Credit Guarantees, Moral Hazard, and the Optimality of Public Reserves *

Matthias Doepke† — UCLA
Robert M. Townsend‡ — University of Chicago

August 2001

Abstract

In this paper we show that public reserves with a low return and a partial credit-guarantee scheme can be optimal if banks face a moral hazard problem with both hidden actions and hidden information. In our model, banks face uncertain returns on their loans or investments, and both the level of investment and the actual returns are unobservable to anyone but the bank itself. We formulate the problem of providing optimal incentive-compatible credit insurance to the banks, and find that the optimal contract has the feature that low-return public reserves are used. This occurs even though public reserves are dominated in return by other investments, and would not be used in a full-information environment. In order to compute the optimal partial insurance scheme, we develop general recursive methods to solve for optimal contracts in dynamic principal-agent models with hidden income and hidden actions.

*We are grateful to Harold Cole, Hugo Hopenhayn, Ruilin Zhou, and seminar participants at Chicago, MIT, Boston University, IIES, and the SED meeting in San Jose for helpful comments.

†Department of Economics, UCLA, Box 951444, Los Angeles, CA 90095. E-mail address: doepke@econ.ucla.edu.

‡Department of Economics, The University of Chicago, 1126 East 59th Street, Chicago, IL 60637. E-mail address: rtownsen@midway.uchicago.edu.

1 Introduction

We consider the problem of providing optimal incentive-compatible insurance in an environment where banks face uncertain returns on their loans or investments, and the level of investment and the actual returns on investment are unobservable to anyone but the bank itself. The model that we use to analyze the role of public reserves is thus a representative of a wider class of models with unobserved states and unobserved actions. In such environments, optimal allocations cannot be computed with standard methods, since the number of incentive constraints becomes too large. We develop recursive methods to solve for optimal incentive-compatible contracts that require only a relatively small number of constraints.

In dynamic principal-agent models with unobserved income shocks the principal generally needs some control over the intertemporal rate of substitution of the agent in order to induce the agent to report his income truthfully. Controlling the intertemporal rate of substitution becomes difficult, however, if the agent can take unobserved actions that influence his future consumption possibilities. Indeed, Allen (1985) shows that if the agent has unobserved access to borrowing and lending in a perfect credit market, the principal cannot offer any additional insurance. The reason is that any agent will choose the reporting scheme that yields the highest net discounted transfer regardless of what the actual endowment is. The planner is no longer able to distinguish agents with low and high endowments.

For real-world applications, an important question is what the optimal insurance scheme looks like if the agent can take hidden actions that are different from accessing a perfect credit market. Examples of such actions include risky investments, or storage at a return lower than the credit-market interest rate. In this paper, we present methods to compute the optimal insurance contract in a wide class of such environments. We apply our methods to a banking model where bankers face uncertain returns on their investments, and both the level of investment and the realized return are observed only by the bankers themselves. We show that in this environment an insurance scheme can arise in which high-yielding private investments and low-yielding public reserves coexist. In other words, our model provides a new rationale for bank reserves. Reserves arise as a feature of the optimal insurance contract in an environment characterized by moral hazard. The traditional justification for public reserves, the need to provide liquidity, does not play a role in our formulation.

The setup of our banking model is as follows. A large number of bank owners or investors are trying to decide on their own investment strategies and whether or not to adopt some kind of collective credit guarantee scheme. Each bank owner/investor, given his own unobserved income, has to decide on his current consumption compensation, how much to take out of the bank, as well as the amount to invest in a variety of risky loans, assets that have an uncertain payoff in income next period. Other things equal, each bank owner/investor would like to smooth out next period's income fluctuations.

There are two ways to do that smoothing. One way is for all bank owner/investors to remain on their own, to respond to higher income this period by consuming more and also saving more, that is investing more in risky loans, though the return on those loans next period remains uncertain. A second way is for bank owner/investors to get together and collectively insure one another against loan default, that is, levying premia when portfolios are successful and paying indemnities against the low yields of some of the projects financed with bank loans. In effect this is like a deposit insurance scheme, an attempt to stabilize returns on investment against the risk in loan portfolios, but here it is the bankers' money that is at risk and the loan return the risk that is insured. The possibility of such a scheme is limited, however, if the level of investment and realized investment returns are unobserved. The bank then becomes a 'black box', with all the inner workings hidden to outsiders. More formally, the classic moral hazard problem arises: A low current-period level of investment makes high returns less likely next period, so if loan default were insured completely, and levels of investment were not known, no banker would invest or be diligent and all would claim indemnities. Recall again that investment comes at the expense of current consumption levels, and private, unobserved investment cannot be required.

Yet an alternative partial insurance or loan guarantee scheme is possible. Bank owners can set up a separate public reserve fund financed with premia but with more limited possibilities for indemnities. Money placed in this fund is under collective control and can be put into low-yield but safe assets. The advantage of the fund is that there is presumed to be no moral hazard problem, so next periods relatively low returns can be used to compensate bank owners with bad loans. For a variety of parameter values it turns out to be optimal for bank owners to to pay premia into the public fund. While the bank owners still invest in loans, they are doing less internal investing than they would if they were entirely on their own. Surprisingly, the insurance gain from the public reserve fund is so large that it can under certain parameter values dominate what might seem to be more efficient investment. That is, at an optimum, the average yield from more liberal internal investment among bank owners can dominate the yield-safe asset, but still it is optimal to set up the public reserve. The reason is that eliciting the information about actual loan yields under the more liberal policy would be too costly in terms of incentives. Also, since bank owners invest less in risky loans when the reserve fund is used, the variability of their income in the next period is reduced, further mitigating the original insurance problem. In effect, we have created a rational for reserves.

This policy only disappears when the economy's internal rate of return (i.e., the average return on investments) is so high that it makes sense for the public authority to issue its own bond, that is, borrow from outsiders, and lend to the bankers. Even here, however, some insurance is still possible in the sense that the bank with successful loan portfolios pay back more to the public authority than do those with loan defaults. In effect, though all bankers might appear collectively to be borrowing at the outside rate, a risk contingency kicks in for those with loan defaults.

On the other hand, when the economy's internal rate of return is low (though possibly still higher than the rate of return on public reserves) in the optimal insurance scheme bankers do not invest at all in risky assets with unobserved returns. All intertemporal smoothing is provided through safe investments with observed returns by the fund. Here we observe the seemingly perverse situation that the utility of the bankers depends negatively on the return of investment projects. Even though in the optimum the investment projects are not used at all, their rate of return still influences the optimal allocation. With a high rate of return, it is more difficult to provide incentives for truthful reporting, and consequently less insurance is possible.

Our results here are related to a literature on the optimality of public insurance, or ironically, how public insurance can be non-optimal. The idea in Attanasio and Rios-Rull (2000) is that the provision of public insurance for aggregate shocks can undercut self-sustaining private insurance against idiosyncratic shocks. Most of the models feature limited commitment, rather than private information, and people continue to participate in private schemes only as they weigh the benefit of future participation, insurance against a bad shock, against the current cost, e.g., of paying a premium now. Of course, such insurance is most valuable when one needs it most, particularly when aggregate shocks are perverse. Thus publicly engineered transfers to those adversely affected, which ceteris paribus can be a good thing, can undercut some of the benefits of private transfers, and this latter effect can, depending on parameter values, be large enough as to cause a net welfare loss. The result is related to Perri and Krueger (1999) who show that progressive taxation can have perverse effects, and to Ligon, Thomas, and Worrall (2000) who show that private storage in an insurance context can have the same perverse effect.

Our model here concentrates on moral hazard and hidden information. We show that high-yielding private investments can coexist with low-yielding public investments in the optimum, and that higher returns on private uncontrolled investment that would be welfare improving in autarky can have a perverse welfare effect to the information constrained optimum. However, one thing which sets our paper apart from the literature is that we do not suppose that the public aspect of insurance are imposed exogenously and suboptimally onto the privately optimizing agents. Rather, we solve a mechanism design problem which derives the optimal public policy. In effect there is little distinction between the public and the private sector. We might equally well expect the reserve fund schemes we described to be adopted by the private sector or to be imposed by an optimizing government. We only know that some kind of collectivity is required.

Our results appear related to several private and public schemes we have observed in the past or observe contemporaneously. In typical contemporary deposit insurance schemes the public puts money in demand deposits and financial institutions and the funds are, in turn, put via loans into long-term and potentially illiquid projects. There is a risk that deposits will need to be withdrawn, and that a bank will find itself unable to meet the demand. To guard against that, financial institutions are required to put some of their assets into liquid if low-yielding reserves. Banks also pay premia into an insurance fund. In turn, when faced with excessive withdrawals, the bank can draw on reserves and if necessary get an indemnity from the fund. Though reserves and deposit insurance funds are nominally used to guard against liquidity, not insolvency problems, the boundary in practice is less than clear. In any event, in our setup, we do not make a distinction between short term low yield assets, on the one hand, and long term high yield assets on the other. Neither do we make a distinction between liquidity and insolvency. The banks of our model do on occasion suffer from genuinely low portfolio yields, and subject to moral hazard and hidden information problems, it is precisely this risk that they wish to and are able to partially insure.

Other historically observed schemes seem to have similar group insurance aspects. Much work has been done lately in understanding the Suffolk Bank system. Groups of private banks in the New England area with circulating notes agreed voluntarily to put some of their assets into the Suffolk Bank. In turn the Suffolk Bank would redeem individual bank notes at par, even when a given bank was under stress or failed. Apparently notes thus circulated at par over a wide geographic area, whereas prior to the Suffolk Bank system, notes circulated at an increasing discount the further from the issuing bank.

While our emphasis is on credit guarantees, the methods developed in this paper can be applied to a more general class of models with hidden endowments and hidden actions. Other examples include hidden storage with an arbitrary return, including storage of grain at a loss, or holding money in an inflationary world. In such environments, optimal allocations cannot be computed with standard methods, since the number of incentive constraints becomes too large.

On the technical side, the main contribution of this paper is to develop recursive methods to solve for optimal incentive-compatible contracts that require only a relatively small number of constraints and can be computed using current computer technology. The key feature of our formulations is that they allow the planner to specify behavior off the equilibrium path, which leads to a dramatic reduction in the number of constraints that need to be imposed when computing the optimal contract. Some of our computational techniques derive from Prescott (1997). Specifically, Prescott introduced off-path utility bounds to reduce the dimensionality of private-information problems. While he concentrates on a static moral-hazard model, we extend the method to a dynamic framework.

Our paper is also related to Fernandes and Phelan (2000). Fernandes and Phelan develop recursive methods to deal with dynamic incentive problems with links between the periods, and as we do, they use a vector of utility promises as the state variable. The main difference is that Fernandes and Phelan do not consider environments in which both actions and states are unobservable to the planner, which is the main source of complexity in our setup. Werning (2001) develops a recursive first-order approach to compute a dynamic moral hazard problem with storage. Werning's method applies to a more restricted setting in which output is observed, the return on storage does not exceed the credit-market return, and storage is not subject to random shocks. We regard his methods as complementary to ours, since the computational cost of the first-order approach is lower as long as it is justified. Another feature of our approach is that we start from a general setup which allows randomization and unrestricted communication, and show from first principles that our recursive formulations are equivalent to the original formulation.

A related theoretical paper is Cole and Kocherlakota (2001). Cole and Kocherlakota consider the special case of an economy in which the agent can only save, but not borrow.

They find that as long as the return on storage is sufficiently high (the gross return has to be at least as high as the inverse of the discount factor), the optimal allocation is equivalent to access to a perfect credit market. As in Allen's result, no insurance is possible beyond self-insurance using borrowing and lending. Cole and Kocherlakota interpret this outcome as a microfoundation for missing security markets, since the optimal allocation can be decentralized using one-period bonds only. Our methods can be used to compute optimal outcomes in a similar environment with a lower return on savings, or in environments with a wedge between borrowing and lending rates. The issue of unobserved access to credit markets is also discussed in Fudenberg, Holmstrom, and Milgrom (1990), who derive conditions under which optimal dynamic incentive contracts can be implemented as a sequence of short-term contracts, one of the conditions being that the principal and agent can access the credit market on equal terms.

The paper is organized as follows. In Section 2 we introduce the economic environment that underlies our mechanism design problem. Section 3 develops four different recursive formulations for the planner's problem. All formulations are linear programs, but they differ in their computational needs (that is, the number of variables and constraints). In Sections 4 and 5 we apply our computational methods to the banking environment described above, and show how public reserves arise as the optimal solution to a mechanism design problem. Section 6 discusses a number of additional numerical examples. Section 7 concludes.

All proofs are contained in the mathematical appendix. Section A.1 introduces a general formulation of our mechanism design problem with arbitrary message spaces and full history dependence. Sections A.2 to A.5 derive the recursive representations that were introduced in Section 3 from this general version, and show that all formulations are equivalent.

2 A General Principal-Agent Model with Hidden States and Hidden Actions

In the following sections we develop a number of recursive formulations for a general mechanism design problem, with the aim to provide a formulation that lends itself easily to numerical computation. For maximum generality, when deriving the different recursive formulations we concentrate on the case of infinitely many periods with unobserved endowments and actions in every period. With little change in notation, the formulations can be adapted to models with finitely many periods and/or partially observable endowments and actions.

The physical setup is identical for all programs that we consider. At the beginning of each period the agent receives an endowment e from a finite set E. The endowment 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, the action is unobservable for the planner. In most examples below we will concentrate on positive a and interpret them as storage, but without any changes in the setup we could also allow for a to be negative, which can be interpreted as borrowing. The interpretation is the usual small-economy one, with unrestricted access to outside credit markets. The agent consumes the amount $e+\tau-a$, and enjoys period utility $u(e+\tau-a)$. Our methods do not require any specific assumptions on the utility function $u(\cdot)$, apart from it being real-valued.

The action a influences the probability distribution over the endowment in the next period. Probability $p(e|a)$ denotes the probability of endowment e if the agent took action a in the previous period. For tractability, we assume that all states occur with positive probability, regardless of the action:

Assumption 1 *The probability distribution over the endowment e satisfies $p(e|a) > 0$ for all $e \in E$, all $a \in A$.*

In the first period the probability $p(e)$ of endowment e does not depend on any prior actions.

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 endowment, and the planner has to be able to send a recommendation for the action.

In what follows Q is the discount factor of the planner, and β is the discount factor of the agent. The planner is risk-neutral and minimizes the expected discounted transfer, while the agent maximizes discounted utility. Another 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. The discount factor Q is given by $Q = \frac{1}{1+r}$, where r is taken to be the outside credit-market interest rate for borrowers and lenders. When the discounted surplus is zero, we have attained a Pareto optimum. 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 are bigger than zero, because utility will still be well defined.

3 Recursive Formulations

In this section, we present a series of recursive formulations of the general problem. The main issue that we have to deal with is that the presence of hidden actions and hidden states implies a large number of incentive constraints. The most straightforward formulations of the problem are therefore computationally infeasible once grids get large.

From the outset, we work with a recursive formulation that uses a vector of promised utilities as the beginning-of-period state variable. *In the Mathematical Appendix we provide a formal justification of this setup.* The planner promises to deliver utility $w(e)$ in case state e gets realized. The entire vector of promised utilities is denoted $\mathbf{w}$. The planner has a beginning-of-period value function $V(\mathbf{w}, e)$ that gives the planner's expected discounted surplus if the vector of promised utilities is $\mathbf{w}$ and state e is realized.

We assume that there are finite grids for all choice variables. $\#E$ is the number of grid points for the endowment, $\#T$ 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 $\#\mathbf{W}$. We restrict attention to finite grids in order to be able to solve the model on the computer. Infinite sets can be approximated by making the number of points in the grid large.

The choice of the set $\mathbf{W}$ is not arbitrary, because there are utility vectors that cannot be implemented by the planner. For example, because both endowment and action are unobserved, the planner will never be able to assign higher utilities to lower endowments than to higher endowments. On the other hand, we can be sure that at least some utility vectors are feasible, because the planner can always decide to make the transfer not depend on the endowment, and assign the utilities that result from this policy. In practice, we determine the set $\mathbf{W}$ numerically. We start with a fine grid, and throw out all points that turn out to be infeasible.

3.1 A Direct Approach

We start with the most straightforward formulation of the problem. The agent comes into the period with a vector of promised utilities $\mathbf{w}$. At the beginning of the period, the state e is realized, and the agent makes a report about the state to the planner. Conditional on this report, the planner delivers a contract consisting of lotteries over transfer τ, recommended storage a, and promised utilities for the next period $\mathbf{w}'$. The planner maximizes surplus subject to promise-keeping and incentive-compatibility constraints. The maximization problem of the planner defines a functional equation in the value function $V(\mathbf{w}, e)$.

Program 1:
The maximization problem of a planner who has promised vector $\mathbf{w}$ and *has received report* e is:

$$V(\mathbf{w}, e) = \max_{\pi \geq 0} \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[-\tau + Q \sum\nolimits_E p(e'|a) V(\mathbf{w}', e')\right] \tag{1}$$

subject to the constraints (2) to (5) below. The first constraint is that the $\pi(\cdot)$ sum to one to form a probability measure:

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) = 1. \tag{2}$$

Second, the contract has to deliver the utility that was promised for state e:

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[u(e + \tau - a) + \beta \sum\nolimits_E p(e'|a) w'(e')\right] = w(e). \tag{3}$$

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

$$\forall \tau, a, \hat{a} \neq a: \quad \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(e + \tau - \hat{a}) + \beta \sum\nolimits_E p(e'|\hat{a}) w'(e')\right]$$
$$\leq \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(e + \tau - a) + \beta \sum\nolimits_E p(e'|a) w'(e')\right]. \quad (4)$$

Finally, the agent needs incentives to tell the truth, so that no agent with endowment $\hat{e}$ 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 endowment $\hat{e}$ 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 $\hat{e}$ 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:

$$\forall \hat{e}, \delta: \quad \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \delta(\tau, a)) + \beta \sum\nolimits_E p(e'|\delta(\tau, a)) w'(e')\right] \leq w(\hat{e}). \quad (5)$$

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 (3). Of course, there are $\#E$ problem 1's to solve.

The problem with this version is that the number of truth-telling constraints is very large. 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 (2) and one promise-keeping constraint (3). The number of obedience constraints (4) is $\#T \times \#A \times (\#A - 1)$, and the number of truth-telling constraints (5) is $\#E \times (\#A)^{(\#T \times \#A)}$. 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 this example and a given vector of utility promises $\mathbf{w}$ and realized state e Program 1 is a linear program with 200 variables and 2,097,174 constraints. If we increase the number of possible actions a to ten, the number of truth-telling constraints alone exceeds 10^{100}, which is well above current estimates of the number of atoms in the universe. 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 a method to deal with this specific problem. Afterwards, we will develop formulations that are also capable of computing solutions in environments with many actions.

Before we proceed to the next formulation, we need to address what happens in the very first period. At the beginning of time the agent does not yet have a utility vector assigned. Instead, we want to trace out the Pareto frontier by varying an initial utility promise W_0, a scalar, to the agent. By varying W_0 we can compute the utility of the planner $V(W_0)$, and in this way trace out the entire ex ante Pareto frontier. The maximization problem of the planner in period zero is:

$$V(W_0) = \max_{\pi \geq 0} \sum_{E,T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|W_0, e)p(e) \left[-\tau + Q \sum\nolimits_E p(e'|a)V(\mathbf{w}', e')\right] \tag{6}$$

subject to the constraints (7) to (10) below. The first constraint is that the contracts have to be probability measures for each e:

$$\forall e: \quad \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|W_0, e) = 1. \tag{7}$$

Second, the contracts have to deliver promised utility W_0:

$$\sum_{E,T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|W_0, e)p(e) \left[u(e + \tau - a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right] = W_0. \tag{8}$$

Third, the agent has to be obedient in case he told the truth about e. For each state e, transfer τ, recommended action a, and alternative action $\hat{a} \neq a$ we require:

$$\begin{aligned}\forall e, \tau, a, \hat{a} \neq a: \quad & \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|W_0, e) \left[u(e + \tau - \hat{a}) + \beta \sum\nolimits_E p(e'|\hat{a})w'(e')\right] \\ & \leq \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|W_0, e) \left[u(e + \tau - a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right]. \end{aligned} \tag{9}$$

Finally, the agent has to tell the truth and be obedient. For each actual state e, counterfactual state $\hat{e}$ and each function $\delta : T \times A \rightarrow A$ we require:

$$\forall e, \hat{e}, \delta : \quad \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|W_0, \hat{e}) \left[u(e + \tau - \delta(\tau, a)) + \beta \sum\nolimits_E p(e'|\delta(\tau, a))w'(e')\right]$$
$$\leq \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|W_0, e) \left[u(e + \tau - a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right]. \qquad (10)$$

This method of handling the first period is especially useful if we restrict attention to a two-period model. $V(\mathbf{w}, e)$ on the right-hand side of (6) is then the value function for the second period as in the left-hand side of (1), with the future, third period suppressed. Given the value function for the second period, we only have to compute the above single program to find the optimal contract in the first period for a given promised utility W_0. One does not have to iterate on a functional equation to find $V(\mathbf{w}, e)$ in (1), and hence the optimal $\pi(\cdot)$.

Models with a finite number of periods can be solved in the same way with a finite number of iterations. In an infinite-period setting we do need to find the value function $V(\mathbf{w}, e)$, that is, $V(\mathbf{w}, e)$ in (1) is a solution to a functional equation, found in practice by iteration to convergence. It is therefore more efficient in period zero to use that function $V(\mathbf{w}, e)$ to trace out the Pareto frontier. This can be done by letting 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 is:

$$V(W_0) = \max_{\pi \geq 0} \sum_{\mathbf{W}} \pi(\mathbf{w}|W_0) \ \left[\sum\nolimits_E p(e)V(e, \mathbf{w})\right] \qquad (11)$$

subject to a probability and a promise-keeping constraint:

$$\sum_{\mathbf{W}} \pi(\mathbf{w}|W_0) = 1. \qquad (12)$$

$$\sum_{\mathbf{W}} \pi(\mathbf{w}|W_0) \left[\sum\nolimits_E p(e) \ w(e)\right] = W_0. \qquad (13)$$

This method is simpler in the infinite-period setting, and it is the one that we use with all other formulations below.

We will now develop a different formulation of the problem with a much smaller number of incentive constraints.

3.2 Double Reporting

The basic idea of this section is to divide the period into two parts and to let the agent report the endowment a second time after the transfer is received, but before a recommendation for the action is received. On the equilibrium path the agent will make the correct report twice and follow the recommended action, and the optimal allocation will be the same as in the first formulation. On the other hand, the second report allows the planner to specify behavior off the equilibrium path, because outcomes are determined even if the second report does not coincide with the first. We will see that this possibility leads to a significant reduction in the number of incentive constraints. *A formal proof of the equivalence of this formulation to the first one is contained in the Appendix.*

As before, the agent comes into the period with a vector of promised utilities $\mathbf{w}$. At the beginning of the period, the agent observes the state e and makes a report to the planner. Then the planner delivers the transfer τ, and afterwards the agent reports the endowment e again. The incentive-compatibility constraints will ensure that this second report be correct, even if the first report was false. Because now the planner receives a report after the transfer, the number of possible transfers does not affect the number of truth-telling constraints, as it did in (5).

Program 2:
The optimization problem of a planner who promised utility vector $\mathbf{w}$ and *has already received first report* e is:

$$V(\mathbf{w}, e) = \max_{\pi \geq 0} \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, e) \; \left[-\tau + Q \sum\nolimits_{E} p(e'|a) V(\mathbf{w}', e')\right] \tag{14}$$

subject to constraints (15)-(20) below. Notice that the contract $\pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, e)$ is conditioned on two reports e, unlike in (1). The first constraint, much like (2), is that the $\pi(\cdot)$ form a probability measure for any second report $\hat{e}$:

$$\forall \hat{e} : \quad \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, \hat{e}) = 1. \tag{15}$$

Since the second report is made *after* the transfer, we have to enforce that the transfer does not depend on the second report. For all $\hat{e} \neq e$ and all τ, we require:

$$\forall \hat{e} \neq e, \tau : \quad \sum_{A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, e) = \sum_{A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, \hat{e}). \tag{16}$$

Given that the agent told the truth twice, the contract has to deliver the promised utility $w(e)$ for state e from vector $\mathbf{w}$. That is, much like (3) above:

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, e) \left[u(e + \tau - a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right] = w(e). \tag{17}$$

Next, the agent needs to be obedient. Given that the second report is true, it has to be optimal for the agent to follow the recommended action. For each true state $\hat{e}$, transfer τ, recommended action a, and alternative action $\hat{a} \neq a$ we require much like (4):

$$\begin{aligned} \forall \hat{e}, \tau, a, \hat{a} \neq a: \quad & \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, \hat{e}) \left[u(\hat{e} + \tau - \hat{a}) + \beta \sum\nolimits_E p(e'|\hat{a})w'(e')\right] \\ & \leq \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, \hat{e}) \left[u(\hat{e} + \tau - a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right]. \end{aligned} \tag{18}$$

We also have to ensure that the agent prefers to tell the truth at the second report, no matter what he reported the first time around. *Since the transfer is already known at the time the second report is made, the number of deviations from the recommended actions that we have to consider does not depend on the number of possible transfers*. For each actual $\hat{e}$, transfer τ, second report $\hat{\hat{e}} \neq \hat{e}$, and action strategy $\delta : A \to A$, we require:

$$\begin{aligned} \forall \hat{e}, \tau, \hat{\hat{e}} \neq \hat{e}, \delta: \quad & \sum_{A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, \hat{\hat{e}}) \left[u(\hat{e} + \tau - \delta(a)) + \beta \sum\nolimits_E p(e'|\delta(a))w'(e')\right] \\ & \leq \sum_{A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, \hat{e}) \left[u(\hat{e} + \tau - a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right]. \end{aligned} \tag{19}$$

Finally, we also have to ensure that the first report e be correct. That is, an agent who is truly at state $\hat{e}$ and should get $w(\hat{e})$, but made a counterfactual first report e, cannot get more utility than was promised for state $\hat{e}$. For all $\hat{e} \neq e$ we require:

$$\forall \hat{e} \neq e: \quad \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e, \hat{e}) \left[u(\hat{e} + \tau - a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right] \leq w(\hat{e}). \tag{20}$$

Notice that these latter truth-telling constraints do not involve deviations in the action a. At the time of the first report the agent knows that the second report will be correct and that he will take the recommended action, because constraints (18) and (19) hold.

The number of variables in this formulation is $\#E \times \#T \times \#A \times \#\mathbf{W}$. Thus, the number of variables increased relative to the first version, since π now also depends on the second report $\hat{e}$. Similarly, there are $\#E$ probability constraints (15), $(\#E-1)\times\#T$ independence constraints (16), and there is one promise-keeping constraint (17). The total number of obedience constraints (18) is $\#E \times \#T \times \#A \times (\#A-1)$. The number of truth-telling constraints for the second report (19) is $\#E \times \#T \times (\#E-1) \times (\#A)^{(\#A)}$, and the number of truth-telling constraints for the first report (20) is $\#E-1$. Going back to our example, consider a program with two states e, ten transfers τ, two actions a, and ten utility vectors $\mathbf{w}'$. For this example Program 2 has 400 variables and 134 constraints. Compared to Program 1, the number of variables increases by 200, but the number of constraints decreases by more then two million. This makes it possible to solve Program 2 on a standard personal computer. Program 2 does less well if the number of actions is large. If we increase the number of actions a to ten, Program 2 has 2000 variables and more than 10^{11} truth-telling constraints. This is a lot smaller than Program 1, but still too big to be handled by standard computer hardware. The key advantage of Program 2 relative to Program 1 is that the number of constraints does not increase exponentially with the number of possible transfers τ. As long as the number of possible actions a is small, this formulation allows computation with fine grids for the other variables. However, the number of constraints still increases exponentially with the number of actions. In the next section we will address methods to deal with this problem.

The equivalence of this formulation to the original Program 1 is not immediate, since the formulation builds on behavior off the equilibrium path that was left unspecified in Program 1. In the appendix, we derive both Program 1 and Program 2 from first principles, and show that the two formulations are equivalent.

3.3 Assigning Off-Path Utility Bounds

We saw already in the last section that specifying behavior off the equilibrium path can lead to a reduction in the number of incentive-compatibility constraints. We will now exploit this idea in a way similar to Prescott (1997) in order to reduce the number of truth-telling constraints. The choice variables in the new formulation include utility bounds $v(\cdot)$ that specify the maximum utility an agent can get when lying about the endowment and receiving a certain recommendation. Specifically, for a given reported endowment e, $v(\hat{e}, e, \tau, a)$ is an upper bound for the utility of an agent who actually has endowment $\hat{e} \neq e$, reported endowment e nevertheless, and received transfer τ and recommendation a. 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 this utility of saying e when being at state $\hat{e}$ is no larger than the utility promise $w(\hat{e})$ for $\hat{e}$.

Program 3:
The optimization problem of the planner in this formulation given report e and promised utility vector $\mathbf{w}$ is:

$$V(\mathbf{w}, e) = \max_{\pi \geq 0, v} \sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[-\tau + Q \sum\nolimits_E p(e'|a)V(\mathbf{w}', e')\right] \tag{21}$$

subject to the constraints (22)-(26) below. Apart from the addition of the utility bounds $v(\cdot)$ the objective function (21) is identical to (1). The first constraint is the probability measure constraint, identical with (2):

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) = 1. \tag{22}$$

The second constraint is the promise-keeping constraint, identical with (3):

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[u(e + \tau - a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right] = w(e). \tag{23}$$

We have to ensure that the agent is obedient and follows the recommendations of the planner, given that the report is true. For each transfer τ, recommended action a, and alternative action $\hat{a} \neq a$, we require as in (4):

$$\begin{aligned} \forall \tau, a, \hat{a} \neq a : \quad & \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[u(e + \tau - \hat{a}) + \beta \sum\nolimits_E p(e'|\hat{a})w'(e')\right] \\ & \leq \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[u(e + \tau - a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right]. \end{aligned} \tag{24}$$

Next, 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 that $v(\hat{e}, e, \tau, a)$, 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:

$$\forall \hat{e} \neq e, \tau, a, \hat{a}: \quad \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[u(\hat{e} + \tau - \hat{a}) + \beta \sum\nolimits_{E} p(e'|\hat{a}) w'(e')\right] \leq v(\hat{e}, e, \tau, a). \tag{25}$$

Finally, 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:

$$\forall \hat{e} \neq e: \quad \sum_{T,A} v(\hat{e}, e, \tau, a) \leq w(\hat{e}). \tag{26}$$

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 (22) and one promise-keeping constraint (23). The number of obedience constraints (24) is $\#T \times \#A \times (\#A - 1)$. There are $(\#E - 1) \times \#T \times (\#A)^2$ constraints (25) to implement the utility bounds, and $(\#E - 1)$ truth-telling constraints (26). 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 3, 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, which is even smaller than Program 2. For the first time, the program is still computable if we increase the number of actions a to ten. In that case, Program 3 has 1100 variables and 1903 constraints, which is still sufficiently small to be solved on a personal computer.

As pointed out by Prescott (1997), when implementing formulations like this numerically one needs to pay attention to the scaling of utility. Linear programming routines usually search only over nonnegative values for the choice variables. Therefore utility needs to be rescaled so that the utility bounds are positive. Given the finite grids for endowment and transfer, it is always possible to rescale utility this way.

3.4 Subdividing the Problem

In the last formulation, the number of constraints gets large if both the grids for transfer τ and action a are made very fine. In practice, this leads to memory problems when computing. In this section we develop a formulation in which *the transfer and the recommendation are assigned at two different stages.* At each stage, the number of variables and constraints is relatively small. This enables us to compute problems with fine grids for actions and transfers.

As in Section 3.2 on double reporting, the period is divided into two parts. In the first subperiod the agent reports the endowment, and the planner assigns the transfer, an interim utility when the agent is telling the truth, as well as a vector of utility bounds in case the agent was lying. In the second subperiod the planner assigns an action and a vector of promised utilities for the next period.

The solution to the problem in the second subperiod is computed for each combination of endowment e, transfer τ, interim utility $w_m(e)$ (m for "middle" or interim) along the truth-telling path, and vector of utility bounds for lying, $\bar{\mathbf{w}}_m(\hat{e}, e)$. Here $\bar{w}_m(\hat{e}, e)$ is an upper bound on the utility *an agent can get who has endowment $\hat{e}$, but reported e nevertheless.* In order to make the exposition more transparent, we write $w_m(e)$ and the vector $\bar{\mathbf{w}}_m(\hat{e}, e)$ as a function of e to indicate the state for which the interim utility is assigned. $\bar{\mathbf{w}}_m(\hat{e}, e)$ is the vector of $\bar{w}_m(\hat{e}, e)$ with components running over endowments $\hat{e} \neq e$. The choice variables in the second subperiod are lotteries over the action a and the vector of promised utilities $\mathbf{w}'$ for the next period.

We use $V_m[e, \tau, w_m(e), \bar{\mathbf{w}}_m(\hat{e}, e)]$ to denote the utility of the planner if the true state is e, the transfer is τ, and $w_m(e)$ and $\bar{\mathbf{w}}_m(\hat{e}, e)$ are the assigned interim utility and utility bounds for lying. The function $V_m(\cdot)$ is determined in the second subperiod (Program 4b below), but as is typical of dynamic programs, we now take that function $V_m(\cdot)$ as given.

We will analyze the first subperiod. In the first subperiod the planner assigns transfers, interim utilities, and utility bounds. When choosing the utility assignments, the planner is restricted to assignments that can actually be implemented in the second subperiod. We use $\mathcal{W}(e, \tau)$ to denote the set of feasible utility assignments for a given state e and transfer τ. For now we will take this set as given, and define it precisely below when we turn to the second subperiod.

The agent comes into the first subperiod with a vector of promised utilities $\mathbf{w}$, to be effected depending on the realized state e. The planner assigns a transfer τ, on-path interim utilities $w_m(e)$, and off-path utility bounds $\bar{\mathbf{w}}_m(\hat{e}, e)$ subject to promise-keeping and truth-telling constraints.

Program 4a:

The maximization problem of the planner given reported endowment e and promised utility vector $\mathbf{w}$ is:

$$V(\mathbf{w},e) = \max_{\pi \geq 0} \sum_{T} \sum_{\mathcal{W}(e,\tau)} \pi(\tau, w_m(e), \bar{\mathbf{w}}_m(\hat{e},e)|\mathbf{w},e) \; V_m[e,\tau,w_m(e),\bar{\mathbf{w}}_m(\hat{e},e)] \tag{27}$$

subject to the constraints (28)-(30) below. The first constraint is the probability constraint:

$$\sum_{T} \sum_{\mathcal{W}(e,\tau)} \pi(\tau, w_m(e), \bar{\mathbf{w}}_m(\hat{e},e)|\mathbf{w},e) = 1. \tag{28}$$

Promises have to be kept:

$$\sum_{T} \sum_{\mathcal{W}(e,\tau)} \pi(\tau, w_m(e), \bar{\mathbf{w}}_m(\hat{e},e)|\mathbf{w},e) \; w_m(e) = w(e). \tag{29}$$

Finally, truth telling has to be observed, so that this e branch is not made too tempting. An agent who is at state $\hat{e}$ should not be able to gain by claiming to be at state e. For all $\hat{e} \neq e$, we require:

$$\forall \hat{e} \neq e: \quad \sum_{T} \sum_{\mathcal{W}(e,\tau)} \pi(\tau, w_m(e), \bar{\mathbf{w}}_m(\hat{e},e)|\mathbf{w},e) \; \bar{w}_m(\hat{e},e) \leq w(\hat{e}). \tag{30}$$

The number of variables in this program is $\sum_T \#\mathcal{W}(e,\tau)$, and there are $1 + (\#E)$ constraints: one probability constraint (28), one promise-keeping constraint (29), and (#E-1) truth-telling constraints (30).

We now turn to the second subperiod. As before in Section 3.3, apart from the lotteries over actions and utilities, the program in the second subperiod also assigns utility bounds. $v(\hat{e},e,\tau,a)$ is an upper bound on the utility an agent can get who has true endowment $\hat{e}$, reported endowment e, and receives transfer τ and recommendation a. These utility bounds are weighted by the probability of receiving recommendation a.

Program 4b:
The following program, given endowment e, transfer τ, and interim utilities and utility bounds, determines $V_m[\cdot]$:

$$V_m[e,\tau,w_m(e),\bar{\mathbf{w}}_m(\hat{e},e)] = \max_{\pi \geq 0, v} \sum_{A,\mathbf{W}'} \pi(a,\mathbf{w}') \left[-\tau + Q \sum\nolimits_E p(e'|a) V(\mathbf{w}',e')\right] \tag{31}$$

subject to constraints (32)-(36) below. The first constraint is the usual probability constraint:

$$\sum_{A,\mathbf{W}'} \pi(a,\mathbf{w}') = 1. \tag{32}$$

The promise-keeping constraint requires that the interim utility $w_m(e)$ that was promised is actually delivered:

$$\sum_{A,\mathbf{W}'} \pi(a,\mathbf{w}') \left[u(e+\tau-a) + \beta \sum\nolimits_E p(e'|a) w'(e')\right] = w_m(e). \tag{33}$$

The obedience constraints ensure that given that the reported endowment is correct, the agent carries out the recommended action. For all recommended a and alternative actions $\hat{a} \neq a$, we require:

$$\begin{aligned} \forall a, \hat{a} \neq a: \quad & \sum_{\mathbf{W}'} \pi(a,\mathbf{w}') \left[u(e+\tau-\hat{a}) + \beta \sum\nolimits_E p(e'|\hat{a}) w'(e')\right] \\ & \leq \sum_{\mathbf{W}'} \pi(a,\mathbf{w}') \left[u(e+\tau-a) + \beta \sum\nolimits_E p(e'|a) w'(e')\right]. \end{aligned} \tag{34}$$

The next set of constraints implements, or respects, the utility bounds. The utility of having reported e, being at state $\hat{e}$, having received transfer τ and recommendation a cannot be larger than $v(\hat{e},e,\tau,a)$. For all $\hat{e} \neq e$, all recommendations a, and all actions $\hat{a}$ we require:

$$\forall \hat{e} \neq e, a, \hat{a}: \quad \sum_{\mathbf{W}'} \pi(a,\mathbf{w}') \left[u(\hat{e}+\tau-\hat{a}) + \beta \sum\nolimits_E p(e'|\hat{a}) w'(e')\right] \leq v(\hat{e},e,\tau,a), \tag{35}$$

where as usual e and τ are givens for Program 4b. Finally, the utility bounds for misreporting the endowment have to be observed. For each endowment $\hat{e} \neq e$ we require:

$$\forall \hat{e} \neq e: \quad \sum_A v(\hat{e},e,\tau,a) \leq \bar{w}_m(\hat{e},e). \tag{36}$$

The number of variables in this program is $\#A \times \#\mathbf{W}$ under $\pi(\cdot)$ plus $(\#E-1) \times \#A$ under $v(\cdot)$, where again e and τ are fixed at this stage of the problem. There is one probability constraint (32) and one promise-keeping constraint (33), and there are $\#A \times (\#A-1)$ obedience constraints (34). The number of constraints (35) to implement the utility bounds is $(\#E-1) \times (\#A)^2$, and the number of truth-telling constraints (36) is $(\#E-1)$. Notice that the number of constraints and variables does not depend on the grid for the transfer.

More possible transfers will *increase the number of programs to be computed* at this second stage (Program 4b), but *the programs will not change in size*. In our example with two endowments e, ten transfers τ, ten actions a, and ten utility vectors $\mathbf{w}'$, Program 4b has 110 variables and 193 constraints, while Program 4a has 3 constraints. The number of variables in Program 4a and the number of Programs 4b that need to be computed depends on the grid for interim utilities. For example, if the grid for interim utilities for each τ and e has ten values, Program 4a has 100 variables, and 1000 Programs 4b need to be computed. In practice, it is generally faster to use Program 3 as long as it is feasible to do so. Large programs can only be computed using Program 4a and 4b, since the individual programs are smaller and require less memory than Program 3.

We still have to define the set $\mathcal{W}(e,\tau)$ of feasible interim utility assignments in Program 4a. It is the set of all assignments for which Program 4b has a solution. With $\mathbf{W}_m$ being the utility grid that is used for $w_m(e)$ and $\bar{\mathbf{w}}_m(\hat{e},e)$, define:

$$\mathcal{W}(e,\tau) = \left\{(w_m(e), \bar{\mathbf{w}}_m(\hat{e},e)) \in (\mathbf{W}_m)^{(\#E)} \mid V_m[e,\tau,w_m(e),\bar{\mathbf{w}}_m(\hat{e},e)] \text{ is defined}\right\} \qquad (37)$$

In other words, we vary utility assignments as parameters or states in Program 4b and rule out the ones for which there is no feasible solution.

4 Risky Banking and the Optimality of Public Reserves

In this section, we will apply the methods developed above to compute optimal incentive-compatible contracts in a *two-period* banking model. For simplicity, suppose that banks (the agents) enter the first period with an observed initial endowment or net worth (a special version of the general setup). The initial endowment can be withdrawn from the bank as profit and consumed by the owner, or it can be invested in varying quantities of loans. Investment is unobserved. The return on that investment, or the endowment in the second period, is random, and can be observed only by the bank itself. Depending on the level of investment in the first period, a high endowment in the second period is more or less likely. Still, the return on loans is uncertain. It is possible that the projects fail and that even a bank that invested at the highest possible level receives the lowest possible endowment in the second period.

Bankers are assumed to be risk-averse, so that other things being equal, they would like to smooth out income fluctuations and insure against the uncertain returns of their investments. If the bankers are on their own, their only mechanism for smoothing income between the two periods is to invest in risky loans. They have no possibility of insuring against the uncertain returns on their investments. But in addition to the banks, there is also a planner who has access to a risk-free storage technology with zero net yield. We solve the problem of providing optimal incentive-compatible insurance to the banks, subject to the requirement that the planner receives a zero discounted surplus. The planner could be identified with the central bank. An alternative assumption would be that there is a continuum of banks who face idiosyncratic, but not aggregate, risk. In that case, the outcome of the planning problem can be interpreted as a mutual insurance scheme run by the banks themselves. The results do not depend on specific institutional assumptions, that is, the optimal contract can be implemented in different ways.

Figures 1 to 4 show features of the optimal contract for a number of variations of our model. In all cases, the known endowment in the first period is .89, a simplification of the general setup. In the second period the endowment can be low or high, .4 or .8. Three investment levels (*in addition to zero investment*) are possible, .1, .2, and .3, and the expected incremental returns per unit invested are 1.31, 1.01, and .90, respectively. In other words, investing .1 increases the expected endowment in the second period by .131, or 1.31 per unit (.131/.1), investing .2 increases the expected endowment by an additional .101 or 1.01 per unit, and so on. The grids for each transfer have 60 elements, and the grid for the promised utility vector which is assigned in period one has 3600 elements.[1] Storage by the planner has zero net return (i.e., the discount factor of the planner equals one, $Q = 1$). The discount factor β of the bankers equals one as well. This assumption is made for simplicity, the results do not change if the discount factor of the bankers is less than one. Notice that the return on private investments exceeds the return on public storage up to an investment level of .2, but is lower after that. If the aim were to maximize the expected endowment in the second period as in a neoclassical world, we therefore would expect the bank to invest .2.

The graphs show outcomes for two different assumptions on the matrix P, which determines the probability of each endowment in period two (down the column) as a function of investment in period 1 (across the row). For Figures 1 and 2, with zero investment the

[1]However, the program dismisses utility vectors which turn out to infeasible, so the dimension of the linear programs which are actually computed is smaller than the one implied by the grid sizes.

probability of the low endowment is $.975$:

$$P = \begin{pmatrix} .9750 & .6475 & .3950 & .1700 \\ .0250 & .3525 & .6050 & .8300 \end{pmatrix}.$$

In other words, with low investment income from investment is likely to be low, but by the same token the variance or riskiness of output in the second period is low as well. If investment is increased, the higher income is more likely, but the variance of output increases as well. The assumption of an increasing variance of output with increasing level of investment is a natural one, but we will later also consider the opposite case to illustrate the relationship of risk to public reserves.

Figure 1 shows the policy functions under full information and full commitment for the case that investment increases the variance of output. Suppose all actions and endowments are observed by the planner. The policy functions are shown as a function of the initial utility promise W_0 to the agent, and the vertical line in all graphs marks the initial utility promise at which the expected surplus of the planner is exactly zero. Since full insurance is possible, investment is at the first-best level of .2 throughout (see the graph in the upper right-hand corner). The observed idiosyncratic shocks across bank-specific investments can be completely smoothed away, leaving the highest social return. The endowments were chosen such that the transfer in the first period (upper-left corner) is exactly zero if the expected surplus of the planner is zero. In other words, there are no public reserves created at the initial date (recall that $Q = 1$). In the second period, the planner implements insurance and transfers resources from high-output bankers with successful projects to low-output bankers with failed projects (lower-left corner). In summary, under full information, public storage is not used, and investment is at the efficient level. The planner merely insures completely against income fluctuations in the second period by transferring resources between banks with successful loans and those with failed loans.

Figure 2 shows the policy functions for the same economy under the assumption of private information. Now both the level of investment in the first period and the return on investment in the second period are unobserved by the planner, and the planner has to provide incentives for obedience (investing properly) and truthful reporting. Investment by the agent (upper-right corner) falls to .1 throughout. Note that private investment is still positive and has a high marginal rate of return above the outside interest rate. At zero expected surplus, the planner taxes or takes in premia in the first period, and establishes a social fund invested at the outside interest rate. In the second period, he returns re-

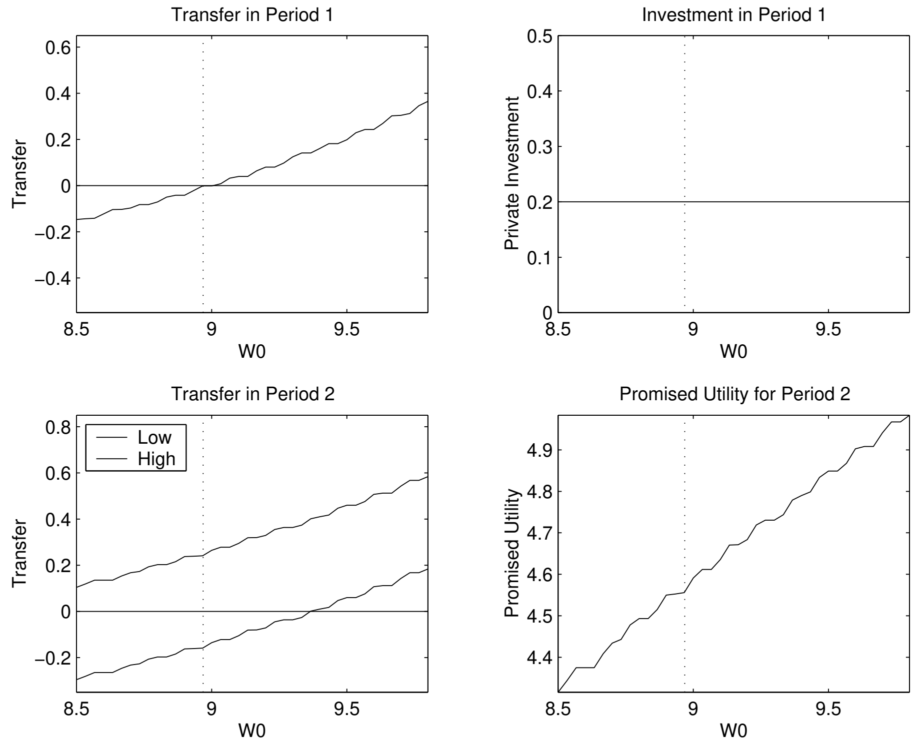

Figure 1: Policy Functions, Unconstrained Solution. Investment Increases Variance of Output in Period 2.

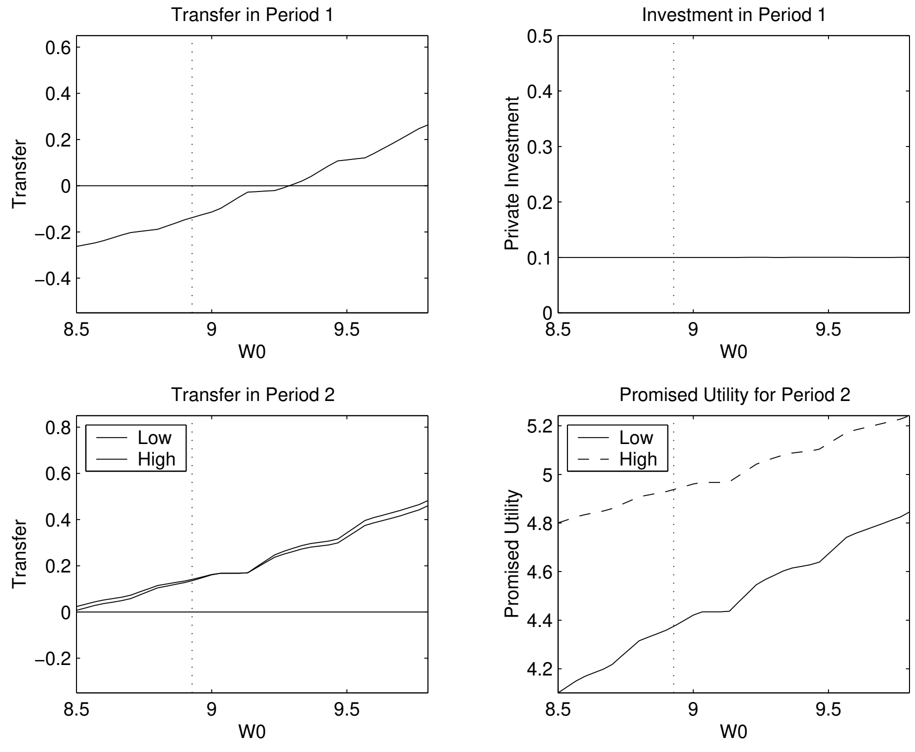

Figure 2: Policy Functions, Constrained Solution. Investment Increases Variance of Output in Period 2.

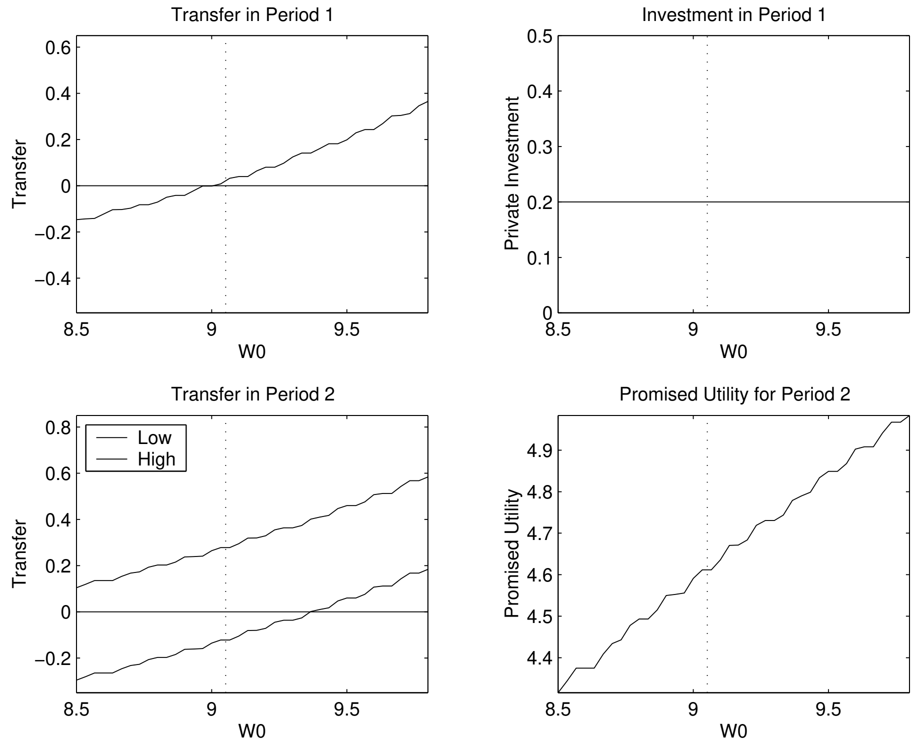

Figure 3: Policy Functions, Unconstrained Solution. Investment Decreases Variance of Output in Period 2.

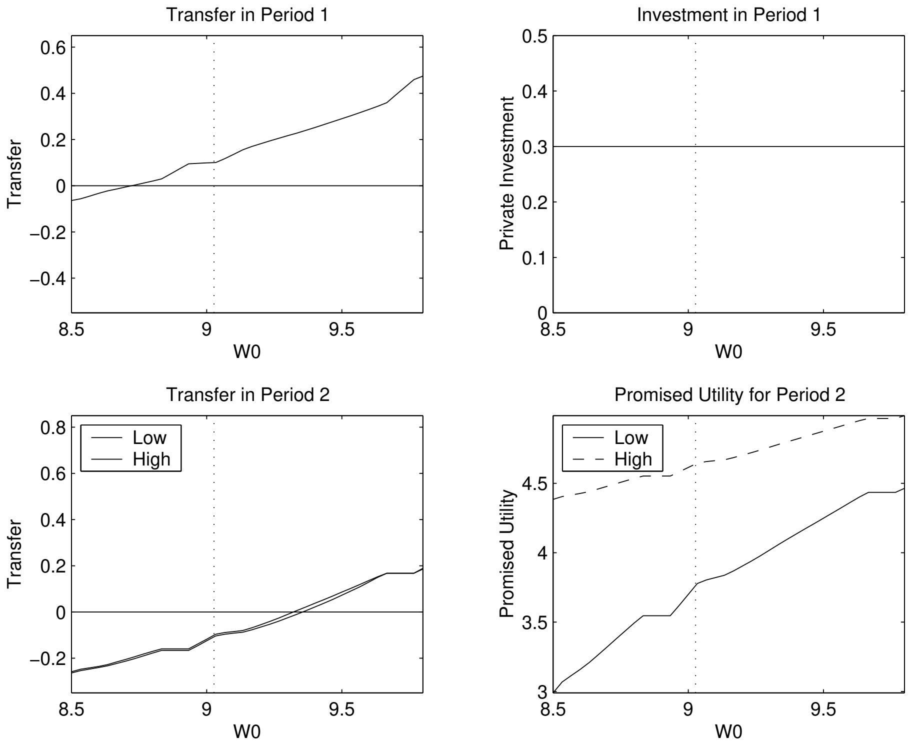

Figure 4: Policy Functions, Constrained Solution. Investment Decreases Variance of Output in Period 2.

sources from indemnities and liquidation of the social fund. The planner now uses public storage to smooth the income of the banks. Transfers in the second period vary only little as a function of the realized endowment. Since the endowment is unobserved and incentives for truthful reporting have to be provided, only little additional insurance against income fluctuations in the second period is possible (due to lotteries and differences in risk aversion).

In summary, under moral hazard we observe inefficient use of private investment in the sense that if there were full information the level of investment would be higher. We also observe the coexistence of high-yielding private investment with low-yielding public reserves. What is the intuition for these results? The key assumption that generates this outcome is that increasing investment increases the variance of output, and under private information, fluctuations in output are not well insured. The choice between public and private investment exhibits the usual risk-return tradeoff. While this assumption seems reasonable, we also computed the model under the opposite assumption to demonstrate how the optimal contract varies with the relationship between investment and risk.

In Figures 3 and 4, at zero investment, the probability of getting the low endowment is reduced from .975 to .83. As investment increases to .3, the probability of getting the low endowment falls to 0.025:

$$P = \begin{pmatrix} .8300 & .5025 & .2500 & .0250 \\ .1700 & .4975 & .7500 & .9750 \end{pmatrix}.$$

This implies that going from zero investment to full investment actually decreases the variance of output in the second period. Apart from this difference, the expected return on investments and other aspects of the environments underlying Figures 1 and 2 and Figures 3 and 4 are identical.

Figure 3 shows the policy functions under full information for the case that investment decreases the variance of output. As in Figure 1, investment is at the first-best level of .2 throughout. In fact, the only difference to Figure 1 is that the vertical line indicating zero surplus for the planner has moved slightly to the right, since the expected endowment in the second period is higher. In other words, at zero surplus the agent receives a higher initial expected utility W_0.

Figure 4 shows the policy functions for the same economy (investment decreases the variance of output) with private information. Investment by the agent (upper-right corner) now rises to .3 (notice that the marginal return is lower than the outside market rate) instead of falling to .1 as it did in Figure 2. Evidently, the inability to control private investment has done some damage. At zero expected surplus, the planner or community fund provides positive transfers to everyone in the first period by borrowing from the outside market, and requires repayment in the second period. No public reserves are held in the low-interest asset at any time. This is exactly the opposite of Figure 2.

In summary, we find that public reserves with a low return can be optimal if banks face a moral-hazard problem. The key underlying assumption is that increased lending or investment increases risk, as seems natural. Using public reserves lowers overall risk and is therefore optimal in the information-constrained solution, even though public storage would be inefficient in the first-best solution. This conclusion would seem to rely on the absence of insurance, that is, bankers who do the private investing bear virtually all the risk arising from random returns. Below, we show that public storage is still optimal in environments where the planner can provide partial incentive-compatible insurance against income shocks.

5 Risky Banking in a Three-Period Setting

In the two-period banking examples in the preceding section there was little scope for insurance against the income shock in the second period (beyond the provision of risk-free investment through the planner), since the world ends after that period. Consequently, the transfer in the second period is (almost) the same across both income realizations, and all risk is born by the bankers. In this section, we show that the basic result of coexisting risky private and low-return public investment holds up in environments with additional periods, where more insurance is possible.

Figures 5 and 6 show an extension of the environment described above to three periods.[2] The initial endowment is known, but in periods 2 and 3 the income of the bankers is unknown and depends on previous investment. Again, this is a modification of the general setup which facilitates computations. As before, the planner can borrow or lend at zero

[2]Even with the methods described in this paper, computing this model is not a trivial task. The main difficulty in the banking example is that public and private investment are close substitutes, so that inaccuracies which arise from using grids for transfers and utility promises can lead to jumps in the policy functions for private and public investment. When the two forms of investment are close substitutes, fine grids for transfers and utility promises need to be chosen, and computation time is high.

net yield. The known endowment in the first period is 1.135. In the second period, the possible income levels are 0.9 and 1.5, and in the third period income can be 0.7 or 1.3. The incomes are chosen such that with investment at the efficient level expected income is constant across the periods. As in the two-period environment, four levels of private investment (0, .1, .2, and .3) are possible. As before, the expected incremental marginal returns for investing at level .1, .2, and .3 are given by 1.31, 1.01, and .9. The probabilities over the two income levels depending on investment are given by the matrix P:

$$P = \begin{pmatrix} .9950 & .7767 & .6083 & .4583 \\ .0050 & .2233 & .3917 & .5417 \end{pmatrix}.$$

The first column corresponds to zero investment, the second column to an investment of .1, and so on. With zero investment, the probability of getting the low income level is .995, and the probability for the high income is .005. The income variance increases with investment in this example. The grids for the transfers in period one and two have 30 elements, and there are 70 grid points for the transfer in the third period. The grid for the promised utility vector which is assigned in period one has 3600 elements, and the grid for the utility assignment in the second period has 3600 elements.

Figure 5 shows the outcome under full information (income and investment level is observed in each period). The outcome is computed using the same grids on transfers and utility promises that were used for the constrained solution below. Private investment is at the efficient level of .2 both in period 1 and period 2 (see graphs in the middle column). Through income-contingent transfers, the planner provides full insurance, i.e., in periods 2 and 3 the difference between the transfer for low- and high-income bankers equals their difference in income in that period. At zero expected surplus for the planner, the average transfer turns out to be zero in every period.

Figure 6 shows the results under private information (income and investment are unobserved). The policies are shown as a function of the realized income. The transfers in the second and third period and the promised utilities for the second period turn out to be contingent on income in the second period, and the utility promises for the third period are contingent on income in periods 2 and 3 ('LH' means low income in 2 and high income in 3, and so on). The optimal policies for private and public investment mirror the results in the two-period environment. Private investment is only at .1 throughout. At zero expected surplus, the transfer is negative in the first period, which implies that the community is creating a safe investment fund. In the second period, the average trans-

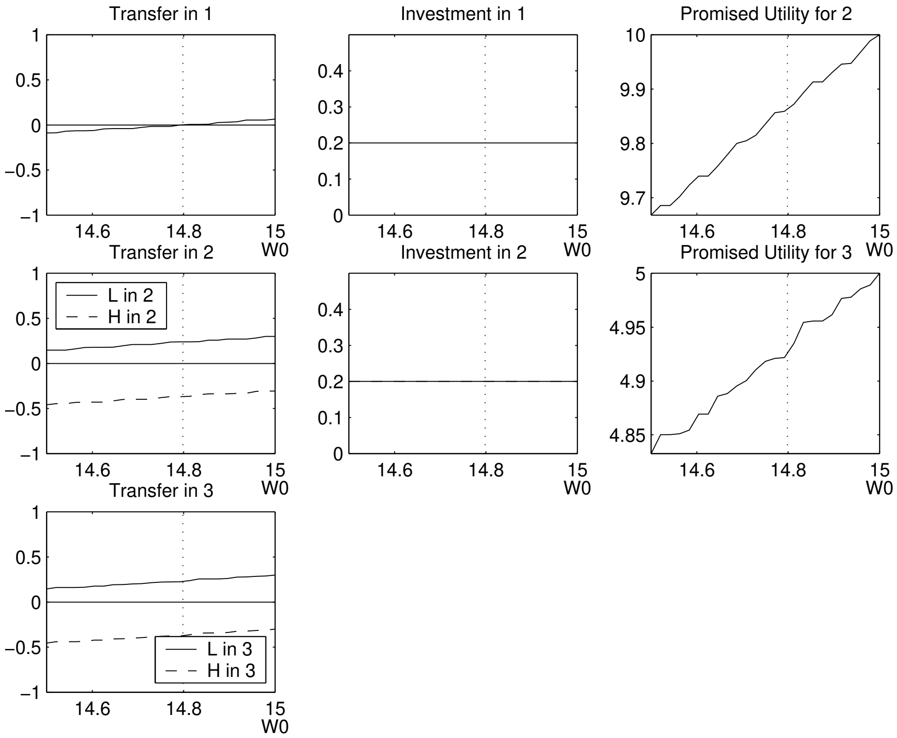

Figure 5: Policy Functions, Unconstrained Solution.

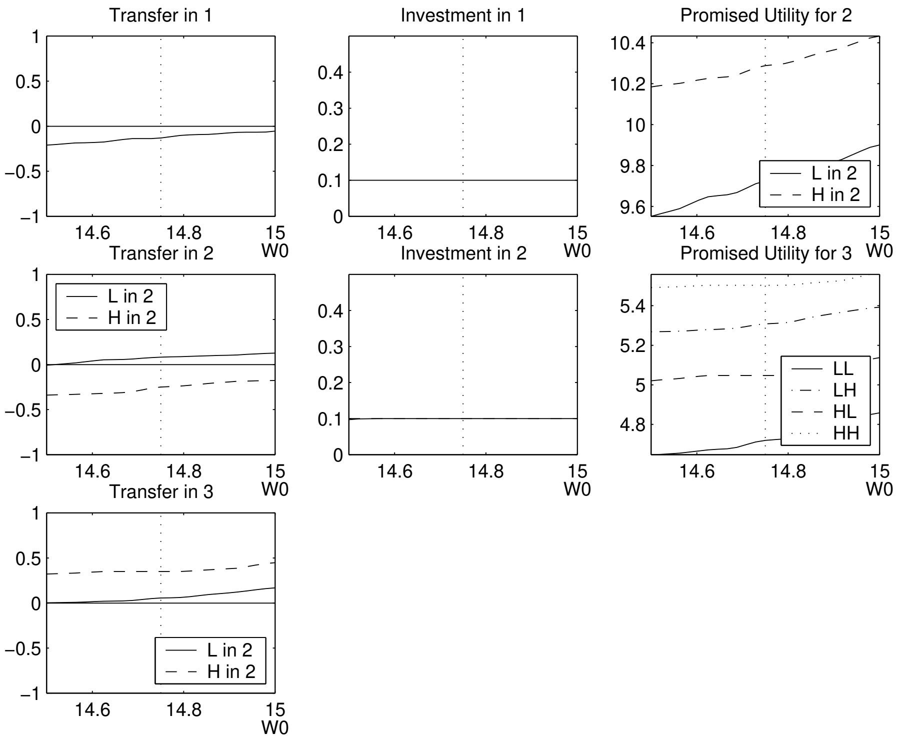

Figure 6: Policy Functions, Constrained Solution.

fer is close to zero, meaning that the fund is maintained. In the third period the average transfer is positive, so the community is returning the safe investment to the bankers.

In the three-period environment, the planner can provide a substantial amount of insurance against the income shock in the second period. The difference between the transfer for bankers with high and low income (solid and dotted line) is about .3, or half of the income difference. Of course, perfect insurance is not possible. The planner can provide this insurance by trading of transfers in the second period against utility promises for the third period (bankers with low income in the second period receive a higher transfer, but a lower utility promise for the third period). Notice that the transfer in the third period is lower for agents who received low income *in the second period*. This form of insurance was not possible in the two-period economy. But even though a substantial amount of insurance is provided, it is still optimal to replace high-yield private investment with low-yield, low-risk public reserves.

The income shock in the third period cannot be insured, since the absence of another period makes it impossible for the planner to trade off transfers against future promised utilities. Notice that unlike in the full-information environment, the transfer in the third period is not contingent on income in that same period.

6 Other Numerical Examples

In the banking model, we concentrated on the situation of private investment with a high return and public storage with a low return. Also, the return on private investment varied with the level of investment. In this section, we analyze some additional numerical examples concentrating on information-constrained insurance. In these examples, the return to private investment is constant, independent of the level of investment, but again different from the public return. We consider both cases in which the return on private investment exceeds the return on public storage, and cases in which the private return is lower than the public return. Here we observe the seemingly perverse effect that *lowering* the return on private investment can actually *increase* the utility of the agent. The reason is that once the return is so low that private investment is not used in the constrained optimum, lowering the return even further makes it easier for the planner to provide incentives for truth-telling and obedience, which results in more insurance. In each case

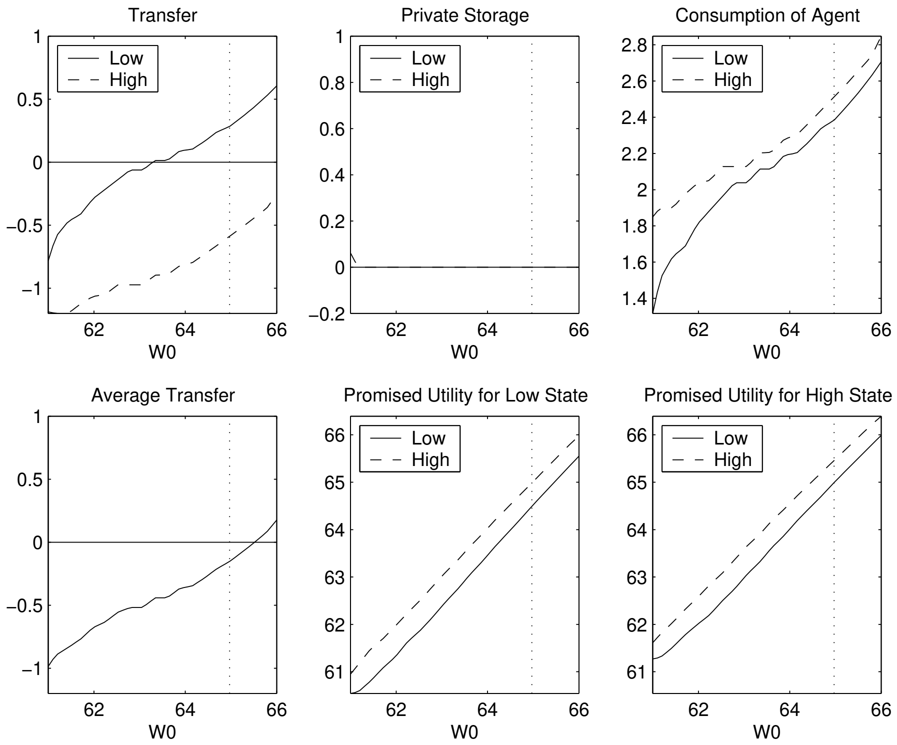

Figure 7: Policy Functions, R=1

we used Program 3 for our computations. The investment technology of the agent will be interpreted as storage.

In all example environments the agent has logarithmic utility. There are two possible endowments e (2.1 or 3.1), three storage levels a (from 0 to 0.8), and thirty transfers τ (from -1.2 to 1) in every period. Planner and agent have the same discount factor $\beta = Q = .952$. The discount factor of the planner corresponds to a risk-free interest rate of 5%. The utility grids have fifty elements. In the initial period, each endowment occurs with probability .5. In all other periods, the probability distribution over endowments depends on the investment level in the previous period. If the investment level is zero, the high endowment occurs with probability .5. The probability of the high endowment increases linearly with the investment level. The expected return on storage R determines how fast the probability of the high endowment increases. For example, for $R = 1$ an increase in storage of 0.1 results in an increase in the expected endowment of 0.1. Since the difference between the two endowment 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 endowment of 3.1 by 0.1, and an equal decrease in the probability of the low endowment of 2.1.

Figure 7 shows the policy functions for $R = 1$ *in an infinite-period setting* as a function of the initial utility promise W_0. We computed the policy functions by iterating on the Bellman operator until convergence, and then solving the initial planning problem (11). It turns out that the choices of the planner are similar for all values of R that are lower than the risk-free interest rate. Therefore we use Figure 7 at $R = 1$ to represent general features of the optimal incentive-compatible contract for $R < 1.05$. The choices of the planner are shown as a function of the initial utility promise to the agent. In all graphs a vertical line marks the initial utility level at which the expected discounted surplus of the planner is zero. As before, if there is a continuum of agents, the planner can be interpreted as a mere programming device to determine the optimal allocation for the community, and the zero-surplus utility level corresponds to the constrained Pareto-optimal allocation for a continuum of agents.

As long as $R < 1.05$ (the market rate), the planner never recommends a positive investment or storage level ever; only the outside credit market is used. The planner uses transfers and utility promises to provide incentive-compatible insurance every period. The transfer for the low endowment is higher than for the high endowment, and at zero expected surplus high-income agents today pay premia, while low-income agents receive indemnities. The difference in consumption between low- and high-endowment agents is only about .2, even though the endowment differs by 1. In order to insure incentive compatibility (i.e., induce the high-endowment agent to report truthfully), the planner promises higher future utility to agents with the high endowment. Both consumption and utility promises increase with initial promised utility.

We can use the optimal policy functions to compute the long-run distribution of consumption and utility in an economy with a continuum of agents. Atkeson and Lucas, Jr. (1992) found that in a class of dynamic incentive problems the efficient allocation has the property that the distribution of consumption diverges over time, i.e., a diminishing fraction of the population receiving an increasing fraction of resources. Phelan (1998) examines the long-run implications of incentive problems more generally. He finds that the limiting distribution for consumption and utility crucially depends on features of the utility function, namely whether or not marginal utility is bounded away from zero as consumption approaches infinity. In our setup, the set of possible consumption and utility values is bounded, since finite grids are used. Figures 8 and 9 show that the long-run distribution of consumption and utility is twin-peaked. Consumption disperses over time, and after a series of good or bad shocks either corner for consumption can be reached.

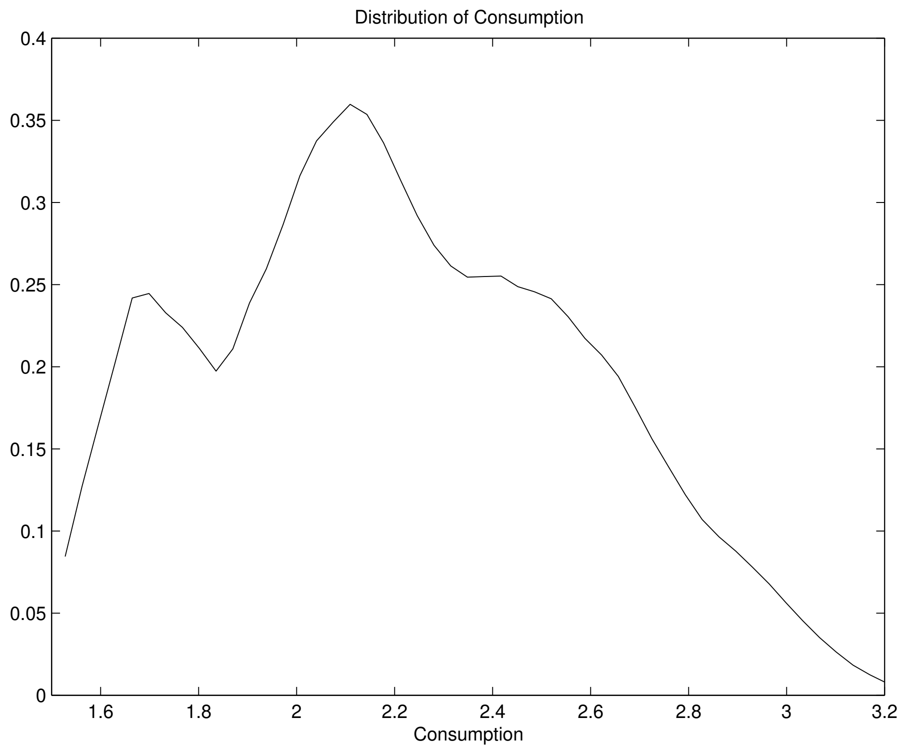

Figure 8: The Long-Run Distribution of Consumption

The corners are not absorbing states, however. For example, agents who consume minimum consumption are thrown back into the distribution once a positive shock hits. Our interpretation for the twin-peaks distribution is that the lower peak is accounted for by agents who were deflected away from the lower bound for consumption.

An interesting question is how high the return on storage needs to be before the planner starts to recommend the use of storage. It turns out that if the planner is nearly indifferent between recommending storage or not, numerical inaccuracies due to finite grids can have large impact on the result. For that reason we made additional computations in a two-period setting, which allowed us to use very fine grids. Figure 10 shows the policy functions in the two-period model when the return on storage equals the credit-market return, and Figure 11 shows the policy functions for $R = 1.07$, i.e., the return on storage slightly exceeds the credit market return. We see that even if the return on storage is slightly higher than the credit-market return, the planner still recommends zero storage. Obviously, under full information the planner would recommend the maximum storage level whenever the return on storage exceeds the credit-market return. The reason that this does not occur in our computations are the incentive constraints. Here, keeping the

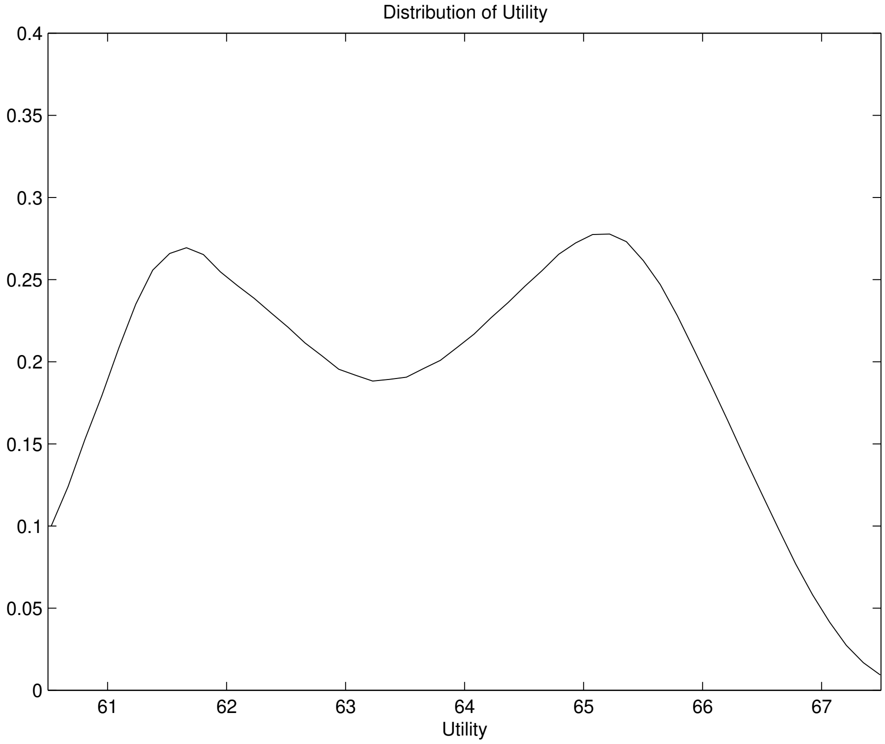

Figure 9: The Long-Run Distribution of Utility

agents at the zero corner for storage makes it easier to provide incentives for truth-telling and obedience.

Figure 12 shows the policy functions for $R = 1.10$. Since the risk-free interest rate is 5%, storage is much more productive than lending in the credit market, the net return is twice as high. The effect of the high return starts to dominate, and consequently the planner recommends the maximum storage level of 0.8 to both agents. Notice that the expected transfer to both agents is positive over the entire range of initial utility promise. In effect, the planner is lending money to the agents, who invest it into the productive storage technology. Other than the different use of storage, the planner provides insurance against income shocks in a similar way to the case with a lower R. Agents with a low endowment get a higher transfer, but lower utility promises.

Figure 13 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 is zero. The utility of the agent is shown for values of R between 0 and 1.1. The top line is the utility the agent gets under full information, that is, when both endowment and storage are observable by the planner. In this case the planner provides full insurance, consumption and utility

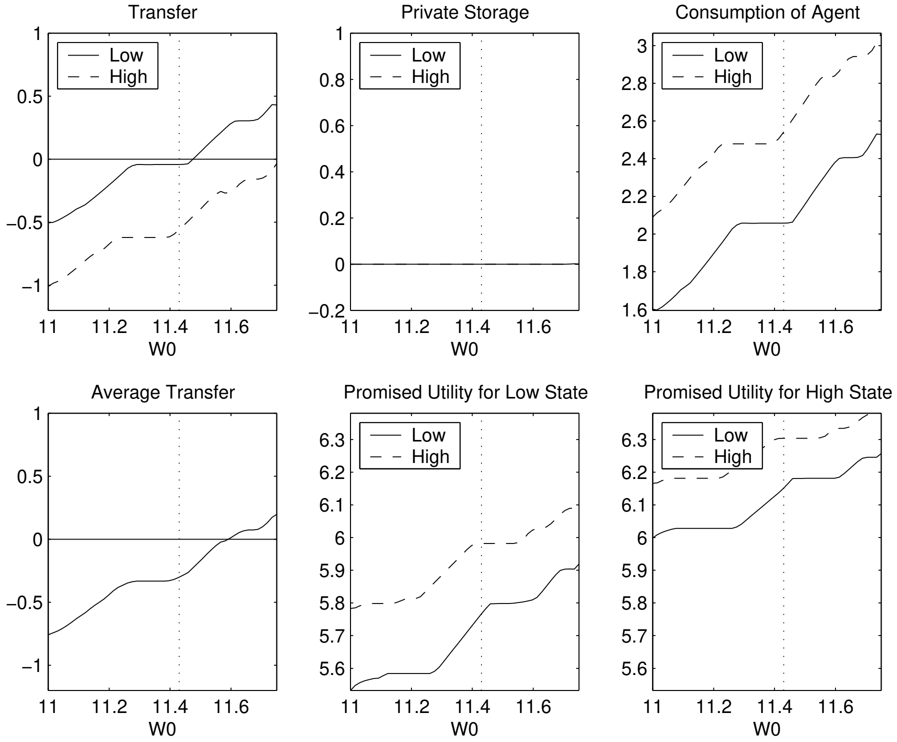

Figure 10: Policy Functions, Two-Period Model, R=1.05

promises are independent of the endowment. It turns out that 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, raising the return on storage does not extend the unconstrained Pareto frontier. Storage is never used, and the utility of the agent does not depend on the return on storage. When R exceeds 1.05, the Pareto frontier expands and the utility of the agent increases. The lower line shows the utility of the agent under autarky. When the return on storage is sufficiently high, 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 middle line shows the utility of the agent with hidden endowments and hidden action. Once the return on storage is sufficiently high, 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 storage levels, a higher return on storage has no positive effects. On the contrary, with a high return on storage 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. Therefore

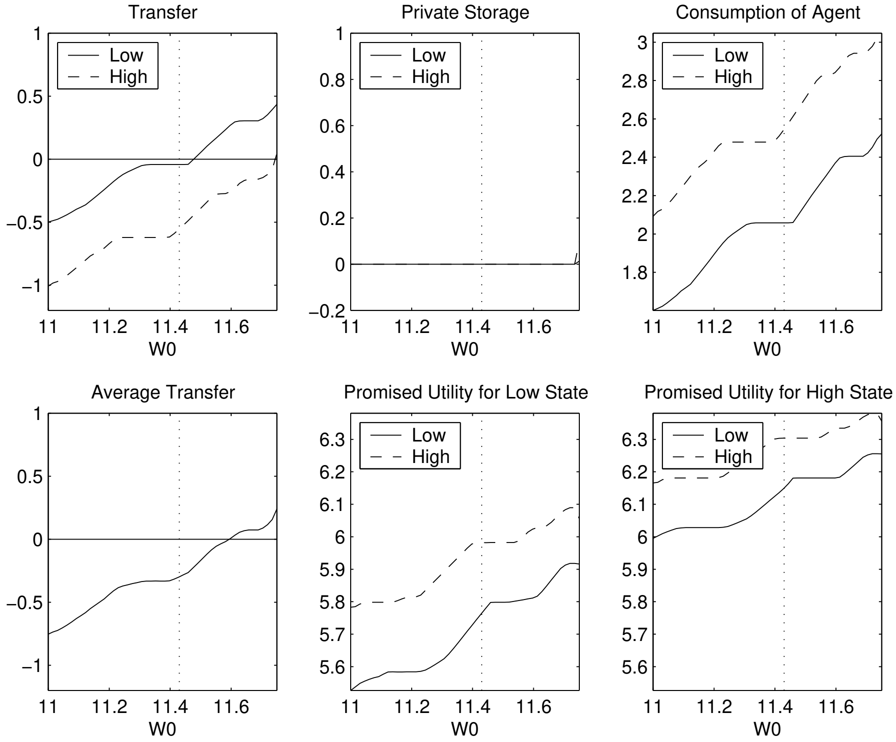

Figure 11: Policy Functions, Two-Period Model, R=1.07

raising the return on storage shrinks the set of feasible allocations, and the utility of the agent falls. Once R exceeds the credit-market return by a sufficient margin, however, private storage is used in the constrained solution, and consequently further increases in R raise utility.

For $R = 1.05$, when public and private returns are equal, the constrained-optimal allocation and autarky yield similar, but not identical, utility. Recall that private agents in the model do not have perfect access to credit markets in the sense that there is an upper bound on borrowing. But we can also see that once the return on investment is somewhat lower than the credit-market interest rate, the gains from optimal incentive-compatible insurance become sizable. Already at $R = 1$ the utility gain from going from autarky to incentive-compatible insurance is as large as the utility advantage of the full-information solution relative to the incentive-compatible solution. At $R = 1$, the difference in utility between autarky and the constrained optimum is equivalent to a gain of 1.4% of consumption in the first period. At $R = 0.95$, the gain is already 1.93%, and at $R = 0.7$ the gain is 2.3%. Thus for all but the highest values for return on storage the gains from optimal insurance are sizable.

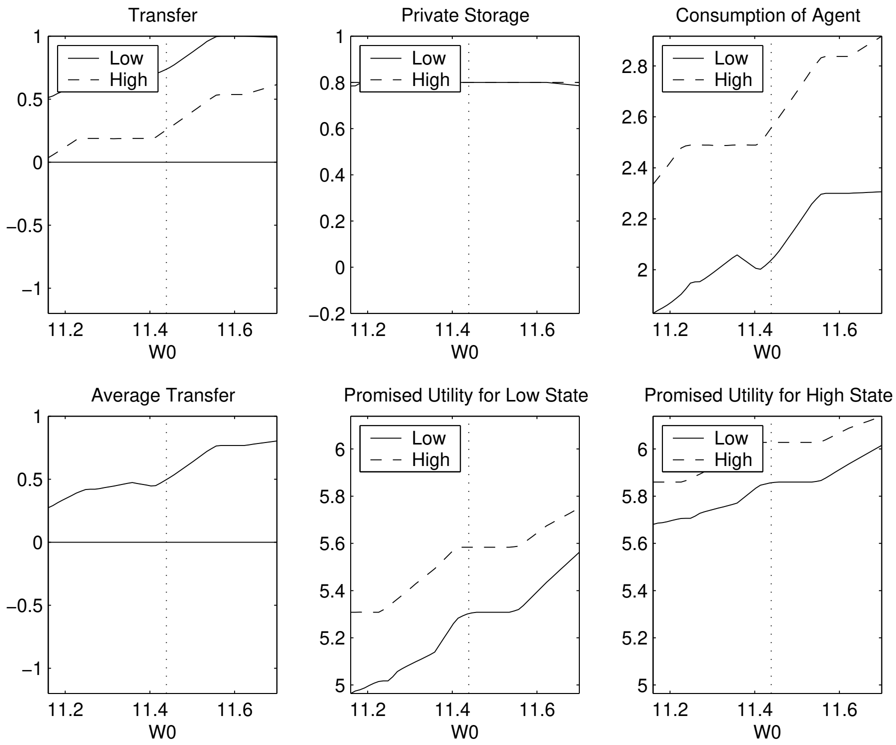

Figure 12: Policy Functions, Two-Period Model, R=1.10

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 endowment. What happens if only one type of constraint is introduced? If we remove the truth-telling constraints (i.e., we assume that the actual endowment 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 of income shocks are fully insured. Conversely, if only the obedience constraints are removed, as if storage could be observed, full insurance cannot be achieved. Since the planner will recommend zero storage, the outcome is the same as in the situation where the storage technology has zero gross return, as in Figure 13 for $R = 0$. Without the obedience constraints, more insurance is possible, but it still falls well short of full insurance.

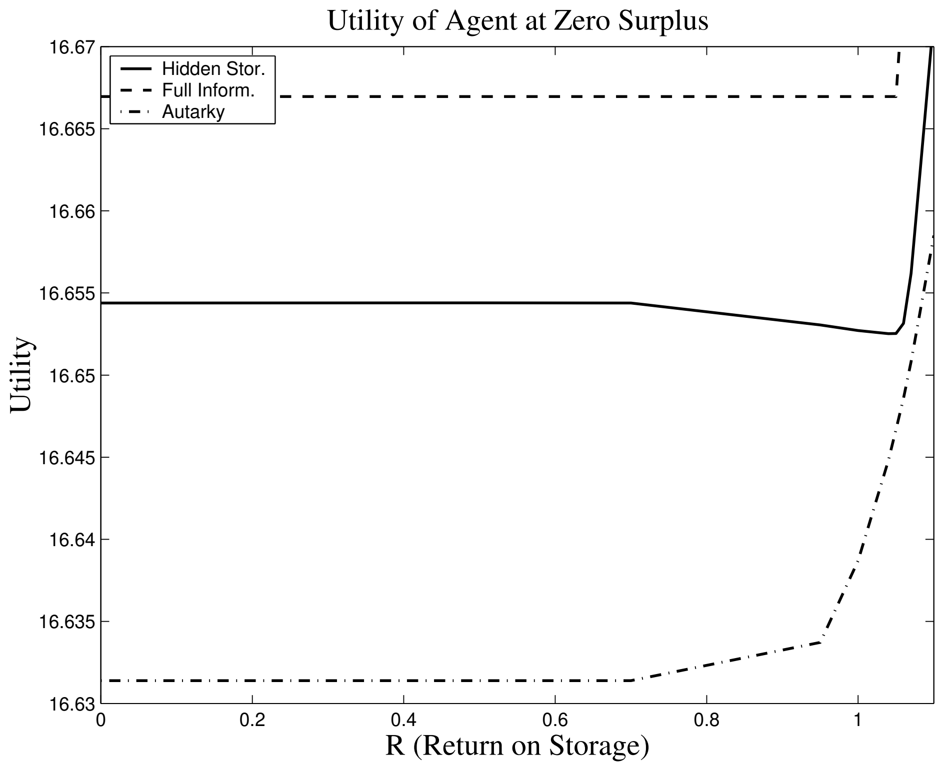

Figure 13: Utility of the Agent with Full Information, Private Information, and Autarky

7 Conclusions

In this paper we show how a dynamic mechanism design problem with hidden states and hidden actions can be formulated in a way that is suitable for computation. The main theme of the methods developed in the paper is to allow the planner to specify behavior off the equilibrium path. In Program 2 this happens because the agent reports his endowment more than once, so that the planner has to specify what happens if there are conflicting reports. In Programs 3 and 4 the planner chooses bounds that limit the utility the agent can get on certain branches off the equilibrium path. By applying these methods, we can solve problems on a standard PC that otherwise would be impossible to compute on any computer.

We use our methods to solve for the optimal credit insurance scheme in a model in which banks face uncertain returns on their investments and are subject to moral hazard. We find that in the constrained optimum low-yielding public reserves and high-yielding private investments can coexist, even though public reserves would not be used in the first-best solution. To this end, our model provides a rationale for bank reserves that relies on moral hazard only, without reference to traditional justifications for reserves like the need for liquidity.

In another application, we examine how the optimal insurance contract varies with the return of the private investment technology. Specifically, we observe that in certain cases the utility of the agent can actually decrease as the return of the private investment 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.

A Mathematical Appendix

In this Appendix we derive all Programs discussed in the paper from first principles. We start with a general setup with arbitrary message spaces, no truth-telling constraints, and strategies that are allowed to be fully history dependent. We then show that the general setup can be reduced to Program 1. The first step is to reprove the revelation principle. We thus limit ourselves to a single message space which is equal to the space of endowments, and show that the agent can be required to tell the truth. Likewise, the message space for the planner is the space of actions, and the agent can be required to follow the recommended action. In a second step we introduce continuation utilities as a state variable to limit history dependence and arrive at the recursive version of Program 1. Once Program 1 is derived, we follow the same procedure to show that Program 2 is also equivalent to the general setup (and therefore to Program 1). The remaining theorems show the equivalence of Program 3 and Program 4 to Program 1 and Program 2. We conclude that all Programs are equivalent, and none of them is more restrictive then the general communication game that we take as our starting point.

A.1 The General Setup

We start with a general communication game with arbitrary message spaces and full history-dependence.

As before, at the beginning of each period the agent realizes an endowment e. Then the agent sends a message 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. Then the agent sends a second message m_2, where m_2 is in some finite set M_2. Now the planner sends a message $r \in R$ to the agent, and R is finite as well. Finally the agent takes an action $a \in A$.

We will use h_t to denote all these choices and the realized endowment in period t :

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

We denote the space of all possible h_t by H_t. H_t is given by:

$$H_t \equiv \{(e_t, m_{1t}, m_{2t}, \tau_t, r_t, a_t) | e_t \in E, m_{1t} \in M_1, m_{2t} \in M_2, \tau_t \in T, r_t \in R, a_t \in A\}.$$

The history of the game up to time t will be denoted by h^t:

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

The set of all possible histories up to time t is denoted by H^t and is 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 of the game up to time $t-1$. On the other hand, the planner sees neither the true endowment nor the true action. We will use g_t and g^t to denote the history of the planner. We therefore have:

$$g_t \equiv \{m_{1t}, m_{2t}, \tau_t, r_t\}.$$

The planner's history of the game up to time t will be denoted by g^t:

$$g^t \equiv \{g_{-1}, g_0, g_1, \ldots, g_t\}.$$

The set G^t of all histories up to 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 $g^t(h^t)$ of the history h^t of the agent. There is no information present at the beginning of time, and consequently we define $h_{-1} \equiv g_{-1} \equiv \emptyset$.

The choices by the planner are described by a pair of outcome functions $\pi(\tau_t|m_{1t}, g^{t-1})$ and $\pi(r_t|m_{1t}, m_{2t}, \tau_t, g^{t-1})$ which map the history up to the last period and events that occurred already in the current period into a probability distribution over transfer τ_t and a report r_t. To save on notation, we will write the choice object of the planner as a joint probability distribution $\pi(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1})$ that depends only on the history and the reports by the agent. Since the transfer is assigned before the second report, we have to enforce that the transfer does not depend on the second report. Therefore the planner can only choose outcome functions that satisfy the condition:

$$\sum_R \pi(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}) = \sum_R \pi(\tau_t, r_t|m_{1t}, \bar{m}_{2t}, g^{t-1}) \tag{38}$$

for all $\tau_t, g^{t-1}, m_{1t}, m_{2t}$, and $\bar{m}_{2t}$.

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 and the endowment into a probability distribution over the first report m_{1t}, a function $\sigma(m_{2t}|e_t, m_{1t}, \tau, h^{t-1})$ which determines a probability distribution over the second report m_{2t}, and a function $\sigma(a_t|e_t, m_{1t}, m_{2t}, \tau, r_t, h^{t-1})$ which determines the action. Again, to save on notation it will be useful to join the last two elements of the strategy into a joint probability distribution $\sigma(m_{2t}, a_t|e_t, m_{1t}, \tau, r_t, h^{t-1})$. Since the second report m_{2t} cannot depend on the report r_t by the planner, strategies are chosen subject to the constraint:

$$\sum_A \sigma(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) = \sum_A \sigma(m_{2t}, a_t|e_t, m_{1t}, \tau_t, \bar{r}_t, h^{t-1}) \tag{39}$$

for all m_{2t}, e_t, m_{1t}, τ_t, h^{t-1}, r_t, and $\bar{r}_t$.

We use $p(h^t)$ to denote the probability of history h^t. This probability is a function of the outcome function and the strategy, but we use $p(h^t)$ simply as a shorthand to avoid writing out the entire sequence of underlying probabilities. In the initial period, for $h^0 = \{e_0, m_{10}, m_{20}, \tau_0, r_0, a_0\}$ we have:

$$p(h^0) = p(e_0)\sigma(m_{10}|e_0, h^{-1})\pi(\tau_0, r_0|m_{10}, m_{20}, g^{-1})\sigma(m_{20}, a_0|e_0, m_{10}, \tau_0, r_0, h^{-1}), \tag{40}$$

where $p(e_0)$ is the exogenously given probability of state e_0 in the initial period. In subsequent periods, the probability distribution over the endowment depends on the action in the previous period. Since the action in the last period is a part of the history, we can write this probability as $p(e_t|h^{t-1})$. The probability of history h^t is then given by:

$$p(h^t) = p(h^{t-1})p(e_t|h^{t-1})\sigma(m_{1t}|e_t, h^{t-1})\pi(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1})\sigma(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}). \tag{41}$$

Here h^{t-1} is the history up to time $t-1$ that is contained in history h^t. For $t \geq k$, $p(h^t|h^k)$ denotes the probability of history h^t, conditional that history h^k occurred.

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

$$U(\pi, \sigma) \equiv \sum_{t=0}^{\infty} \beta^t \left[\sum_{H^{t-1}} p(h^{t-1}) \left[\sum_E p(e|h^{t-1}) \left[\sum_{M_1, M_2, T, R, A} \sigma(m_{1t}|e_t, h^{t-1})\pi(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}) \right.\right.\right.$$
$$\left.\left.\left. \sigma(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1})u(e_t + \tau_t - a_t) \right]\right]\right]. \tag{42}$$

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 state h^k through k has been realized is given by:

$$U(\pi,\sigma|h^k) \equiv \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{h^{t-1}|h^k} p(h^{t-1}|h^k) \Bigg[\sum_{E} p(e|h^{t-1}) \Bigg[\sum_{M_1,M_2,T,R,A} \sigma(m_{1t}|e_t,h^{t-1}) \\ \pi(\tau_t,r_t|m_{1t},m_{2t},g^{t-1})\sigma(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg]\Bigg]\Bigg]. \quad (43)$$

We now want to define an *equilibrium strategy* 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 given outcome function $\pi^\star$, for any alternative strategy $\hat{\sigma}$ and any history h^k, an equilibrium strategy $\sigma^\star$ has to satisfy:

$$\forall \hat{\sigma}, h^k: \quad \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{h^{t-1}|h^k} p(h^{t-1}|h^k) \Bigg[\sum_{E} p(e|h^{t-1}) \Bigg[\sum_{M_1,M_2,T,R,A} \sigma^\star(m_{1t}|e_t,h^{t-1}) \\ \pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1})\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg]\Bigg]\Bigg] \\ \geq \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{h^{t-1}|h^k} \hat{p}(h^{t-1}|h^k) \Bigg[\sum_{E} p(e|h^{t-1}) \Bigg[\sum_{M_1,M_2,T,R,A} \hat{\sigma}(m_{1t}|e_t,h^{t-1}) \\ \pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1})\hat{\sigma}(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg]\Bigg]\Bigg], \quad (44)$$

where $\hat{p}(h^{t-1}|h^k)$ is the probability generated under $\hat{\sigma}$. Inequality (44) imposes optimization from any history h^k on. In addition, we also require that the strategy be optimal at any node that starts after an arbitrary first report in a period is made, i.e., even if in any period $k+1$ the first report was generated by a strategy $\hat{\sigma}$, it is optimal to revert to $\sigma^\star$ from the second report in period $k+1$ on. For any alternative strategy $\hat{\sigma}$ and any history

h^k, an equilibrium strategy $\sigma^\star$ therefore also has to satisfy:

$$\forall \hat{\sigma}, h^k : \quad \sum_E p(e|h^k) \Bigg[\sum_{M_1, M_2, T, R, A} \hat{\sigma}(m_{1k+1}|e_{k+1}, h^k)$$

$$\pi^\star(\tau_{k+1}, r_{k+1}|m_{1k+1}, m_{2k+1}, g^k)\sigma^\star(m_{2k+1}, a_{k+1}|e_{k+1}, m_{1k+1}, \tau_{k+1}, r_{k+1}, h^k)u(e_{k+1}+\tau_{k+1}-a_{k+1})\Bigg]$$

$$+ \sum_{t>k+1}^{\infty} \beta^t \Bigg[\sum_{h^{t-1}|h^k} p(h^{t-1}|h^k) \Bigg[\sum_E p(e|h^{t-1}) \Bigg[\sum_{M_1, M_2, T, R, A} \sigma^\star(m_{1t}|e_t, h^{t-1})$$

$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1})\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1})u(e_t + \tau_t - a_t)\Bigg]\Bigg]\Bigg]$$

$$\geq \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{h^{t-1}|h^k} \hat{p}(h^{t-1}|h^k) \Bigg[\sum_E p(e|h^{t-1}) \Bigg[\sum_{M_1, M_2, T, R, A} \hat{\sigma}(m_{1t}|e_t, h^{t-1})$$

$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1})\hat{\sigma}(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1})u(e_t + \tau_t - a_t)\Bigg]\Bigg]\Bigg]. \quad (45)$$

This condition is not restrictive, since by (44) even without this condition the agent chooses the second report optimally conditional on any first report that occurs with positive probability. The only additional effect of condition (45) is to force the agent to choose the second report and action optimally even conditional on first reports that occur with zero probability. Requiring optimal behavior off the equilibrium path will help us later in deriving recursive formulations of the general game that are open to efficient computation.

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

Definition 1 *Given an outcome function $\pi^\star$, an equilibrium strategy is a strategy such that inequalities* (44) *and* (45) *are 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 at time zero (42).

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

$$\sum_{t=0}^{\infty}\beta^t\Bigg[\sum_{H^{t-1}}p(h^{t-1})\Bigg[\sum_{E}p(e|h^{t-1})\Bigg[\sum_{M_1,M_2,T,R,A}\sigma(m_{1t}|e_t,h^{t-1})$$

$$\pi(\tau_t,r_t|m_{1t},m_{2t},g^{t-1})\sigma(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg]\Bigg]\Bigg]=W_0. \quad (46)$$

Definition 2 *A feasible contract* $\{\pi,\sigma\}$ *is an outcome function* π *together with a corresponding equilibrium strategy* σ *such that the agent realizes his initial reservation utility.*

Definition 3 *A feasible allocation is a probability distribution over endowments, transfers and actions that is implied by a feasible contract.*

The set of feasible contracts is characterized by the promise-keeping constraint (46), by the optimality conditions (44) and (45), and a number of adding-up constraints that ensure that both outcome function and strategy consist of probability measures.

The objective function of the planner is:

$$V(\pi,\sigma)\equiv\sum_{t=0}^{\infty}Q^t\Bigg[\sum_{H^{t-1}}p(h^{t-1})\Bigg[\sum_{E}p(e|h^{t-1})\Bigg[\sum_{M_1,M_2,T,R,A}\sigma(m_{1t}|e_t,h^{t-1})$$

$$\pi(\tau_t,r_t|m_{1t},m_{2t},g^{t-1})\sigma(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})(-\tau_t)\Bigg]\Bigg]\Bigg] \quad (47)$$

The planner's problem is to choose a feasible contract that maximizes (47). By construction, this will be Pareto optimal.

Definition 4 *An optimal contract is a feasible contract that solves the planner's problem.*

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

Proof: We need to show that for some W_0 the set of feasible contracts is nonempty and compact, and the objective function is continuous. To see that the set of feasible contracts 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 (46), (44), and (45) 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 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. □

A.2 Derivation of Program 1

The Revelation Principle

Now we want to use the revelation principle to reduce the setup to Program 1. The first step is to to show that without loss of generality we can restrict attention to the case were there is just one message space each for the agent and the planner. The message space of the agent is equal to the space of endowments E, and the agent will be induced to tell the truth. The message space for the planner is 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.* Our strategy is to start with a feasible allocation in the general setup, and then to derive a feasible outcome function in the restricted setup that implements the same allocation. In order to do that, we will start with a precise definition of our new restricted setup.

The planner chooses an outcome function $\pi(\tau_t, a_t|e_t, s^{t-1})$ that determines the transfer and the recommended action as a function of the endowment of the agent and the history up to the last period. In this setup, the history as known by the planner consists of all reported endowments, transfers, and recommended actions up to the last period. The contemporary part of history at time t is denoted s_t,

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

the set of all possible s_t is $S_t = E \times T \times A$, the history up through time t is s^t,

$$s^t = \{s_{-1}, s_0, \ldots, s_t\},$$

and the set of all histories up through time t is denoted S^t and is given by:

$$S^t \equiv S_{-1} \times S_0 \times S_1 \times \ldots \times S_t.$$

The planner chooses an outcome function subject to a number of constraints. 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, endowments and histories, and that:

$$\forall e_t, s^{t-1}: \quad \sum_{T,A} \pi(\tau_t, a_t|e_t, s^{t-1}) = 1. \tag{48}$$

Given an outcome function, we define probabilities $p(s^t)$ over histories and $p(e_t|s^{t-1})$ over endowments conditional on histories, based on recommended actions, in the obvious way. Given these probabilities, the outcome function has to deliver reservation utility W_0 to the agent, provided that the agent reports truthfully and takes the recommended actions:

$$\sum_{t=0}^{\infty} \beta^t \left[\sum_{S^{t-1},E,T,A} p(s^{t-1}) p(e_t|s^{t-1}) \pi(\tau_t, a_t|e_t, s^{t-1}) u(e_t + \tau_t - a_t) \right] = W_0. \tag{49}$$

Finally, it has to be optimal for the agent to tell the truth and follow the recommended action. We write a possible deviation strategy as a function $\delta_e(s^{t-1}, e_t)$ that determines the reported endowment as a function of the history and the true endowment, and a function

$\delta_a(s^{t-1}, e_t, \tau_t, a_t)$ that determines the action as a function of history, endowment, transfer, and recommended action. Since the action may be different, 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 $\hat{p}$. We require that the actions of the agent be optimal from any history s^k on. Then for every possible deviation δ_e, δ_a and any history s^k, the outcome function has to satisfy:

$$\forall \delta_e, \delta_a, s^k: \quad \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s^k)\hat{p}(e_t|s^{t-1})\pi\big(\tau_t, a_t|\delta_e(s^{t-1}, e_t), s^{t-1}\big)u\big(e_t+\tau_t-\delta_a(s^{t-1}, e_t, \tau_t, a_t)\big)\Bigg]$$
$$\leq \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} p(s^{t-1}|s^k)p(e_t|s^{t-1})\pi(\tau_t, a_t|e_t, s^{t-1})u(e_t + \tau_t - a_t)\Bigg]. \quad (50)$$

Definition 5 *An outcome function is feasible under truth telling and obedience if it satisfies the constraints* (48), (49) *and* (50) *above.*

Definition 6 *A feasible allocation in the truth-telling mechanism is a probability distribution over endowments, transfers and actions that is implied by a feasible outcome function.*

We are now ready to state the main result of this section.

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

Proof: 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 $\{\pi^\star, \sigma^\star\}$. We will now define an outcome function for the truth-telling mechanism that implements the same allocation. To complete the proof, then we will have to show that this outcome function satisfies constraints (48) to (50).

We will define the outcome function such that the allocation is the one implemented by $(\pi^\star, \sigma^\star)$ along the equilibrium path. To do that, let $H^t(s^t)$ be the set of histories in the general game that coincide in terms of the endowment, transfer, and action sequence with 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:

$$p(h^t|s^t) \equiv \frac{p(h^t)}{\sum_{h^t \in H^t(s^t)} p(h^t)} \quad (51)$$

If s^t has zero probability (that is, if the sequence s^t of endowments, transfers, and actions occurs with probability zero in the allocation implemented by $\{\pi^\star, \sigma^\star\}$), the definition of $p(h^t|s^t)$ is irrelevant, and is therefore left unspecified. Using (51), we also have the conditional probability of h^t given s^k for $t > k$:

$$p(h^t|s^k) = \sum_{h^k \in H^k(s^k)} p(h^t|h^k)p(h^k|s^k). \tag{52}$$

Now we define an outcome function for the truth-telling mechanism by:

$$\pi(\tau_t, a_t|e_t, s^{t-1}) \equiv \sum_{h^t \in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1, M_2, R} \sigma^\star(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) \tag{53}$$

Basically, the outcome function is gained by integrating out the message spaces M_1, M_2, and R.

We now have to verify that with this choice of an outcome function conditions (48) to (50) above are satisfied. In showing this, we can make use of the fact that $\{\pi^\star, \sigma^\star\}$ are probability measures and satisfy (44) and (45).

We start with the probability-measure constraint (48):

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

Using definition (53) on the left-hand side yields:

$$\sum_{T,A} \Bigg[\sum_{h^{t-1} \in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1, M_2, R} \sigma^\star(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) \Bigg] = 1$$
$$\sum_{h^{t-1} \in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1, T, M_2, R, A} \sigma^\star(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) = 1$$

by a change in the order of summation, and since we sum over probability measures $\sigma^\star$ and π, we get:

$$\sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) = 1 \text{ or:}$$

$$1 = 1,$$

because of definition (51), thus (48) is satisfied.

The promise-keeping constraint (49) is next:

$$\sum_{t=0}^{\infty}\beta^t\Bigg[\sum_{S^{t-1},E,T,A} p(s^{t-1})p(e_t|s^{t-1})\pi(\tau_t,a_t|e_t,s^{t-1})u(e_t+\tau_t-a_t)\Bigg] = W_0.$$

Using definition (53) in the left-hand side yields:

$$\sum_{t=0}^{\infty}\beta^t\Bigg[\sum_{S^{t-1},E,T,A} p(s^{t-1})p(e_t|s^{t-1})\sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1})\sum_{M_1,M_2,R}\sigma^\star(m_{1t}|e_t,h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg] = W_0$$

$$\sum_{t=0}^{\infty}\beta^t\Bigg[\sum_{S^{t-1}}\sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(s^{t-1})p(h^{t-1}|s^{t-1})\sum_{E} p(e_t|h^{t-1})\sum_{M_1,M_2,T,R,A}\sigma^\star(m_{1t}|e_t,h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg] = W_0$$

by a change in the order of summation and since $p(e_t|h^{t-1}) = p(e_t|s^{t-1})$ if $h^t \in H^t(s^t)$,

$$\sum_{t=0}^{\infty}\beta^t\Bigg[\sum_{H^{t-1}} p(h^{t-1})\sum_{E} p(e_t|h^{t-1})\sum_{M_1,M_2,T,R,A}\sigma^\star(m_{1t}|e_t,h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg] = W_0,$$

where the last step used definition (51). The last equation is (46) and therefore satisfied. Thus (49) is satisfied.

The last step is to show that (50) is satisfied for any δ_e, δ_a, and s^k:

$$\sum_{t>k}^{\infty}\beta^t\Big[\sum_{S^{t-1},E,T,A}\hat{p}(s^{t-1}|s^k)\hat{p}(e_t|s^{t-1})\pi\big(\tau_t,a_t|\delta_e(s^{t-1},e_t),s^{t-1}\big)u\big(e_t+\tau_t-\delta_a(s^{t-1},e_t,\tau_t,a_t)\big)\Big]$$
$$\leq \sum_{t>k}^{\infty}\beta^t\Big[\sum_{S^{t-1},E,T,A}p(s^{t-1}|s^k)p(e_t|s^{t-1})\pi(\tau_t,a_t|e_t,s^{t-1})u(e_t+\tau_t-a_t)\Big].$$

Using definition (53) on both sides gives:

$$\sum_{t>k}^{\infty}\beta^t\Big[\sum_{S^{t-1},E,T,A}\hat{p}(s^{t-1}|s^k)\hat{p}(e_t|s^{t-1})\sum_{h^{t-1}\in H^{t-1}(s^{t-1})}p(h^{t-1}|s^{t-1})\sum_{M_1,M_2,R}\sigma^\star(m_{1t}|\delta_e(s^{t-1},e_t),h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|\delta_e(s^{t-1},e_t),m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-\delta_a(s^{t-1},e_t,\tau_t,a_t))\Big]$$
$$\leq \sum_{t>k}^{\infty}\beta^t\Big[\sum_{S^{t-1},E,T,A}p(s^{t-1}|s^k)p(e_t|s^{t-1})\sum_{h^{t-1}\in H^{t-1}(s^{t-1})}p(h^{t-1}|s^{t-1})\sum_{M_1,M_2,R}\sigma^\star(m_{1t}|e_t,h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Big]$$

The deviation of the agent defines a deviation strategy $\hat{\sigma}$, given by:

$$\hat{\sigma}(m_{1t}|e_t,h^{t-1})\equiv\sigma^\star(m_{1t}|\delta_e(s^{t-1},e_t),h^{t-1})$$

and:

$$\hat{\sigma}(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})\equiv\sigma^\star(m_{2t},a_t|\delta_e(s^{t-1},e_t),m_{1t},\tau_t,r_t,h^{t-1}).$$

Changing the notation for the agent's strategy in this way we get:

$$\sum_{t>k}^{\infty}\beta^t\Big[\sum_{S^{t-1},E,T,A}\hat{p}(s^{t-1}|s^k)\hat{p}(e_t|s^{t-1})\sum_{h^{t-1}\in H^{t-1}(s^{t-1})}p(h^{t-1}|s^{t-1})\sum_{M_1,M_2,R}\hat{\sigma}(m_{1t}|e_t,h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\hat{\sigma}(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Big]$$
$$\leq \sum_{t>k}^{\infty}\beta^t\Big[\sum_{S^{t-1},E,T,A}p(s^{t-1}|s^k)p(e_t|s^{t-1})\sum_{h^{t-1}\in H^{t-1}(s^{t-1})}p(h^{t-1}|s^{t-1})\sum_{M_1,M_2,R}\sigma^\star(m_{1t}|e_t,h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Big]\quad(54)$$

By a change in the order of summation and since $p(e_t|h^{t-1}) = p(e_t|s^{t-1})$ if $h^t \in H^t(s^t)$, this is:

$$\sum_{t>k}^{\infty} \beta^t \Big[\sum_{S^{t-1}} \sum_{h^{t-1} \in H^{t-1}(s^{t-1})} \hat{p}(s^{t-1}|s^k) p(h^{t-1}|s^{t-1}) \sum_{E} \hat{p}(e_t|h^{t-1}) \sum_{M_1,M_2,T,R,A} \hat{\sigma}(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\hat{\sigma}(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) u(e_t + \tau_t - a_t) \Big]$$
$$\leq \sum_{t>k}^{\infty} \beta^t \Big[\sum_{S^{t-1}} \sum_{h^{t-1} \in H^{t-1}(s^{t-1})} p(s^{t-1}|s^k) p(h^{t-1}|s^{t-1}) \sum_{E} p(e_t|h^{t-1}) \sum_{M_1,M_2,T,R,A} \sigma^\star(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) u(e_t + \tau_t - a_t) \Big].$$

This in turn can also be written as:

$$\sum_{t>k}^{\infty} \beta^t \Big[\sum_{H^{t-1}} \hat{p}(h^{t-1}|s^k) \sum_{E} \hat{p}(e_t|h^{t-1}) \sum_{M_1,M_2,T,R,A} \hat{\sigma}(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\hat{\sigma}(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) u(e_t + \tau_t - a_t) \Big]$$
$$\leq \sum_{t>k}^{\infty} \beta^t \Big[\sum_{H^{t-1}} p(h^{t-1}|s^k) \sum_{E} p(e_t|h^{t-1}) \sum_{M_1,M_2,T,R,A} \sigma^\star(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) u(e_t + \tau_t - a_t) \Big].$$

In the last equation probabilities are still conditioned on s^k. We can use (52) to condition probabilities on h^k instead:

$$\sum_{H^k(s^k)} p(h^k|s^k) \Big[\sum_{t>k}^{\infty} \beta^t \Big[\sum_{H^{t-1}} \hat{p}(h^{t-1}|h^k) \sum_{E} \hat{p}(e_t|h^{t-1}) \sum_{M_1,M_2,T,R,A} \hat{\sigma}(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\hat{\sigma}(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) u(e_t + \tau_t - a_t) \Big] \Big]$$
$$\leq \sum_{H^k(s^k)} p(h^k|s^k) \Big[\sum_{t>k}^{\infty} \beta^t \Big[\sum_{H^{t-1}} p(h^{t-1}|h^k) \sum_{E} p(e_t|h^{t-1}) \sum_{M_1,M_2,T,R,A} \sigma^\star(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1}) u(e_t + \tau_t - a_t) \Big] \Big].$$

This inequality can be obtained by summing over both sides of (44) and therefore holds, so that (50) is satisfied. Thus the specified contract (53) satisfied all conditions for a feasible contract, which completes the proof. □

The Recursive Representation

We now have a version of our general game that imposes 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 current game to a recursive version with a vector of promised utilities as the state variable. The result will be Program 1, which stands at the beginning of our discussion in the paper.

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 the endowment e_k. It will be useful to consider a problem in which the planner has to deliver a vector of reservation utilities $\mathbf{w}_0$, depending on the endowment in the initial period. The original planners problem can then be cast as choosing the vector of initial utility assignments $\mathbf{w}_0$ which yields the highest expected surplus for the planner, given the initial probability distribution over states $e \in E$.

As before, the choice object of the planner is given by $\pi(\tau_t, a_t|e_t, s^{t-1})$. However, we picture the planner as picking the contract conditional on the initial report e_0 by the agent. The constraints are the same as before, namely (48) to (50), but we will write them down separately for each initial report e_0. The first set of constraints are the probability-measure constraints (48). For all e_t, s^{t-1}, we require:

$$\forall e_t, s^{t-1}: \quad \sum_{T,A} \pi(\tau_t, a_t|e_t, s^{t-1}) = 1. \tag{55}$$

There is now a promise-keeping constraint like (49) below for each possible initial endowment. For all e_0, we require:

$$\begin{aligned}\forall e_0: \quad & \sum_{T_0,A_0} \pi(\tau_0, a_0|e_0, s^{-1}) \Big[u(e_0 + \tau_0 - a_0) \\ & + \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1},E,T,A} p(s^{t-1}|s_0) p(e_t|s^{t-1}) \pi(\tau_t, a_t|e_t, s^{t-1}) u(e_t + \tau_t - a_t) \Big] = w_0(e_0). \end{aligned} \tag{56}$$

Next, akin to (50), we need to guarantee that the agent is obedient, assuming that the initial report e_0 was correct. A deviation consists of an alternative reporting strategy

$\delta_e(e_t, s^{t-1})$ *from tomorrow on* and an alternative action strategy $\delta_a(s^t)$. We use $\delta(s^t)$ to denote the history that was reported under the deviation if the true history is s^t, and $\hat{p}(\cdot)$ are the adjusted probabilities of states s^t and endowments e_t. For all e_0 and all deviations, we require:

$$\forall e_0, \delta_e, \delta_a : \quad \sum_{T_0, A_0} \pi(\tau_0, a_0|e_0, s^{-1}) \Big[u(e_0 + \tau_0 - \delta_a(s^0))$$

$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1}, E, T, A} \hat{p}(s^{t-1}|s_0) \hat{p}(e_t|s^{t-1}) \pi(\tau_t, a_t|\delta_e(e_t, s^{t-1}), \delta(s^{t-1})) u(e_t + \tau_t - \delta_a(s^t)) \Big] \leq w_0(e_0). \tag{57}$$

Finally, an agent needs an incentive to tell the truth already in the initial period. For all reported e_0, all actual $\hat{e}_0 \neq e_0$, and all deviations, we require:

$$\forall e_0, \hat{e}_0 \neq e_0, \delta_e, \delta_a : \quad \sum_{T_0, A_0} \pi(\tau_0, a_0|e_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - \delta_a(s^0))$$

$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1}, E, T, A} \hat{p}(s^{t-1}|s_0) \hat{p}(e_t|s^{t-1}) \pi(\tau_t, a_t|\delta_e(e_t, \delta(s^{t-1})), \delta(s^{t-1})) u(e_t + \tau_t - \delta_a(s^t)) \Big] \leq w_0(\hat{e}_0). \tag{58}$$

Constraints (55) to (58) are merely a different of way of writing down our original constraints (48) to (50), in a way that will be useful for our recursive representation. Since the vector of promised utilities $\mathbf{w}$ will serve as our state variable, it will be important to show that the set of all feasible utility vectors has nice properties.

Definition 7 *The set* $\mathbf{W}$ *is given by all vectors* $\mathbf{w} \in R^{\#E}$ *that satisfy constraints* (55) *to* (58) *for some outcome function* $\pi(\tau_t, a_t|e_t, s^{t-1})$.

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

Proof: (Outline) 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 the endowment, 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. Finally, we can show that $\mathbf{W}$ is closed by a contradiction argument. Assume that $\mathbf{W}$ is not closed. Then there exists a converging sequence $\mathbf{w}_n$ such that each element of the sequence is in $\mathbf{W}$, but its limit $\mathbf{w}$ is not. Corresponding to each $\mathbf{w}_n$ there is a contract $\pi(\tau_t, a_t|e_t, s^{t-1})_n$ satisfying constraints (55) to (58). 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 is easy to show that $\mathbf{w}$ must satisfy (55) to (58) when $\pi(\tau_t, a_t|e_t, s^{t-1})$ is the chosen contract. $\square$

Now we consider the problem of a planner who has promised utility vector $\mathbf{w} \in \mathbf{W}$ after the agent has delivered report e_0.

Problem P:

$$V(\mathbf{w}_0, e_0) = \max_{\pi} \sum_{T_0,A_0} \pi(\tau_0, a_0|e_0, s^{-1}) \Big[-\tau_0 + \sum_{t=1}^{\infty} Q^t \sum_{S^{t-1},E,T,A} p(s^{t-1}|s_0) p(e_t|s^{t-1}) \pi(\tau_t, a_t|e_t, s^{t-1})(-\tau_t) \Big]$$

subject to constraints (55) to (58) above, for a given $e_0 \in E$ and $\mathbf{w} \in \mathbf{W}$. We want to show that this problem has a recursive structure.

To do this, we need to define on-path future utilities. For all s^{k-1}, e^k, let:

$$\begin{aligned} w(e_k, s^{k-1}) = & \sum_{T_k,A_k} \pi(\tau_k, a_k|e_k, s^{k-1}) \Big[u(e_k + \tau_k - a_k) \\ & + \sum_{t=k+1}^{\infty} \beta^{t-k} \sum_{S^{t-1},E,T,A} p(s^{t-1}|e_k, s^{k-1})) p(e_t|s^{t-1}) \pi(\tau_t, a_t|e_t, s^{t-1}) u(e_t + \tau_t - a_t) \Big], \end{aligned} \quad (59)$$

and let $\mathbf{w}(s^{k-1})$ be the vector of these utilities over all e_k. We can now show the following result:

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 is an optimal contract for Problem P with $e_0 = e_k$ and $\mathbf{w}_0 = \mathbf{w}(s^{k-1})$.*

Proof: The proof is by construction. We have shown earlier that an optimal contract exists. Let $\tilde{\pi}$ be an optimal contract from time zero, and π_k an optimal contract for $e_0 = e_k$

and $\mathbf{w}_0 = \mathbf{w}(s^{k-1})$, with the elements of vector $\mathbf{w}(s^{k-1})$ defined in (59). Now construct a new contract $\pi^\star$ that is equal to π_k from (e_k, s^{k-1}) on, and equals $\tilde{\pi}$ until time k and on all future branches other than e_k, s^{k-1}. We claim that $\pi^\star$ is also an optimal contract. To show this, we have to demonstrate that $\pi^\star$ satisfies constraints (55) to (58), 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 $\mathbf{w}_0 = \mathbf{w}(s^{k-1})$ as the starting point also constrain the choices of the planner in the original program from (e_k, s^{k-1}) on. 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 $\tilde{\pi}$. Since $\tilde{\pi}$ is assumed to be an optimal contract, if $\pi^\star$ satisfies (55) to (58), it must be optimal as well. Thus we only have to show that (55) to (58) 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 (55) 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 (56) is satisfied since the new contract delivers the same on-path utilities by construction. We still have to show that the incentive constraints (57) and (58) are satisfied. We will do this by contradiction. Suppose that (57) is not satisfied by contract $\pi^\star$. Then there is a deviation $\delta_e(e_t, s^{t-1})$, $\delta_a(s^t)$ such that:

$$\sum_{T_0,A_0} \pi^\star(\tau_0, a_0|e_0, s^{-1}) \Big[u(e_0 + \tau_0 - \delta_a(s^0))$$
$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s_0)\hat{p}(e_t|s^{t-1})\pi^\star(\tau_t, a_t|\delta_e(e_t, s^{t-1}), \delta(s^{t-1}))u(e_t + \tau_t - \delta_a(s^t)) \Big] > w_0(e_0). \tag{60}$$

Now use $w(e_k, s^{k-1}, \delta)$ to denote the continuation utility of the agent from time k on under the deviation strategy. Now we can rewrite (60) as:

$$\sum_{T_0,A_0} \pi^\star(\tau_0, a_0|e_0, s^{-1}) \Big[u(e_0 + \tau_0 - \delta_a(s^0))$$
$$+ \sum_{t=1}^{k-1} \beta^t \sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s_0)\hat{p}(e_t|s^{t-1})\pi^\star(\tau_t, a_t|\delta_e(e_t, s^{t-1}), \delta(s^{t-1}))u(e_t + \tau_t - \delta_a(s^t))$$
$$+ \beta^k \sum_{E,s^{k-1}} \hat{p}(s^{k-1}|s_0)\hat{p}(e_k|s^{k-1})w(e_k, s^{k-1}, \delta) \Big] > w_0(e_0). \tag{61}$$

Notice that for s^{k-1} that are reached with positive probability under the deviation we have $w(e_k, s^{k-1}, \delta) \leq w(e_k, \delta(s^{k-1}))$, where $\delta(s^k)$ is the history as seen by the planner (reported endowments, delivered transfers, and recommended actions) under the deviation strategy. Otherwise, either $\tilde{\pi}$ or π_k would violate incentive constraints. Therefore (61) implies:

$$\sum_{T_0,A_0} \pi^\star(\tau_0, a_0|e_0, s^{-1}) \Big[u(e_0 + \tau_0 - \delta_a(s^0))$$
$$+ \sum_{t=1}^{k-1} \beta^t \sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s_0)\hat{p}(e_t|s^{t-1})\pi^\star(\tau_t, a_t|\delta_e(e_t, s^{t-1}), \delta(s^{t-1}))u(e_t + \tau_t - \delta_a(s^t))$$
$$+ \beta^k \sum_{E,s^{k-1}} \hat{p}(s^{k-1}|s_0)\hat{p}(e_k|s^{k-1})w(e_k, \delta(s^{k-1}))\Big] > w_0(e_0). \quad (62)$$

But now the left-hand side of (62) is the utility that the agent gets under plan $\tilde{\pi}$ from following the deviation strategy until time k, and following the recommendations of the planner afterwards. Thus (62) contradicts the incentive compatibility of $\tilde{\pi}$. We obtain a contradiction, $\pi^\star$ actually satisfies (57). The proof for constraint (58) follows the same lines. This shows that plan $\pi^\star$ is within the constraints of the original problem. Since $\pi^\star$ yields at least as much surplus as $\tilde{\pi}$ and $\tilde{\pi}$ is an optimal contract, $\pi^\star$ must be optimal as well. □

Given this result, we know that the maximized surplus of the planner can be written as:

$$V(\mathbf{w}_0, e_0) = \sum_{A,T} \pi^\star(\tau_0, a_0|e_0, s^{-1}) \left[-\tau_0 + Q \sum_{E} p(e_1|s^0)V(e_1, \mathbf{w}(s^0)) \right]. \quad (63)$$

In light of (63), we can cast the planner's problem as choosing transfers and actions in the present period, and choosing continuation utilities from the set $\mathbf{W}$ from tomorrow on conditioned on history $s^0 = \{e_0, \tau_0, a_0\}$. We will write the choices of the planner as a function of the vector of promised utilities $\mathbf{w}$ that has to be delivered in the current period, and the current state e. The choices of the planner are therefore functions $\pi(\tau, a|\mathbf{w}, e)$ and $\Pi(\mathbf{w}'|\mathbf{w}, e, \tau, a)$, where Π is a probability measure over set $\mathbf{W}$, and $\mathbf{w}'$ are the promised utilities from tomorrow on. We still need to determine which constraints need to be placed on these choices in order to guarantee that the implied contract satisfies (55) to (58). In order to reproduce (55), we need to impose:

$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) = 1. \tag{64}$$

The promise-keeping constraint (56) will be satisfied if we impose:

$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) \left[u(e + \tau - a) + \beta \int_{\mathbf{W}} [\sum_{E} p(e'|a) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \tau, a) \right] = w(e). \tag{65}$$

The incentive constraints are more subtle. We first require that the agent cannot gain by following another action strategy $\delta_a(\tau, a)$, assuming that the reported endowment was correct:

$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) \left[u(e + \tau - \delta_a(\tau, a)) + \beta \int_{\mathbf{W}} [\sum_{E} p(e'|\delta_a(\tau, a)) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \tau, a) \right] \leq w(e). \tag{66}$$

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:

$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) \left[u(\hat{e} + \tau - \delta_a(\tau, a)) + \beta \int_{\mathbf{W}} [\sum_{E} p(e'|\delta_a(\tau, a)) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \tau, a) \right] \leq w(\hat{e}). \tag{67}$$

The constraints above 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. The constraints therefore imply that (57) and (58) hold for one-time deviations. We still have to show that (66) and (67) are sufficient to prevent deviations in multiple periods. 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 (66) and (67) 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-time 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 problem of the planner as solving:

$$V(\mathbf{w},e)=\max_{\pi\geq 0}\sum_{A,T}\pi(\tau,a|\mathbf{w},e)\left[-\tau+\beta\int_{\mathbf{W}}\sum_{E}p(e'|a)V(\mathbf{w}',e')\Pi(\mathbf{w}'|\mathbf{w},e,\tau,a)\right] \tag{68}$$

by choice of π and Π, subject to constraints (64) to (67) above. Program 1 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 (68) defines an operator T that maps functions $V(\mathbf{w},e)$ into $TV(\mathbf{w},e)$. It is easy to show that T maps bounded continuous functions into bounded continuos 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.

Computing the Value Set

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, 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 8 $\mathbf{w}\in B(\mathbf{W}')$ *if for all* $e\in E$ *there exists a measure* π *and a probability measure* Π *on* $\mathbf{W}'$ *such that the following equations hold:*

$$\sum_{T,A}\pi(\tau,a|\mathbf{w},e)=1, \tag{69}$$

$$\sum_{T,A}\pi(\tau,a|\mathbf{w},e)\left[u(e+\tau-a)+\beta\int_{\mathbf{W}'}[\sum_{E}p(e'|a)w'(e')]\Pi(\mathbf{w}'|\mathbf{w},e,\tau,a)\right]=w(e), \tag{70}$$

$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) \left[u(e + \tau - \delta_a(\tau, a)) + \beta \int_{\mathbf{W}'} [\sum_E p(e'|\delta_a(\tau, a)) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \tau, a) \right] \leq w(e), \tag{71}$$

and:

$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e) \left[u(\hat{e} + \tau - \delta_a(\tau, a)) + \beta \int_{\mathbf{W}'} [\sum_E p(e'|\delta_a(\tau, a)) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \tau, a) \right] \leq w(\hat{e}). \tag{72}$$

Notice that constraints (69) to (72) are almost identical to (64) to (67), the only difference being that integration occurs over $\mathbf{W}'$ instead of $\mathbf{W}$. Intuitively, the set $B(\mathbf{W}')$ consists of all utility vectors $\mathbf{w}$ that are feasible today (observing all incentive constraints), given that utility vectors from tomorrow on 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 5

- *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 $\bar{\mathbf{W}}$.*

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 on. 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.

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

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 (55) to (58) 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 (64) to (67). 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 (55) to (58), which we have done already earlier in the discussion of Proposition 4.

A.3 Derivation of Program 2

The Revelation Principle Once Again

We now will follow the same steps as above to derive Program 2. In Program 2, the agent gets to report the endowment twice, once before the transfer is received, and another time before the planner recommends an action. While on the equilibrium path the agent will tell the truth twice, having two reports allows us to specify off-path behavior, which in turn simplifies computations.

We start by applying the revelation principle. We define a restricted version of the game in which there are two messages from the agent and a recommendation by the planner. The agent will be required both to tell the truth and to be obedient. Both message spaces of the agent are equal to the space of endowments E. The message space for the planner is equal to the space of actions A, and the agent will be required to follow the recommended action. The point again is that requiring or inducing honesty and obedience can be done without loss of generality.

The planner chooses an outcome function at t $\pi(\tau_t, a_t|e_{1t}, e_{2t}, s^{t-1})$ that determines the transfer and the action as a function of the two reported endowments by the agent and the history up to period $t-1$. Notice that since this function is also specified for the case that the two reports differ, the planner in effect specifies behavior off the equilibrium path. The history as known by the planner consists of all reported endowments, transfers, and recommended actions up to the last period. As before, the contemporary history at time t is denoted s_t,

$$s_t \equiv \{e_{1t}, e_{2t}, \tau_t, a_t\},$$

the set of all possible s_t is S_t, the history up through and including time t is s^t,

$$s^t = \{s_{-1}, s_0, \ldots, s_t\},$$

and the set of all histories up through time t is denoted S^t and is given by:

$$S^t \equiv S_{-1} \times S_0 \times S_1 \times \ldots \times S_t.$$

The planner chooses an outcome function subject to a number of constraints. 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, endowments and histories, and that:

$$\forall e_{1t}, e_{2t}, s^{t-1}: \quad \sum_{T,A} \pi(\tau_t, a_t|e_{1t}, e_{2t}, s^{t-1}) = 1. \tag{73}$$

Next, the transfer τ cannot depend on the second report e_{2t}, since the second report is made only after the agent receives the transfer. This implies the following condition:

$$\forall \tau_t, e_{1t}, e_{2t}, \bar{e}_{2t}, s^{t-1}: \quad \sum_A \pi(\tau_t, a_t | e_{1t}, e_{2t}, s^{t-1}) = \sum_A \pi(\tau_t, a_t | e_{1t}, \bar{e}_{2t}, s^{t-1}). \tag{74}$$

Given an outcome function, probabilities $p(s^t)$ and $p(e_t|s^{t-1})$ of histories and endowments are defined in the obvious way. Given these probabilities, the outcome function has to deliver reservation utility W_0 to the agent, provided that the agent reports truthfully twice and takes the recommended actions:

$$\sum_{t=0}^{\infty} \beta^t \left[\sum_{S^{t-1},E,T,A} p(s^{t-1}) p(e_t|s^{t-1}) \pi(\tau_t, a_t | e_t, e_t, s^{t-1}) u(e_t + \tau_t - a_t) \right] = W_0. \tag{75}$$

Also, it has to be optimal for the agent to tell the truth each time and follow the recommended action. We will write a possible deviation strategy as a function $\delta_{e1}(s^{t-1}, e_t)$ that determines the first reported endowment as a function of the history and and the true endowment, a function $\delta_{e2}(s^{t-1}, e_t)$ that determines the second reported endowment as a function of the history s^{t-1} and and the true endowment e_t, and a function $\delta_a(s^{t-1}, e_t, \tau_t, a_t)$ that determines the action a as a function of history s^{t-1}, endowment e_t, transfer τ_t, and recommended action a_t. Since the action may be different, 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 $\hat{p}$.

We require that truth-telling and obedience be optimal for the agent, starting in any period and from any history. In addition, we also require that the prescribed actions be optimal for the agent at the node in the middle of a period, i.e., after the transfer has been assigned, regardless of what the first report was. Thus even if the agent misreports the endowment at the beginning of the period, it has to be optimal to report truthfully at the second report. As we will see later on, this requirement leads to major simplifications in the computation of the optimal outcome function.

The first requirement is that the prescribed reports and actions be optimal, given any history h^k that has been already realized, inducing honesty and obedience in the future:

$$\forall h^k, \delta_{e1}, \delta_{e2}, \delta_a : \quad \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s^k)\hat{p}(e_t|s^{t-1})$$
$$\pi(\tau_t, a_t|\delta_{e1}(s^{t-1}, e_t), \delta_{e2}(s^{t-1}, e_t), s^{t-1})u(e_t + \tau_t - \delta_a(s^t))\Bigg]$$
$$\leq \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} p(s^{t-1}|s^k)p(e_t|s^{t-1})\pi(\tau_t, a_t|e_t, e_t, s^{t-1})u(e_t + \tau_t - a_t). \quad (76)$$

In fact, we place an even more strict restriction on the outcome function. Given history h^k, we require that even if the first report in period $k+1$ was wrong, it is optimal for the agent to tell the truth at the second report in $k+1$ and follow the recommended action, and to be honest and obedient in the future. This implies the following condition:

$$\forall h^k, \delta_{e1}, \delta_{e2}, \delta_a : \quad \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1})\hat{p}(e_t|s^{t-1})$$
$$\pi\big((\tau_t, a_t|\delta_{e1}(s^{t-1}, e_t), \delta_{e2}(s^{t-1}, e_t), s^{t-1})\big)u\big(e_t + \tau_t - \delta_a(s^t)\big)\Bigg]$$
$$\leq \beta^{k+1}\Bigg[\sum_{E,T,A} p(e_t|s^k)\pi\big(\tau_{k+1}, a_{k+1}|\delta_{e1}(s^k, e_{k+1}), e_{k+1}, s^k\big)u(e_{k+1} + \tau_{k+1} - a_{k+1})\Bigg]$$
$$+ \sum_{t>k+1}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} p(s^{t-1}|s^k)p(e_t|s^{t-1})\pi(\tau_t, a_t|e_t, e_t, s^{t-1})u(e_t + \tau_t - a_t)\Bigg] \quad (77)$$

Definition 9 *An outcome function is feasible under truth telling and obedience if it satisfies the constraints* (73), (74), (75), (76), *and* (77) *above.*

Definition 10 *A feasible allocation is a probability distribution over endowments, transfers and actions that is implied by a feasible outcome function.*

We are now going to show that this setup does not restrict the set of feasible allocations.

Proposition 7 *For any message spaces M_1, M_2, and R, any allocation that is feasible in the general mechanism is also feasible in the truth-telling mechanism.*

Proof: Corresponding to any feasible allocation in the general setup there is a feasible contract that implements this allocation. Fix this contract $(\pi^\star, \sigma^\star)$. We will now define an outcome function for the truth-telling mechanism that implements the same allocation. To complete the proof, we then will have to show that this outcome function satisfies constraints (73) to (77).

We will define the outcome function such that the allocation is the one implemented by $(\pi^\star, \sigma^\star)$ along the equilibrium path. To do that, let $H^t(s^t)$ be the set of histories in the general game that coincide in terms of the transfer and action sequence with history s^t in the restricted game, and such that the endowment sequence in h^t matches the sequence of second reports in h^t. At this point, we restrict attention to those s^t in which the agent makes the same report twice. Define $p(h^t|s^t)$ is the probability of history h^t conditional on s^t, i.e:

$$p(h^t|s^t) \equiv \frac{p(h^t)}{\sum_{h^t \in H^t(s^t)} p(h^t)} \tag{78}$$

If s^t occurs with zero probability under $(\pi^\star, \sigma^\star)$, we set $p(h^t|s^t)$ to zero. We still need to define probabilities for the case when the agent makes differing reports with a period. The probabilities are chosen to resemble the case in which the agent was mistaken about the endowment when making the first report, but correct at the second report:

$$\begin{aligned} p(h^t|s^t) \equiv\ & p(h^{t-1}|s^{t-1}) p(e_{2t}|h^{t-1}) \\ & \sum_{M_1, M_2, R} \sigma(m_{1t}|e_{1t}, h^{t-1}) \pi(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}) \sigma(m_{2t}, a_t|e_{2t}, m_{1t}, \tau_t, r_t, h^{t-1}). \end{aligned} \tag{79}$$

Now we define an outcome function for the truth-telling mechanism by:

$$\begin{aligned} \pi(\tau_t, a_t|e_{1t}, e_{2t}, s^{t-1}) \equiv & \sum_{h^{t-1} \in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1, M_2, R} \sigma^\star(m_{1t}|e_{1t}, h^{t-1}) \\ & \pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1})) \sigma^\star(m_{2t}, a_t|e_{2t}, m_{1t}, \tau_t, r_t, h^{t-1}). \end{aligned} \tag{80}$$

We now have to verify that this outcome function satisfies the conditions (73) to (77) above. Using definition (80) in (73) yields:

$$
\begin{aligned}
\sum_{T,A} \pi(\tau_t, a_t|e_{1t}, e_{2t}, s^{t-1}) =& \sum_{T,A} \Bigg[\sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \sigma^\star(m_{1t}|e_{1t}, h^{t-1}) \\
& \pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_{2t}, m_{1t}, \tau_t, r_t, h^{t-1}) \Bigg] \\
=& \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,T,M_2,R,A} \sigma^\star(m_{1t}|e_{1t}, h^{t-1}) \\
& \pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_{2t}, m_{1t}, \tau_t, r_t, h^{t-1})
\end{aligned}
$$

by a change in the order of summation, and since we sum over probability measures, we get:

$$
\begin{aligned}
=& \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \\
=& 1,
\end{aligned}
$$

thus (73) is satisfied. Using definition (80) in (74) gives:

$$
\begin{aligned}
& \sum_{A} \Bigg[\sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \sigma^\star(m_{1t}|e_{1t}, h^{t-1}) \\
& \quad \pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_{2t}, m_{1t}, \tau_t, r_t, h^{t-1}) \Bigg] \\
& \quad = \sum_{A} \Bigg[\sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \sigma^\star(m_{1t}|e_{1t}, h^{t-1}) \\
& \qquad \pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|\bar{e}_{2t}, m_{1t}, \tau_t, r_t, h^{t-1}) \Bigg].
\end{aligned}
$$

Using (38) and changing the order of summation, this is:

$$\sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \sigma^\star(m_{1t}|e_{1t},h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1})) \sum_{M_2,A} \Big[\sigma^\star(m_{2t},a_t|e_{2t},m_{1t},\tau_t,r_t,h^{t-1})\Big]$$
$$= \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,R} \sigma^\star(m_{1t}|e_{1t},h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1})) \sum_{M_2,A} \Big[\sigma^\star(m_{2t},a_t|\bar{e}_{2t},m_{1t},\tau_t,r_t,h^{t-1})\Big],$$

or:

$$\sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \sigma^\star(m_{1t}|e_{1t},h^{t-1})\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))$$
$$= \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,R} \sigma^\star(m_{1t}|e_{1t},h^{t-1})\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1})).$$

Both sides are identical now, thus (74) is satisfied. Using definition (80) in (75) yields:

$$\sum_{t=0}^{\infty}\beta^t\Big[\sum_{S^{t-1},E,T,A} p(s^{t-1})p(e_t|s^{t-1})\pi(\tau_t,a_t|e_t,e_t,s^{t-1})u(e_t+\tau_t-a_t)\Big]$$
$$= \sum_{t=0}^{\infty}\beta^t\Big[\sum_{S^{t-1},E,T,A} p(s^{t-1})p(e_t|s^{t-1}) \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \sigma^\star(m_{1t}|e_t,h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Big]$$
$$= \sum_{t=0}^{\infty}\beta^t\Big[\sum_{S^{t-1}} \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(s^{t-1})p(h^{t-1}|s^{t-1}) \sum_{E} p(e_t|h^{t-1}) \sum_{M_1,M_2,T,R,A} \sigma^\star(m_{1t}|e_t,h^{t-1})$$
$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Big]$$

by a change in the order of summation and since $p(e_t|h^{t-1}) = p(e_t|s^{t-1})$ if $h^t \in H^t(s^t)$,

$$= \sum_{t=0}^{\infty} \beta^t \Bigg[\sum_{H^{t-1}} p(h^{t-1}) \sum_{E} p(e_t|h^{t-1}) \sum_{M_1,M_2,T,R,A} \sigma^\star(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1})u(e_t + \tau_t - a_t)\Bigg]$$

by using definition (51),

$$= W_0$$

by (46). Thus (75) is satisfied. Next, using definition (80) on both sides of (76) gives:

$$\sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s^k)\hat{p}(e_t|s^{t-1}) \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \sigma^\star(m_{1t}|\delta_{e1}(s^{t-1}, e_t), h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|\delta_{e2}(s^{t-1}, e_t), m_{1t}, \tau_t, r_t, h^{t-1})u(e_t + \tau_t - \delta_a(s^t))\Bigg]$$
$$\leq \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} p(s^{t-1}|s^k)p(e_t|s^{t-1}) \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \sigma^\star(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1})u(e_t + \tau_t - a_t)\Bigg]$$

By a change in notation for the agent's strategy this can be written as:

$$\sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s^k)\hat{p}(e_t|s^{t-1}) \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \hat{\sigma}(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\hat{\sigma}(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1})u(e_t + \tau_t - a_t)\Bigg]$$
$$\leq \sum_{t>k}^{\infty} \beta^t \Bigg[\sum_{S^{t-1},E,T,A} p(s^{t-1}|s^k)p(e_t|s^{t-1}) \sum_{h^{t-1}\in H^{t-1}(s^{t-1})} p(h^{t-1}|s^{t-1}) \sum_{M_1,M_2,R} \sigma^\star(m_{1t}|e_t, h^{t-1})$$
$$\pi^\star(\tau_t, r_t|m_{1t}, m_{2t}, g^{t-1}(h^{t-1}))\sigma^\star(m_{2t}, a_t|e_t, m_{1t}, \tau_t, r_t, h^{t-1})u(e_t + \tau_t - a_t)\Bigg].$$

This inequality is identical to (54), and it was shown above that (54) is equivalent to (44) and therefore holds. Finally, we will show that (77) holds by applying much the same procedure. Using definition (80) on both sides of (77) gives:

$$\sum_{t>k}^{\infty}\beta^t\Bigg[\sum_{S^{t-1},E,T,A}\hat{p}(s^{t-1}|s^k)\hat{p}(e_t|s^{t-1})\sum_{h^{t-1}\in H^{t-1}(s^{t-1})}p(h^{t-1}|s^{t-1})\sum_{M_1,M_2,R}\sigma^\star(m_{1t}|\delta_{e1}(s^{t-1},e_t),h^{t-1})$$

$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|\delta_{e2}(s^{t-1},e_t),m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-\delta_a(s^t))\Bigg]$$

$$\leq \beta^{k+1}\Bigg[\sum_{E,T,A}p(e_t|s^k)\sum_{H^k(s^k)}p(h^k|s^k)\sum_{M_1,M_2,R}\sigma^\star(m_{1k+1}|\delta_{e1}(s^k,e_{k+1}),h^k(h^{k+1}))$$

$$\pi^\star(\tau_{k+1},r_{k+1}|m_{1k+1},m_{2k+1},g^k(h^{k+1}))\sigma^\star(m_{2k+1},a_{k+1}|e_t,m_{1k+1},\tau_{k+1},r_{k+1},h^k(h^{k+1}))u(e_{k+1}+\tau_{k+1}-a_{k+1})\Bigg]$$

$$+\sum_{t>k+1}^{\infty}\beta^t\Bigg[\sum_{S^{t-1},E,T,A}p(s^{t-1}|s^k)p(e_t|s^{t-1})\sum_{h^{t-1}\in H^{t-1}(s^{t-1})}p(h^{t-1}|s^{t-1})\sum_{M_1,M_2,R}\sigma^\star(m_{1t}|e_t,h^{t-1})$$

$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg]$$

By a change in notation for the agent's strategy this can be written as:

$$\sum_{t>k}^{\infty}\beta^t\Bigg[\sum_{S^{t-1},E,T,A}\hat{p}(s^{t-1}|s^k)\hat{p}(e_t|s^{t-1})\sum_{h^{t-1}\in H^{t-1}(s^{t-1})}p(h^{t-1}|s^{t-1})\sum_{M_1,M_2,R}\hat{\sigma}(m_{1t}|e_t,h^{t-1})$$

$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\hat{\sigma}(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg]$$

$$\leq \beta^{k+1}\Bigg[\sum_{E,T,A}p(e_t|s^k)\sum_{H^k(s^k)}p(h^k|s^k)\sum_{M_1,M_2,R}\hat{\sigma}(m_{1k+1}|e_{k+1},h^k(h^{k+1}))$$

$$\pi^\star(\tau_{k+1},r_{k+1}|m_{1k+1},m_{2k+1},g^k(h^{k+1}))\sigma^\star(m_{2k+1},a_{k+1}|e_t,m_{1k+1},\tau_{k+1},r_{k+1},h^k(h^{k+1}))u(e_{k+1}+\tau_{k+1}-a_{k+1})\Bigg]$$

$$+\sum_{t>k+1}^{\infty}\beta^t\Bigg[\sum_{S^{t-1},E,T,A}p(s^{t-1}|s^k)p(e_t|s^{t-1})\sum_{h^{t-1}\in H^{t-1}(s^{t-1})}p(h^{t-1}|s^{t-1})\sum_{M_1,M_2,R}\sigma^\star(m_{1t}|e_t,h^{t-1})$$

$$\pi^\star(\tau_t,r_t|m_{1t},m_{2t},g^{t-1}(h^{t-1}))\sigma^\star(m_{2t},a_t|e_t,m_{1t},\tau_t,r_t,h^{t-1})u(e_t+\tau_t-a_t)\Bigg],$$

and this equation can be transformed into (45) with exactly the same steps that were used above. □

The Recursive Representation

As before with Program 1, the next step is to reduce the current game to a recursive version with a vector of promised utilities as the state variable. The result will be Program 2.

Again, we consider a problem in which the planner has to deliver a vector of reservation utilities $\mathbf{w}_0$, depending on the endowment in the first period.

The choice object of the planner is given by $\pi(\tau_t, a_t|e_{1t}, e_{2t}, s^{t-1})$. The planner picks the contract conditional on the initial report e_{10} (the first report in period 0) by the agent. The constraints are the same as before, but we will write them down separately for each initial report e_{10}. The first set of constraints are the probability-measure constraints. For all e_t, s^{t-1}, we require:

$$\forall e_{1t}, e_{2t}, s^{t-1}: \quad \sum_{T,A} \pi(\tau_t, a_t|e_{1t}, e_{2t}, s^{t-1}) = 1. \tag{81}$$

As before, we need to impose a consistency constraint to ensure that the transfer τ does not depend on the second report e_{2t}:

$$\forall \tau_t, e_{1t}, e_{2t}, \bar{e}_{2t}, s^{t-1}: \quad \sum_{A} \pi(\tau_t, a_t|e_{1t}, e_{2t}, s^{t-1}) = \sum_{A} \pi(\tau_t, a_t|e_{1t}, \bar{e}_{2t}, s^{t-1}). \tag{82}$$

There is a promise-keeping constraint for each possible initial endowment. For all e_0, we require:

$$\begin{aligned}\forall e_0: \quad & \sum_{T_0,A_0} \pi(\tau_0, a_0|e_0, e_0, s^{-1}) \Big[u(e_0 + \tau_0 - a_0) \\ & + \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1},E,T,A} p(s^{t-1}|s_0) p(e_t|s^{t-1}) \pi(\tau_t, a_t|e_t, e_t, s^{t-1}) u(e_t + \tau_t - a_t) \Big] = w_0(e_0). \end{aligned} \tag{83}$$

Next, we need to guarantee that the agent is obedient, assuming that the second report in period zero $\hat{e}_0$ was correct. A deviation consists of an alternative reporting strategy $\delta_{e1}(e_t, s^{t-1})$, $\delta_{e2}(e_t, s^{t-1})$ *from the second period on* and an alternative action strategy $\delta_a(s^t)$. We use $\delta(s^t)$ to denote the history that was reported under the deviation if the true history is s^t, and $\hat{p}(\cdot)$ are the adjusted probabilities of states s^t and endowments e_t. For all e_0, $\hat{e}_0$, and all deviations, we require:

$$\forall e_0, \hat{e}_0, \delta_{e1}, \delta_{e1}, \delta_a : \quad \sum_{T_0, A_0} \pi(\tau_0, a_0 | e_0, \hat{e}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - \delta_a(s^0))$$

$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1}, E, T, A} \hat{p}(s^{t-1}|s_0) \hat{p}(e_t|s^{t-1}) \pi(\tau_t, a_t | \delta_{e1}(e_t, s^{t-1}), \delta_{e2}(e_t, s^{t-1}), \delta(s^{t-1})) u(e_t + \tau_t - \delta_a(s^t)) \Big]$$

$$\leq \sum_{T_0, A_0} \pi(\tau_0, a_0 | e_0, \hat{e}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - a_0)$$

$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1}, E, T, A} \hat{p}(s^{t-1}|s_0) \hat{p}(e_t|s^{t-1}) \pi(\tau_t, a_t | e_t, e_t, s^{t-1}), s^{t-1}) u(e_t + \tau_t - a_t) \Big] \quad (84)$$

We also have to make sure that the agent prefers to report the correct endowment $\hat{e}_0$ at the second report:

$$\forall e_0, \hat{e}_0, \hat{\hat{e}}_0, \delta_{e1}, \delta_{e1}, \delta_a : \quad \sum_{T_0, A_0} \pi(\tau_0, a_0 | e_0, \hat{\hat{e}}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - \delta_a(s^0))$$

$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1}, E, T, A} \hat{p}(s^{t-1}|s_0) \hat{p}(e_t|s^{t-1}) \pi(\tau_t, a_t | \delta_{e1}(e_t, s^{t-1}), \delta_{e2}(e_t, s^{t-1}), \delta(s^{t-1})) u(e_t + \tau_t - \delta_a(s^t)) \Big]$$

$$\leq \sum_{T_0, A_0} \pi(\tau_0, a_0 | e_0, \hat{e}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - a_0)$$

$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1}, E, T, A} \hat{p}(s^{t-1}|s_0) \hat{p}(e_t|s^{t-1}) \pi(\tau_t, a_t | e_t, e_t, s^{t-1}), s^{t-1}) u(e_t + \tau_t - a_t) \Big] \quad (85)$$

Finally, an agent needs incentives to tell the truth already at the first report. Since the constraints above ensure that from the second report on the agent will not deviate, deviations in actions do not need to be considered:

$$\forall e_0, \hat{e}_0 \neq e_0, : \quad \sum_{T_0, A_0} \pi(\tau_0, a_0 | e_0, \hat{e_0}, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - a_0)$$

$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1}, E, T, A} p(s^{t-1}|s_0) p(e_t|s^{t-1}) \pi(\tau_t, a_t | e_t, e_t, s^{t-1}) u(e_t + \tau_t - a_t) \Big] \leq w_0(\hat{e}_0). \quad (86)$$

Constraints (81) to (86) are merely a different of way of writing down our original constraints, in a way that will be useful for our recursive representation. Since the vector of promised utilities $\mathbf{w}$ will serve as our state variable, it will be important to show that the set of all feasible utility vectors has nice properties.

Definition 11 *The set* $\mathbf{W}$ *is given by all vectors* $\mathbf{w} \in R^{\#E}$ *that satisfy constraints* (81) *to* (86) *for some outcome function* $\pi(\tau_t, a_t | e_t, e_t, s^{t-1})$.

Proposition 8 *The set* $\mathbf{W}$ *is nonempty and compact.*

Proof: (See proof for Proposition 3) □

Now we consider the problem of a planner who has promised utility vector $\mathbf{w} \in \mathbf{W}$, and who has received the initial report e_0 (the first report in the initial period).

Problem P:

$$V(\mathbf{w}_0, e_0) = \max_{\pi} \sum_{T_0, A_0} \pi(\tau_0, a_0 | e_0, e_0, s^{-1}) \Big[-\tau_0 + \sum_{t=1}^{\infty} Q^t \sum_{S^{t-1}, E, T, A} p(s^{t-1}|s_0) p(e_t|s^{t-1}) \pi(\tau_t, a_t | e_t, e_t, s^{t-1})(-\tau_t) \Big]$$

Subject to constraints (81) to (86) above, for a given $e_0 \in E$ and $\mathbf{w} \in \mathbf{W}$. We want to show that this problem has a recursive structure.

To do this, we need to define on-path future utilities. For all s^{k-1}, e^k, let:

$$\begin{aligned} w(e_k, s^{k-1}) = &\sum_{T_k, A_k} \pi(\tau_k, a_k | e_k, e_k, s^{k-1}) \Big[u(e_k + \tau_k - a_k) \\ &+ \sum_{t=k+1}^{\infty} \beta^{t-k} \sum_{S^{t-1}, E, T, A} p(s^{t-1} | e_k, s^{k-1})) p(e_t | s^{t-1}) \pi(\tau_t, a_t | e_t, e_t, s^{t-1}) u(e_t + \tau_t - a_t) \Big], \end{aligned} \tag{87}$$

and let $\mathbf{w}(s^{k-1})$ be the vector of these utilities over all e_k. We can now show the following result:

Proposition 9 *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 *is an optimal contract for Problem P with* $e_0 = e_k$ *and* $\mathbf{w}_0 = \mathbf{w}(s^{k-1})$.

Proof: The proof is by construction. We have shown earlier that an optimal contract exists. Let $\tilde{\pi}$ be an optimal contract from time zero, and π_k an optimal contract for $e_0 = e_k$ and $\mathbf{w}_0 = \mathbf{w}(s^{k-1})$. Now construct a new contract $\pi^\star$ that is equal to π_k from (e_k, s^{k-1})

on, and equals $\tilde{\pi}^\star$ until time k and on all branches other than e_k, s^{k-1}. We claim that $\pi^\star$ is an optimal contract. To show this, we have to demonstrate that $\pi^\star$ satisfies constraints (81) to (86), 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 $\mathbf{w}_0 = \mathbf{w}(s^{k-1})$ as the starting point also constrain the choices of the planner in the original program from (e_k, s^{k-1}) on. By reoptimizing at (e_k, s^{k-1}) 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 $\tilde{\pi}$. Since $\tilde{\pi}$ is assumed to be an optimal contract, if $\pi^\star$ satisfies (81) to (86), it must be optimal as well. Thus we only have to show that (81) to (86) 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 and consistency constraints (81) and (82) 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 (83) is satisfied since the new contract delivers the same on-path utilities by construction. We still have to show that the incentive constraints (84), (85), and (86) are satisfied. We will do this by contradiction.

Suppose that (84) is not satisfied by contract $\pi^\star$. Then there is a deviation $\delta_{e1}(e_t, s^{t-1})$, $\delta_{e2}(e_t, s^{t-1})$, $\delta_a(s^t)$ such that:

$$\sum_{T_0,A_0} \pi^\star(\tau_0, a_0|e_0, \hat{e}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - \delta_a(s^0))$$
$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s_0)\hat{p}(e_t|s^{t-1})\pi^\star(\tau_t, a_t|\delta_{e1}(e_t, s^{t-1}), \delta_{e2}(e_t, s^{t-1}), \delta(s^{t-1}))u(e_t+\tau_t-\delta_a(s^t)) \Big]$$
$$> \sum_{T_0,A_0} \pi^\star(\tau_0, a_0|e_0, \hat{e}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - a_0)$$
$$+ \sum_{t=1}^{\infty} \beta^t \sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s_0)\hat{p}(e_t|s^{t-1})\pi^\star(\tau_t, a_t|e_t, e_t, s^{t-1}), s^{t-1})u(e_t + \tau_t - a_t) \Big] \quad (88)$$

Now use $w(e_k, s^{k-1}, \delta)$ to denote the continuation utility of the agent from time k on under the deviation strategy. Now we can rewrite (89) as:

$$\sum_{T_0,A_0} \pi^\star(\tau_0, a_0|e_0, \hat{e}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - \delta_a(s^0))$$
$$+\sum_{t=1}^{k-1} \beta^t \sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s_0)\hat{p}(e_t|s^{t-1})\pi^\star(\tau_t, a_t|\delta_{e1}(e_t, s^{t-1}), \delta_{e2}(e_t, s^{t-1}), \delta(s^{t-1}))u(e_t+\tau_t-\delta_a(s^t))$$
$$+ \beta^k \sum_{E,s^{k-1}} \hat{p}(s^{k-1}|s_0)\hat{p}(e_k|s^{k-1})w(e_k, s^{k-1}, \delta)\Big]$$
$$> \sum_{T_0,A_0} \pi^\star(\tau_0, a_0|e_0, \hat{e}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - a_0)$$
$$+\sum_{t=1}^{k-1} \beta^t \sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s_0)\hat{p}(e_t|s^{t-1})\pi^\star(\tau_t, a_t|e_t, e_t, s^{t-1}), s^{t-1})u(e_t + \tau_t - a_t)$$
$$+ \beta^k \sum_{E,s^{k-1}} \hat{p}(s^{k-1}|s_0)\hat{p}(e_k|s^{k-1})w(e_k, s^{k-1})\Big] \quad (89)$$

Notice that for s^{k-1} that are reached with positive probability under the deviation we have $w(e_k, s^{k-1}, \delta) \leq w(e_k, \delta(s^{k-1}))$, where $\delta(s^k)$ is the history as seen by the planner (reported endowments, delivered transfers, and recommended actions) under the deviation strategy. Otherwise, either $\tilde{\pi}$ or π_k would violate incentive constraints. Therefore (89) implies:

$$\sum_{T_0,A_0} \pi^\star(\tau_0, a_0|e_0, \hat{e}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - \delta_a(s^0))$$
$$+\sum_{t=1}^{k-1} \beta^t \sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s_0)\hat{p}(e_t|s^{t-1})\pi^\star(\tau_t, a_t|\delta_{e1}(e_t, s^{t-1}), \delta_{e2}(e_t, s^{t-1}), \delta(s^{t-1}))u(e_t+\tau_t-\delta_a(s^t))$$
$$+ \beta^k \sum_{E,s^{k-1}} \hat{p}(s^{k-1}|s_0)\hat{p}(e_k|s^{k-1})w(e_k, s^{k-1})\Big]$$
$$> \sum_{T_0,A_0} \pi^\star(\tau_0, a_0|e_0, \hat{e}_0, s^{-1}) \Big[u(\hat{e}_0 + \tau_0 - a_0)$$
$$+\sum_{t=1}^{k-1} \beta^t \sum_{S^{t-1},E,T,A} \hat{p}(s^{t-1}|s_0)\hat{p}(e_t|s^{t-1})\pi^\star(\tau_t, a_t|e_t, e_t, s^{t-1}), s^{t-1})u(e_t + \tau_t - a_t)$$
$$+ \beta^k \sum_{E,s^{k-1}} \hat{p}(s^{k-1}|s_0)\hat{p}(e_k|s^{k-1})w(e_k, s^{k-1})\Big] \quad (90)$$

But now the left-hand side of (90) is the utility that the agent gets under plan $\tilde{\pi}$ from following the deviation strategy until time k, and following the recommendations of the planner afterwards. Thus (90) contradicts the incentive compatibility of $\tilde{\pi}$. We obtain a contradiction, $\pi^\star$ actually satisfies (84). The proof for constraints (85) and (86) follows the same lines. This shows that plan $\pi^\star$ is within the constraints of the original problem. Since $\pi^\star$ yields at least as much surplus as $\tilde{\pi}$ and $\tilde{\pi}$ is an optimal contract, $\pi^\star$ must be optimal as well. □

Given this result, we know that the maximized surplus of the planner can be written as:

$$V(\mathbf{w}_0, e_0) = \sum_{A,T} \pi^\star(\tau_0, a_0 | e_0, e_0, s^{-1}) \left[-\tau_0 + Q \sum_E p(e_1|s^0) V(e_1, \mathbf{w}(s^0)) \right]. \tag{91}$$

In light of (91), we can cast the planner's problem as choosing transfers and actions in the present period, and choosing continuation utilities from the set $\mathbf{W}$ from tomorrow on. We will write the choices of the planner as a function of the vector of promised utilities $\mathbf{w}$ that has to be delivered in the current period, and the current state e. The choices of the planner are therefore functions $\pi(\tau, a|\mathbf{w}, e_1, e_2)$ and $\Pi(\mathbf{w}'|\mathbf{w}, e_1, e_2, \tau, a)$, where Π is a probability measure over set $\mathbf{W}$, and $\mathbf{w}'$ are the promised utilities from tomorrow on. We still need to determine which constraints need to be placed on these choices in order to guarantee that the implied contract satisfies (81) to (86). In order to reproduce (81), we need to impose:

$$\forall e_2: \quad \sum_{T,A} \pi(\tau, a|\mathbf{w}, e, e_2) = 1. \tag{92}$$

The consistency constraint (82) is satisfied if we impose:

$$\forall \tau, e_2, \bar{e}_2: \quad \sum_A \pi(\tau, a|\mathbf{w}, e, e_2) = \sum_A \pi(\tau, a|\mathbf{w}, e, \bar{e}_2). \tag{93}$$

The promise-keeping constraint (83) will be satisfied if we impose:

$$\sum_{T,A} \pi(\tau, a|\mathbf{w}, e, e) \left[u(e + \tau - a) + \beta \int_{\mathbf{W}} [\sum_E p(e'|a) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \tau, a) \right] = w(e). \tag{94}$$

We now get to the incentive constraints. The equivalent for constraint (84) can be written separately for each recommended action a:

$$\forall \hat{e}, \tau, a, \hat{a}: \quad u(\hat{e}+\tau-\hat{a}) + \beta \int_{\mathbf{W}} [\sum_E p(e'|\hat{a}) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \hat{e}, \tau, a)$$
$$\leq u(\hat{e}+\tau-a) + \beta \int_{\mathbf{W}} [\sum_E p(e'|a) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \hat{e}, \tau, a). \tag{95}$$

In order to insure that the second report is correct we need to impose:

$$\forall \hat{e}, \tau, \hat{\hat{e}} \neq \hat{e}, \delta: \quad \sum_{A,\mathbf{W}'} \pi(\tau, a, |\mathbf{w}, e, \hat{e}) \left[u(\hat{e}+\tau-\delta(a)) + \beta \int_{\mathbf{W}} [\sum_E p(e'|\delta(a)) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \hat{\hat{e}}, \tau, a) \right]$$
$$\leq \sum_{A,\mathbf{W}'} \pi(\tau, a|\mathbf{w}, e, \hat{e}) \left[u(\hat{e}+\tau-a) + \beta \int_{\mathbf{W}} [\sum_E p(e'|a) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \hat{e}, \tau, a) \right] \tag{96}$$

Finally, to make sure that the first report is correct, we require:

$$\forall \hat{e} \neq e: \quad \sum_{T,A} \pi(\tau, a, |\mathbf{w}, e, \hat{e}) \left[u(\hat{e}+\tau-a) + \beta \int_{\mathbf{W}} [\sum_E p(e'|a) w'(e')] \Pi(\mathbf{w}'|\mathbf{w}, e, \hat{e}, \tau, a) \right] \leq w(\hat{e}). \tag{97}$$

The constraints above rule out that the agent can gain from misreporting or disobedience in any period, given that he goes back to truth-telling and obedience from the next period on. The constraints therefore imply that (84), (85), and (86) hold for one-time deviations. It can be shown with the same methods used for Program 1 that (95) to (97) imply (84) to (86) even for a finite or infinite number of deviations.

Thus we are justified to pose the problem of the planner as solving:

$$V(\mathbf{w}, e) = \max_{\pi \geq 0} \sum_{A,T} \pi(\tau, a|\mathbf{w}, e, e) \left[-\tau + \beta \int_{\mathbf{W}} \sum_E p(e'|a) V(\mathbf{w}', e') \Pi(\mathbf{w}'|\mathbf{w}, e, e, \tau, a) \right] \tag{98}$$

by choice of π and Π, subject to constraints (92) to (97) above. Program 2 is a version of this problem with a discrete grid for promised utilities. Of course, we now assume that the function $V(\mathbf{w}, e)$ is already known. $V(\mathbf{w}, e)$ can be computed with standard dynamic programming techniques. Specifically, the right-hand side of (98) defines an operator T that maps functions $V(\mathbf{w}, e)$ into $TV(\mathbf{w}, e)$. It is easy to show 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.

A.4 Equivalence of Programs 1 and 3

We now want to show that Program 3 is equivalent to Program 1. In both programs, the planner chooses lotteries over transfer, action, and promised utilities. Even though in Program 3 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 10 *The allocations that can be implemented in Program 3 and 1 are identical.*

Proof: We want to show that constraints (22)-(26) in Program 3 place the same restrictions on the outcome function $\pi(\cdot)$ as the constraints (2)-(5) of Program 1. The probability constraints (2) and (22), the promise-keeping constraints (3) and (23), and the obedience constraints (4) and (24) are identical. This leaves us with the truth-telling constraints. Let us first assume we have found a lottery $\pi(\tau, a, \mathbf{w}'|\mathbf{w}, e)$ that satisfies the truth telling constraint (5) 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 (25) and (26) in Program 3. For each $\hat{e}$, τ, and a, define $v(\hat{e}, e, \tau, a)$ as the maximum of the left hand side of (25) over all $\hat{a}$:

$$v(\hat{e}, e, \tau, a) \equiv \max_{\hat{a}} \left\{ \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[u(\hat{e} + \tau - \hat{a}) + \beta \sum\nolimits_{E} p(e'|\hat{a}) w'(e') \right] \right\}. \tag{99}$$

Then clearly (25) is satisfied, since the left-hand side of (25) 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 (25):

$$\hat{\delta}(\tau, a) \equiv \text{argmax}_{\hat{a}} \left\{ \sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[u(\hat{e} + \tau - \hat{a}) + \beta \sum\nolimits_{E} p(e'|\hat{a}) w'(e') \right] \right\}. \tag{100}$$

Since $\pi(\tau, a, \mathbf{w}'|\mathbf{w}, e)$ satisfies (5) for any function $\delta(\cdot)$ by assumption, we have for our particular $\hat{\delta}(\cdot)$:

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \hat{\delta}(\tau, a)) + \beta \sum\nolimits_E p(e'|\hat{\delta}(\tau, a))w'(e')\right] \leq w(\hat{e}). \tag{101}$$

By the way we chose $\hat{\delta}(\cdot)$ and the $v(\cdot)$, we have from (25):

$$\sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \ \left[u(\hat{e} + \tau - \hat{\delta}(\tau, a)) + \beta \sum\nolimits_E p(e'|\hat{\delta}(\tau, a))w'(e')\right] = v(\hat{e}, e, \tau, a). \tag{102}$$

Substituting the left-hand side into (101), we get:

$$\sum_{T,A} v(\hat{e}, e, \tau, a) \leq w(\hat{e}). \tag{103}$$

which is (26).

Conversely, suppose we have found a lottery $\pi(\tau, a, \mathbf{w}'|\mathbf{w}, e)$ that satisfies (25) and (26) in Program 3 for some choice of $v(\hat{e}, e, \tau, a)$. By (25), we have then for any $\hat{e}$ and $\hat{a}$ and hence any $\delta : T \times S \to A$:

$$\sum_{\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \delta(\tau, a)) + \beta \sum\nolimits_E p(e'|\delta(\tau, a))w'(e')\right] \leq v(\hat{e}, e, \tau, a). \tag{104}$$

Substituting the left-hand side of (104) into the assumed (26) for the $v(\hat{e}, e, \tau, a)$, we maintain the inequality:

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \left[u(\hat{e} + \tau - \delta(\tau, a)) + \beta \sum\nolimits_E p(e'|\delta(\tau, a))w'(e')\right] \leq w(\hat{e}). \tag{105}$$

But this is (5) in Program 1. Therefore the sets of constraints are equivalent, which proves that Program 1 and Program 3 are equivalent. □

A.5 Equivalence of Programs 3 and 4

We now will show that Program 3 and Program 4 are equivalent in the sense that the same allocations are feasible in both programs. By allocation we mean a probability distribution over transfer, storage, and promised future utilities. In order to prove this claim, we will start with an allocation that satisfies the constraints in Program 3 and then show that this allocation is also feasible in Program 4. Afterwards, we will start with an allocation that is feasible in Program 4, and then show that it is also feasible in Program 3. Since in both Programs the planner has the same objective of maximizing surplus, this result also implies that the maximizing allocation will be the same in both programs, as will be the utility of the planner.

Proposition 11 *The allocations that can be implemented in Program 3 and 4 are identical.*

Proof: Assume that for a given vector of promised utilities $\mathbf{w}$ and a given true state e we have found an allocation $\pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)$ that satisfies the constraints (22)-(26) of Program 3 for some specification of utility bounds $v^\star(\hat{e}, e, \tau, a)$. Here and below the * denotes a solution to Program 3. We now need to find contracts and utility bounds for Programs 4a and 4b that implement the same allocation. The contracts that we choose have the property that the planner does not randomize over intermediate utilities in Program 4a. Given a transfer τ, one specific vector of intermediate utilities is assigned in Program 4a. Given this transfer and intermediate utility vector, in Program 4b the correct distribution over storage and future utilities is implemented. We will first describe which lotteries over storage and promised utilities to choose in stage 4b, and then move backwards to Program 4a.

At stage 4b, the planner recommends an action and assigns a vector of promised utilities, conditional on the transfer and the vector of interim utilities that were assigned at 4a. Since we will choose contracts at stage 4a that do not randomize over interim utilities, for each transfer τ we have to specify only one contract at stage 4b. Fix the transfer τ. We want to find a contract $\pi(a, \mathbf{w}')$ that implements the same distribution over a and $\mathbf{w}'$, conditional on τ, as $\pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)$. In order to do this, define:

$$\pi(a, \mathbf{w}') \equiv \frac{\pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)}{\sum_{A \times \mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)}. \tag{106}$$

We also have to define the utility bounds. Our choices are:

$$v(\hat{e}, e, \tau, a) \equiv \frac{v^\star(\hat{e}, e, \tau, a)}{\sum_{A \times \mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)}. \tag{107}$$

These choices also correspond to a specific vector of interim utilities. We choose these utilities to ensure that all constraints in Program 4b are satisfied. We write these utilities as a function of τ so that we can identify which interim utilities correspond to which transfers in Program 4a. Specifically, we let:

$$w_m(e)(\tau) \equiv \sum_{A,\mathbf{W}'} \pi(a, \mathbf{w}') \left[u(e+\tau-a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right], \tag{108}$$

where the $\pi(a, \mathbf{w}')$ are defined in (106). For each $\hat{e}$, we define

$$\bar{w}_m(\hat{e}, e)(\tau) \equiv \sum_A v(\hat{e}, e, \tau, a), \tag{109}$$

where the $v(\hat{e}, e, \tau, a))$ are defined in (107). We now have to show that all constraints Program 4b are satisfied. We can use the fact that $\pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)$ and the $v^\star(\hat{e}, e, \tau, a)$ satisfy all constraints in Program 3.

We will start with the probability constraint (32). Substituting (106) into the left-hand side of (32), we get:

$$\frac{\sum_{A\times\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)}{\sum_{A\times\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)} = 1, \tag{110}$$

which on the left-hand side clearly equals unity. Therefore (32) is satisfied. The promise-keeping constraint (33) is identical to (108) and therefore also satisfied. We now move to the obedience constraint (34). Substituting our definition (106) into what one would like to show, namely (34), we need to establish:

$$\begin{aligned}\sum_{\mathbf{W}'} \frac{\pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)}{\sum_{A\times\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)} & \left[u(e+\tau-a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right] \\ \geq \sum_{\mathbf{W}'} \frac{\pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)}{\sum_{A\times\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)} & \left[u(e+\tau-\hat{a}) + \beta \sum\nolimits_E p(e'|\hat{a})w'(e')\right],\end{aligned}$$

or, equivalently, by pulling out the common constant in the denominator:

$$\begin{aligned}\sum_{\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e) & \left[u(e+\tau-a) + \beta \sum\nolimits_E p(e'|a)w'(e')\right] \\ \geq \sum_{\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e) & \left[u(e+\tau-\hat{a}) + \beta \sum\nolimits_E p(e'|\hat{a})w'(e')\right].\end{aligned} \tag{111}$$

This is (24) in Program 3 and therefore by assumption satisfied. Next, substituting our definitions (106) and (107) into what we would like to show, namely (35), and multiplying by $\sum_{A\times\mathbf{W}'}\pi^\star(\tau,a,\mathbf{w}'|\mathbf{w},e)$ on both sides we need to establish:

$$\sum_{\mathbf{W}'}\pi^\star(\tau,a,\mathbf{w}'|\mathbf{w},e)\ \left[u(\hat{e}+\tau-\hat{a})+\beta\sum\nolimits_{E}p(e'|\hat{a})w'(e')\right]\le v^\star(\hat{e},e,\tau,a), \tag{112}$$

But this is (25) and therefore by assumption true. Finally, the constraints (36) are satisfied because we defined the utility bounds in (109) accordingly. Therefore in Program 4b all constraints are satisfied.

We now move to stage 4a. As mentioned earlier, for a given τ, the contract we use does not randomize over interim utilities. Since we have already determined which interim utilities correspond to which transfer τ, the only thing left to do is to assign probability mass to different transfers and their associated interim utilities. We choose the contract such that the probability of each transfer equals the probability of that transfer in contract $\pi^\star(\cdot)$. For each τ, we have:

$$\pi(\tau,w_m(e)(\tau),\bar{\mathbf{w}}_m(\hat{e},e)(\tau))\equiv\sum_{A,\mathbf{W}}\pi^\star(\tau,a,\mathbf{w}'|\mathbf{w},e). \tag{113}$$

We set the probability of all other combinations of transfer and interim utilities to zero. Clearly, by choosing this contract we implement the same allocation over transfer, storage, and promised utilities as $\pi^\star(\cdot)$ does. We still have to check whether the constraints (28)-(30) of Program 4a are satisfied. The first constraint (28) is the probability constraint. We have:

$$\begin{aligned}\sum_T\sum_{\mathcal{W}(e,\tau)}\pi(\tau,w_m(e),\bar{\mathbf{w}}_m(\hat{e},e)) &= \sum_T\pi(\tau,w_m(e)(\tau),\bar{\mathbf{w}}_m(\hat{e},e)(\tau)) \\ &= \sum_{T,A,\mathbf{W}'}\pi^\star(\tau,a,\mathbf{w}'|\mathbf{w},e).\end{aligned} \tag{114}$$

since the right-hand side of (114) is just a sum over the left-hand side of (113). Since the $\pi^\star(\tau,a,\mathbf{w}'|\mathbf{w},e)$ satisfy (22) we have, as required:

$$=1. \tag{115}$$

For the promise-keeping constraint (29) we start on its left-hand side and use (106) and (108) to get:

$$\begin{aligned}
\sum_T \sum_{\mathcal{W}(e,\tau)} \pi(\tau, w_m(e), \bar{\mathbf{w}}_m(\hat{e}, e))\; w_m(e) &= \sum_T \pi(\tau, w_m(e)(\tau), \bar{\mathbf{w}}_m(\hat{e}, e)(\tau))\; w_m(e)(\tau) \\
&= \sum_{T,A,\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)\; w_m(e)(\tau) \\
&= \sum_{T,A,\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)\; \left[u(e+\tau-a) + \beta \textstyle\sum_E p(e'|a) w'(e')\right] \\
&= w(e), \qquad (116)
\end{aligned}$$

where the last equality follows because the $\pi^\star(\cdot)$ satisfy (23). Finally, for the truth-telling constraint (30), using (113) and starting on its left-hand side, given any $\hat{e} \neq e$ we get:

$$\begin{aligned}
\sum_T \sum_{\mathcal{W}(e,\tau)} \pi(\tau, w_m(e), \bar{\mathbf{w}}_m(\hat{e}, e))\; \bar{w}_m(\hat{e}, e) &= \sum_T \pi(\tau, w_m(e)(\tau), \bar{\mathbf{w}}_m(\hat{e}, e)(\tau))\; \bar{w}_m(\hat{e}, e)(\tau) \\
&= \sum_{T,A,\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)\; \bar{w}_m(\hat{e}, e)(\tau) \\
&= \sum_{T,A,\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)\; \left[\sum_A v(\hat{e}, e, \tau, a)\right]
\end{aligned}$$

by (109),

$$= \sum_{T,A,\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)\; \left[\frac{\sum_A v^\star(\hat{e}, e, \tau, a)}{\sum_{A,\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)}\right]$$

because of (107). The $\pi(\cdot)$ will cancel, that is:

$$\begin{aligned}
&= \sum_T \left[\sum_{A,\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)\; \left[\frac{\sum_A v^\star(\hat{e}, e, \tau, a)}{\sum_{A,\mathbf{W}'} \pi^\star(\tau, a, \mathbf{w}'|\mathbf{w}, e)}\right]\right] \\
&= \sum_{T,A} v^\star(\hat{e}, e, \tau, a) \\
&\leq w(\hat{e}). \qquad (117)
\end{aligned}$$

The last inequality follows because the $v^\star(\cdot)$ by assumption satisfy (26). Thus all constraints are satisfied. Therefore an allocation that is feasible in Program 3 can always be implemented in Program 4.

We will now show conversely that an allocation that is feasible in Program 4 is also feasible in Program 3. For a given vector of promised utilities $\mathbf{w}$ and a given true endowment, assume that we have found contracts and utility bounds $\pi^\star(\tau, w_m, \bar{\mathbf{w}}_m)$, $\pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m)$, and $v^\star(\hat{e}, e, \tau, a)(\tau, w_m, \bar{\mathbf{w}}_m)$ that satisfy all constraints in Programs 4a and 4b. We use * to denote a solution to Program 4a and 4b and write the contracts and utility bounds for Program 4a as a function of transfer and interim utilities to indicate to which transfer and interim utility vector the contracts applies to. In this section, we drop the arguments from the interim utilities in order to fit the equations on the page. We now have to find contracts and utility bounds for Program 3 that implement the same allocation and that satisfy constraints (22) to (26). For each τ, a, $\mathbf{w}'$, define as the obvious guess:

$$\pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \equiv \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m)\ \pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m). \tag{118}$$

Basically, we integrate out the interim utilities. We guess the following utility bounds:

$$v(\hat{e}, e, \tau, a) \equiv \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m)\ v^\star(\hat{e}, e, \tau, a). \tag{119}$$

Clearly, the chosen contract implements the same allocation. We have to show that the contract and the utility bounds satisfy the constraints (22)-(26) of Program 3.

We start with the probability constraint (22). Using definition (118), we get:

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) = \sum_{T,A,\mathbf{W}} \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m)\ \pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m).$$

Rewriting the order of summation:

$$= \sum_{T} \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m)\ \left[\sum_{A,\mathbf{W}'} \pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m) \right].$$

Using the probability constraint (32) in Program 4b gives:

$$= \sum_{T} \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m) = 1, \tag{120}$$

where, because of the probability constraint (28) in Program 4a, the left-hand side of (120) equals unity, and thus (22) is satisfied. We now will show that the promise-keeping constraint (23) holds. Starting on its left-hand side and using definition (118) we get:

$$\sum_{T,A,\mathbf{W}'} \pi(\tau, a, \mathbf{w}'|\mathbf{w}, e) \; \left[u(e+\tau-a) + \beta \textstyle\sum_E p(e'|a)w'(e')\right]$$
$$= \sum_{T,A,\mathbf{W}'} \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m) \; \pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m) \; \left[u(e+\tau-a) + \beta \textstyle\sum_E p(e'|a)w'(e')\right].$$

Rearranging the order of summation gives:

$$= \sum_T \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m) \; \left[\sum_{A,\mathbf{W}'} \pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m) \; \left[u(e+\tau-a) + \beta \textstyle\sum_E p(e'|a)w'(e')\right]\right].$$

Using the promise-keeping constraint (33) from Program 4b gives:

$$= \sum_T \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m) \; w_m = w(e), \tag{121}$$

where for the last step we used the promise-keeping constraint (29) from Program 4a.

Thus promise keeping is satisfied as well. Substituting definition (118) into both sides of what we hope will be the obedience constraint (24), we get:

$$\sum_{\mathbf{W}'} \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m) \; \pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m) \; \left[u(e+\tau-a) + \beta \textstyle\sum_E p(e'|a)w'(e')\right]$$
$$\geq \sum_{\mathbf{W}'} \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m) \; \pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m) \; \left[u(e+\tau-\hat{a}) + \beta \textstyle\sum_E p(e'|\hat{a})w'(e')\right]. \tag{122}$$

Rearranging the order of summation yields:

$$\sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m) \; \left[\sum_{\mathbf{W}'} \pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m) \; \left[u(e+\tau-a) + \beta \textstyle\sum_E p(e'|a)w'(e')\right]\right]$$
$$\geq \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau, w_m, \bar{\mathbf{w}}_m) \; \left[\sum_{\mathbf{W}'} \pi^\star(a, \mathbf{w}')(w_m, \bar{\mathbf{w}}_m) \; \left[u(e+\tau-\hat{a}) + \beta \textstyle\sum_E p(e'|\hat{a})w'(e')\right]\right]. \tag{123}$$

This inequality holds because of (34) term by term before summing over $\mathcal{W}(e,\tau)$. Next, substituting our definitions (118) and (119) into what we hope will be constraint (25), we want to establish that:

$$\sum_{\mathbf{W}'}\sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau,w_m,\bar{\mathbf{w}}_m)\;\pi^\star(a,\mathbf{w}')(w_m,\bar{\mathbf{w}}_m)\;\Big[u(\hat{e}+\tau-\hat{a})+\beta\sum\nolimits_E p(e'|\hat{a})w'(e')\Big]$$
$$\leq \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau,w_m,\bar{\mathbf{w}}_m)\;v^\star(\hat{e},e,\tau,a). \tag{124}$$

Rearranging gives:

$$\sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau,w_m,\bar{\mathbf{w}}_m)\;\left[\sum_{\mathbf{W}'}\pi^\star(a,\mathbf{w}')(w_m,\bar{\mathbf{w}}_m)\;\Big[u(\hat{e}+\tau-\hat{a})+\beta\sum\nolimits_E p(e'|\hat{a})w'(e')\Big]\right]$$
$$\leq \sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau,w_m,\bar{\mathbf{w}}_m)\;v^\star(\hat{e},e,\tau,a), \tag{125}$$

which holds because by (35) it is assumed to hold pointwise in Program 4b. Finally, starting from the left-hand side of the hoped for (26), we use definition (119) to get:

$$\sum_{T,A} v(\hat{e},e,\tau,a) = \sum_{T,A}\sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau,w_m,\bar{\mathbf{w}}_m)\;v^\star(\hat{e},e,\tau,a).$$

Changing the order of summation gives:

$$= \sum_{T}\sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau,w_m,\bar{\mathbf{w}}_m)\;\left[\sum_{A} v^\star(\hat{e},e,\tau,a)\right],$$

because of the assumed (36) in Program 4b this is:

$$\leq \sum_{T}\sum_{\mathcal{W}(e,\tau)} \pi^\star(\tau,w_m,\bar{\mathbf{w}}_m)\;\bar{w}_m(\hat{e},e)$$
$$\leq w(\hat{e}), \tag{126}$$

where the last inequality follows from (30). Thus all constraints are satisfied, which completes the proof. Since the feasible allocations are identical in both Programs, and the objective function is the same, this also implies that Program 3 and 4 have the same solution. $\square$

Corollary 1 *Program 1, Program 2, Program 3, and Program 4 are all equivalent.*

Proof: We showed in Appendix A2 and A3 that all allocations that are feasible in the general mechanism are also feasible in Program 1 and Program 2. We also showed in A4 that Program 1 and Program 3 are equivalent, and Proposition 9 here in A5 establishes that Program 3 and Program 4 are equivalent. Since the feasible allocations are identical and the objective function is the same, all programs have the same solutions. □

References

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

Atkeson, Andrew and Robert E. Lucas, Jr. 1992. "On Efficient Distribution with Private Information." *Review of Economic Studies* 59:427–453.

Attanasio, Orazio and Jose-Victor Rios-Rull. 2000. "Consumption Smoothing in Island Economies: Can Public Insurance Reduce Welfare?" *European Economic Review* 44 (7): 1225–58.

Cole, Harold L. and Narayana Kocherlakota. 2001. "Efficient Allocations with Hidden Income and Hidden Storage." *Review of Economic Studies* 68 (3): 523–42.

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.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. 2000. "Mutual Insurance, Individual Savings, and Limited Commitment." *Review of Economic Dynamics* 3 (2): 216–246.

Perri, Fabrizio and Dirk Krueger. 1999. "Risk Sharing: Private Insurance Markets or Redistributive Taxes?" Federal Reserve Bank of Minneapolis Staff Report 262.

Phelan, Christopher. 1998. "On the Long Run Implications of Repeated Moral Hazard." *Journal of Economic Theory* 79:174–91.

Prescott, Edward S. 1997. "Computing Private Information Problems with Dynamic Programming Methods." Federal Reserve Bank of Richmond.

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