$$

Note that similar constraints are written for the $\hat{e}$ problem, so that agents with $\hat{e}$ receive $w(\hat{e})$ from a constraint like (21). For a given vector of utility promises, there are $\#E$ Program 1's to solve.

Program 1 allows us to numerically solve the auxiliary planning problem for a given vector of utility promises by using linear programming and iteration on the value function. To recover the original planning problem with a scalar utility promise W_0, we let the planner offer a lottery $\pi(\mathbf{w}|W_0)$ over utility vectors $\mathbf{w}$ before the first period starts and before e is known. The problem of the planner at this initial stage is:

$$V(W_0) = \max_{\pi \geq 0} \sum_{\mathbf{W}} \pi(\mathbf{w}|W_0) \left[\sum\nolimits_E \mu(e) V(e, \mathbf{w})\right] \tag{24}$$

subject to a probability and a promise-keeping constraint:

$$\sum_{\mathbf{W}} \pi(\mathbf{w}|W_0) = 1. \tag{25}$$

$$\sum_{\mathbf{w}} \pi(\mathbf{w}|W_0) \left[\sum\nolimits_E \mu(e)\, w(e)\right] \geq W_0. \tag{26}$$

The same methods can be used for computing models with finitely many periods. With finitely many periods, the value functions carry time subscripts. The last period T would be computed first by solving Program 1 with all terms involving utility promises omitted. The computed value function $V_T(\mathbf{w}, e)$ for period T is then an input in the computation of the value function for period $T-1$. Moving backward in time, the value function for the initial period is computed last.

An important practical limitation of the approach outlined thus far is that the number of truth-telling constraints in Program 1 is very large, which makes computation practically infeasible even for problem with relatively small grids. For each state $\hat{e}$ there is a constraint for each function $\delta : T \times A \to A$, and there are $(\#A)^{(\#T \times \#A)}$ such functions. Unless the grids for τ and a are rather sparse, memory problems make the computation of this program infeasible. The total number of variables in this formulation, the number of objects under $\pi(\cdot)$, is $\#T \times \#A \times \#\mathbf{W}$. There is one probability constraint (20) and one promise-keeping constraint (21). The number of obedience constraints (22) is $\#T \times \#A \times (\#A - 1)$, and the number of truth-telling constraints (23) is $(\#E - 1) \times (\#A)^{(\#T \times \#A)}$. Thus, the number of constraints grows exponentially with the product of the grid sizes for actions and transfers.

As an example, consider a program with two states e, ten transfers τ, two actions a, and ten utility vectors $\mathbf{w}'$. With only ten points, the grids for transfers and utility promises are rather sparse. Still, for a given vector of utility promises $\mathbf{w}$ and reported state e Program 1 is a linear program with 200 variables and 1,048,598 constraints. If we increase the number of possible actions a to ten, the number of truth-telling constraints alone is 10^{100}. Clearly, such programs will not be computable now or any time in the future. It is especially harmful that the grid size for the transfer causes computational problems, as it does here because of the dimensionality of $\delta(\tau, a)$. One can imagine economic environments in which there are only a small number of options for actions available, but it is much harder to come up with a reason why the planner should be restricted to a small set of transfers. In the next section we present an alternative formulation, equivalent to the one developed here, that can be used for computing solutions in environments with many transfers and many actions.

5 Off-Path Utility Bounds and Efficient Computation

In this section we develop an alternative recursive formulation of the planning problem which is equivalent to Program 1, but requires a much smaller number of constraints. The key method for reducing the number of constraints in the program is to allow the planner to specify utility promises off the equilibrium path, as proposed by Prescott (2003).

The choice variables in the new formulation include utility bounds $v(\cdot)$ that specify the maximum utility (that is, current utility plus expected future utility) that an agent can get when lying about income and receiving a certain recommendation. Specifically, for a given reported income e, $v(\hat{e}, e, \tau, a)$ is an upper bound for the utility of an agent who actually has income $\hat{e} \neq e$, reported income e nevertheless, and received transfer τ and recommendation a. An intermediate step would be to assign exact utilities for $\hat{e}, e, \tau, a$, in which case we have something like the threat keeping in Fernandes and Phelan (2000). But as in Prescott (2003), we do better by only imposing utility bounds. This utility bound is already weighted by the probability of receiving transfer τ and recommendation a. Thus, in order to compute the total expected utility that can be achieved by reporting e when the true state is $\hat{e}$, we simply have to sum the $v(\hat{e}, e, \tau, a)$ over all possible transfers τ and recommendations a. The truth-telling constraint is then that the utility of saying e when being at state $\hat{e}$ is no larger than the utility promise $w(\hat{e})$ for $\hat{e}$.

The planner's optimization problem **Program 2** given report e and promised utility vector **w** is:

$$V(\mathbf{w}, e) = \max_{\pi \geq 0, v} \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[-\tau + Q \sum\nolimits_E \mu(e'|a) V(\mathbf{w}', e')\right] \tag{27}$$

subject to the probability measure constraint (20), the promise-keeping constraint (21), the obedience constraints (22), and the constraints (28) and (29) below. The first new constraint requires that the utility bounds have to be observed. An agent who reported state e, is in fact at state $\hat{e}$, received transfer τ, and got the recommendation a, cannot receive more utility than $v(\hat{e}, e, \tau, a)$, where again $v(\hat{e}, e, \tau, a)$ incorporates the probabilities of transfer τ and recommendation a. For each state $\hat{e} \neq e$, transfer τ, recommendation a, and all possible actions $\hat{a}$ we require:

$$\forall \hat{e} \neq e, \tau, a, \hat{a}: \quad \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \hat{a}) + \beta \sum\nolimits_E \mu(e'|\hat{a}) w'(e')\right] \leq v(\hat{e}, e, \tau, a). \tag{28}$$

Also, the truth-telling constraints are that the utility of an agent who is at state $\hat{e}$ but reports e cannot be larger than the utility promise for $\hat{e}$. For each $\hat{e} \neq e$ we require:

$$\forall \hat{e} \neq e: \quad \sum_{T,A} v(\hat{e}, e, \tau, a) \leq w(\hat{e}). \tag{29}$$

Thus, the new constraints (28) and (29) in Program 2 replace the truth-telling constraints (23) in Program 1. The number of variables in this problem is $\#T \times \#A \times \#\mathbf{W}$ under $\pi(\cdot)$ plus $(\#E-1) \times \#T \times \#A$, where the latter terms reflect the utility bounds $v(\cdot)$ that are now choice variables. There is one probability constraint (20) and one promise-keeping constraint (21). The number of obedience constraints (22) is $\#T \times \#A \times (\#A - 1)$. There are $(\#E - 1) \times \#T \times (\#A)^2$ constraints (28) to implement the utility bounds, and $(\#E-1)$ truth-telling constraints (29). Notice that the number of constraints does not increase exponentially in any of the grid sizes. The number of constraints is approximately quadratic in $\#A$ and approximately linear in all other grid sizes. This makes it possible to compute models with a large number of actions. In Program 2, our example with two states e, ten transfers τ, two actions a, and ten utility vectors $\mathbf{w}'$ is a linear program with 220 variables and 63 constraints. The program is still computable if we increase the number of actions a to ten. In that case, Program 2 has 1100 variables and 1903 constraints, which is sufficiently small to be easily solved on a personal computer.

We now want to show that Program 2 is equivalent to Program 1. In both programs, the planner chooses lotteries over transfer, action, and promised utilities. Even though in Program 2 the planner also chooses utility bounds, in both programs the planner's utility depends only on the lotteries, and not on the bounds. The objective functions are identical. In order to demonstrate that the two programs are equivalent, it is therefore sufficient to show that the set of feasible lotteries is identical. We therefore have to compare the set of constraints in the two programs.

Proposition 5 *Program 1 and Program 2 are equivalent.*

The proof (contained in the Appendix) consists of showing that constraints (28) and (29) in Program 2 place the same restrictions on the outcome function $\pi(\cdot)$ as the constraints (23) of Program 1. The probability constraint (20), the promise-keeping constraint (21), and the obedience constraints (22) are imposed in both programs. Therefore one only needs to show that for any $\pi(\cdot)$ that satisfies the incentive constraints (23) in Program 1, one can find utility bounds such that the same outcome function satisfies (28) and (29) in

Program 2 (and vice versa). Since the objective function is identical, it then follows that the programs are equivalent.

In Program 2, the number of constraints gets large if both the grids for transfer τ and action a are made very fine. In practice, this may lead to memory problems when computing. Further reductions in the number of constraints are possible in a formulation in which the transfer and the recommendation are assigned at two different stages. Here we subdivide the period completely and assign interim utility promises as an additional state variable. This procedure yields even smaller programs. The reduction comes at the expense of an increase in the number of programs that needs to be computed. The programs for both stages are described in Doepke and Townsend (2002). The main remaining limitation of our approach is that we are restricted to relatively small spaces for the state e (for example, the examples in Section 7 have two possible states). If there are many states, the number of feasible utility vectors in $\mathbf{W}'$ becomes large, which increases both the number of constraints and the number of programs that need to be computed.

Even in the relatively efficient Program 2, the linear programs generally have the feature that the number of constraints far exceeds the number of variables, which implies that a large number of constraints are not binding at the optimum. One may wonder whether there are situations where it can be determined beforehand which constraints are not going to be binding, which would lead to further computational simplifications. In some applications such simplification may indeed be possible. For example, Thomas and Worrall (1990) show in an environment similar to ours, but without lotteries and without hidden actions, that the truth-telling constraint is binding for the high-income agent, but not for the low-income agent. Under sufficiently restrictive assumptions on utility and technology, similar results could be obtained for our environment.

In terms of reducing the size of the linear programs, however, it would be far more valuable to show that it is sufficient to check only "local" deviations in terms of the action (since the number of constraints is asymptotically linear in the number of incomes, but quadratic in the number of actions). Here we find that under standard assumptions (in particular, concave utility) there is only limited scope for further improvements. To see why, consider an environment with two incomes (high and low), concave utility, and an investment technology that is linear in storage (i.e., $\mu(e|a)$ is linear in a). Consistent with the storage interpretation, A is assumed to be a finite set of non-negative values that includes zero storage. Finally, assume that it can be shown in advance that the optimal contract will assign zero storage to the agent in every state, as is typically the case of the return on storage is lower that the planner's credit market return. Will it be possible to implement the optimal contract by checking only the "local" constraints, i.e., constraints that involve deviations to the lowest positive storage level? As far as the obedience constraints (22) are concerned, the answer is yes. In (22), the marginal gain from deviating to an $\hat{a} > 0$ (higher future utility) is constant regardless of $\hat{a}$ (since $\mu(e|\hat{a})$ is linear), whereas the marginal loss from deviating (lower current utility) is increasing in $\hat{a}$ (since the utility function is concave). If (22) is satisfied for a given $\hat{a}^\star$, it is therefore also satisfied for all $\hat{a} > \hat{a}^\star$, so checking the local constraint indeed suffices to implement the obedience constraints.

The same argument does not work, however, for the truth-telling constraints. One might think that the constraints (28) once again only need to be imposed for the lowest possible $\hat{a}$. However, the problem is that we do not know a priori which $\hat{a}$ maximizes the left-hand side of (28), since (28) describes behavior off the equilibrium path. The optimal storage level of the equilibrium path will generally not be equal to zero. Consider a rich agent who lies and pretends to be poor. Under the optimal contract, relative to telling the truth such an agent would receive a higher transfer today, but lower future utility. Thus, current marginal utility would decrease, while future marginal utility would increase, which makes storage particularly attractive. In other words, the setup is such that a *joint* deviation in report and action is more attractive than a single deviation in the report. This complementarity between the two types of deviation makes it impossible to simplify the program by considering only local constraints when implementing truth-telling.

While this may seem like an unfortunate outcome, this result underlines the strengths of our methods relative to the first-order approach. By construction, in the first-order approach only local deviations are checked (the difference being that the size of the deviation is infinitesimal). If, however, there are complementarities between different types of deviations, the local constraints are not sufficient, and the first-order approach is not applicable. Kocherlakota (2004) points out that such complementarities generically arise in a number of environments, and the discussion above shows that the same concern is important in the hidden-storage world considered here. Our methods provide a tractable approach for solving this class of problems.

6 Multiple Actions and Observable States

Up to this point, we have restricted our analysis to an environment with a single state (the beginning-of-period income) and a single action (the storage decision at the end of the period). In many applications, there are multiple states realized throughout the period, some of which may be observed, and there are multiple actions. Consider, for example, the canonical unemployment insurance problem with hidden savings, versions of which are analyzed by Werning (2001) and Abraham and Pavoni (2002). There are two states, namely beginning-of-period assets and the outcome of the job search. It is typically assumed that the employment status is observable, while assets are not. There are also two unobserved actions, namely the job-search intensity and the savings or storage decision. In this section, we show how our methods can be naturally extended to richer environments of this kind. As we will see, additional observable states do not add any substantial complications. Multiple unobserved actions can be handled by introducing multiple *layers* of utility bounds, each corresponding to a different action. In effect, at the the first stage the planner makes "promises on promises."

Figure 2 sums up the timing of events within the period. In the extended model, there is still a beginning-of-period state e. After e is realized, the agent takes an unobserved action a_1. Next, an observable state y which depends on e and a_1 through probability function $f(y|e, a_1)$ is realized. In the unemployment insurance model, a_1 is typically interpreted as search effort, and y is the outcome of the search (success or failure). After y is realized, the timing is as in our original setup. The planner hands out a transfer τ, the agent takes a second unobserved action a_2 (which can be interpreted as storage), and finally tomorrow's state e' is determined by the probability function $g(e'|e, a_1, y, \tau, a_2)$, which has full support over tomorrow's state space E. Notice that we allow income to be autocorrelated and to depend on all variables realized within the period. The period utility function is more general as well and takes the form $u(e, a_1, y, \tau, a_2)$. As before, we assume that every income and every state always occurs with positive probability:

Assumption 3 *The probability distribution over the state e satisfies $g(e'|e, a_1, y, \tau, a_2) > 0$ for all $e \in E$, $a_1 \in A_1$, $y \in Y$, $\tau \in T$, and $a_2 \in A_2$. The probability distribution over the state y satisfies $f(y|e, a_1) > 0$ for all $e \in E$ and $a_1 \in A_1$.*

In this more general environment, versions of Proposition 2 (revelation principle) and Propositions 3 and 4 (recursive mechanism) can be proven, with little change in nota-

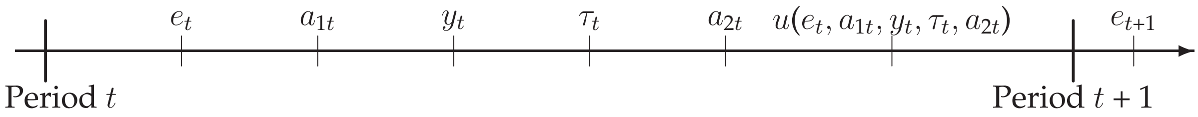

Figure 2: The Sequence of Events in Period t

tion, by following the same steps as in the single-action environment. We therefore limit ourselves to outlining the main steps here, without going through every detail. After introducing message spaces corresponding to the state and the two actions, in the double-action environment the history at time t can be expressed at:

$$h_t \equiv \{e_t, m_{1t}, m_{2t}, a_{1t}, y_t, \tau_t, m_{3t}, a_{2t}\},$$

where m_{1t} is a message sent by the agent which serves as a report of the state e, whereas m_{2t} and m_{3t} are messages sent by the planner which serve as a recommendations for the two actions. The part of the history observable by the planner is given by:

$$s_t \equiv \{m_{1t}, m_{2t}, y_t, \tau_t, m_{3t}\}.$$

The histories h^t and s^t up to time t, strategies σ and outcome functions π, and probabilities over histories are defined as in Section 3. The definitions of optimal strategies, equilibria, and optimal equilibria (Definitions 1 to 3 above) also carry over to the more general case. As in Section 4, the next step is then to define a more restricted direct mechanism with truth-telling and obedience. In the restricted mechanism, message spaces are equal to the spaces that each report corresponds to, so that we have $M_1 = E$, $M_2 = A_1$, and $M_3 = A_2$. The contemporary part of the history of the planner is then given by:

$$s_t \equiv \{e_t, a_{1t}, y_t, \tau_t, a_{2t}\},$$

where e_t is interpreted as a report, and a_{1t} and a_{2t} are recommendations. A mechanism with truth-telling and obedience can be written as $\pi(a_{1t}, y_t, \tau_t, a_{2t}|e_t, s^{t-1})$, and as in Section 4 this mechanism has to satisfy a probability constraint, a promise-keeping constraint, and a set of incentive constraints. Since y_t is drawn by nature, there is an additional set of constraints which requires that the choices of the planner are consistent with the probability function $f(y|e, a_1)$. The revelation principle is proven, as in Proposition 2, by taking an arbitrary pair of an outcome function and a corresponding equilibrium strategy, and "translating" this pair into a corresponding outcome function $\pi^\star$ under truth-telling and obedience. In our case, this translation is given by:

$$\pi^\star(a_{1t}, y_t, \tau_t, a_{2t}|e_t, s^{t-1}) \equiv \sum_{H^{t-1}(s^{t-1}), M_1, M_2, M_3} p(h_{t-1}|s^{t-1})\, \sigma(m_{1t}|e_t, h^{t-1})\pi(m_{2t}|m_{1t}, s^{t-1})$$
$$\sigma(a_{1t}|e_t, m_{1t}, m_{2t}, h^{t-1})\, f(y_t|e_t, a_{1t})\, \pi(\tau_t|m_{1t}, m_{2t}, y_t, s^{t-1})$$
$$\pi(m_{3t}|m_{1t}, m_{2t}, y_t, \tau_t, s^{t-1})\, \sigma(a_{2t}|e_t, m_{1t}, m_{2t}, a_{1t}, y_t, \tau_t, m_{3t}, h^{t-1}).$$

Since the pair (π, σ) is an equilibrium, it can be shown that $\pi^\star$ satisfies all required constraints as in Proposition 2, so that using the direct mechanism does not impose any loss of generality.

The next step towards a computable version of the planning problem is to show that the problem is recursive in a vector of state-contingent utility promises (Proposition 4). Here the arguments made for the single-action environment in Section 4 carry over almost directly. To be more specific, in our baseline environment a vector of state-contingent utility promises is a sufficient state variable because past actions influence future outcomes *only* through their impact on the future state e'. Once e' is realized, past actions do not have further effects. This is the key property that we exploited in the proof of Proposition 4 to show that it is sufficient to condition utility promises on e'. The extended two-action environment adds some new complications within the period, but it preserves the property that events in any given period affect future outcomes only through their influence on the next period's state e'. Consequently, the same arguments that we made in the proof of Proposition 4 also pertain to the extended environment.

So far, we have argued that the same steps that led to Program 1 in the single-action environment carry over with little changes to our extended double-action environment. The final step of our earlier analysis, leading to Program 2, dealt with the "curse of dimensionality" in terms of the within-period incentive constraints. In particular, we showed how off-path utility bounds could be used to dramatically reduce the number of incentive constraints. While we will see that an extension of this method applies to the double-action environment, it is here where the two environments are the most different. We therefore discuss the last step of the analysis in detail. We start be writing down the planning problem arising in this setup in the standard form, akin to Program 1. The optimization problem **Program 3** given report e and promised utility vector $\mathbf{w}$ is:

$$V(\mathbf{w}, e) = \max_{\pi \geq 0} \sum_{A_1,Y,T,A_2,\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \left[-\tau + Q \sum\nolimits_E g(e'|e, a_1, y, \tau, a_2) V(\mathbf{w}', e')\right] \tag{30}$$

subject to the constraints below. The first constraint is the probability measure constraint:

$$\sum_{A_1,Y,T,A_2,\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) = 1. \tag{31}$$

The second constraint is the promise-keeping constraint:

$$\sum_{A_1,Y,T,A_2,\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \left[u(e, a_1, y, \tau, a_2) + \beta \sum\nolimits_E g(e'|e, a_1, y, \tau, a_2) w'(e')\right] = w(e). \tag{32}$$

Third, the contract has to obey the true probability distribution $f(y|e, a_1)$:

$$\forall \bar{a}_1 \in A_1, \bar{y} \in Y : \quad \frac{\sum_{T,A_2,\mathbf{W}'} \pi(\bar{a}_1, \bar{y}, \tau, a_2, \mathbf{w}'|\mathbf{w}, e)}{\sum_{Y,T,A_2,\mathbf{W}'} \pi(\bar{a}_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e)} = f(\bar{y}|e, \bar{a}_1). \tag{33}$$

Fourth, we have to ensure that the agent prefers to take the correct first action a_1, *conditional on the report being true*:

$$\begin{aligned} \forall a_1, \hat{a}_1 \neq a_1, \delta_2 : \quad & \sum_{Y,T,A_2,\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \frac{f(y|e, \hat{a}_1)}{f(y|e, a_1)} \\ & \left[u(e, \hat{a}_1, y, \tau, \delta_2(a_1, y, \tau, a_2)) + \beta \sum\nolimits_E g(e'|e, \hat{a}_1, y, \tau, \delta_2(a_1, y, \tau, a_2)) w'(e')\right] \\ \leq & \sum_{Y,T,A_2,\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \left[u(e, a_1, y, \tau, a_2) + \beta \sum\nolimits_E g(e'|e, a_1, y, \tau, a_2) w'(e')\right]. \end{aligned} \tag{34}$$

Notice that we have to allow arbitrary deviations in the second action at this stage, since we do not know which second action is optimal if the first action is not the correct one.

Fifth, conditional on having been truthful and obedient up to that point, the agent needs incentives to take the correct second action a_2:

$$\forall a_1, y, \tau, a_2, \hat{a}_2 \neq a_2 :$$
$$\sum_{\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \left[u(e, a_1, y, \tau, \hat{a}_2) + \beta \sum\nolimits_E g(e'|e, a_1, y, \tau, \hat{a}_2) w'(e')\right]$$
$$\leq \sum_{\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \left[u(e, a_1, y, \tau, a_2) + \beta \sum\nolimits_E g(e'|e, a_1, y, \tau, a_2) w'(e')\right]. \quad (35)$$

Sixth, we need to provide incentives for making the correct report, regardless of the action strategies δ_1 and δ_2 to be followed later:

$$\forall \hat{e} \neq e, \delta_1, \delta_2 :$$
$$\sum_{A_1, Y, T, A_2, \mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \frac{f(y|\hat{e}, \delta_1(a_1))}{f(y|e, a_1)}$$
$$\left[u(\hat{e}, \delta_1(a_1), y, \tau, \delta_2(a_1, y, \tau, a_2)) + \beta \sum\nolimits_E g(e'|\hat{e}, \delta(a_1), y, \tau, \delta_2(a_1, y, \tau, a_2)) w'(e')\right] \leq w(\hat{e}). \quad (36)$$

The number of variables in this problem is $\#A_1 \times \#Y \times \#T \times \#A_2 \times \#\mathbf{W}$. The number of constraints is prohibitively high even for small grid sizes: There are $\#E \times \#A_1^{\#A_1} \times \#A_2^{\#A_1 \times \#Y \times \#T \times \#A_2}$ truth telling constraints (36) alone. The key problem is that there are *two* arbitrary deviation functions δ_1 and δ_2 at different stages.

In the single-action environment we were able to substantially reduce the number of constraints by introducing off-path utility bounds, and the same approach is going to be useful here. We will show, however, that to do deal with two actions, two *layers* of utility bounds are required, one corresponding to each action. The first layer, expressed as $v_1(\hat{e}, a_1)$, provides an upper bound on the utility of an agent who is at state $\hat{e}$ but reports e instead, and subsequently receives the recommendation a_1 for the first action. These bounds will be used to simplify the truth-telling constraints as in Program 2 above. The second layer can be written as $v_2(\hat{e}, a_1, \hat{a}_1, y, \tau, a_2)$ and enables us to implement the obedience constraints. Here $v_2(\hat{e}, a_1, \hat{a}_1, y, \tau, a_2)$ bounds the utility of an agent at true state $\hat{e}$ who reports e, receives recommendation a_1, does $\hat{a}_1$, reaches state y, receives transfer τ, and gets the second recommendation a_2.

Using these bounds, the planner's optimization problem **Program 4** given report e and promised utility vector $\mathbf{w}$ is:

$$V(\mathbf{w}, e) = \max_{\pi \geq 0, v_1, v_2} \sum_{A_1, Y, T, A_2, \mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \left[-\tau + Q \sum\nolimits_E g(e'|e, a_1, y, \tau, a_2) V(\mathbf{w}', e')\right] \tag{37}$$

subject to the probability measure constraint (31), the promise-keeping constraint (32), the consistency constraints (33), the obedience constraint for the second action (35), and the new constraints (38)-(41) below. The first new constraint is the obedience constraint for the first action:

$$\forall a_1, \hat{a}_1 \neq a_1: \quad \sum_{Y,T,A_2} v_2(e, a_1, \hat{a}_1, y, \tau, a_2)$$
$$\leq \sum_{Y,T,A_2,\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \left[u(e, a_1, y, \tau, a_2) + \beta \sum\nolimits_E g(e'|e, a_1, y, \tau, a_2) w'(e')\right]. \tag{38}$$

The second constraint provides incentives for making the correct report:

$$\forall \hat{e} \neq e: \quad \sum_{A_1} v_1(\hat{e}, a_1) \leq w(\hat{e}). \tag{39}$$

Next, the first-layer utility bounds actually have to bound utility, regardless of the first action $\hat{a}_1$ actually undertaken:

$$\forall \hat{e} \neq e, a_1, \hat{a}_1: \quad \sum_{Y,T,A_2} v_2(\hat{e}, a_1, \hat{a}_1, y, \tau, a_2) \leq v_1(\hat{e}, a_1). \tag{40}$$

The second-layer bounds have to be true bounds as well:

$$\forall \hat{e}, a_1, \hat{a}_1, y, \tau, a_2, \hat{a}_2:$$
$$\sum_{\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \frac{f(y|\hat{e}, \hat{a}_1)}{f(y|e, a_1)}$$
$$\left[u(\hat{e}, \hat{a}_1, y, \tau, \hat{a}_2) + \beta \sum\nolimits_E g(e'|\hat{e}, \hat{a}_1, y, \tau, \hat{a}_2) w'(e')\right] \leq v_2(\hat{e}, a_1, \hat{a}_1, y, \tau, a_2). \tag{41}$$

The reason why two layers of utility bounds are required is that once the agent lies about the state ($\hat{e} \neq e$) and we are therefore off the equilibrium path, we do not know which first action $\hat{a}_1$ is optimal for a given recommendation a_1. If we did, it would be possible in (39) (the only constraint where the v_1 bounds are used) to sum over the corresponding v_2 bounds directly. Instead, we have to write the constraint (39) using v_1 bounds, and then impose a separate set of constraints (40) and (41) to make sure that the v_1 are true bounds on utility for *any* first action $\hat{a}_1$.

Proposition 6 *Program 3 and Program 4 are equivalent.*

The proof (contained in the Appendix) follows the same approach as the proof of Proposition 5.

In Program 4 the number of constraints is asymptotically proportional to the number of possible actions (at each stage) squared, whereas in the version without utility bounds the number of constraints grows exponentially with the number of actions. The number of variables in Program 4 is $[\#A_1 \times \#Y \times \#T \times \#A_2 \times \#\mathbf{W}] + [(\#E - 1) \times \#A_1] + [\#E \times \#A_1^2 \times \#Y \times \#T \times \#A_2]$, where the last two terms correspond to the utility bounds v_1 and v_2. There is one probability constraint (31) and one promise-keeping constraint (32). There are $\#A_1 \times \#Y$ consistency constraints (33), $\#A_1 \times \#Y \times \#T \times \#A_2 \times (\#A_2 - 1)$ obedience constraints for the second action (35), $\#A_1 \times (\#A_1 - 1)$ obedience constraints for the first action (38), $\#E - 1$ truth telling constraints (39), $(\#E - 1) \times \#A_1^2$ constraints to implement the first-level utility bounds (40), and $\#E \times \#A_1^2 \times \#Y \times \#T \times \#A_2^2$ constraints to implement the second-level utility bounds (41) (including the cases $\hat{e} = e$ and $\hat{a}_1 = a_1$).

Program 4 is small enough to be computable on a standard PC for moderate grid sizes. If memory is an issue, the size of the programs could be further decreased by subdividing each period in two different parts. Also, the complexity of the resulting programs is sensitive to the precise timing of events in the period. We have assumed that the agent's second action a_2 takes place *after* the planner hands out the transfer τ, which implies that the utility bounds v_2 have to be conditional on τ. If in a given application the timing were such that the planner provides a transfer at the end of the period, the utility bounds would no longer depend on τ, which substantially reduces the size of the programs. The approach presented here could be further generalized by introducing yet more actions or states, where each action would result in a new layer of utility bounds.

7 Computed Examples

7.1 Hidden Storage

In this section, we use our methods to compute optimal allocations in a variant of the single-action environment underlying Programs 1 and 2, where the hidden action is designed to resemble storage at a fixed return. The planner has access to an outside credit market at a fixed interest rate, and wants to provide insurance to an agent who has hidden income shocks. The agent can hide part of the income, and invest in a technology which raises the expected income in the following period.

Our interest in this application stems from its relation to the existing literature: there are two special cases of this environment for which the optimal insurance mechanism can be derived analytically. Allen (1985) shows that if the hidden technology amounts to unobserved access to a perfect credit market at the same interest rate that the planner has access to, the principal cannot offer any insurance beyond self-insurance in the credit market. Cole and Kocherlakota (2001b) and Ljungqvist and Sargent (2003) extend this result to the case where the agent can only store the income, but not borrow. As long as the return on storage is identical to the planner's credit market return, the optimal outcome is once again equivalent to credit market access for the agent. If the results of Allen (1985), Cole and Kocherlakota (2001b), and Ljungqvist and Sargent (2003) held more generally, the implementation of optimal policies would be straightforward: the planner only would have to provide credit market access to the agent, which is much simpler than the history-contingent transfer schemes that arise in environments without hidden actions. However, all of the literature relies on the assumption that planner and agent face the same return. Our recursive methods make it possible to compute outcomes for a variety of returns of the hidden storage technology, both below and above the credit-market interest rate. We find in these computations that once the hidden storage technology has a return sufficiently different from the credit market interest rate, the equivalence of optimal insurance and credit market access breaks down. This is true regardless of whether the return on storage is above or below the interest rate; it is only required that the deviation be sufficiently large.

We illustrate this result in an environment with three periods. Planner and agent have the same discount factor $\beta = Q = 1.05^{-1}$, corresponding to a risk-free interest rate of five percent for the planner. The agent has logarithmic utility. In each period there are two possible incomes e (2.1 or 3.1) and three storage levels a (0, 0.3, or 0.6). In periods one and two, the planner can use a grid of 50 transfers τ (from -1.4 to 1.4). In period three outcomes are easier to compute since there are no further actions utility promises, which allows us to use a grid with 4,000 grid points for transfers. In each period, the utility grids have 100 elements per income, so that (given two incomes) there are 10,000 possible utility vectors. In the initial period, each income occurs with probability 0.5. In all other periods, the probability distribution over incomes depends on the storage level in the previous period. If storage is zero, the high income occurs with probability 0.025. The probability of the high income increases linearly with storage. The expected return on storage R determines how fast the probability of the high income increases.[11] For example, for $R = 1$, an increase in storage of 0.1 results in an increase in the expected income of 0.1. Since the difference between the two incomes equals one $(3.1 - 2.1)$, with $R = 1$ increasing the storage level a by 0.1 results in an increase in the probability of the high income of 3.1 by 0.1, and an equal decrease in the probability of the low income of 2.1.

We computed optimal policy functions for a variety of returns R as a function of the (scalar) initial utility promise W_0. Program 2 was used for all computations.[12] Figure 3 shows the policy functions for $R = 1.1$ as a function of the initial utility promise W_0 and the realized income in the first period (low or high, represented by solid and dashed lines in the graph). In all graphs a vertical line marks the initial utility level at which the expected discounted surplus of the planner is zero. At $R = 1.1$, the net expected return on storage is 10 percent, which exceeds the credit market interest rate of 5 percent.

[11] Notice that our storage technology is non-standard in the sense that it is stochastic. Storage increases the probability of a high income realization within a given distribution, as opposed to shifting the entire distribution upwards. However, comparing our computations to existing theoretical results indicates that this choice does not have a major effect on the results (see below).

[12] The computations were carried out on a PC with an Intel P4 2.1 GHz CPU, 512 MB of RAM, and running Windows XP. The matrices were set up in MATLAB, and the optimization steps were carried out using the ILOG CPLEX 7.5 linear program solver. In the last period (period 3), each linear program had 4,000 variables and 3 constraints. Each of these programs took about 0.1 seconds to compute. In period 2, the programs had 34,800 variables and 753 constraints, resulting in a computing time of 0.25 seconds per program, and in period 1 the programs had 200,100 variables and 753 constraints, with a computation time of 1.8 seconds per program. The variation in the number of variables between periods one and two stems from differences in the number of feasible utility vectors. At each stage, a separate program needs to be computed for each income and each feasible utility vector. The actual number of programs to be computed varies, since some infeasible utility vectors are disregarded from the outset. On average, 1500 programs had to be computed for period 3, and 3000 programs for periods 1 and 2. The total computation time averaged 3 minutes for period 3, 45 minutes for period 2, and 280 minutes for period 1. In all computed solutions to the initial planning problem (24), the surplus of the planner $V(W_0)$ is a strictly decreasing function of the utility promise W_0, implying that the promise-keeping constraint (26) is always binding.

Consequently, the storage technology is used to the maximum of 0.6 regardless of the realized income. The planner uses transfers and utility promises to provide incentive-compatible insurance to the agent with the low income realization. The transfer for the low income is higher than for the high income, and at zero expected surplus high-income agents today are taxed, while low-income agents receive a transfer. The difference in consumption between low- and high-income agents is only about 0.3, even though the income differs by 1. In other words, the planner is able to insure about seventy percent the income risk in the first period. In order to ensure incentive compatibility (i.e., induce the high-income agent to report truthfully), the planner promises a schedule of higher future utility from tomorrow onward to agents with the high income.

The optimal insurance scheme for returns R below 1.05 (the market rate) is similar to the outcome for R = 1.1, with the key difference that the private storage technology is not used. Interestingly, the planner still does not recommend the use of private storage if the return on storage is slightly above the credit market return. Even though using storage would increase the expected income, there would also be added uncertainty since the storage technology is stochastic. Using private storage pays off only if the return is sufficiently far above the credit-market return to compensate for the added risk.

Figure 4 shows how the expected utility of the agent varies with the return on storage, subject to the requirement that the discounted surplus of the planner be zero. The utility of the agent is shown for values of R between 0.6 and 1.1 (below $R = 0.6$, utilities no longer change with R). The top line is the utility the agent gets under full information, that is, when both income and storage are observable by the planner. In this case the planner provides full insurance, and consumption and utility promises are independent of income. The utility of the agent does not vary with the return on storage as long as $R \leq 1.05$. This is not surprising, since 1.05 is the risk-free interest rate. As long as the return on storage does not exceed the credit-market return, storage is never used. When R exceeds 1.05, the Pareto frontier expands and the utility of the agent increases dramatically. In effect, when $R > 1.05$ an arbitrage opportunity arises which is only limited by the upper bound on storage.

The lower line shows the utility of the agent under autarky. When the return on storage is sufficiently high (above about 0.91, still a relatively dismal return) the agent uses storage to self-insure against the income shocks. Since under autarky storage is the only insurance for the agent, utility increases with the return on storage.

The solid line shows the utility of the agent under optimal incentive-constrained insurance with hidden income and hidden actions. Once the return on storage exceeds some critical value (about 0.7), the utility of the agent decreases with the return on storage, instead of increasing as it does under autarky. As long as the planner never recommends positive private storage, a higher return on storage has no positive effects. On the contrary, it becomes harder to satisfy the obedience constraints which require that the agent does not prefer a positive storage level when the planner recommends zero storage. Raising the return on storage shrinks the set of feasible allocations, and the utility of the agent falls. Hence, the properties of the storage technology influence the optimal insurance contract even if private storage is never used in equilibrium.

The dotted line represents the utility of an agent who has access to a perfect credit market under the same conditions as the planner. The agent can also use the private storage technology in addition to borrowing and lending in the credit market. This utility is computed by solving for the optimal saving and borrowing decisions explicitly, i.e., no finite grids on consumption or assets are imposed. As long as the return on storage does not exceed the credit market interest rate, once again zero storage is optimal, and the utility is independent of the return on storage. When the return on storage exceeds the interest rate, the agent may gain by using the storage technology in addition to smoothing income in the credit market.

We know that, in principle, the information-constrained solution should be no worse that what the agent can achieve independently when given access to an outside credit market, since any borrowing-and-lending plan is incentive compatible. In other words, the optimal insurance scheme can reproduce anything that can be achieved in the credit market, so that the credit-market solution should provide a lower bound on the information-constrained solution. In Figure 4, the optimal constrained solution does dip below the credit-market solution for a small range of R, which can be attributed to the finite grids used to compute the information-constrained regime. Even so, the figure provides a measure of the accuracy of our code, as the difference between the solution and the lower bound is very small.

Thus, within the bounds of approximation our computations bear out the prediction by Cole and Kocherlakota (2001b) that the credit-market utility should equal the optimal-insurance utility at the point where the interest rate and return on storage are equal, despite the fact that our storage technology is stochastic. For returns sufficiently below $R = 1.05$, optimal information-constrained insurance does significantly better than credit market access. For $R < 1.05$, the utility of the agent increases with R under both regimes, but the information-constrained solution once again dominates for sufficiently high returns.

In sum, Figure 4 shows that if the return on private storage is sufficiently different from the outside interest rate, the agent does significantly better with optimal insurance than with borrowing and lending. What matters here is the ability of the agent to circumvent the insurance scheme offered by the planner. The planner provides incentives for truthful reporting by exploiting differences in the intertemporal rate of substitution between agents with high and low income realizations. When the return on storage is close to the credit-market interest rate, agents have the possibility of smoothing income on their own, so that differences in rates of substitution are small. Consequently, little additional insurance can be offered. When there is a large difference between return on storage and interest rate, however, it is costly for agents to deviate from the recommended storage scheme (no storage at low returns, maximum storage at high returns). The planner can use the resulting differences in rates of substitution to provide incentives for truthful reporting, which results in more insurance.[13]

The optimal allocations in this application are computed subject to two types of incentive constraints: obedience constraints force the agent to take the right action, and truth-telling constraints ensure that the agent reports the actual income. What happens if only one type of constraint is introduced? If we remove the truth-telling constraints (i.e., we assume that the actual income can be observed by the planner) and assume that the return on storage does not exceed the credit-market return, the situation is simple: The planner can provide full insurance against income shocks. In this case the obedience constraints are not binding since the optimal storage level is zero, which is also what the agent prefers if income shocks are fully insured. Conversely, if only the obedience constraints are removed, as if storage could be observed, full insurance cannot be achieved. The outcome is essentially the same as in the situation where storage is hidden, but the storage technology has zero gross return. The utilities are as in Figure 4 for $R = 0.6$.

[13]The exact gain from optimal insurance depends on the return on storage. At low returns, the quantitative gains relative to credit-market access are small. The maximum utility differential is equivalent to a permanent increase in consumption of slightly above one-twentieth of one percent. If the return to storage is so low that the storage technology is never used, self-insurance in a credit market achieves 95 percent of the utility gain of optimal insurance relative to autarky. The relative gains from optimal insurance are higher when storage is actually used. At the maximum return of $R = 1.1$, the credit-market solution yields only 65 percent of the utility gain relative to autarky that can be attained by optimal insurance. These numbers, of course, are particular to our example. Generally, the absolute gains achieved by optimal insurance increase with the amount of risk inherent in income realizations and the agent's degree of risk aversion.

	$e = 0, a_1 = 0$	$e = 0, a_1 = 1$	$e = 1, a_1 = 0$	$e = 1, a_1 = 1$
$P(y = 1)$	0.9	0.8	0.9	0.1
$P(y = 10)$	0.1	0.2	0.1	0.9

Table 1: Probability Distribution over Output y

7.2 Unobserved Effort with Hidden Productivity and Training

We now present another example to put our methods for computing the double-action environment in Section 6 to work. Since in this environment there are a number of different types of incentive constraints, we chose an extension of a relatively standard unobserved-effort economy to demonstrate how these different constraints interact. The state e in this economy will be labeled as "productivity." The first action a_1 is interpreted as effort, and a_1 together with the state e jointly determine the probability distribution over output y. The second unobserved action a_2 is interpreted as "training" or "skill acquisition," and determines the probability distribution over future productivity e'.

The agent enters the period with a productivity state e that can be either high or low (normalized to $e = 1$ and $e = 0$). The first unobserved action ("effort") a_1 can also be either high or low ($a_1 = 1$ or $a_1 = 0$). Next, (observable) income y is realized. Once again just two values are possible, $y = 10$ or $y = 1$. Table 1 shows the probability distribution over output as a function of e and a_1. Notice that when effort a_1 is low, the probability distribution over output does not depend on productivity e. A productive agent, on the other hand, gets a much larger increase in the probability of getting high output by exerting high effort than an unproductive agent. After output is realized, the planner hands out a transfer τ. The second unobserved action a_2 ("training") takes place at the end of the period, and can take the two values $a_2 = 1$ and $a_2 = 0$. For a worker who does the training ($a_2 = 1$), the probability of being productive in the next period ($e' = 1$) is 0.9, while without training ($a_2 = 0$) the probability is only 0.1. The utility function of the agent is given by:

$$u(c, a_1, a_2) = \sqrt{c} - \alpha_1 a_1 - \alpha_2 a_2,$$

where $\alpha_1 = 1$ and $\alpha_2 = 0.2$. Thus, preferences are separable in consumption and the two effort choices (this simplifies interpreting the results; our computational method does not require separability assumptions).

In summary, we are considering a fairly standard unobserved effort model, which is enriched by the second unobserved action "training" and the unobserved state "productivity."[14] To map out how the optimal solution of the planning problem is affected by the different incentive problems, we computed solutions under different assumptions on observability, ranging from full information to the case where e, a_1, and a_2 are all unobserved. For simplicity, a two-period model was computed, and we used Program 4 for all computations (while omitting some of the incentive constraints when some or all of e, a_1, and a_2 are observable).[15] In the first period, each productivity state occurs with probability 0.5.

As a reference, Figure 5 shows the optimal policy functions in the first period under full information (e, a_1, and a_2 all observed) as a function of the reservation W_0 utility of the agent. Displayed are the two action choices in the first period as a function of the state (productive or unproductive), consumption as a function of output (separate graphs for each productivity state), and the utility promises conditional on tomorrow's productivity (again, there are separate graphs for each productivity state in the first period). Given the effect of the effort a_1 on output y, it turns out to be optimal to require high effort only from productive agents. Since utility is concave in consumption, consumption is not conditional on type or output. At first sight, it may seem surprising that the future utility promise is lower for the high state than the low state. The reason is that in the second period consumption is independent of type or output once again, whereas only the productive agent will be required to put in high effort, which explains the lower promise for productive agents.

We now start to move towards the information-constrained solution. Figure 6 shows the outcome if the state e (productivity) is unobserved in both periods, but the two actions continue to be public information. The planner now needs to provide incentives to the agents to reveal the true state. As the upper-left panel shows, for most of the range of initial utility promises only the productive agent is required to exert high effort in the first period. Since effort is costly in terms of utility, the productive agent is compensated for this effort with higher consumption, lower required training, and higher future utility promises. Notice, however, that consumption does *not* depend on realized output: since

[14]Other interpretations of this environment are possible, such as two-stage crop production where the first action is "harvest" and the second action is "sowing" for the next season.

[15]We used the same technology as described in Footnote 12. The grid for the transfer τ had 50 values in both periods, and the grid for utility promises had 30 elements per state in income (resulting in 900 possible combinations, not all of which where feasible). 1800 programs had to be solved in period 2. In the version with e, a_1, and a_2 all unobserved, each linear program had 200 variables and 12 constraints and took on average 0.007 seconds to compute. In the first (initial) period, each program had 99602 variables and 3607 constraints, which resulted in an average computation time of 10.78 seconds on average. The programs where some or all of e, a_1, and a_2 are observed have fewer constraints and compute correspondingly faster.

effort is observable, the planner does not need to provide incentives to exert high effort once the productivity state e has been revealed.

Over most of the range of initial utility promises, unlike in the full-information solution the utility promises for period 2 turn out not to depend on the future state e'. In the full information solution, productive agents receive lower future utility promises, but now the planner needs to provide incentives for truthful reporting of the state. In particular, the key difficulty is to prevent a productive agent from claiming to be unproductive, since in the first-best solution productive agents receive lower utility. Even when the state e' is hidden, it is possible to deliver more utility to the unproductive agent in the second period, namely by prescribing high effort even when low productivity is reported, and then punishing high output. This hurts the productive agent more, as the productive agent is more likely to experience high output. Assigning high effort thus allows the planner to distinguish productive from unproductive agents, because then the outcome y' provides a precise signal on the realization of e'. However, using this scheme turns out to be efficient only for low utility promises W_0 (the range for which the utility promise for low e' exceeds the promise for high e'), and the scheme is only used in the second period.

In the next version, displayed in Figure 7, both e and a_1 are unobserved, while a_2 is public information. The planner now needs to provide incentives to put in effort a_1, which is mainly accomplished by making consumption depend on y. Only productive agents are affected by this change, since unproductive agents are not asked to put in high effort. Second-period consumption (not displayed) now depends on the output realization in the first period as well. A second key difference to Figure 6 is that the planner is no longer able to promise higher utility for the low state, even for low reservation utilities. The reason is that in Figure 6 the two productivity types could be distinguished only by assigning high effort $a_1 = 1$ to both productive and unproductive agents in the second period. This is no longer incentive compatible once effort a_1 is unobserved as well. Thus, the combination of the two sets of incentive constraints forces the planner to offer identical utility promises to productive and unproductive agents, which is not the case if either set of constraints is imposed in isolation.

To see how unobservability of training a_2 affects the results, it is best to start with a version where only a_2 is unobserved, while e and a_1 are public information. The results are displayed in Figure 8. To provide incentives for training, the planner has to promise higher utility for the high future productivity state. Thus, in the region where positive training is assigned, the ranking of the utility promises across the two states is reversed relative to the full-information solution. Finally, Figure 9 shows the fully incentive-constrained version where e, a_1, and a_2 are all unobserved. In the range where the recommended a_2 is positive, the incentive constraints compel the planner to offer more utility promise for the future productive state. Even if the recommended a_2 is zero, the planner has to offer at least as much utility at the productive state compared to the unproductive state to induce truthful reporting of e (as in Figure 7). Finally, consumption depends on y if $a_1 = 1$ is recommended, so that high effort is incentive compatible.

To summarize, the results demonstrate that the features of the optimal contract are quite sensitive to the assumptions on exactly what can be observed by the planner. In Figure 6, for example, the planner wants to provide incentives for reporting the true endowment e. For low utility promises, this is accomplished by imposing high effort on the agent, because then the outcome y provides a precise signal on the realization of e. The insurance scheme that deals with the hidden state e therefore relies on the action a_1 being controlled by the planner. If a_1 is hidden as well, it not incentive-compatible for the agent to put in high effort given that high output is punished subsequently, so the original insurance scheme is undercut. We conjecture that similar outcomes will arise in a variety of environments with multiple hidden actions and states.

8 Conclusions

In this paper we show how general dynamic mechanism design problems with hidden states and hidden actions can be formulated in a way that is suitable for computation. Our methods can be applied to wide variety of models, including applications with income shocks or productivity shocks, and actions that could be interpreted as storage, investment, work effort, search effort, or training. We start from a general planning problem which allows arbitrary message spaces, lotteries, and history dependence. The planning problem is reduced to a recursive version which imposes truth-telling and obedience and uses vectors of utility promises as state variables. We also develop methods to reduce the number of constraints that need to be imposed when computing the optimal contract. The main idea of these methods is to allow the planner to put structure on outcomes off the equilibrium path. In particular, the planner chooses bounds that limit the utility the agent can get on certain off-equilibrium branches that are reached if only if the agent misreports the state at the beginning of the period. If there are multiple hidden actions within the period, we show that multiple *layers* of such utility bounds can be used to implement the optimal contract efficiently.

In an application to a model with hidden storage, we use our methods to examine how the optimal insurance contract varies with the return of the private storage technology. The only existing results for this environment concern special cases where the return on storage is equal to the credit-market return faced by the planner (see Allen 1985, Cole and Kocherlakota 2001b, and Ljungqvist and Sargent 2003). We observe that in certain cases the utility of the agent can actually decrease as the return of the private technology rises. This occurs when the return on public storage is sufficiently high so that in the constrained optimum only public storage is used. If now the return to the private technology rises, it becomes more difficult for the planner to provide incentives for truth-telling and obedience, and consequently less insurance is possible. Thus the effect of a rising return of private investment on the agent's utility in the constrained optimum is exactly the opposite of the effect that prevails under autarky, where a higher return raises utility.

In an application to a hidden effort model with productivity and training, we show how the optimal contract varies depending on which events and actions can be observed by the planner. In particular, we demonstrate how the different types of incentive constraints interact, in the sense that the optimal outcome with multiple unobserved states and actions looks very different from a combination of the outcomes with each state or action unobserved individually. The intuition for this finding is that insurance for a single information problem often works by making outcomes contingent on an observed state or action, as in the standard unobserved effort model where incentives are provided by making consumption depend on output, even though the agent is risk averse. If there are multiple unobserved states and actions, the ability of the planner to impose such contingencies is undercut, and the insurance cannot be provided in the standard form.

A Mathematical Appendix

A.1 Proofs for all Propositions

Proposition 1 *There are reservation utilities $W_0 \in R$ such that an optimal equilibrium exists.*

Proof of Proposition 1 We need to show that for some W_0 the set of equilibria is nonempty and compact, and that the objective function is continuous. To see that the set of equilibria is nonempty, notice that the planner can assign a zero transfer in all periods, and always send the same message. If the strategy of the agent is to choose the actions that are optimal under autarky, clearly all constraints are trivially satisfied for the corresponding initial utility W_0. The set of all contracts that satisfy the probability-measure constraints is compact in the product topology since π and σ are probability measures on finite support. Since only equalities and weak inequalities are involved, it can be shown that the constraints (3) and (4) define a closed subset of this set. Since closed subsets of compact sets are compact, the set of all feasible contracts is compact. We still need to show that the objective function of the planner is continuous. Notice that the product topology corresponds to pointwise convergence, i.e., we need to show that for a sequence of contracts that converges pointwise, the surplus of the planner converges. This is easy to show since we assume that the discount factor of the planner is smaller than one, and that the set of transfers is bounded. Let π_n, σ_n, be a sequence of contracts that converges pointwise to π, σ, and choose $\epsilon > 0$. We have to show that there is an N such that $|V(\pi_n, \sigma_n) - V(\pi, \sigma)| < \epsilon$. Since the transfer τ is bounded and $Q < 1$, there is an T such that the discounted surplus of the planner from time T on is smaller than $\epsilon/2$. Thus we only have to make the difference for the first T periods smaller than $\epsilon/2$, which is the usual Euclidian finite-dimensional case. □

Proposition 2 (Revelation Principle) *For any message spaces M_1 and M_2, any allocation that is feasible in the general mechanism is also feasible in the truth-telling mechanism.*

Proof of Proposition 2 (Outline) Corresponding to any feasible allocation in the general setup there is a feasible contract that implements this allocation. Fix a feasible allocation and the corresponding contract $\{\pi, \sigma\}$. We will now define an outcome function for the truth-telling mechanism that implements the same allocation. To complete the proof, we then have to show that this outcome function satisfies constraints (6) to (8).

We will define the outcome function such that the allocation is the one implemented by (π, σ) along the equilibrium path. To do that, let $H^t(s^t)$ be the set of histories h^t in the general game such that the sequence of incomes, transfers, and actions in h^t coincides with the sequence of reported incomes, transfers, and recommended actions in history s^t in the restricted game. Likewise, define $p(h^t|s^t)$ as the probability of history h^t conditional on s^t:

$$p(h^t|s^t) \equiv \frac{p(h^t)}{\sum_{H^t(s^t)} p(h^t)} \tag{42}$$

If s^t has zero probability (that is, if the sequence s^t of incomes, transfers, and actions occurs with probability zero in the allocation implemented by $\{\pi, \sigma\}$), the definition of $p(h^t|s^t)$ is irrelevant, and is therefore left unspecified. Now we define an outcome function for the truth-telling mechanism by:

$$\pi^\star(\tau_t, a_t|e_t, s^{t-1}) \equiv \sum_{H^{t-1}(s^{t-1}),M_1,M_2} p(h^{t-1}|s^{t-1})\, \sigma(m_{1t}|e_t, h^{t-1})\, \pi(\tau_t|m_{1t}, s^{t-1}(h^{t-1}))$$
$$\pi(m_{2t}|m_{1t}, \tau_t, s^{t-1}(h^{t-1}))\, \sigma(a_t|e_t, m_{1t}, \tau_t, m_{2t}, h^{t-1}). \quad (43)$$

Basically, the outcome function is gained by integrating out the message spaces M_1, , and M_2 and prescribing the outcomes that occur on the equilibrium path.

We now have to verify that with this choice of an outcome function conditions (6) to (8) above are satisfied. In showing this, we can make use of the fact that $\{\pi, \sigma\}$ are probability measures and satisfy (3) and (4). The proof proceeds by substituting (43) into (3) and (4), and showing that the resulting equations imply (6) to (8). This is carried out in detail in Doepke and Townsend (2002). □

Proposition 3 *The set* $\mathbf{W}$ *is nonempty and compact.*

Proof of Proposition 3 To see that $\mathbf{W}$ is nonempty, notice that the planner can always assign a zero transfer in every period and recommend the optimal action that the agent would have chosen without the planner. For the $w_0(e)$ that equals the expected utility of the agent under autarky under state e, all constraints are satisfied. To see that $\mathbf{W}$ is bounded, notice that there are finite grids for income, the transfer, and the action. This implies that in every period consumption and therefore utility is bounded from above and from below. Since the discount factor β is smaller than one, total expected utility is also bounded. Since each $w_0(e)$ has to satisfy a promise-keeping constraint with equality, the set $\mathbf{W}$ must be bounded. To see that $\mathbf{W}$ is closed, assume there exists a converging sequence $\mathbf{w}_n$ such that $\mathbf{w}_n \in \mathbf{W}$ for all n. Corresponding to each $\mathbf{w}_n$ there is a contract $\pi(\tau_t, a_t|e_t, s^{t-1})_n$ satisfying constraints (6), (7), and (8). Since the contracts are within a compact subset of R^∞ with respect to the product topology, there is a convergent subsequence with limit $\pi(\tau_t, a_t|e_t, s^{t-1})$. It then follows that w must satisfy (6), (7), and (8) when $\pi(\tau_t, a_t|e_t, s^{t-1})$ is the chosen contract. □

Proposition 4 *For all* $\mathbf{w}_0 \in \mathbf{W}$ *and* $e_0 \in E$*, and for any* s^{k-1} *and* e_k*, there is an optimal contract* $\pi^\star$ *such that the remaining contract from* s^{k-1} *and* e_k *on is an optimal contract for the auxiliary planning problem with* $e_0 = e_k$ *and* $\mathbf{w}_0 = \mathbf{w}(s^{k-1}, \pi^\star)$.

Proof of Proposition 4 We will first construct $\pi^\star$ from a contract which is optimal from time zero and another contract which is optimal starting at s^{k-1} and e_k. We will then show by a contradiction argument that $\pi^\star$ is an optimal contract from time zero as well.

We have shown earlier that an optimal contract exists. Let π be an optimal contract from time zero, and π_k an optimal contract for $e_0 = e_k$ and $\mathbf{w}_0 = \mathbf{w}(s^{k-1}, \pi)$, with the elements of vector $\mathbf{w}(s^{k-1}, \pi)$ defined in (12). Now construct a new contract $\pi^\star$ that is equal to π_k from (e_k, s^{k-1}) on, and equals π until time k and on all future branches other than e_k, s^{k-1}. First, notice that by the way $\pi^\star$ is constructed, we have $\mathbf{w}(s^{k-1}, \pi) = \mathbf{w}(s^{k-1}, \pi^\star)$, and since $\pi^\star$ equals π_k from (e_k, s^{k-1}), $\pi^\star$ fulfills the reoptimization requirement of the proposition. We now claim that $\pi^\star$ is also an optimal contract. To show this, we have to demonstrate that $\pi^\star$ satisfies constraints (6), (7), and (8), and that it maximizes the surplus of the planner subject to these constraints. To start, notice that the constraints that are imposed if we compute an optimal contract taking $e_0 = e_k$ and $\mathbf{w}_0 = \mathbf{w}(s^{k-1}, \pi)$ as the starting point also constrain the choices of the planner in the original program from (e_k, s^{k-1}) onward. By reoptimizing at (e_k, s^{k-1}) in period k as if the game were restarted, the planner clearly cannot lower his surplus since no additional constraints are imposed. Therefore the total surplus from contract $\pi^\star$ cannot be lower than the surplus from π. Since π is assumed to be an optimal contract, if $\pi^\star$ satisfies (6), (7), and (8), it must be optimal as well. Thus we only have to show that (6), (7), and (8) are satisfied, or in other words, that reoptimizing at e_k, s^{k-1} does not violate any constraints of the original problem.

The probability constraints (6) are satisfied by contract $\pi^\star$, since the reoptimized contract is subject to the same probability constraints as the original contract. The promise-keeping constraint (7) is satisfied since the new contract delivers the same on-path utilities by construction. We still have to show that the incentive constraints (8) are satisfied. We will do this by contradiction. Suppose that (8) is not satisfied by contract $\pi^\star$. Then there is a deviation δ such that for some s^l, e_{l+1}:

$$\sum_{t=l+1}^{\infty} \beta^t \left[\sum_{H^t} p(h^t|\pi^\star, \delta, s^l, e_{l+1}) u(e_t + \tau_t - \delta_a(h^{t-1}, e_t, \tau_t, a_t)) \right] > \sum_{t=l+1}^{\infty} \beta^t \left[\sum_{S^t} p(s^t|\pi^\star, s^l, e_{l+1}) u(e_t + \tau_t - a_t) \right]. \quad (44)$$

Consider first the case $l + 1 \geq k$. On any branch that starts at or after time k, contract $\pi^\star$ is entirely identical to either π or π_k. But then (44) implies that either π or π_k violates incentive-compatibility (8), a contradiction. Consider now the case $l + 1 < k$. Here the contradiction is not immediate, since the remaining contract is a mixture of π and π_k. Using $w(e_k, s^{k-1}, \delta)$ to denote the continuation utility of the agent from time k on under the deviation strategy, we can rewrite (44) as:

$$\sum_{t=l+1}^{k-1} \beta^t \left[\sum_{H^t} p(h^t|\pi^\star, \delta, s^l, e_{l+1}) u(e_t + \tau_t - \delta_a(h^{t-1}, e_t, \tau_t, a_t)) \right] + \\ \beta^k \sum_{E, s^{k-1}} p(e_k, s^{k-1}|\pi^\star, \delta, s^l, e_{l+1}) w(e_k, s^{k-1}, \delta) \\ > \sum_{t=l+1}^{k-1} \beta^t \left[\sum_{S^t} p(s^t|\pi^\star, s^l, e_{l+1}) u(e_t + \tau_t - a_t) \right] + \beta^k \sum_{E, s^{k-1}} p(e_k, s^{k-1}|\pi^\star, s^l, e_{l+1}) w(e_k, s^{k-1}). \tag{45}$$

Notice that for s^{k-1} that are reached with positive probability under the deviation we have:

$$w(e_k, s^{k-1}, \delta) \leq w(e_k, \delta(s^{k-1})), \tag{46}$$

where $\delta(s^k)$ is the history as seen by the planner (reported incomes, delivered transfers, and recommended actions) under the deviation strategy. Otherwise, either π or π_k would violate incentive constraints. To see why, assume that e_k, s^{k-1} is a branch after which $\pi^\star$ is identical to π. If we had $w(e_k, s^{k-1}, \delta) > w(e_k, \delta(s^{k-1}))$, an agent under contract π who reached history $\delta(s^{k-1})$ could gain by following the deviation strategy δ afterwards. This cannot be the case since π is assumed to be an optimal contract, and therefore deviations are never profitable. Using (46) in (45) gives:

$$\sum_{t=l+1}^{k-1} \beta^t \left[\sum_{H^t} p(h^t|\pi^\star, \delta, s^l, e_{l+1}) u(e_t + \tau_t - \delta_a(h^{t-1}, e_t, \tau_t, a_t)) \right] + \\ \beta^k \sum_{E, s^{k-1}} p(e_k, s^{k-1}|\pi^\star, \delta, s^l, e_{l+1}) w(e_k, \delta(s^{k-1})) \\ > \sum_{t=l+1}^{k-1} \beta^t \left[\sum_{S^t} p(s^t|\pi^\star, s^l, e_{l+1}) u(e_t + \tau_t - a_t) \right] + \beta^k \sum_{E, s^{k-1}} p(e_k, s^{k-1}|\pi^\star, s^l, e_{l+1}) w(e_k, s^{k-1}). \tag{47}$$

The outcome function $\pi^\star$ enters (47) only up to time $k-1$. Since up to time $k-1$ the outcome function $\pi^\star$ is identical to π, and since by construction of $\pi^\star$ continuation utilities at time k are the same under $\pi^\star$ and π, we can rewrite (47) as:

$$\sum_{t=l+1}^{k-1} \beta^t \left[\sum_{H^t} p(h^t|\pi, \delta, s^l, e_{l+1}) u(e_t + \tau_t - \delta_a(h^{t-1}, e_t, \tau_t, a_t)) \right] + \beta^k \sum_{E, s^{k-1}} p(e_k, s^{k-1}|\pi, \delta, s^l, e_{l+1}) w(e_k, \delta(s^{k-1}))$$
$$> \sum_{t=l+1}^{k-1} \beta^t \left[\sum_{S^t} p(s^t|\pi, s^l, e_{l+1}) u(e_t + \tau_t - a_t) \right] + \beta^k \sum_{E, s^{k-1}} p(e_k, s^{k-1}|\pi, s^l, e_{l+1}) w(e_k, s^{k-1}). \quad (48)$$

But now the left-hand side of (48) is the utility that the agent gets under plan π from following the deviation strategy until time k and following the recommendations of the planner afterwards. Thus (48) contradicts the incentive compatibility of π. We obtain a contradiction, thus $\pi^\star$ actually satisfies (8). This shows that plan $\pi^\star$ is within the constraints of the original problem. Since $\pi^\star$ yields at least as much surplus as π and π is an optimal contract, $\pi^\star$ must be optimal as well. □

Proposition 5 *Program 1 and Program 2 are equivalent.*

Proof of Proposition 5 We want to show that constraints (28) and (29) in Program 2 place the same restrictions on the outcome function $\pi(\cdot)$ as the truth-telling constraints (23) of Program 1. Let us first assume we have found a lottery $\pi(\tau, a, \mathbf{w}'|\mathbf{w}, e)$ that satisfies the truth telling constraint (23) of Program 1 for all $\hat{e}$ and $\delta : T \times S \to A$. We have to show that there exist utility bounds $v(\hat{e}, e, \tau, a)$ such that the same lottery satisfies (28) and (29) in Program 2. For each $\hat{e}$, τ, and a, define $v(\hat{e}, e, \tau, a)$ as the maximum of the left hand side of (28) over all $\hat{a}$:

$$v(\hat{e}, e, \tau, a) \equiv \max_{\hat{a}} \left\{ \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \hat{a}) + \beta \sum\nolimits_E \mu(e'|\hat{a}) w'(e') \right] \right\}. \quad (49)$$

Then clearly (28) is satisfied, since the left-hand side of (28) runs over $\hat{a}$. Now for each τ and a, define $\hat{\delta}(\cdot)$ by setting $\hat{\delta}(\tau, a)$ equal to the $\hat{a}$ that maximizes the left-hand side of (28):

$$\hat{\delta}(\tau, a) \equiv \text{argmax}_{\hat{a}} \left\{ \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \hat{a}) + \beta \sum\nolimits_E \mu(e'|\hat{a}) w'(e') \right] \right\}. \quad (50)$$

Since $\pi(\tau, a, \mathbf{w}'|\mathbf{w}, e)$ satisfies (23) for any function $\delta(\cdot)$ by assumption, we have for our particular $\hat{\delta}(\cdot)$:

$$\sum_{T, A, \mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \hat{\delta}(\tau, a)) + \beta \sum\nolimits_E \mu(e'|\hat{\delta}(\tau, a)) w'(e') \right] \leq w(\hat{e}). \quad (51)$$

By the way we chose $\hat{\delta}(\cdot)$ and the $v(\cdot)$, we have from (28):

$$\sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[u(\hat{e} + \tau - \hat{\delta}(\tau, a)) + \beta \sum\nolimits_E \mu(e'|\hat{\delta}(\tau, a))w'(e')\right] = v(\hat{e}, e, \tau, a). \tag{52}$$

Substituting the left-hand side into (51), we get:

$$\sum_{T,A} v(\hat{e}, e, \tau, a) \leq w(\hat{e}). \tag{53}$$

which is (29).

Conversely, suppose we have found a lottery $\pi(\tau, a, \mathbf{w}'|\mathbf{w}, e)$ that satisfies (28) and (29) in Program 2 for some choice of $v(\hat{e}, e, \tau, a)$. By (28), we have then for any $\hat{e}$ and $\hat{a}$ and hence any $\delta : T \times S \to A$:

$$\sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \delta(\tau, a)) + \beta \sum\nolimits_E \mu(e'|\delta(\tau, a))w'(e')\right] \leq v(\hat{e}, e, \tau, a). \tag{54}$$

Substituting the left-hand side of (54) into the assumed (29) for the $v(\hat{e}, e, \tau, a)$, we maintain the inequality:

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \delta(\tau, a)) + \beta \sum\nolimits_E \mu(e'|\delta(\tau, a))w'(e')\right] \leq w(\hat{e}). \tag{55}$$

But this is (23) in Program 1. Therefore the sets of constraints are equivalent, which proves that Program 1 and Program 2 are equivalent. □

Proposition 6 *Program 3 and Program 4 are equivalent.*

Proof of Proposition 6 We want to show that constraints (38)-(41) in Program 4 place the same restrictions on the outcome function $\pi(\cdot)$ as the constraints (34) and (36) of Program 3. Let us first assume we have found a lottery $\pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e)$ that satisfies the constraints (34) and (36) of Program 3 for all $\hat{e}$, δ_1, and δ_2. We have to show that there exist utility bounds $v(\hat{e}, e, \tau, a)$ such that the same lottery satisfies (39)-(41) in Program 4. For each $\hat{e}$, a_1, $\hat{a}_1$, y, τ, and a_2 define $v_2(\hat{e}, a_1, \hat{a}_1, y, \tau, a_2)$ as the maximum of the left hand side of (41) over all $\hat{a}_2$:

$$v_2(\hat{e}, a_1, \hat{a}_1, y, \tau, a_2) \equiv \max_{\hat{a}_2} \left\{ \sum_{\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \frac{f(y|\hat{e}, \hat{a}_1)}{f(y|e, a_1)} \right.$$
$$\left. \left[u(\hat{e}, \hat{a}_1, y, \tau, \hat{a}_2) + \beta \sum\nolimits_E g(e'|\hat{e}, \hat{a}_1, y, \tau, \hat{a}_2)w'(e')\right] \right\}. \tag{56}$$

This choice guarantees that (38) and (41) are satisfied. Next, for all $\hat{e}$ and a_1, define $v_1(\hat{e}, a_1)$ as the maximum over the left-hand side of (40):

$$v_1(\hat{e}, a_1) \equiv \max_{\hat{a}_1} \left\{ \sum_{Y,T,A_2} v_2(\hat{e}, a_1, \hat{a}_1, y, \tau, a_2) \right\}. \tag{57}$$

This choice guarantees that (40) is satisfied. Finally, (36), (40), and (41) jointly imply that (39) is satisfied (if it were violated, we could construct a deviation strategy that violates (36)).

Conversely, suppose we have found a lottery $\pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e)$ that satisfies (38)-(41) in Program 4 for some choice of utility bounds. Because of (41), we have for any δ_2:

$$\forall \hat{e}, a_1, \hat{a}_1, y, \tau, a_2 :$$
$$\sum_{\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \frac{f(y|\hat{e}, \hat{a}_1)}{f(y|e, a_1)}$$
$$\left[u(\hat{e}, \hat{a}_1, y, \tau, \delta_2(a_1, y, \tau, a_2)) + \beta \sum\nolimits_E g(e'|\hat{e}, \hat{a}_1, y, \tau, \delta_2(a_1, y, \tau, a_2)) w'(e')\right]$$
$$\leq v_2(\hat{e}, a_1, \hat{a}_1, y, \tau, a_2). \tag{58}$$

Plugging the left-hand side of this equation into (38) for the $v_2(\hat{e}, a_1, \hat{a}_1, y, \tau, a_2)$ maintains the inequality, and therefore gives (34). We also have for any given δ_1 and δ_2:

$$\forall \hat{e}, a_1, y, \tau, a_2 :$$
$$\sum_{\mathbf{W}'} \pi(a_1, y, \tau, a_2, \mathbf{w}'|\mathbf{w}, e) \frac{f(y|\hat{e}, \delta_1(a_1))}{f(y|e, a_1)}$$
$$\left[u(\hat{e}, \delta_1(a_1), y, \tau, \delta_2(a_1, y, \tau, a_2)) + \beta \sum\nolimits_E g(e'|\hat{e}, \delta_1(a_1), y, \tau, \delta_2(a_1, y, \tau, a_2)) w'(e')\right]$$
$$\leq v_2(\hat{e}, a_1, \delta_1(a_1), y, \tau, a_2). \tag{59}$$

Substituting the left-hand side of (40) for the given δ_1 into (39) gives:

$$\forall \hat{e} \neq e : \quad \sum_{A_1,Y,T,A_2} v_2(\hat{e}, a_1, \delta a_1, y, \tau, a_2) \leq w(\hat{e}). \tag{60}$$

Now we can replace the $v_2(\hat{e}, a_1, \delta a_1, y, \tau, a_2)$ in this equation by the left-hand side of (59) while maintaining the inequality:

$$\forall \hat{e} \neq e:$$
$$\sum_{A_1,Y,T,A_2,\mathbf{W}'} \pi(a_1,y,\tau,a_2,\mathbf{w}'|\mathbf{w},e)\frac{f(y|\hat{e},\delta_1(a_1))}{f(y|e,a_1)}$$
$$\Big[u(\hat{e},\delta_1(a_1),y,\tau,\delta_2(a_1,y,\tau,a_2)) + \beta\sum\nolimits_E g(e'|\hat{e},\delta_1(a_1),y,\tau,\delta_2(a_1,y,\tau,a_2))w'(e')\Big] \leq w(\hat{e}), \quad (61)$$

But this is (36) in Program 3 for our particular choice of δ_1 and δ_2. Since the choice of δ_1 and δ_2 was arbitrary, (36) is satisfied, which shows that Program 1 and Program 2 are equivalent. □

A.2 Computing the Value Set

Our analysis was based on the assumption that the set $\mathbf{W}$ of feasible utility vectors is known in advance. In practice, $\mathbf{W}$ is not known and needs to be computed alongside the value function $V(\mathbf{w},e)$. $\mathbf{W}$ can be computed with the dynamic-programming methods described in detail in Abreu, Pierce, and Stachetti (1990), henceforth APS. An outline of the method follows.

We start by defining an operator B that maps nonempty compact subsets of $\mathbf{R}^{\#E}$ into nonempty compact subsets of $\mathbf{R}^{\#E}$. Let $\mathbf{W}'$ be a nonempty compact subset of $\mathbf{R}^{\#E}$. Then $B(\mathbf{W}')$ is defined as follows:

Definition 6 *A utility vector* $\mathbf{w} \in B(\mathbf{W}')$ *if there exist probabilities* $\pi(\tau,a|\mathbf{w},e)$ *and future utilities* $w'(\mathbf{w},e,\tau,a) \in \mathbf{W}'$ *such that* (14) *to* (17) *hold.*

The key point is that utility promises are chosen from the set $\mathbf{W}'$ instead of the true value set $\mathbf{W}$. Intuitively, $B(\mathbf{W}')$ consists of all utility vectors $\mathbf{w}$ that are feasible today (observing all incentive constraints), given that utility vectors from tomorrow onward are drawn from the set $\mathbf{W}'$. The fact that B maps compact set into compact sets follows from the fact that all constraints are linear and involve only weak inequalities. Clearly, the true set of feasible utility vectors $\mathbf{W}$ satisfies $\mathbf{W} = B(\mathbf{W})$, thus $\mathbf{W}$ is a fixed point of B. The computational approach described in APS consists of using B to define a shrinking sequence of sets that converges to $\mathbf{W}$.

To do this, we need to start with a set $\mathbf{W}_0$ that is known to be larger than $\mathbf{W}$ a priori. In our case, this is easy to do, since consumption is bounded and therefore lifetime utility is bounded above and below. We can choose $\mathbf{W}_0$ as an interval in $\mathbf{R}^{\#E}$ from a lower bound that is lower than the utility from receiving the lowest consumption forever to a number that exceeds utility from consuming the highest consumption forever. We can now define a sequence of sets $\mathbf{W}_n$ by defining $\mathbf{W}_{n+1}$ as $\mathbf{W}_{n+1} = B(\mathbf{W}_n)$. We have the following results:

Proposition 7

- *The sequence* $\mathbf{W}_n$ *is shrinking, i.e., for any* n, $\mathbf{W}_{n+1}$ *is a subset of* $\mathbf{W}_n$.
- *For all* n, $\mathbf{W}$ *is a subset of* $\mathbf{W}_n$.
- *The sequence* $\mathbf{W}_n$ *converges to a limit* $\bar{\mathbf{W}}$, *and* $\mathbf{W}$ *is a subset of* $\bar{\mathbf{W}}$.

Proof of Proposition 7 To see that $\mathbf{W}_n$ is shrinking, we only need to show that $\mathbf{W}_1$ is a subset of $\mathbf{W}_0$. Since $\mathbf{W}_0$ is an interval, it suffices to show that the upper bound of $\mathbf{W}_1$ is lower than the upper bound of $\mathbf{W}_0$, and that the lower bound of $\mathbf{W}_1$ is higher than the lower bound of $\mathbf{W}_0$. The upper bound of $\mathbf{W}_1$ is reached by assigning maximum consumption in the first period and the maximum utility vector in $\mathbf{W}_0$ from the second period onward. But the maximum utility vector $\mathbf{W}_0$ by construction corresponds to consuming more than maximum consumption every period, and since utility is discounted, the highest utility vector in $\mathbf{W}_1$ therefore is smaller than the highest utility vector in $\mathbf{W}_0$.

To see that $\mathbf{W}$ is a subset of all $\mathbf{W}_n$, notice that by the definition of B, if $\mathbf{C}$ is a subset of $\mathbf{D}$, $B(\mathbf{C})$ is a subset of $B(\mathbf{D})$. Since $\mathbf{W}$ is a subset of $\mathbf{W}_0$ and $\mathbf{W} = B(\mathbf{W})$, we have that $\mathbf{W}$ is a subset of $\mathbf{W}_1 = B(\mathbf{W}_0)$, and correspondingly for all the other elements.

Finally, $\mathbf{W}_n$ has to converge to a nonempty limit since it is a decreasing sequence of compact sets, and the nonempty set $\mathbf{W}$ is a subset of all elements of the sequence. $\square$

Up to this point, we know that $\mathbf{W}_n$ converges to $\bar{\mathbf{W}}$ and that $\mathbf{W}$ is a subset of $\bar{\mathbf{W}}$. What we want to show is that $\bar{\mathbf{W}}$ and $\mathbf{W}$ are actually identical. What we still need to show, therefore, is that $\bar{\mathbf{W}}$ is also a subset of $\mathbf{W}$.

Proposition 8 *The limit set* $\bar{\mathbf{W}}$ *is a subset of the true value set* $\mathbf{W}$.

Proof of Proposition 8 The outline of the proof is as follows. To show that an element $\mathbf{w}$ of $\bar{\mathbf{W}}$ is in $\mathbf{W}$, we have to find $\pi(\tau_t, a_t|e_t, s^{t-1})$ that satisfy constraints (6), (8), and (10) for $\mathbf{w}$. These $\pi(\tau_t, a_t|e_t, s^{t-1})$ can be constructed period by period from the π that are implicit in the definition of the operator B. Notice that in each period continuation utilities are drawn from the same set $\bar{\mathbf{W}}$, since $\bar{\mathbf{W}}$ as the limit of the sequence $\mathbf{W}_n$ satisfies $\bar{\mathbf{W}} = B(\bar{\mathbf{W}})$. By definition of B, the resulting $\pi(\tau_t, a_t|e_t, s^{t-1})$ satisfy the period-by-period constraints (14) to (17). We therefore need to show that satisfying the period-by-period constraints (with a given set of continuation utilities) is equivalent to satisfying the original constraints (6), (8), and (10), which we have done above in Section 4.2. $\square$

For numerical implementation of this algorithm, we have to deal with the additional restriction that the set of possible future utilities needs to be finite, so that linear programming can be used. Instead of starting with the set $\mathbf{W}_0$ which is an interval in $\mathbf{R}^{\#E}$, we therefore start out with a finite subset $\mathbf{W}_0^D$ of $\mathbf{W}_0$, where the utility promises for each type lie on a discrete grid. To apply the APS algorithm, an operator B^D is now defined as mapping subsets of $\mathbf{W}_0^D$ into subsets of $\mathbf{W}_0^D$, that is, both current utilities and future utility promises are drawn from the same set:

Definition 7 *For any $\mathbf{W}' \subseteq \mathbf{W}_0^D$, the set $B^D(\mathbf{W}')$ consists of all $\mathbf{w} \in \mathbf{W}_0^D$ such that for all $e \in E$ there exist probabilities $\pi(\tau, a, \mathbf{w}'|\mathbf{w}, e)$ such that* (20) *to* (23) *hold.*

Analogously to the above, we can now define a sequence $\mathbf{W}_n^D$ by setting $\mathbf{W}_{n+1}^D = B^D(\mathbf{W}_n^D)$. This sequence has the following properties:

Proposition 9

- *The sequence $\mathbf{W}_n^D$ is shrinking, i.e., for any n, $\mathbf{W}_{n+1}^D$ is a subset of $\mathbf{W}_n^D$.*
- *For all n, $\mathbf{W}_n^D$ is a subset of $\mathbf{W}_n$.*
- *The sequence $\mathbf{W}_n^D$ converges to a limit $\bar{\mathbf{W}}^D$, and $\bar{\mathbf{W}}^D$ is a subset of $\mathbf{W}$.*

Proof of Proposition 9 To show that $\mathbf{W}_n^D$ is shrinking, it suffices to show that $\mathbf{W}_1^D$ is a subset of $\mathbf{W}_0^D$, which is true by definition of the operator B^D. The second claim is true since $\mathbf{W}_0^D$ is a subset of $\mathbf{W}_0$, and the operator B^D is more restrictive than the operator B, since current utilities are drawn from a smaller set. $\mathbf{W}_n^D$ converges since it is a decreasing sequence of compact sets. Finally, since $\mathbf{W}_n^D \subseteq \mathbf{W}_n$ for all n, we must have $\bar{\mathbf{W}}^D \subseteq \bar{\mathbf{W}} = \mathbf{W}$. □

Notice that the proof does not rule out that the limit set $\bar{\mathbf{W}}^D$ may be the empty set; that may well be the case if the initial set $\mathbf{W}_0^D$ is too coarse. We do know, however, that any element of $\bar{\mathbf{W}}^D$ is a member of the true value set. In the language of Judd, Yeltekin, and Conklin (2003), $\bar{\mathbf{W}}^D$ is an *inner approximation* of the true value set $\mathbf{W}$.

The question arises whether it is possible to assess how well the true value set is being approximated. In the context of dynamic games, Judd, Yeltekin, and Conklin (2003) approach this problem by complementing inner approximations of the value set with outer approximations, i.e., sets that are known to be larger than the true value set. The difference between inner and outer approximation then provides bounds on the quality of the approximation. The same approach is not applicable here, however, since it requires the ability to compute tight outer approximations. Our linear programming approach requires the use of discrete value sets, which generally cannot be used to provide outer approximations.

There exist a number of general results on the approximation of infinite-dimensional linear programming problems by programs with finitely many variables. If we used the true value set $\mathbf{W}$ as the set of possible utility promises, the resulting linear program would have an infinite number of variables (though still a finite number of constraints), since probability weight could be put on any of the generally infinitely many members of $\mathbf{W}$. We can think of our finite linear programs as versions of the general infinite-dimensional linear programs under additional restrictions, namely that zero probability be assigned to all but a given finite set of future utility promises. Hernández-Lerma and Lasserre (1998) provide conditions under which the values of a sequence of such restricted linear programs converge to the value of the true infinite-dimensional program (see their Proposition 2.2 and Theorem 3.2). Translated into our application, if $\mathbf{W}_k$ is an increasing sequence of finite approximations of $\mathbf{W}$ such that $\mathbf{W}_k \subset \mathbf{W}$, the key condition is that $\cup_{k=0}^{\infty}\mathbf{W}_k$ is dense in $\mathbf{W}$. It may seem that these results should apply easily to our environment, since $\mathbf{W}$ is a subset of an interval in $\mathbf{R}^{\#E}$. However, notice that the $\mathbf{W}_k$ cannot be chosen directly, but would have to be generated as a fixed point of the B^D operator for each k, so without additional assumptions the properties of $\mathbf{W}_k$ cannot be directly inferred.

As a practical matter, it is sometimes possible to judge the quality of approximations if there are special cases with known solutions. The hidden storage application in Section 7.1 provides such an example, since there are theoretical results that apply to the case when the planner and the agent face the same credit-market return.

References

Abraham, Arpad and Nicola Pavoni. 2002. "Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending." Unpublished Manuscript, UCL and LBS.

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–1063.

Allen, Franklin. 1985. "Repeated Principal-Agent Relationships with Lending and Borrowing." *Economics Letters* 17:27–31.

Cole, Harold L. and Narayana Kocherlakota. 2001a. "Dynamic Games with Hidden Actions and Hidden States." *Journal of Economic Theory* 98 (1): 114–26.

———. 2001b. "Efficient Allocations with Hidden Income and Hidden Storage." *Review of Economic Studies* 68 (3): 523–42.

Doepke, Matthias and Robert M. Townsend. 2002. "Dynamic Mechanism Design with Hidden Income and Hidden Actions: Technical Appendix." UCLA Department of Economics Working Paper No. 819.

Fernandes, Ana and Christopher Phelan. 2000. "A Recursive Formulation for Repeated Agency with History Dependence." *Journal of Economic Theory* 91 (2): 223–47.

Fudenberg, Drew, Bengt Holmstrom, and Paul Milgrom. 1990. "Short-Term Contracts and Long-Term Agency Relationships." *Journal of Economic Theory* 51:1–31.

Green, Edward J. 1987. "Lending and the Smoothing of Uninsurable Income." In *Contractual Arrangements for Intertemporal Trade*, edited by E. Prescott and N. Wallace, 3–25. Minneapolis: University of Minnesota Press.

Harris, Milton and Robert M. Townsend. 1977. "Allocation Mechanisms for Asymetrically Informed Agents." Manuscript presented at the NBER-CME Conference Honoring Herb Simon.

———. 1981. "Resource Allocation under Asymmetric Information." *Econometrica* 49 (1): 33–64.

Hernández-Lerma, Onésimo and Jean B. Lasserre. 1998. "Approximation Schemes for Infinite Linear Programs." *SIAM Journal on Optimization* 8 (4): 973–88.

Judd, Kenneth L., Sevin Yeltekin, and James Conklin. 2003. "Computing Supergame Equilibria." *Econometrica* 71 (4): 1239–54.

Kocherlakota, Narayana R. 2004. "Figuring Out the Impact of Hidden Savings on Optimal Unemployment Insurance." *Review of Economic Dynamics* 7 (3): 541–554.

Ljungqvist, Lars and Thomas J. Sargent. 2003. "Recursive Macroeconomic Theory." Manuscript for 2nd Edition.

Myerson, Roger B. 1979. "Incentive Compatibility and the Bargaining Problem." *Econometrica* 47 (1): 61–73.

——— 1982. "Optimal Coordination Mechanisms in Generalized Principal Agent Problems." *Journal of Mathematical Economics* 10 (1): 67–81.

——— 1986. "Multistage Games with Communication." *Econometrica* 54 (2): 323–58.

Phelan, Christopher. 1995. "Repeated Moral Hazard and One-Sided Commitment." *Journal of Economic Theory* 66 (2): 488–506.

Phelan, Christopher and Robert M. Townsend. 1991. "Computing Multi-Period, Information-Constrained Optima." *Review of Economic Studies* 58:853–881.

Prescott, Edward S. 2003. "Communication in Models with Private Information: Theory and Computation." *The Geneva Papers on Risk and Insurance Theory* 28 (2): 105–130.

Spear, S. and S. Srivastava. 1987. "On Repeated Moral Hazard with Discounting." *Review of Economic Studies* 54:599–618.

Thomas, Jonathan and Tim Worrall. 1990. "Income Fluctuations and Asymmetric Information: An Example of a Repeated Principal-Agent Problem." *Journal of Economic Theory* 51 (2): 367–390.

Townsend, Robert M. 1979. "Optimal Contracts and Competitive Markets with Costly State Verification." *Journal of Economic Theory* 21 (2): 265–93.

——— 1982. "Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information." *Journal of Political Economy*, no. 6:1166–86.

——— 1988. "Information Constrained Insurance: The Revelation Principle Extended." *Journal of Monetary Economics* 21 (2/3): 411–50.

Wang, Cheng. 1995. "Dynamic Insurance with Private Information and Balanced Budgets." *Review of Economic Studies* 62 (4): 577–95.

Werning, Iván. 2001. "Repeated Moral Hazard with Unmonitored Wealth: A Recursive First-Order Approach." Unpublished Manuscript, University of Chicago.

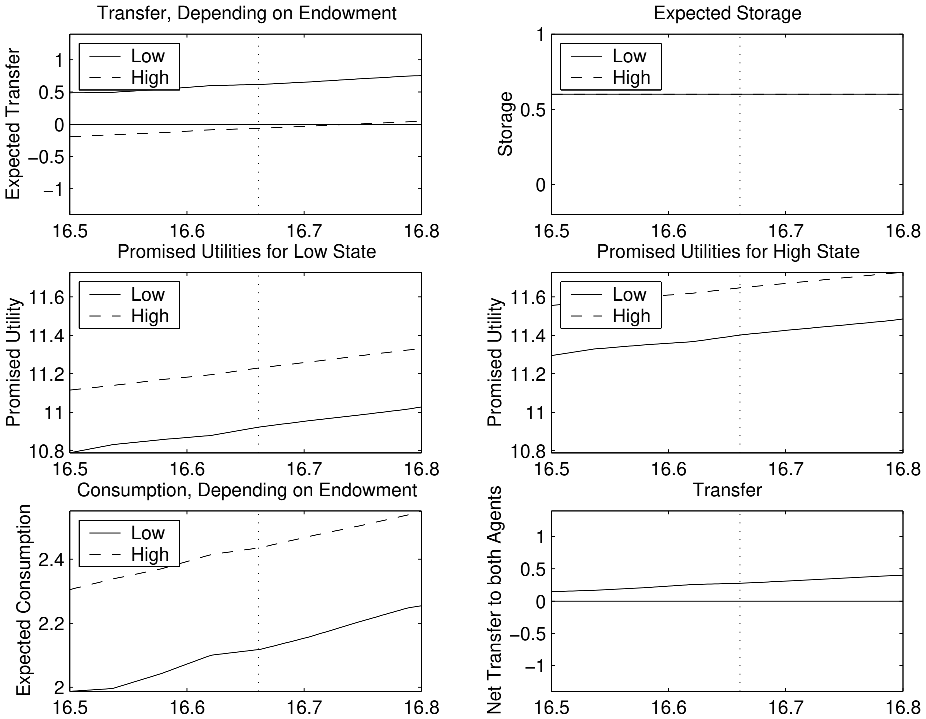

Figure 3: Policy Functions, R=1.1

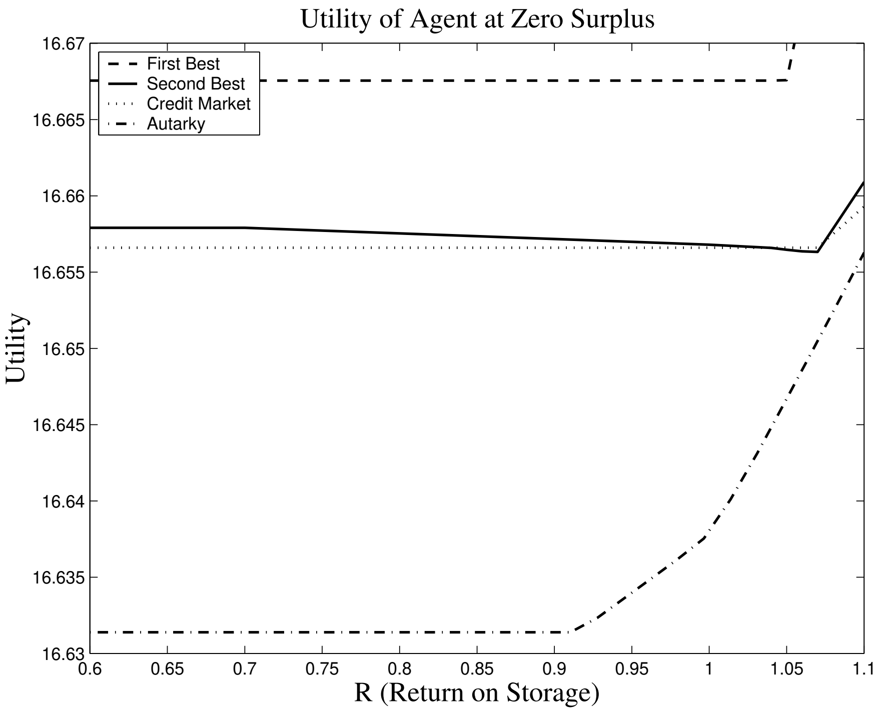

Figure 4: Utility of the Agent with Full Information, Private Information, Credit-Market Access, and Autarky

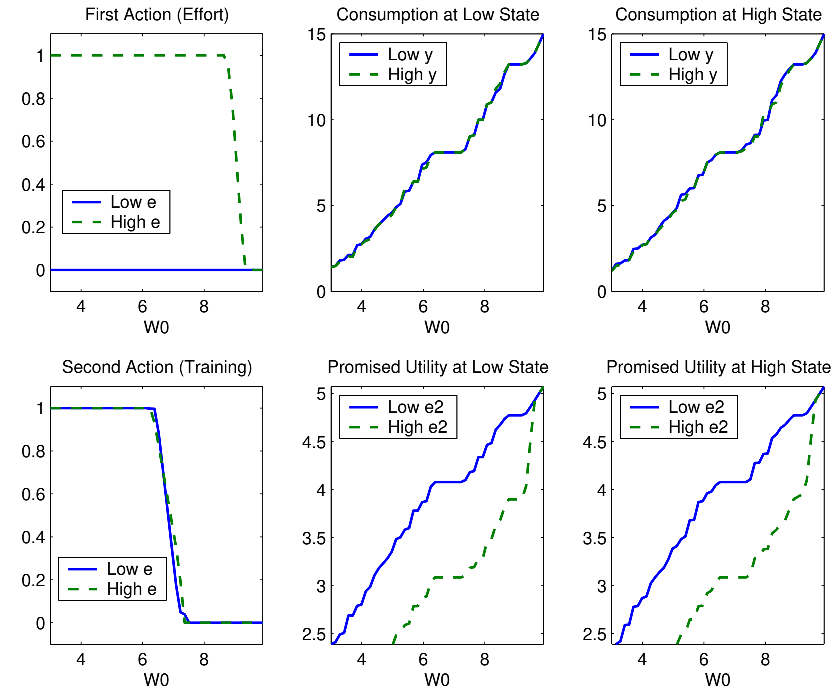

Figure 5: First-Period Policy Functions, Full Information

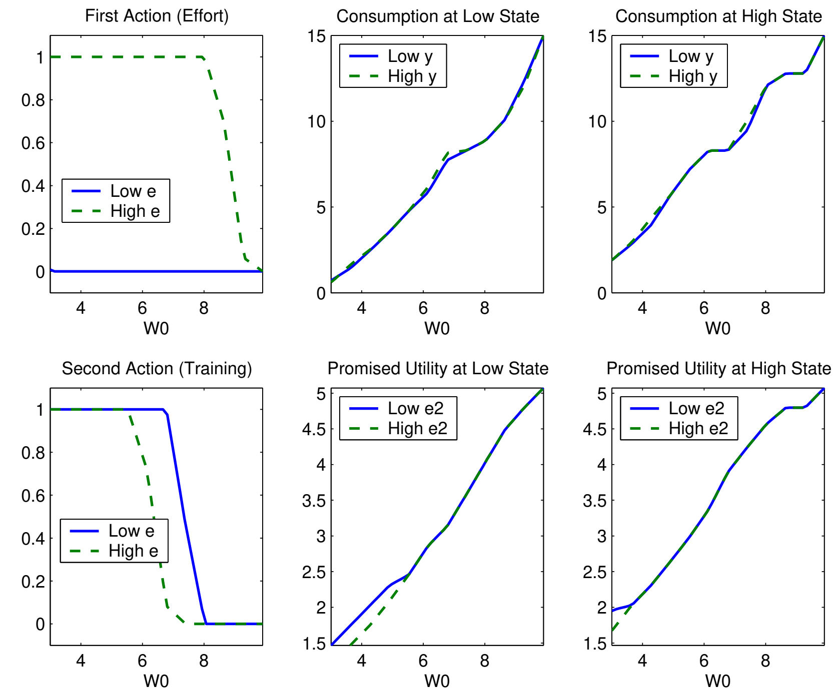

Figure 6: First-Period Policy Functions, e Private Information

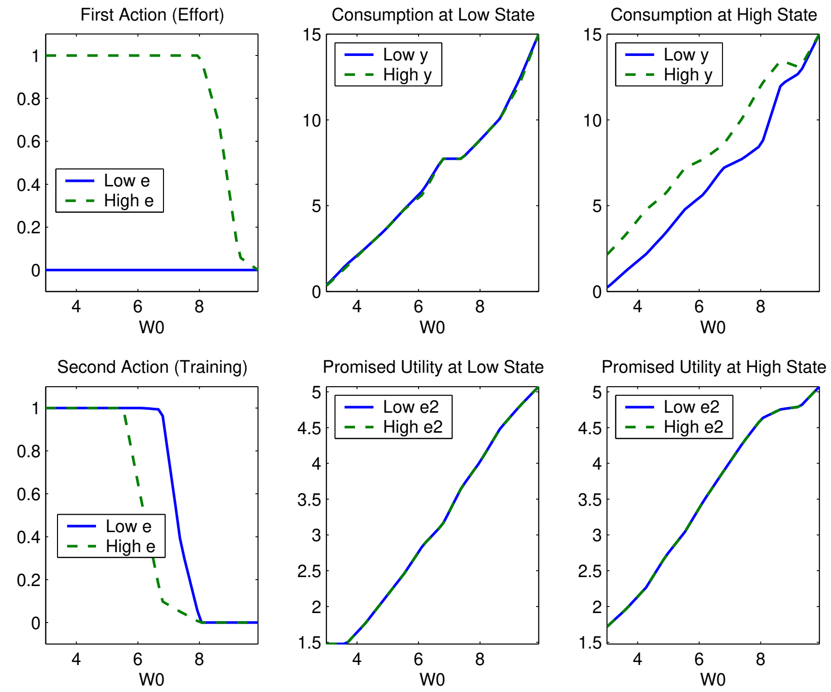

Figure 7: First-Period Policy Functions, e, a_1 Private Information

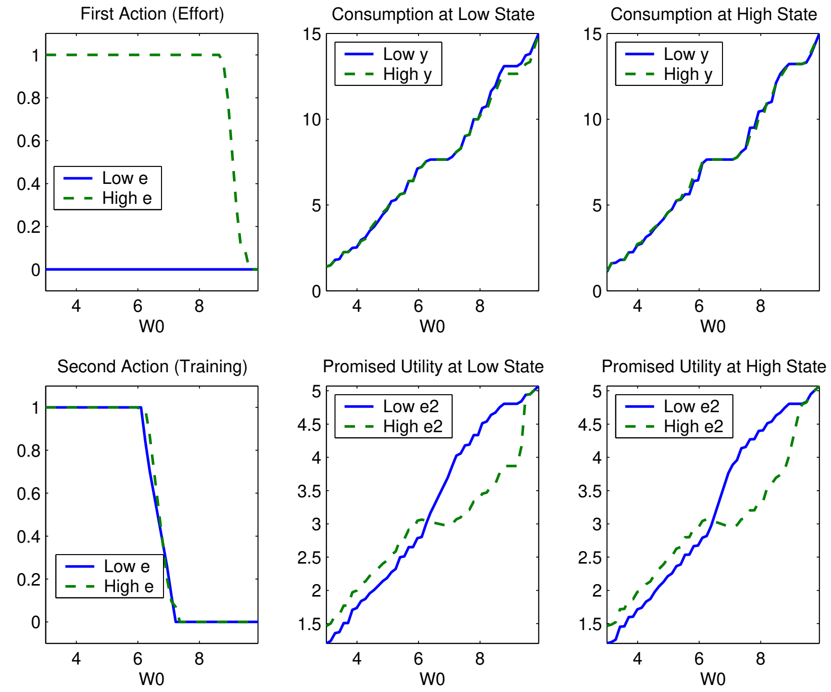

Figure 8: First-Period Policy Functions, a_2 Private Information

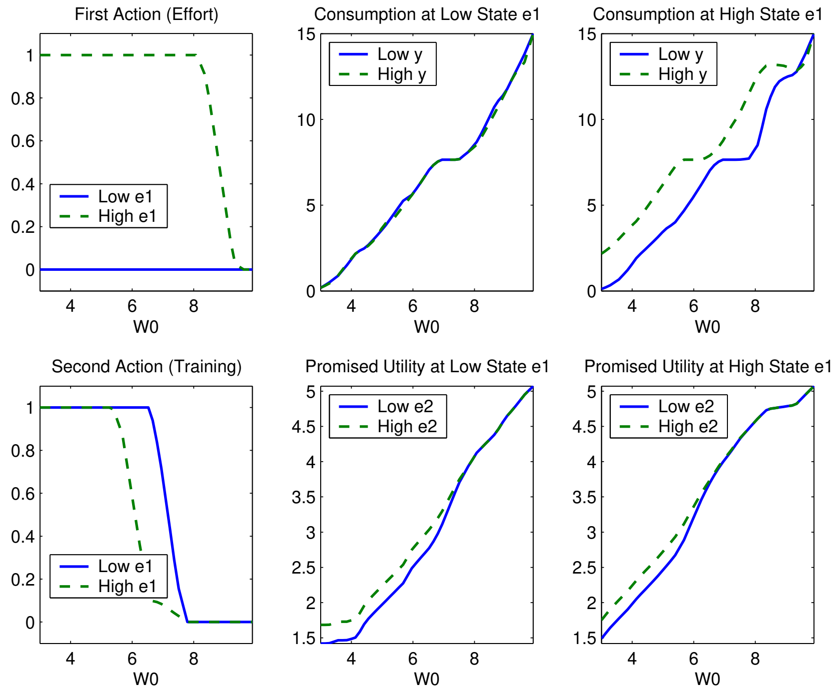

Figure 9: First-Period Policy Functions, e, a_1, and a_2 Private Information