

Dynamic Risk Sharing With Private Information And Costly Verification of Storage

Urs Haegler¹

Department of Economics, Royal Holloway, University of London, and Financial Markets Group, London School of Economics

e-mail: u.haegler@rhbnc.ac.uk

November 1999

Abstract

This paper studies the role of dynamic risk sharing in a world where agents have private information about their incomes, and in which their storage activities can be monitored at a cost. A principal offers long-term, full-commitment contracts to the agents, promising them consumption smoothing to some extent, as well as insurance against low incomes. The study examines the choice of optimal reporting and verification strategies and their effect on the efficiency of the contracts. The main results are that allowing for costly verifiability of storage may enable the principal to offer income-contingent transfer schemes that are Pareto-superior to pure borrowing and lending, and that only agents reporting very low incomes may be investigated.

¹The author would like to thank Harold Cole, Jeff Frank, Douglas Gale, Martin Hellwig, Anil Kashyap, Nobu Kiyotaki, Bob Nobay and David Webb for helpful discussions.

1 Introduction

This paper studies the role of dynamic risk sharing in a world where agents have private information about their incomes, and in which their storage activities can be monitored at a cost only. An agent's income risk is shared with a principal, who may be interpreted as a government (benefit) agency that is obliged to have a balanced budget.¹ The principal offers long-term contracts to individuals who seek not only intertemporal consumption smoothing but also intratemporal insurance against low incomes. The purpose of the study is to examine the choice of optimal reporting and verification strategies and their effect on the efficiency of long-term income-insurance contracts.

Most institutions providing such insurance services, public or private, require information regarding an agent's income or wealth. Usually it is the agents themselves who are asked to provide this type of information. However, very often they have an incentive to misrepresent data, because the transfers they will receive or be asked to pay depend on their reports. On the other hand, many of these institutions are entitled to check the truthfulness of reports made. Townsend (1979) analyses a model with deterministic verification strategies. However, as such auditing activity is costly, random verification strategies may be preferable, as shown by Mookherjee and Png (1989)) within a static framework. Indeed, verification in real-world situations usually occurs on the basis of random samples rather than being an action taken in each case. Therefore, we adapt the use of random verification strategies to our dynamic context.

There is now a substantive number of papers dealing with efficient allocations under informational asymmetries in a dynamic framework. This literature was pioneered by Radner (1981), Townsend (1982), Rubinstein and Yaari (1983), with most of the later models building on the approach put forward by Green (1987). These contributions show that optimal truth-telling contracts may generate consumption allocations that are Pareto-superior to those achieved by pure lending and borrowing, i.e. in a permanent-income-hypothesis (PIH) framework according to (Friedman (1957)). This improvement is achieved by invoking long-term contracts which promise to an agent a sequence of transfers with an expected net present value of transfers that is decreasing in his income report. Hence, such

¹Alternatively, one could think of the principal as as competitively behaved financial intermediary.

contracts offer some insurance against bad income shocks. In order to establish a truth-telling equilibrium, the intertemporal structure of these transfers, however, must provide agents with incentives to report their incomes truthfully.

One important feature common to all these papers is that there is no possibility of intertemporal storage. This is justified by assuming that either the physical properties of (consumption) goods or the technology available do not allow for storage, or that inventories are perfectly (costlessly) monitorable by the principal.

Neither of these assumptions is particularly convincing, though, which is why other studies have abandoned them. Contributions by Allen (1985) and, more recently, Cole and Kocherlakota (1997) have shown that the efficient allocation may then be ‘forced back’ to coincide with a pure-debt securities (PIH) equilibrium. In other words, if assets can be hidden, then efficient allocations are identical to what can be achieved in an incomplete-markets economy where all trade in contingent claims is shut down (Aiyagari and Williamson (1997)). The problem is that any dynamic transfer plan that would provide some insurance can be undone by agents using a storage policy that suits their intertemporal consumption preferences. This in turn undermines the incentives to report their incomes truthfully in the first place.²

These negative results cast severe doubts on the usefulness of the private-information approach to dynamic risk sharing. The present paper re-assesses this situation by considering a framework where storage can be concealed only if the principal abstains from the costly investigation of an agent’s inventory, and in which the former commits to a monitoring policy that induces truth-telling by the latter. The model provides a generalised framework of dynamic risk sharing that nests the non-storability and the non-verifiability settings described above as two extreme cases. The verifiability cost is implicitly assumed to be zero in the former and going to infinity in the latter.

Two results are central to the study. Firstly, it is shown that introducing

²This finding appears to be fairly robust. The introduction of lotteries, for instance, cannot be used to overturn the results as long as preferences exhibit non-increasing absolute risk aversion (an assumption that is often imposed on utility functions in this type of literature [see, e.g., Thomas and Worrall (1990)]). It also seems to hold in a world with bilateral asymmetric information and balanced budgets (Wang (1995)), as well as in a framework put forward more recently by Fernandes and Phelan (1999) which examines cases in which the privately observed shocks to income or preferences are serially correlated.

monitorability of storage at a finite cost goes some way towards re-establishing the Pareto-superiority of income-contingent contracts, provided the verification cost is sufficiently low and the penalty to be imposed on an agent that has failed to comply with the contractual storage prescriptions can be made severe enough.

Secondly, whilst there is never an equilibrium in which the storage of an agent reporting the highest possible income is investigated, there are parameter constellations for which there is no verification for the next lower income reports either. Only agents that declare a sufficiently low income have their storage checked with positive probability. As a consequence, those with higher incomes first make a lump-sum transfer payment in the reporting period. After that, they expect a (future) consumption allocation that corresponds to that obtained through pure borrowing-and-lending arrangements, i.e. characterised by a net present value of zero. This is due to the fact that, with certainty, their storage will not be verified. Thus, if the transfer scheme did promise them anything else but optimal consumption smoothing they can always respond to that by storing optimally, as there will be no monitoring and hence no risk of being penalised.

In contrast, agents that report a low income have their storage investigated with positive probability. They expect to receive a sequence of transfers that has positive net present value but does not allow them to smooth consumption optimally. This is similar to what agents who earn less than the highest income experience in Green and Oh (1991), for example.

The main reason for the assumption that it is the storage levels that can be verified at a cost, rather than the shocks to the privately observed variable (income) as in papers by Smith and Wang (1997) and Wang and Williamson (1998), e.g., is the wider applicability of such an approach. Atkeson and Lucas (1992), for instance, model an economy in which there are shocks to tastes, and not incomes. The scenario of a principal verifying intangible preferences seems less plausible than that of monitoring tangible assets. Moreover, whilst verifying wealth certainly has its difficulties, monitoring an agent's income is not necessarily more straightforward. An individual, for instance, may obtain unobserved income through transfers by relatives or through 'moonlighting' activities.

The study bears some relevance to the debate about means testing in social welfare and pension schemes. In many industrialised countries there are political discussions on whether to move away from a system of universal benefits towards a

system of means-tested social benefits. In the former everyone makes contributions and receives payments according to their incomes such that the present value of the two flows match each other more or less. In the latter, transfers are based more specifically on current income and wealth.

To some extent, universal benefits can be viewed as merely redistributing incomes intertemporally. Means-tested benefits, on the other hand, usually incorporate a substantive degree of intratemporal redistribution of incomes and wealth from the richer to the poorer, thereby providing income insurance.³

A number of features described in the analysis are consistent with observations from real-world welfare and tax systems. Individuals reporting medium or high incomes make transfers (pay taxes) in relation to these reports. Apart from that, they appear to rely on various forms of pure lending and borrowing to smooth intertemporal consumption (e.g. through pension schemes). They are usually not explicitly participating in the system of welfare benefits. The latter pays transfers with a positive net present value to households reporting low incomes, but often tries to monitor wealth (and sometimes income) of benefit applicants.

The paper is organised as follows. In Section 2 we describe the physical framework and set up the basic problem to be solved by the contractual parties. In Section 3 we note a number of general results which help us to simplify the analysis of the model of Section 2. We identify a number of features of the optimal contract, putting particular emphasis on incentive compatibility, and we discuss penalties, rewards and verification strategies. In Section 4 we illustrate the general model with a two-period example, in which there are three income levels. In particular, we compare the equilibrium allocation with those found in models without storability and in models with infinite verification costs. Finally, we make a number of concluding remarks in Section 5.

³It should be mentioned, however, that one of the main disadvantages of means-tested benefits are considered to be the disincentives to work and save that they provide (see, e.g., ?). The trade-off between this negative effect and the improved insurance feature of means testing will not be examined here.

2 The Environment and the Model

We consider an economy with discrete time, as represented by the set of time periods $\{0\} \cup \mathcal{T}$, where $\mathcal{T} \equiv \{1, 2, \dots, T\}$ and $T < \infty$. There is a continuum of *a priori* identical agents, distributed over the unit interval $[0, 1]$ and alive for $T + 1$ periods. In each time period $t \in \mathcal{T}$, each agent receives a privately observed random endowment y_t of a single variety of a storable consumption good. The realisation of y_t is drawn from a finite set $\Theta = \{\theta_1, \theta_2, \dots, \theta_n\}$. We assume that, for each agent, $0 \leq \theta_1 < \theta_2 < \dots < \theta_n$, and that the y_t are identically and independently distributed with $\Pr[y_t = \theta_i] = p_i > 0$, $\forall t \in \mathcal{T}$ and $\theta_i \in \Theta$, $\sum_i p_i = 1$. The probability distribution of the agents' endowments is common knowledge.

The utility a representative agent assigns to a consumption stream $c_t^T \equiv \{c_\tau\}_{\tau=t}^T$, at the beginning of period $t \in \mathcal{T}$, is given by

$$U(c_t^T) = \sum_{\tau=t}^T \beta^{\tau-t} u(c_\tau), \quad (2.1)$$

where the instantaneous-utility function $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ allows for non-negative consumption only, and is assumed to be strictly increasing and strictly concave, and the discount factor $\beta \in (0, 1)$. Each agent's consumption is private information as well. An agent's history of endowments up to and including period t is denoted by $y^t = (y_1, y_2, \dots, y_t)$, and the set of all possible such histories by Θ^t .

A government authority offers budget-neutral contracts (transfer mechanisms) to the agents. That is, the (date-0) net present value of transfers is zero.⁴

Throughout this paper we assume full commitment, i.e. contracts signed are binding and enforceable for the entire duration T of a relationship between principal and agent. A version of the Revelation Principle applies in our environment. We can therefore focus on direct mechanisms to implement efficient allocations.

⁴Alternatively, we may assume that there exists a large but finite number $N < \infty$ of risk-neutral financial intermediaries that offer contracts to the agents in a perfectly competitive environment. Agents' search costs are assumed to be zero, and therefore all intermediaries will offer identical contracts in equilibrium. Due to perfect competition, the intermediaries' expected profits from these contracts are zero. Hence, we could focus, without loss of generality, on the case where each intermediary serves $\frac{1}{N}$ of the population of agents. Note that due to our assumption of the set of agents constituting a continuum, the law of large numbers applies in the sense that each intermediary is faced with an identical cross-section of the population of agents almost surely.

The contracts specify (positive or negative) transfer payments from the intermediary to an agent. The transfers are contingent on that agent's track record, which consists of the history reported incomes and the history of investigations, as compiled by the principal.

At each point in time $t \in \mathcal{T}$ an agent sends a report \hat{y}_t on his income to the principal.⁵

For a given income history y^{t-1} up to and including period $t - 1$, we denote by y_t^τ an income history up to and including period τ , $\tau \geq t$, that has y^{t-1} as a subhistory.

Each agent has access to a private storage technology, the use of which can be observed by an outsider only at a cost. If the principal never made an effort to verify agents' storage activities, any ex-ante desired allocation which provides some degree of insurance, could be undone ex post through storage. In order to prevent such actions, the principal will have to set appropriate incentives. The latter consist of storage prescriptions specified in the contract, and the threat to monitor certain agents' storage activities.

More specifically, the contract will prescribe storage levels $s_t \geq 0$ as well as a random verification strategy $\pi_t \in [0, 1]$ for each period $t \in \mathcal{T}$. In other words, the intermediary can commit to probabilities π_t of checking whether an agent's inventory is in line with the contractual prescription. In contrast to, e.g., Townsend (1979) who restricts attention to deterministic audit strategies, we follow Mookherjee and Png (1989) in allowing for randomised verification strategies. The difficulty for an intermediary is that he can observe an agent's period- t storage decision $\hat{s}_t \geq 0$ only at the very beginning of the subsequent period, and more importantly, at a monitoring cost $\gamma > 0$ only.

The storage prescription and the verification probability at a particular point in time are contingent on an agent's track record to date. The track record contains the principal's relevant information about an agent. As mentioned previously, it also determines the (possibly negative) transfer payment received by the latter. Generally speaking, if the inventory is found to be out of line with the storage levels prescribed in the contract, i.e. if $\hat{s}_t \neq s_t$, the contract may foresee that a

⁵Of course, there are a number of real-world situations where a principal relies on agents' information about privately observed income. For instance, lenders and tax collectors base their actions on reports received by borrowers and taxpayers, respectively.

lower transfer be made than otherwise. On the other hand, if the investigation shows that the agent has chosen the appropriate level of storage, i.e. if $\hat{s}_t = s_t$, the contract may reward him for his honesty by making the transfer higher than it would have been without an investigation.

Formally, we define a track record up to date $t+1$ as a vector $\hat{h}^t \equiv (\hat{h}_1, \hat{h}_2, \dots, \hat{h}_t)$, where

$\hat{h}_\tau \equiv (\hat{y}_\tau, \iota_\tau, \hat{s}_{\tau-1}^{\iota_\tau})$, $\tau = 1, 2, \dots, t$. Thus, the track record consists of (i) the history of income reports, (ii) the history of verification efforts, as represented by the sequence of indicator values, $\iota^t \equiv (\iota_1, \iota_2, \dots, \iota_t)$ (with $\iota_\tau \in \{0, 1\} \forall \tau \in \mathcal{T}$), and (iii) the history of investigative findings. If an investigation occurs in τ , $\iota_\tau = 1$ and $\hat{s}_{\tau-1}^{\iota_\tau} = \hat{s}_{\tau-1}$; otherwise $\iota_\tau = 0$ and, for convenience, we set $\hat{s}^{\iota_\tau} = s_{\tau-1}$, $\tau = 1, 2, \dots, t$.⁶ (Note that what is investigated in period τ is the inventory carried over from period $\tau - 1$, $\hat{s}_{\tau-1}$.) Let \hat{H}^t represent the set of all such verification histories \hat{h}^t up to period t .

Consumption is required to be non-negative, $c_t \geq 0$ for all t . Hence, if an agent is found to have violated the contractual storage agreement, an intermediary cannot punish an agent more than by confiscating his entire wealth as well as present and future income.⁷

The contract is written in period 0 and specifies for each period the (possibly negative) transfer from the intermediary to the agent, the private storage level, and the probability of investigation. All of them are functions of the track record. A transfer system is a sequence of functions $\{b_t\}_{t=1}^T$ with $b_t : \hat{H}^t \rightarrow \mathbb{R}$. The storage prescriptions are a sequence of functions $\{s_t\}_{t=1}^T$ such that $s_t : \hat{H}^t \rightarrow \mathbb{R}_+$, which assign admissible levels of storage to the agent. The investigation probabilities are a sequence of functions $\{\pi_t\}_{t=1}^T$ with $\pi_t : \hat{H}^{t-1} \rightarrow [0, 1]$. Thus, an agent who has had a track record \hat{h}^{t-1} up to period $t - 1$, is investigated at the beginning of t with probability $\pi_t(\hat{h}^{t-1})$.

⁶In other words, if no investigation occurs the intermediary simply assumes that the agent has complied with the storage prescription s_τ .

⁷Alternatively, one could assume that an investigating intermediary will never find more than a certain fraction of the total amount secretly stored, which would limit the feasible punishment even further. Whilst such a scheme may be somewhat more realistic, the expected additional insights from adopting it do not warrant the increase in modeling costs. As long as the fraction of the storage that can be found is sufficiently large, the results found here should not be affected qualitatively.

For simplicity we assume that the agent does not know whether he has been monitored until after he has made his income report in period t . The principal's investigative (in)activity and the agent's income report then lead to the updated track record \hat{h}^t , on the basis of which the agent receives $b_t(\hat{h}^t)$ and is required to store $s_t(\hat{h}^t)$ in t . An agent's period- t net endowment, $y_t + b_t(\hat{h}^t)$ can either be consumed or be stored at a constant gross rate of return on storage of $R > 0$.⁸

To summarise the description of the model we give the following definition of a contract.

DEFINITION: A contract Γ is a mechanism consisting of a transfer system $\{b_t\}_{t=1}^T$, a sequence of storage prescriptions $\{s_t\}_{t=1}^T$, and contingent probabilities of investigating $\{\pi_t\}_{t=1}^T$.

Each period the principal can store a non-negative amount of the consumption good, S_t , at the gross rate of return $R > 0$. We do not restrict the storage of intermediaries to be non-negative.

The period- t consumption of an agent with a track record \hat{h}^t up to period t can be expressed as

$$c_t(y^t, \hat{h}^t) = y_t + b_t(\hat{h}^t) - \hat{s}_t(y^t, \hat{h}^t) + R\hat{s}_{t-1}(y^{t-1}(y^t), \hat{h}^{t-1}(\hat{h}^t)), \quad (2.2)$$

where $y^{t-1}(y^t)$ and $\hat{h}^{t-1}(\hat{h}^t)$ stand for the sub-histories of income history y^t and track record \hat{h}^t , respectively.

The following two definitions are useful for the analysis below.

DEFINITION: For a given \hat{h}^{t-1} , $\hat{h}_{t-1}^{t_t} \equiv (\hat{h}^{t-1}, (\hat{y}_t, \iota_t, \hat{s}_{t-1}^{t_t}))$ is a track record up to t with ι_t as the verification indicator in period t .

DEFINITION: We label h^{t_t} a full-compliance track record if it contains $\hat{y}_\tau = y_\tau$ and $\hat{s}_{\tau-1}^{t_\tau} = s_{\tau-1} \forall \tau = 1, \dots, t$.

Thus, a full-compliance track record involves both truth-telling and not being caught breaching the storage agreement in any period. The set of all full-compliance

⁸Note that due to the possibility of being investigated and penalised, the actual rate of return on storage may be uncertain, and its expected value therefore smaller than R .

track records is denoted by $H^t \subset \hat{H}^t$.

Furthermore, given a track record \hat{h}^{t-1} up to period t denote by \hat{h}_t^τ the track record up to τ that consists of \hat{h}^{t-1} and the continuation track record from t onwards, i.e. $\hat{h}_t^\tau \equiv (\hat{h}^{t-1}, \hat{h}_t, \dots, \hat{h}_\tau)$.

Finally, define $p(y^t) \equiv \prod_{\tau=1}^t p(y_\tau)$ as the probability of endowment history y^t occurring.

A cost-minimising principal offering a contract (transfer mechanism) that promises the agent an expected utility U solves the optimisation problem (OP)

$$\min_{\{b_t(\cdot), s_t(\cdot), \pi_t(\cdot), S_t(\cdot)\}_{t=1}^T} \sum_{t=1}^T \sum_{y^t \in \Theta^t} \sum_{\iota_t \in \{0,1\}} \left[\pi^{\iota_t}(h^{t-1}) R^{-(t-1)} [b_t(h^{\iota_t}) + \iota_t \gamma] \right] p(y^t) \quad (2.3)$$

subject to the set of constraints

$$\sum_{t=1}^T \sum_{y^t \in \Theta^t} \sum_{\iota_t \in \{0,1\}} \beta^{t-1} \pi^{\iota_t}(h^{t-1}) u[y_t + b_t(h^{\iota_t}) + R s_{t-1}(y^{t-1}, h^{t-1}) - s_t(y^t, h^{\iota_t})] p(y^t) \geq U; \quad (2.4)$$

$$\sum_{t=1}^T \sum_{y^t \in \Theta^t} \sum_{\iota_t \in \{0,1\}} \beta^{t-1} \pi^{\iota_t}(h^{t-1}) u[y_t + b_t(h^{\iota_t}) + R s_{t-1}(y^{t-1}, h^{t-1}) - s_t(y^t, h^{\iota_t})] p(y^t) \geq \quad (2.5)$$

$$\max_{\{\hat{s}_t, \hat{y}_t\}_{t=1}^T} \sum_{t=1}^T \sum_{y^t \in \Theta^t} \sum_{\iota_t \in \{0,1\}} \beta^{t-1} \pi^{\iota_t}(\hat{h}^{t-1}) u[y_t + b_t(\hat{h}^{\iota_t}) + R s_{t-1}(y^{t-1}, \hat{h}^{t-1}) - s_t(y^t, \hat{h}^{\iota_t})] p(y^t)$$

subject to

$$y_t + b_t(\hat{h}_t) + R s_{t-1}(y^{t-1}, h^{t-1}) - s_t(y^t, \hat{h}_t) \geq 0, \forall t \in \mathcal{T},$$

$$\hat{s}_0 = 0, \hat{s}_t(y^t, \hat{h}_t) \geq 0, \forall t \in \mathcal{T};$$

$$y_t + b_t(h_t) + R s_{t-1}(y^{t-1}, h^{t-1}) - s_t(y^t, h_t) \geq 0, \quad \forall t \in T, \quad (2.6)$$

$$\sum_{t=1}^T \sum_{y^t \in \Theta^t} \sum_{\iota_t \in \{0,1\}} [b_t(h^{\iota_t}) + S_t(h^{\iota_t}) - R S_{t-1}(h^{t-1}) + \pi_t(h^{t-1}) \gamma] p(y^t) \geq 0, \forall t \in \mathcal{T}, \quad (2.7)$$

$$s_t \geq 0, \quad \forall t \in \mathcal{T}, \quad (2.8)$$

$$S_0 = s_0 = 0. \quad (2.9)$$

$$S_T \geq 0. \quad (2.10)$$

Thus, the principal seeks to offer contracts that minimise the sum of discounted transfer payments and investigation costs, subject to a number constraints. The promise-keeping constraint (2.4) guarantees the agent an ex-ante expected utility level U . The incentive-compatibility constraint (2.5) not only induces agents to report their incomes truthfully, but also requires them comply with the contractual storage prescriptions. Expression (2.6) excludes negative consumption, and (2.7) is the principal's feasibility constraint. Given that we do not impose a non-negativity constraint on the principal's storage (except for the final period, as indicated by (2.10)), (2.7) can be trivially satisfied for periods 1 through $T - 1$ by simply storing negative amounts if necessary. Requirement (2.8) constrains agent's storage to be non-negative, and (2.9) says that neither the principals nor the agents have any goods to begin with.

3 Incentive Compatibility

Optimisation problem OP presented at the end of the previous section appears cumbersome. In this section, we will characterise incentive-compatibility constraints in a way that makes the problem more tractable.

To this end we will first adapt a result from Cole and Kocherlakota (1997) to show that we do not lose any generality by focusing on contracts that prescribe zero storage for agents in each period. Assume a given equilibrium implements a consumption allocation that prescribes positive storage for some agents in some periods. By adjusting transfers appropriately we can then establish an identical consumption allocation in which the storage level are always prescribed to be zero.⁹

PROPOSITION 3.1: *For any contract Γ^0 that solves OP there is another contract Γ^1 which also solves OP and has $s_t^1 = 0$ (zero functions) $\forall t \in \mathcal{T}$.*

Proof: Let $\Gamma^0 \equiv \{b_t^0, s_t^0, \pi_t^0, S_t^0\}_{t=1}^T$ be a solution to the original problem. Consider a modified mechanism with transfer, private and public storage, as well as investigation probability functions

$$b_t^1 \equiv b_t^0 - s_t^0 + R s_{t-1}^0,$$

⁹The solution to the program in the previous section is generally not unique because there may be many efficient allocations with positive private storage yielding the same consumption stream.

$$s_t^1 \equiv 0,$$

$$S_t^1 \equiv S_t^0 + s_t^0,$$

and

$$\pi_t^1 \equiv \pi_t^0,$$

$\forall t \in \mathcal{T}$, respectively.

Thus, for each contingency transfers are reduced by the net storage prescribed by the old contract, and all new prescribed storage levels are set to zero. Investigation probabilities remain unchanged and new public storage are simply augmented by old private savings.

Define as $\hat{s}_t^1 \equiv \hat{s}_t^0 - s_t^0$ the agent's actual period- t storage under the new mechanism. It is then easily checked that replacing $\{b_t^0, s_t^0, \pi_t^0, S_t^0\}_{t=1}^T$ with the new mechanism $\{b_t^1, s_t^1, \pi_t^1, S_t^1\}_{t=1}^T$ in the original problem does not change the principal's profit, nor does it affect the constraints (2.4) and (2.6) through (2.10). As for the incentive-compatibility constraint (2.5), note that the left-hand side is not altered either. The optimisation problem on the right-hand side of (2.5) takes into account two types of deviations from the contractual agreement, i.e. the misreporting of income and the non-adherence to prescribed storage. Assume that, when switching from Γ^0 to Γ^1 , there is a period $\tau \in \mathcal{T}$ in which some deviation of either type increases the right-hand side of (2.5). That is, for a given $\hat{h}^{\tau-1}$, there is an announcement/storage choice $(\hat{y}_\tau^1, \hat{s}_\tau^1)$ such that the expression

$$\sum_{t=1}^T \sum_{y^t \in \Theta^t} \sum_{\iota_t \in \{0,1\}} \beta^{t-1} \pi^{\iota_t}(\hat{h}^{t-1}) u[y_t + b_t(\hat{h}^{\iota_t}) + R\hat{s}_{t-1}(y^{t-1}, \hat{h}^{t-1}) - \hat{s}_t(y^t, \hat{h}^{\iota_t})] p(y^t)$$

is higher than under the original mechanism. But $(\hat{y}_\tau^0, \hat{s}_\tau^0)$ could have been chosen to equal $(\hat{y}_\tau^1, s_\tau^1 + \hat{s}_\tau^1)$. Then the same levels of consumption and therefore the same instantaneous utilities would have been achieved under the old mechanism. Hence, $\{b_t^0, s_t^0, \pi_t^0, S_t^0\}_{t=1}^T$ cannot have been a solution to the original problem, which contradicts our initial assumption. \square

In contrast to most other models within this literature, the present problem does not allow for a formulation that is recursive in the continuation utility of the agents. Unless monitoring of inventories occurs with certainty there is always the possibility that agents store goods secretly. Thus, their preferences over

future transfer streams are not common knowledge. This problem is caused by the presence of a hidden state variable we simplify the problem in a different fashion.¹⁰

Income misrepresentation may induce an agent to deviate from the contractually prescribed zero-storage level, both in the current as well as in consecutive periods. This is because an agent under-reporting his income will receive a transfer too high to be consumed immediately. Instead he prefers to smooth out expected consumption across time periods. In order to do so it will be necessary for him to store part of the resources. The presence of inventories affects the way future income and transfer streams are valued by an agent.

Before proceeding, we need to introduce some additional notation: For any given track record \hat{h}^{t-1} , denote as

$$h_{t-1}^{\iota_\tau} \equiv \left(\hat{h}^{t-1}, (y_t, \iota_t, 0), (y_{t+1}, \iota_{t+1}, 0), \dots, (y_\tau, \iota_\tau, 0) \right)$$

a track record consisting of \hat{h}^{t-1} and the continuation that has $\hat{y}_n = y_n$ and $\hat{s}_{n-1} = s_{n-1} = 0$, $n = t, \dots, \tau$. In this case, there is truth-telling and adherence to zero-storage prescriptions *from t onwards*. Define further, for period- t income reports $y_t \in \Theta$,

$$\hat{h}_{t_0}^{\iota_\tau} \equiv \left(\hat{h}^{t-1}, (\hat{y}_t, \iota_t, 0) \right),$$

$\iota_t = 0, 1$, which refers to the case where an agent stores zero in $t - 1$ ($\hat{s}_{t-1} = 0$);

$$\hat{h}_{t_+}^0 \equiv \left(\hat{h}^{t-1}, (\hat{y}_t, 0, 0) \right),$$

which deals with the case where the agent decides to store a strictly positive amount in $t - 1$ but is not investigated in t ($\hat{s}_{t-1} > 0$ and $\hat{s}_{t-1}^t = 0$); and

$$\hat{h}_{t_+}^1 \equiv \left(\hat{h}^{t-1}, (\hat{y}_t, 1, \hat{s}_{t-1}) \right),$$

the track record when the agent decides to store a strictly positive amount in $t - 1$ and is caught in t ($\hat{s}_{t-1} > 0$ $\hat{s}_{t-1}^t = \hat{s}_{t-1}$).

¹⁰The problem of a hidden state variable also appears when there is serial correlation in income shocks. Nevertheless, as long as the memory of the stochastic process is finite, the resulting problem can, in principle, still be formulated recursively. See Fernandes and Phelan (1999) for details.

For convenience we also define

$$U_{t+1}(\hat{h}^t, \hat{s}_t) \equiv \max_{\{\hat{s}_\tau\}_{\tau=t+1}^T} \sum_{\tau=t+1}^T \sum_{y_{t+1}^\tau \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} \beta^{\tau-(t+1)} \pi^{\iota_\tau}(\hat{h}^{\tau-1}) \times \\ u \left[y_\tau + b_\tau(\hat{h}_t^{\tau-1}, (\hat{y}_\tau, \iota_\tau, \hat{s}_{\tau-1}^{\iota_\tau})) + R\hat{s}_{\tau-1} - \hat{s}_\tau \right] \frac{p(y_t^\tau)}{p(y^t)},$$

where $\hat{h}_t^t \equiv \hat{h}^t$.

This expression represents the continuation utility to be expected from the contract from $t + 1$ onwards, given a track record \hat{h}^t up to t , and actual period- t storage \hat{s}_t , conditional on the agent telling the truth about all income realisations from $t + 1$ onwards, but storing optimally (i.e. not necessarily complying with storage prescriptions).

The definition enables us to write the incentive-compatibility constraint (2.5) as a sequence of temporary incentive-compatibility constraints. That is, $\forall y^{t-1}$, \hat{h}^{t-1} , and \hat{s}_{t-1} ,

$$\sum_{\tau=t}^T \sum_{y_t^\tau \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} \beta^{\tau-t} \pi^{\iota_\tau}(\hat{h}^{\tau-1}) u \left[y_\tau + b_\tau(h_{t-1}^{\iota_\tau}) \right] \frac{p(y_{t-1}^\tau)}{p(y^{t-1})} \geq \quad (3.11)$$

$$\max_{\hat{s}_t, \hat{y}_t} \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t}(\hat{h}^{t-1}) u \left[y_t + b_t(\hat{h}^{\iota_t}) + R\hat{s}_{t-1} - \hat{s}_t \right] + \beta U_{t+1}(\hat{h}^t, \hat{s}_t).$$

This sequence of constraints implies that truth-telling combined with zero storage in both current and future periods must not be dominated by lying today and truth-telling in the future, combined with individually optimal storage in all future periods.

Lemma 3.2 in the Appendix shows that the sequence of constraints, (3.11) is equivalent to the overall constraint (2.5).

We reduce the complexity of the set of incentive-compatibility constraints further by identifying a number of equilibrium properties of the transfer functions. Specifically, we show that their expected value is decreasing in the current report, which is the basis for demonstrating that investigation probabilities are decreasing in income reports as well.

PROPOSITION 3.3: *In equilibrium, given a track record \hat{h}^{t-1} up to period $t-1$, the expected period- t transfer,*

$$\sum_{\iota_t \in \{0,1\}} \pi^{\iota_t}(h^{t-1})b_t(\hat{h}_{t-1}^{\iota_t}),$$

is decreasing in the current income report, \hat{y}_t .

Proof: See Appendix.

Furthermore, for any income history y^{t-1} and track record \hat{h}^{t-1} , the net present value of transfers from t through T cannot be higher for an agent announcing θ_i in t than that of an agent announcing θ_j , if $\theta_i > \theta_j$. In other words, in equilibrium, and with R as the (gross) discount rate,

$$\begin{aligned} \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t}(h^{t-1})b_t(\hat{h}_{ti}^{\iota_t}) + \sum_{\tau=t+1}^T \sum_{y_i^\tau \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} R^{-(\tau-t)} \pi^{\iota_\tau}(\hat{h}_{ti}^{\tau-1})b_\tau(\hat{h}_{ti}^{\iota_\tau}) \frac{p(y_{t-1}^\tau)}{p(y^{t-1})} \geq \\ \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t}(h^{t-1})b_t(\hat{h}_{ti}^{\iota_t}) + \sum_{\tau=t+1}^T \sum_{y_i^\tau \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} R^{-(\tau-t)} \pi^{\iota_\tau}(\hat{h}_{tj}^{\tau-1})b_\tau(\hat{h}_{ti}^{\iota_\tau}) \frac{p(y_{t-1}^\tau)}{p(y^{t-1})}. \end{aligned}$$

If this condition is violated, none of the ‘downward’ incentive-compatibility constraints is binding, as someone who truthfully announces a higher income not only receives higher net present value of transfers but, due to Proposition 3.3, also enjoys smoother consumption than he would through income misrepresentation. But then the principal could increase its expected payoff by exploiting the strict concavity of the per-period utility functions. Lowering all transfers for the relatively higher income announcement and raising those for the relatively lower ones by a somewhat smaller amount would leave the promised utility U unchanged, and would not upset the other constraints either.

In contrast to frameworks in which there is no storage (and, more generally, in standard adverse-selection models), it is not the case here that, in any given period, only the ‘downward adjacent’ incentive-compatibility constraints are binding.

To clarify this point, we will split up the temporary incentive-compatibility constraints (3.11) into two components. For a track record \hat{h}^{t-1} , the first component consists of all the ‘pure truth-telling’ constraints without storage:

$$\sum_{\tau=t}^T \sum_{y_i^\tau \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} \beta^{\tau-t} \pi^{\iota_\tau}(\hat{h}^{\tau-1}) u \left[y_\tau + b_\tau(h_{t-1}^{\iota_\tau}) \right] \frac{p(y_{t-1}^\tau)}{p(y^{t-1})} \geq \quad (3.12)$$

$$\max_{\hat{y}_t} \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[y_t + b_t(\hat{h}^{t-1}, (\hat{y}_t, \iota_t, \hat{s}_{t-1}^{\iota_t})) \right] \frac{p(y_{t-1}^{\iota_t})}{p(y^{t-1})} + \beta U_{t+1}(\hat{h}_{t_0}^{\iota_t}, 0),$$

for each $t \in \mathcal{T}$, and each $y_t \in \Theta$.

The second component is made up of the conditions that prevent misrepresentation of income and *strictly* positive storage from being superior to telling the truth and zero storage:

$$\sum_{\tau=t}^T \sum_{y_\tau^i \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} \beta^{\tau-t} \pi^{\iota_\tau} (\hat{h}^{\tau-1}) u \left[y_\tau + b_\tau(h_{t-1}^{\iota_\tau}) \right] \frac{p(y_{t-1}^{\iota_\tau})}{p(y^{t-1})} \geq \quad (3.13)$$

$$\max_{\hat{y}_t, \hat{s}_t} \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[y_t + b_t(\hat{h}^{t-1}, (\hat{y}_t, \iota_t, \hat{s}_{t-1}^{\iota_t})) + R\hat{s}_{t-1} - \hat{s}_t \right] \frac{p(y_{t-1}^{\iota_t})}{p(y^{t-1})}$$

$$+ \beta U_{t+1}(\hat{h}_{t+}^{\iota_t}, \hat{s}_t),$$

for each $t \in \mathcal{T}$, and each $y_t \in \Theta$.

A deviation from truth-telling entails two different actions by the agent. Either the agent consumes everything immediately and stores nothing, in which case he has nothing to fear from any investigation; or he stores optimally to spread the benefit of additional funds received from the intermediary over several periods, with the risk of losing it if being caught. The first case represents a corner solution to the optimal-storage problem, whereas the latter implies an internal solution to it.

Lemma 3.4 in the Appendix shows that for the first component, (3.12), only the ‘downward adjacent’ incentive-compatibility constraints are binding.

We cannot make the analogous claim for the second component of the set of incentive-compatibility constraints, (3.13). In some cases the constraint preventing an agent with current income θ_i from reporting θ_{i-k} , $k = 2, 3, \dots, i-1$ and storing a *positive* amount, is binding. We will see this possibility illustrated in the numerical example in Section 4.

Finally, we observe that none of the upward constraints is binding. This is true because, as we have seen above, (i) the NPV of transfers to an agent reporting a higher current income cannot exceed the NPV of transfers made to an agent reporting a lower income, and (ii) the expected current-period transfer is decreasing in the income reported. That implies that if an agent reports an income higher

than the true realisation he obtains not only less in an NPV sense, but the transfer stream is such that intertemporal consumption is also less smooth.

3.1 The Investigation Strategy

In this subsection we analyse the optimal investigation policy. We first show that whenever an agent announces the highest possible level as his income realisation, it is optimal for the intermediary not to inspect that agent's storage. The reason is that the transfer in this case is such that the agent would not choose to store (additionally) anyway.¹¹ It would therefore be wasteful for the principal to monitor under those circumstances.

More formally, note that an agent with a track record \hat{h}^{t-1} and (truthfully stated) income θ_n in period t faces an investigation probability of $\pi^1(\hat{h}^{t-1})$ in the next period. An optimal storage level in period t satisfies the first-order condition

$$\sum_{y_{it}^{t+1} \in \Theta^{t+1}} \sum_{\iota_{t+1} \in \{0,1\}} \beta \pi^{\iota_{t+1}}(\hat{h}_{t-1}, (\theta_n, \iota_t, \hat{s}_{t-1}^{\iota_t})) u' [y_{t+1} + b_{t+1}(h_{t-1}^{\iota_{t+1}})] \frac{p(y_{it}^{t+1})}{p(y^{t-1}, \theta_n)} =$$

$$\sum_{\iota_{t+1} \in \{0,1\}} \pi^{\iota_{t+1}}(\hat{h}^{t-1}) u' [\theta_n + b_t(\hat{h}^{t-1}, (\theta_n, \iota_t, \hat{s}_{t-1}^{\iota_t}))] \frac{p(y_{t-1}, \theta_n)}{p(y^{t-1})},$$

where we have used the definition $y_{in}^{t+1} \equiv (y_{t-1}, \theta_n, y_{t+1})$.

We have seen previously, that transfers can always be constructed such that an equilibrium involves zero storage. Assume now that there is such an equilibrium, but with

$$\pi^1(\hat{h}_{t-1}, (\theta_n, 1, \hat{s}_{t-1}^1)) > 0$$

in $t + 1$. Reducing this probability would induce an increase in private storage, as it makes being caught less likely. But this effect can again be undone by an adjustment of transfers, i.e. a reduction in $b_t(\cdot)$ and an increase in $b_{t+1}(\cdot)$. If carried out appropriately, this intertemporal reallocation does not change the NPV of transfers, nor does it upset the incentive-compatibility constraints. The latter is true because for agents with $y_t \in \{\theta_1, \theta_2, \dots, \theta_{n-1}\}$ the new allocation (transfer scheme) for agents with the highest income is even less attractive than the original

¹¹Of course, this does not necessarily mean that such an agent has no storage at all. In principle he may have carried goods into the present from previous periods.

one. At the same time, however, the principals's payoff is increased as his expected cost is lower. Hence, the original contract cannot constitute an equilibrium.

Furthermore, one can show that $\pi^1(\hat{h}_{t-1}, (y_t, 1, \hat{s}_{t-1}^1))$ is nonincreasing in y_t . Consider an agent with a period- t income realisation $y_t = \theta_i$ and recall that

$$b_t(\hat{h}^{t-1}, (\theta_i, \iota_t, \hat{s}_{t-1}^{\iota_t})) < b_t(\hat{h}^{t-1}, (\theta_{i-1}, \iota_t, \hat{s}_{t-1}^{\iota_t})).$$

For a given expected rate of return the willingness to store is therefore stronger when the agent has lied downward (has reported θ_{i-1}) than when he has reported θ_i truthfully. This is because the intertemporal marginal rate of substitution increases in y_t and decreases in \hat{s}_t . But this implies that for the Euler equation (??) to hold, the expected rate of return must not be higher for agents reporting a lower income. This requires that the probability of investigating in $t+1$ be nonincreasing in period- t income.

3.2 Carrots and Sticks

We now address the question how the transfers are affected by the findings of an investigation, if it takes place. If an agent is caught deviating from the zero-storage policy there should be a penalty in the form of a lower transfer to him, i.e. for all \hat{h}^{t-1} and all period- t reports \hat{y}_t ,

$$b_t(\hat{h}^{t-1}, (\hat{y}_t, 1, a)) < b_t(\hat{h}^{t-1}, (\hat{y}_t, 1, 0)),$$

where a is any strictly positive number. Otherwise there would be no reason to comply with the zero-storage restrictions.

It should also be clear that in equilibrium, if an agent is punished, the penalty takes on the highest possible value. According to constraint 2.6, the negative transfer in a particular period is limited to the sum of the storage detected and the minimum income θ_1 . Penalties that go beyond that limit violate the non-negativity condition on consumption.

On the other hand, a penalty set below this limit cannot be optimal. To see this, increase the penalty (i.e. decrease the transfer) by a small amount, and at the same time reduce the investigation probability such that the net effect on the right-hand side of the incentive-compatibility constraint (2.5) is zero. Thus, without altering the feasibility set one can lower the expected investigation cost for the principal and hence increase his profit.

Perhaps a more interesting question is whether an agent, if found to have complied with the zero-storage rule, should be rewarded with a transfer that is higher than in the case in which there is simply no verification at all. Clearly, this point is relevant only for those income reports for which the probability of investigation is positive. An argument similar to the one made above for the case of a penalty shows that when there is a positive probability of investigation, a reward should be given to an agent found to be without storage.

Suppose first that

$$b_t(\hat{h}^{t-1}, (\hat{y}_t, 1, 0)) < b_t(\hat{h}^{t-1}, (\hat{y}_t, 0, 0)),$$

for some \hat{h}^{t-1} and \hat{y}_t , which would imply that the agent is actually punished when found not storing anything. The principal could increase his profit by reducing the risk the agent is exposed to. This is achieved by increasing the LHS by a small amount and decreasing the RHS by a somewhat larger amount, such that the expected utility of the agent is unchanged.

If $b_t(\hat{h}^{t-1}, (\hat{y}_t, 1, 0)) = b_t(\hat{h}^{t-1}, (\hat{y}_t, 0, 0))$, the principal can increase the LHS slightly, at the same time reducing the investigation probability. As long as the changes are appropriate this will not affect the agent's expected utility. It does, however, reduce the expected cost of the principal who is therefore made better off.

3.3 Competitive Behaviour

So far we have not given any consideration to the assumption that the contract is budget-neutral, or that the principal behaves competitively. As we will see in this subsection, this requirement will pin down the ex-ante expected utility attained by the agents.

A solution to the optimisation problem does not necessarily exist for every U . Define \mathcal{U} as the set of values of U for which an equilibrium exists. For each $U \in \mathcal{U}$, solving the above problem yields a specific equilibrium cost for the intermediary. That is, one can specify a function $C : \mathcal{U} \rightarrow \mathbb{R}$ that assigns an expected cost to each utility level promised to the agent. It is easily shown that for such an environment C is an increasing function. Let us assume that \mathcal{U} is convex and has the property that $C(\inf \mathcal{U}) \leq 0$, $C(\sup \mathcal{U}) \geq 0$. Budget neutrality is equivalent to

the requirement that $C(U)$ equals zero. Hence, the agent can obtain the ex-ante expected utility U^* characterised by $C(U^*) = 0$.

4 A Numerical Example

In this section we illustrate the general model by employing a two-period framework similar to that in Green and Oh (1991). We will also compute the efficient allocation in a pure debt-securities equilibrium as well as that in an equilibrium where goods are not storable.

Apart from setting $T = 2$ we simplify matters further by assuming that there are only three (i.i.d.) income realisations in the first period.¹² With probability p^h the income is θ^h , with probability p^m it is θ^m , and with probability $p^l \equiv 1 - p^h - p^m$ income θ^l is obtained, $\theta^h > \theta^m > \theta^l \geq 0$. In a slight but inconsequential deviation from the general model, we assume that there is a single (non-random) income realisation in period 2, denoted by z .

The report-contingent transfers for period 1 are given by b_1^h , b_1^m and b_1^l for income announcements θ^h , θ^m and θ^l , respectively. Given an income report of θ^i , for period 1, $i = h, m, l$, the contingent transfers for period 2 are b_2^i if storage is not verified, and \tilde{b}_2^i if it is verified to be zero. We remind the reader that in the case where an agent is found to have positive storage in the second period, the intermediary can confiscate the agent's entire wealth consisting of storage and income. In other words, the transfer in this case would be $-(s+z)$, with s denoting storage carried over from period 1, such that the period-2 consumption becomes zero.

We specify the instantaneous-utility function as the exponential,

$$u(c_t) = -e^{-rc_t},$$

where the constant r is the Arrow-Pratt measure of absolute risk aversion.

Define $Y^i \equiv e^{-r\theta^i}$ and $Z \equiv e^{-rz}$, and the utility derived from the transfers as $B_1^i \equiv e^{-rb_1^i}$, $B_2^i \equiv e^{-rb_2^i}$, and $\tilde{B}_2^i \equiv e^{-r\tilde{b}_2^i}$, $i = h, m, l$.

¹²A three-period example has also been analysed. It appears that the results do not differ qualitatively from the ones presented here. Therefore, and due to the fact that one has 39 variables, we only present a two-period model here.

We let $\pi^i \in [0, 1]$ denote the probability with which the storage of an agent who has reported a first-period income of θ^i , will be investigated at the beginning of the second period. Moreover, we define as \hat{s} being the amount an agent decides to carry over from period 1 to period 2, and $\hat{S} \equiv e^{-r\hat{s}}$. Note that in equilibrium $\hat{s} = 0$ and therefore $\hat{S} = 1$.

The principal solves the following problem:

For a given utility level U minimise w.r.t. $\{B_1^i, B_2^i, \tilde{B}_2^i, \pi_i\}_{i=h,m,l}$, the expected cost,

$$-\sum_i p^i \left[r^{-1} \left(\ln B_1^i + R^{-1} \left(\pi^i \ln \tilde{B}_2^i + (1 - \pi^i) \ln B_2^i \right) \right) + \pi^i \gamma \right] \quad (4.14)$$

subject to the constraints

$$\sum_i p^i \left[Y^i B_1^i + R^{-1} \left(\pi^i Z \tilde{B}_2^i + (1 - \pi^i) Z B_2^i \right) \right] = -U, \quad (4.15)$$

$$Y^i B_1^i + \beta \left(\pi^i Z \tilde{B}_2^i + (1 - \pi^i) Z B_2^i \right) \leq \quad (4.16)$$

$$\min \left\{ Y^i B_1^i + \beta \left(\pi^k Z \tilde{B}_2^k + (1 - \pi^k) Z B_2^k \right); \min_{\hat{S} > 1} \frac{Y^i B_1^k}{\hat{S}} + \beta \left(\pi^k + (1 - \pi^k) \hat{S}^R Z B_2^k \right) \right\},$$

$k \neq i$. The right-hand side of the downward incentive-compatibility constraint (4.16) again reflects the two possible types of deviations from the contractual prescriptions. The first argument refers to the case where an agent pretends to have received a lower income than he actually did, but does not store any goods. The second argument deals with the case in which a high-income agent not only lies about his income but also stores optimally, thereby risking to get caught in the second period. If income level θ_k has been declared, this would mean losing storage and second-period income z with probability π^k .

The problem of finding an optimal storage level in the latter case is straightforward in our two-period example. The first-order condition for the minimisation problem can be written as

$$(1 - \pi^k) Z B_2^k \hat{S}^R = \frac{Y^i B_1^k}{\hat{S}}.$$

Due to the concavity of the objective function it is both necessary and sufficient

for an optimum. It can be solved for \hat{S} to yield

$$\hat{S} = \left(\frac{Y_1^i B_1^k}{(1 - \pi^k) Z B_2^k} \right)^{\frac{1}{R+1}}.$$

From the first-order condition we can replace

$$(1 - \pi^k) \hat{S}^R Z B_2^k$$

in the second argument on the right-hand side of (4.16) with

$$\frac{Y^i B_1^k}{\hat{S}}.$$

The objective function of the minimisation problem embedded in (4.16) then becomes

$$(1 + \beta) \frac{Y^i B_1^k}{\hat{S}} + \beta \pi^k,$$

which, after using the solution for \hat{S} , becomes

$$(1 + \beta) \left[(Y^i B_1^k)^{\frac{R}{R+1}} ((1 - \pi^k) Z B_2^k)^{\frac{1}{R+1}} \right] + \beta \pi^k.$$

It is convenient to split up (4.16) into its two components, i.e.

$$Y^i B_1^i + \beta (\pi^i Z \tilde{B}_2^i + (1 - \pi^i) Z B_2^i) \leq Y^i B_1^k + \beta (\pi^k Z \tilde{B}_2^k + (1 - \pi^k) Z B_2^k), \quad (4.17)$$

and

$$Y^i B_1^i + \beta (\pi^i Z \tilde{B}_2^i + (1 - \pi^i) Z B_2^i) \leq (1 + \beta) \left[(Y^i B_1^k)^{\frac{R}{R+1}} ((1 - \pi^k) Z B_2^k)^{\frac{1}{R+1}} \right] + \beta \pi^k. \quad (4.18)$$

After having eliminated the optimisation problem for the agent we can minimise, for each given U , (4.14) subject to (4.15), (4.17) and (4.18). We do so numerically, and then pick the value of U for which the expected cost is equal to zero (budget neutrality).

In Table 1 we present the results for this costly-verification model, and compare them with those obtained from the two other frameworks with asymmetric information mentioned in the introduction. The first is the one in which goods are perishable or storage perfectly monitorable ('no storage'), and in the second verification costs are prohibitively high ('no verification').

The following parameter values used:

$r = 1$, $\beta = 0.95$, $\theta^l = 1$, $\theta^m = 2$, $\theta^h = 3$, $z = 2$, $p^h = 0.1$, $p_m = 0.8$, $p_h = 0.1$, $\gamma = 0.8$.

We also assume that the rate of return on storage is equal to the rate of time preference, i.e. $R = \beta^{-1}$.

We make the following observations. As expected, in the model with costly verification of storage the agent reaches an ex-ante utility level that lies between the (higher) utility obtained in the no-storage model and the (lower) utility obtained in the framework without storage verification.

The ex-post utility level for an agent who has received (and truthfully reports) period-1 income θ^i is denoted by U^i . We see that, whilst there is relatively little difference for receivers of the medium income, the low-income type in the no-verification world suffers from a utility level which is lower than that of his equivalents in the other two models. On the other hand the prohibitively high cost of storage verification benefits high-income agents. The reason is that they will not have to contribute towards the intratemporal redistribution of resources (income insurance), because such a scheme cannot be truthfully implemented. The complete absence of insurance cover is also reflected by the fact the *NPV* of transfers is zero for all three income types, which is a feature that is not shared by the other two models.

The verification of storage also has implications for the intertemporal smoothing of consumption. In all three models the consumption of an agent with the highest income is smoothed perfectly. In other words, the Euler equation holds for agents with the highest incomes. In the framework with costly verification this is due to the fact established earlier in the paper that there is no monitoring of storage for someone who reports θ^h . More generally, it is another illustration of the ‘no-(efficiency)-distortion-at-the-top’ result that is well known from adverse-selection models.

In the model with costly verification the medium-income agent also enjoys perfect consumption smoothing. Again, the reason for this is that he is never investigated, which also accounts for the result that both the medium and the high incomes have the same (negative) *NPV* of transfers. Hence, in this model, and for the given parameters, it is an efficient outcome that the burden of ‘subsidising’ the low-income agent is the same for the other two types.

	Costly Verification	No Storage	No Verification
U	-0.26935	-0.26753	-0.27100
U^h	-0.161254	-0.181293	-0.158028
U^m	-0.269293	-0.265939	-0.263906
U^l	-0.377904	-0.366496	-0.440721
b_1^h	-0.5073	-0.6250	-0.4866
b_1^m	-0.0202	0.0106	0
b_1^l	0.7613	0.8383	0.4874
b_2^h	0.4926	0.3760	-0.5122
b_2^m	-0.0265	-0.0202	0
b_2^l	-0.4771	-0.4782	-0.5130
\tilde{b}_2^m	0.1209	NA	NA
\tilde{b}_2^l	-0.3697	NA	NA
c_1^h	2.4927	2.3750	2.5130
c_1^m	1.9798	2.0106	2
c_1^l	1.7613	1.8384	1.4870
c_2^h	2.4926	2.3750	2.5130
c_2^m	1.9798	1.9734	2
c_2^l	1.5228	1.5217	1.4870
\tilde{c}_2^m	NA	NA	NA
\tilde{c}_2^l	1.6303	NA	NA
π^m	0	NA	NA
π^l	0.0524	NA	NA
NPV^h	-0.0394	-0.2678	0
NPV^m	-0.0394	-0.0146	0
NPV^l	0.3133	0.3840	0

Table 1: Comparison of allocations in three different asymmetric-information frameworks

Also note that NPV^l obtained by the agent with θ^l does not appear to be much lower than what he would receive in a world without storage, despite the fact that the chance of being investigated is only just over 5%. Thus, even a small threat of being investigated (albeit with a large penalty in the case of being caught) is sufficient to re-establish a considerable part of the income-insurance feature exhibited by efficient contracts in a world without storage. An additional finding, which we do not report, is that even large variations in the storage cost do not affect the verification probability substantially.

This example shows the possibility that only agents reporting a low income face a positive probability of being investigated. However, this outcome is by no means a certainty. Consider the following change of parameters. We increase each of the probabilities of receiving a high or a low income from 0.1 to 0.2 and reduce the occurrence of medium incomes from 0.8 to 0.6. We leave all other parameters unchanged such that the NPV of incomes is unaffected.

The values taken by the most important variables under the new parameter set are reported in Table 2, where we confine ourselves to the model with costly storage verification.

The most drastic change is that there is now even a positive probability that an agent reporting medium income is investigated.¹³ As a consequence, there is a slight distortion in the intertemporal consumption allocation for the medium-income agent. He is induced to consume more in the first period than in the second period, unless he obtains a reward of 0.0740 with a probability of 0.0236. This relatively small probability of investigation now leads to a sizeable discrepancy between NPV^h and NPV^m .

We conclude this example with a few remarks regarding implied tax rates. In both the costly-verification and the no-storage allocation of Table 1 the high-income and the medium-income agents' incomes are taxed (on an NPV basis), i.e. they receive negative net transfers. It is interesting to compare the tax rates τ^h and τ^m (NPV of tax payments as a share of NPV of total incomes) for the two types. In the no-storage model $\tau^h = 0.055$ and $\tau^m = 0.004$, which implies that efficiency requires progressive taxation. In contrast, for the costly-verification model one obtains $\tau^h = 0.008$ and $\tau^m = 0.010$, which means that in such a world

¹³Note that this probability is, however, smaller than that for a low-income agent, as it should be according to the analysis in Section 3.

U	-0.27410
c_1^h	2.3690
c_1^m	1.9768
c_1^l	1.7708
c_2^h	2.3690
c_2^m	1.9751
c_2^l	1.5181
\tilde{c}_2^m	2.0491
\tilde{c}_2^l	1.5801
π^m	0.0236
π^l	0.05419
NPV^h	-0.2804
NPV^m	-0.0452
NPV^l	0.3162

Table 2: The allocation with costly verification under first-period income distribution $\theta^h = 0.2$, $\theta^m = 0.6$, $\theta^l = 0.2$.

a tax schedule which is regressive in parts can be allocationally efficient.

This does not seem to hold, however, under the changed income distribution relevant for the results in Table 2. There we obtain $\tau^h = 0.057$ and $\tau^m = 0.012$, which suggests that a progressive tax-rate schedule is optimal.

5 Concluding Remarks

We have studied an economy in which agents, who have private information about their incomes and about their use of a commonly available storage technology, are offered contracts by a government authority or competitive intermediaries to insure them against income shocks. If storage were not monitorable at all, agents would be able to ‘undo’ schemes that redistribute income from those with higher current income to those with lower current income. In other words, each individual receives transfers with a net present value of zero, which means that no proper insurance of incomes takes place.

However, making storage verifiable at a sufficiently low cost enables us to find equilibrium contracts where the promised net present value of transfers received is contingent on incomes reported by the agents. This requires that storage verification be carried out with positive probability for at least some income reports, and that penalties are imposed on agents who have been found to deviate from the contractual storage prescriptions.

It is shown that whilst verification takes place with positive probability for agents reporting lower income, this is not the case for those reporting higher incomes. As a consequence, those with higher incomes tend to make a lump-sum transfer in the reporting period and expect to use pure borrowing and lending arrangements in the future.

In contrast, agents that report lower incomes and whose storage is checked, expect to receive a sequence of transfers that has positive net present value, i.e. that goes beyond pure borrowing and lending.

The fact that an equilibrium which involves monitoring of storage is Pareto-superior (ex ante) to a pure debt-securities equilibrium implies that the use of means testing in benefit systems is a welfare-enhancing policy. The reason is that it allows for the implementation of a scheme that provides insurance against income

shocks.

The analysis also highlights the extent to which contracts offered by private-sector financial intermediaries may contribute towards an improved sharing of dynamic risks.

This paper has focused on the case of individuals' income insurance, but the framework presented should be applicable to risk-sharing problems in other contexts. It may, for example, be used in the design of optimal mutual reserve systems for financial intermediaries, such as banks or insurance companies.

One technical aspect to be addressed in future work is renegotiation-proofness of such contracts. Contract theorists are well aware of the fact that arrangements involving the costly verification of a state variable are prone to renegotiation once the assumption of full commitment has to be abandoned. In other words, when it comes to storage monitoring the intermediary may decide not to do so and save the cost that would be incurred. After all he knows that, on the equilibrium path, the agent would have complied with storage prescriptions anyway, and ex post it would therefore be inefficient to investigate. Of course, this inactivity would be foreseen by all agents who would now have less of an incentive to report truthfully, bringing the original equilibrium to a collapse.

Finally, the model could also be extended to incorporate risky assets in addition to the safe storage technology already available to agents. The purpose of such an exercise would be to study the implications of private information on income and storage, as well as the possibility to investigate the latter at a cost, on the prices of such assets.

Appendix

LEMMA 3.2 *The sequence of incentive-compatibility constraints, (3.11) is equivalent to the incentive-compatibility constraint (2.5).*

Proof: That the sequence (3.11) implies the overall constraint (2.5), can be shown by backward induction. Taking into account that, for obvious reasons, $\hat{s}_T = 0$ is the optimal storage choice in the final period, the incentive-compatibility constraint in $t = T$ is, for all \hat{h}^{T-1} , \hat{s}_{T-1} and y_T ,

$$\sum_{\iota_T \in \{0,1\}} \pi^{\iota_T} (\hat{h}^{T-1}) u \left[y_T + b_T(h_{T-1}^{\iota_T}) + R\hat{s}_{T-1} \right] \geq \max_{\hat{y}_T} \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[y_t + b_t(\hat{h}^{\iota_t}) + R\hat{s}_{T-1} \right].$$

This implies that truth-telling and zero-storage in the final period is weakly preferred to any deviation from this policy.

For period $t = T - 1$ the constraint becomes

$$\sum_{\tau=T-1}^T \sum_{y_{T-1}^{\tau} \in \Theta^{\tau}} \sum_{\iota_{\tau} \in \{0,1\}} \beta^{\tau-(T-1)} \pi^{\iota_{\tau}} (\hat{h}^{\tau-1}) u \left[y_{\tau} + b_{\tau}(h_{T-2}^{\iota_{\tau}}) \right] \frac{p(y_{T-1}^{\tau})}{p(y^{T-1})} \geq$$

$$\max_{\hat{s}_{T-1}, \hat{y}_{T-1}} \sum_{\iota_{T-1} \in \{0,1\}} \pi^{\iota_{T-1}} (\hat{h}^{T-2}) u \left[y_{T-1} + b_t(\hat{h}^{\iota_{T-1}}) + R\hat{s}_{T-2} - \hat{s}_{T-1} \right] + \beta U_T(\hat{h}^{T-1}, \hat{s}_{T-1}),$$

which says that truth-telling and zero-storage in both the penultimate and final period is weakly superior compared to lying or storing a positive amount in $T - 1$ and truth-telling plus zero-storage in T . The latter has already been shown to be no worse than lying and/or positive storage in the previous step. Hence, truth-telling and zero-storage in both $T - 1$ and T is at least as good as deviating from this policy in either one or both periods.

This line of argument can be repeated for $t = T - 2$ down to $t = 1$, thereby proving the statement in one direction.

The reverse, that if the overall constraint (2.5) holds the sequence of temporary constraints is satisfied as well, can also be shown. If this were not the case there exists a history after which deviation from truth-telling delivers a higher continuation utility than truth-telling. But this means that the overall constraint cannot be satisfied, thereby contradicting the original assumption. \square

Proof of Proposition 3.3:

Note that lying about one's income without ever storing is one of the strategies prevented by the incentive-compatibility constraint (2.5). Thus, if an allocation does not satisfy the 'no-storage' subset of constraints, it obviously cannot satisfy the larger set (2.5) itself. Incentive compatibility implies that, without storage, $\forall \hat{h}^{t-1}$,

$$\begin{aligned} & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_i + b_t(h_{t-1}^{\iota_t}) \right] \\ & + \sum_{\tau=t+1}^T \sum_{y_{ii}^\tau \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} \beta^{\tau-(t+1)} \pi^{\iota_\tau} (h_{t-1}^{\tau-1}) u \left[y_\tau + b_\tau(h_{t-1}^{\iota_\tau}) \right] \frac{p(y_{ii}^\tau)}{p(y^{ti})} \geq \\ & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_i + b_t(\hat{h}_{t-1}, (\theta_j, \iota_t, \hat{s}_{t-1}^{\iota_t})) \right] \\ & + \sum_{\tau=t+1}^T \sum_{y_{ii}^\tau \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} \beta^{\tau-(t+1)} \pi^{\iota_\tau} (\hat{h}_{t-1}^{\tau-1}) u \left[y_\tau + b_\tau(\hat{h}_{t-1}, (\theta_j, \iota_t, \hat{s}_{t-1}^{\iota_t}), \hat{h}_t^{\iota_\tau}) \right] \frac{p(y_{ii}^\tau)}{p(y^{ti})}, \end{aligned}$$

and

$$\begin{aligned} & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_j + b_t(h_{t-1}^{\iota_t}) \right] \\ & + \sum_{\tau=t+1}^T \sum_{y_{ij}^\tau \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} \beta^{\tau-(t+1)} \pi^{\iota_\tau} (h_{t-1}^{\tau-1}) u \left[y_\tau + b_\tau(h_{t-1}^{\iota_\tau}) \right] \frac{p(y_{ij}^\tau)}{p(y^{tj})} \geq \\ & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_j + b_t(\hat{h}_{t-1}, (\theta_i, \iota_t, \hat{s}_{t-1}^{\iota_t})) \right] \\ & + \sum_{\tau=t+1}^T \sum_{y_{ij}^\tau \in \Theta^\tau} \sum_{\iota_\tau \in \{0,1\}} \beta^{\tau-(t+1)} \pi^{\iota_\tau} (\hat{h}_{t-1}^{\tau-1}) u \left[y_\tau + b_\tau(\hat{h}_{t-1}, (\theta_i, \iota_t, \hat{s}_{t-1}^{\iota_t}), \hat{h}_t^{\iota_\tau}) \right] \frac{p(y_{ij}^\tau)}{p(y^{tj})}. \end{aligned}$$

Combining these two inequalities yields

$$\begin{aligned} & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_i + b_t(h_{t-1}^{\iota_t}) \right] - \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_i + b_t(\hat{h}_{t-1}, (\theta_j, \iota_t, \hat{s}_{t-1}^{\iota_t})) \right] \geq \\ & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_j + b_t(h_{t-1}^{\iota_t}) \right] - \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_j + b_t(\hat{h}_{t-1}, (\theta_i, \iota_t, \hat{s}_{t-1}^{\iota_t})) \right]. \end{aligned}$$

It can be shown that if $\theta_i > \theta_j$, it is necessary that

$$\sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_i + b_t(h_{t-1}^{\iota_t}) \right] < \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_j + b_t(h_{t-1}^{\iota_t}) \right].$$

□

LEMMA 3.4 *Out of the subset of incentive-compatibility constraints that prohibit strategies of lying and zero storage, only the downward adjacent ones are binding.*

Proof: For $\theta_2, \dots, \theta_n$,

$$\begin{aligned} & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_i + b_t(h_{t-1}^{\iota_t}) \right] \\ & + \sum_{\tau=t+1}^T \sum_{y_{ii}^{\tau} \in \Theta^{\tau}} \sum_{\iota_{\tau} \in \{0,1\}} \beta^{\tau-(t+1)} \pi^{\iota_{\tau}} (h_{t-1}^{\tau-1}) u \left[y_{\tau} + b_{\tau}(h_{t-1}^{\iota_{\tau}}) \right] \frac{p(y_{ii}^{\tau})}{p(y^{t_i})} = \\ & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_i + b_t(\hat{h}_{t-1}, (\theta_{i-1}, \iota_t, \hat{s}_{t-1}^{\iota_t})) \right] \\ & + \sum_{\tau=t+1}^T \sum_{y_{ii}^{\tau} \in \Theta^{\tau}} \sum_{\iota_{\tau} \in \{0,1\}} \beta^{\tau-(t+1)} \pi^{\iota_{\tau}} (\hat{h}_{t-1}^{\tau-1}) u \left[y_{\tau} + b_{\tau}(\hat{h}_{t-1}, (\theta_{i-1}, \iota_t, \hat{s}_{t-1}^{\iota_t}), \hat{h}_t^{\iota_{\tau}}) \right] \frac{p(y_{ii}^{\tau})}{p(y^{t_i})}, \end{aligned}$$

but

$$\begin{aligned} & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_i + b_t(h_{t-1}^{\iota_t}) \right] \\ & + \sum_{\tau=t+1}^T \sum_{y_{ii}^{\tau} \in \Theta^{\tau}} \sum_{\iota_{\tau} \in \{0,1\}} \beta^{\tau-(t+1)} \pi^{\iota_{\tau}} (h_{t-1}^{\tau-1}) u \left[y_{\tau} + b_{\tau}(h_{t-1}^{\iota_{\tau}}) \right] \frac{p(y_{ii}^{\tau})}{p(y^{t_i})} > \\ & \sum_{\iota_t \in \{0,1\}} \pi^{\iota_t} (\hat{h}^{t-1}) u \left[\theta_i + b_t(\hat{h}_{t-1}, (\theta_{i-k}, \iota_t, \hat{s}_{t-1}^{\iota_t})) \right] \\ & + \sum_{\tau=t+1}^T \sum_{y_{ii}^{\tau} \in \Theta^{\tau}} \sum_{\iota_{\tau} \in \{0,1\}} \beta^{\tau-(t+1)} \pi^{\iota_{\tau}} (\hat{h}_{t-1}^{\tau-1}) u \left[y_{\tau} + b_{\tau}(\hat{h}_{t-1}, (\theta_{i-k}, \iota_t, \hat{s}_{t-1}^{\iota_t}), \hat{h}_t^{\iota_{\tau}}) \right] \frac{p(y_{ii}^{\tau})}{p(y^{t_i})}, \end{aligned}$$

for $k = 2, 3, \dots, i-1$. If, for any k , the latter relation were to hold as an equality, transfers would have to be such that the expected continuation utility for an agent with current income θ_i is the same regardless of whether he announces truthfully or declares θ_{i-k} . But then there is an agent with income θ_j , such that $\theta_i > \theta_j > \theta_{i-k}$, who values the expected stream of transfers for type θ_{i-k} more highly than those actually assigned to him. Hence, he would have an incentive to mimic an agent with θ_{i-k} . \square

References

- AIYAGARI, R., and S. WILLIAMSON (1997): “Money and Dynamic Credit Arrangements with Private Information,” mimeo.
- ALLEN, F. (1985): “Repeated Principal-Agent Relationships with Lending and Borrowing,” *Economic Letters*, 17, 27–31.
- ATKESON, A., and R. LUCAS (1992): “On Efficient Distribution with Private Information,” *Review of Economic Studies*, 59, 427–453.
- COLE, H., and N. KOCHERLAKOTA (1997): “Efficient Allocations with Hidden Income and Hidden Storage,” *Federal Reserve Bank of Minneapolis, Staff Report*, No. 238.
- FERNANDES, A., and C. PHELAN (1999): “A Recursive Formulation for Repeated Agency with History Dependence,” *Federal Reserve Bank of Minneapolis, Staff Report*, No. 259.
- FRIEDMAN, M. (1957): “A Theory of the Consumption Function,” National Bureau of Economic Research, New York.
- GREEN, E. (1987): “Lending and Smoothing of Uninsurable Income,” in *Contractual Arrangements for Intertemporal Trade*, ed. by E. Prescott, and N. Wallace, pp. 3–25. Minneapolis: University of Minnesota Press.
- GREEN, E., and S. OH (1991): “Can a ‘Credit Crunch’ Be Efficient?,” *Federal Reserve Bank of Minneapolis Quarterly Review*, Fall 1991, 3–17.
- MOOKHERJEE, D., and I. PNG (1989): “Optimal Auditing, Insurance, and Redistribution,” *Quarterly Journal of Economics*, May, 399–415.
- RADNER, R. (1981): “Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship,” *Econometrica*, 49, 1127–1148.
- RUBINSTEIN, A., and M. YAARI (1983): “Repeated Insurance Contracts and Moral Hazard,” *Journal of Economic Theory*, 30, 74–97.
- SMITH, B., and C. WANG (1997): “Repeated Insurance Relationships in a Costly State Verification Model: With an Application to Deposit Insurance,” *Federal Reserve Bank of Minneapolis, Working Paper*, 574D.

- THOMAS, J., and T. WORRALL (1990): “Income Fluctuation and Asymmetric Information: An Example of Repeated Principal-Agent Problem,” *Journal of Economic Theory*, 51, 367–390.
- TOWNSEND, R. (1979): “Optimal Contracts and Competitive Markets with Costly State Verification,” *Journal of Economic Theory*, 22, 265–293.
- TOWNSEND, R. (1982): “Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information,” *Journal of Political Economy*, 90, 1166–1186.
- WANG, C. (1995): “Dynamic Insurance with Private Information and Balanced Budgets,” *Review of Economic Studies*, 62, 577–595.
- WANG, C., and S. WILLIAMSON (1998): “Debt Contracts and Financial Intermediation with Costly Screening,” *Canadian Journal of Economics*, 31, 573–595.