Available online at www.sciencedirect.com

SCIENCE DIRECT®

Review of Economic Dynamics 7 (2004) 541–554

Review of Economic Dynamics

www.elsevier.com/locate/red

Figuring out the impact of hidden savings on optimal unemployment insurance ☆

Narayana R. Kocherlakota

Stanford University, Federal Reserve Bank of Minneapolis, and NBER, USA

Received 7 August 2002; revised 11 October 2003

Available online 25 May 2004

Abstract

In this paper, I consider the problem of optimal unemployment insurance in a world in which the unemployed agent's job-finding effort is unobservable and his level of savings is unobservable. I show that the first-order approach is not always valid for this problem, and I argue that the available recursive procedures are not currently computationally feasible. Nonetheless, for the case in which the disutility of effort is linear, I am able to provide a complete characterization of the optimal contract: the agent's consumption is constant while he is unemployed, and jumps up to a higher constant and history-independent level of consumption when he finds a job.

1. Introduction

In a recent paper, Hopenhayn and Nicolini (1997) study the properties of an optimal insurance arrangement between a risk-neutral insurer (principal) and a risk-averse worker (agent). They assume that the agent begins life unemployed and expends a hidden amount of effort to find a job in each period. His probability of finding a job is increasing in the amount of effort exerted; once he finds a job, he keeps it forever. Importantly, the insurer has complete control over the agent's consumption, because the agent cannot *secretly* transfer consumption from one period to the next.

They find that in an optimal contract between the principal and the agent, the agent's consumption is a decreasing function of his time spent unemployed. This general result

☆ This paper previously circulated under the title "Simplifying optimal unemployment insurance: The role of hidden savings."

E-mail address: nkocher@stanford.edu.

1094-2025/$ – see front matter
doi:10.1016/j.red.2004.01.003

has two consequences. First, an agent who has been unemployed for t periods has a lower consumption than an agent who has been unemployed for $(t-1)$ periods. Second, an agent who finds a job after a long period of unemployment must make a higher payment to the insurer than an agent who finds a job after a short period of unemployment.

As stated above, Hopenhayn and Nicolini assume that the principal can costlessly monitor the agent's savings and condition contractual payments on this variable. One can show that the optimal contract in Hopenhayn and Nicolini's setting has the property that the agent is *savings-constrained* when unemployed: the agent's shadow interest rate is lower than the principal's shadow interest rate. Nor is this feature of the Hopenhayn–Nicolini contract unique to the unemployment insurance problem. Rogerson (1985a) shows that in settings with repeated moral hazard, it is generally optimal to impose a sufficiently severe punishment for poor output performance that the agent ends up being savings-constrained. Intuitively, the agent would like to save so as to mitigate next period's punishment.[1]

It follows that with moral hazard, the optimal dynamic contract is only incentive-compatible under the assumption that the principal is able to costlessly monitor the agent's asset levels. This assumption is somewhat restrictive. After all, there are a number of ways that a person can transfer resources to the future (like foreign bank accounts or by accumulating durables) that may be hard for outsiders to observe. It is therefore important to understand the intertemporal structure of optimal contracts when the agent is allowed to engage in secret asset accumulation.

This paper is a contribution to this general research agenda. I relax the assumption that savings can be monitored by the principal in the Hopenhayn–Nicolini unemployment insurance model, and assume instead that the agent can secretly save at the same rate as the principal. I then look to solve for the optimal insurance contract.[2] Not surprisingly, this problem is generally impossible to solve analytically. Unfortunately, it is also difficult to solve numerically. In a recent paper, Fernandes and Phelan (2000) have described a recursive formulation for a related class of problems. It is not known, though, how to translate their recursive formulation into a practical computational procedure when savings can take on a continuum of values. Werning (2002) and Abraham and Pavoni (2003) attack the problem by using a computationally feasible first-order approach that replaces the agent's incentive constraints with the corresponding first order conditions. However, I show that even in simple examples, the first-order approach may not be valid because the agent's decision problem is intrinsically non-concave in effort and savings.

It is possible, though, to obtain an analytical solution in a particular case, even when the first-order approach is known to be invalid. I assume that the agent's disutility from effort is linear in the probability of his finding a job, and that the principal wants the agent to exert an interior amount of effort while unemployed. Under these assumptions,

[1] In a recent working paper, Shimer and Werning (2003) consider unemployment insurance in a version of the McCall search paradigm. They assume that the insurer cannot observe the wage drawn by the unemployed agent. They show that if the agent has exponential utility, then the optimal unemployment insurance contract is the same whether or not the agent can secretly save and/or borrow.

[2] I search across all incentive-compatible insurance contracts. Abdulkadiroglu et al. (2002) instead consider an incomplete markets economy with a limited set of possible unemployment insurance systems. They numerically characterize the optimal unemployment insurance system in that set.

I prove that the optimal unemployment insurance contract takes an extremely simple form. During the period that an agent is unemployed, his consumption is constant. When he becomes employed, his consumption jumps up to a new constant level that is independent of the duration of the unemployment spell. This structure implies that once the agent's savings level is unobservable, it is optimal for the agent to be *borrowing-constrained* when unemployed.

The intuition behind this result is as follows. The contract has to be designed to punish the agent as severely as possible, given that it must deter the agent from saving. This intuition would seem to lead to the optimal contract's featuring consumption-smoothing, so that the principal and agent have the same shadow interest rate.[3] However, the very fact that the first-order approach fails is a sign that this intuition is wrong. The binding intertemporal incentive constraint is one in which the agent jointly deviates from the optimal contract by *simultaneously* saving more and working less. When the contract is designed to prevent this joint deviation, the agent ends up being borrowing-constrained given that he does work the amount specified by the contract.

In this paper, I assume that the unemployed agent cannot borrow secretly. I have two reasons for this restriction. The first is technical: in the linear disutility case, there are *no* incentive-compatible contracts (including repetition of any static contract) if agents can engage in both hidden borrowing and lending. The second is more substantive. It is much more difficult for individuals to engage in hidden borrowing than hidden saving, because their loans have to be enforced. In contrast, as Cole and Kocherlakota (2001) explicitly model, hidden saving can take the form of physical investment. Physical investment requires no outside enforcement and so is intrinsically more difficult to monitor.

2. The problem

In this section, I describe a variant of the Hopenhayn–Nicolini unemployment insurance model, augmented to allow for hidden savings. The principal has von Neumann–Morgenstern utility function

$$-\sum_{t=1}^{\infty} \beta^{t-1} c_t$$

and the agent has von Neumann–Morgenstern utility function

$$\sum_{t=1}^{\infty} \beta^{t-1}\big[u(c_t) - v(p_t)\big]$$

where, in both utility functions, c_t is the agent's consumption in period t. The variable p_t is the agent's effort in period t, and lies in the set $[0, 1]$. I assume that $u', -u'', v' > 0$, $v'' \geqslant 0$, and that u is bounded from above and from below. I assume that $0 < \beta < 1$.

[3] This intuition is valid in the environment with hidden income and hidden storage studied by Cole and Kocherlakota (2001). The key difference between the two settings is that in Cole and Kocherlakota, the two types of deviations (storing from period t to period $(t + 1)$ and then lying) are not complementary in utility. In contrast, shirking and storing are complementary in the model studied in this paper.

An agent can be employed or unemployed; he begins life unemployed. The choice of p_t affects the probability of becoming employed for an unemployed agent. Specifically, if an agent is unemployed at the end of period $(t-1)$, then the probability of his becoming employed in period t is p_t, and the probability of his staying unemployed is $(1-p_t)$. If an agent is employed at the end of period $(t-1)$, he stays employed in period t with probability one. Thus, being employed is an absorbing state.

The agent's employment status is observable to others, but his choice of p_t is unobservable. As well, the agent can secretly save at rate $1/\beta - 1$. I consider contracts in this economy that specify two sequences $\{c_t^E, c_t^U\}_{t=1}^{\infty}$. Given such a contract, an agent who is unemployed in period t receives compensation c_t^U from the principal. If an agent became employed for the first time in period t, then his compensation from the principal in period $s \geqslant t$ is c_t^E. Thus, once an agent is employed, his compensation is constant over time. (It is simple to show that because the principal and agent have the same discount factor, this smooth compensation is efficient in this economy.)

I assume that the principal wants to (weakly) implement a sequence of effort choices $p^* = \{p_t^*\}_{t=1}^{\infty}$ by the agent when unemployed, where $1 > p_t^* > 0$ for all t. I define an incentive-compatible contract (c^E, c^U) to be one such that:

$$\begin{aligned}\left\{S_t^*, p_t^*\right\}_{t=1}^{\infty} \in \arg\max_{\{S_t, p_t\}_{t=1}^{\infty}} \sum_{t=1}^{\infty} \beta^{t-1} \prod_{s=1}^{t-1}(1-p_s)\big\{p_t u\big(\zeta_t^E\big)/(1-\beta) - v(p_t) \\ +(1-p_t)u\big(\zeta_t^U\big)\big\}\end{aligned}$$

$$\begin{aligned}&\text{s.t.}\quad \zeta_t^E = c_t^E + S_{t-1}(1-\beta)/\beta \quad \text{for all } t,\\ &\zeta_t^U = c_t^U + S_{t-1}/\beta - S_t \quad \text{for all } t,\\ &S_t, p_t, 1-p_t, \zeta_t^E, \zeta_t^U \geqslant 0 \quad \text{for all } t,\\ &S_0 = 0,\end{aligned}$$

so that it is weakly optimal for an unemployed agent to choose p_t^* in all t. Note that if an agent becomes employed in period t with savings S_{t-1}, then his optimally smoothed consumption is $(c_t^E + S_{t-1}(1-\beta)/\beta)$ in every period thereafter.[4]

It is straightforward to show that, given any incentive-compatible contract, there exists a payoff-equivalent contract $(c^{E\prime}, c^{U\prime})$ in which the agent's optimal savings sequence is

[4] I do not formally model *why* the principal desires to implement an interior p^*. It is standard in principal-agent problems to model the principal's objective as being linear in p; this assumption, if the agent's utility function is linear in p, would generically result in the principal's preferring a bang-bang specification for p.

However, in this unemployment insurance problem, the principal may prefer an interior choice for p because of search externalities. Suppose that the principal is contracting with a unit measure of agents, and that a given agent's disutility from choosing a probability p is given by $p\Psi(\overline{p})$, where $\overline{p}$ is the average p chosen by the other agents in the economy. If Ψ is increasing, then there are congestion effects—it becomes harder for a given agent to find a job when other agents are searching a lot. When designing the optimal contract, the principal internalizes the externality implicit in Ψ, and the principal's choice of p will, for a generic class of problems, be interior.

zero. I restrict attention to these contracts that induce zero savings, and look to solve the following cost-minimization problem (*UIP*):

$$\min_{c^E, c^U} \sum_{t=1}^{\infty} \beta^{t-1} \prod_{s=1}^{t-1} \bigl(1 - p_s^*\bigr) \bigl\{ p_t^* c_t^E / (1-\beta) + \bigl(1 - p_t^*\bigr) c_t^U \bigr\}$$

s.t.

$$(0, p^*) \in \arg \max_{S \geqslant 0, 1 \geqslant p \geqslant 0} \sum_{t=1}^{\infty} \beta^{t-1} \prod_{s=1}^{t-1} (1 - p_s) \bigl\{ p_t u\bigl(c_t^E + S_{t-1}(1-\beta)/\beta\bigr)/(1-\beta)$$

$$+ (1 - p_t) u\bigl(c_t^U + S_{t-1}/\beta - S_t\bigr) - v(p_t) \bigr\}$$

$$\text{s.t.} \quad c_t^U + S_{t-1}\beta^{-1} - S_t \geqslant 0 \quad \text{for all } t, \qquad S_0 = 0;$$

$$\sum_{t=1}^{\infty} \beta^{t-1} \prod_{s=1}^{t-1} (1 - p_s) \bigl\{ p_t u\bigl(c_t^E\bigr)/(1-\beta) + (1 - p_t) u\bigl(c_t^U\bigr) \bigr\} \geqslant u_*$$

$$c_t^E, c_t^U \geqslant 0 \quad \text{for all } t.$$

In words: What contracts are the minimal-cost incentive-compatible contracts among all those that provide the agent with ex-ante utility of at least u^*?

3. Difficulties

In this section, I consider two recently developed approaches to solving *UIP*. The first is to make the problem recursive in some fashion. I find that this approach is, at least of this writing, computationally infeasible. The second is to use a version of the first-order approach. I find that this approach will not work if the curvature of v is sufficiently small.

3.1. Can we make the problem recursive?

Much of the recent analysis of dynamic moral hazard problems is based on an insight of Spear and Srivastava (1987). They show that, without hidden savings or other hidden state variables, dynamic moral hazard problems are recursive in the following sense: in each period, the principal chooses current consumption and next period's continuation utility so as to minimize his costs subject to the incentive constraints, and subject to delivering a specified amount of continuation utility to the agent. Hence, the principal-agent problem is recursive with respect to a one-dimensional state variable: continuation utility.

The difficulty in making *UIP* recursive in a similar fashion is that if an agent brings savings into the period, his response to any given contract is different than if he does not bring savings. In other words, the presence of hidden savings essentially introduces an adverse selection problem at each date. Fernandes and Phelan (2000) show how to deal with this kind of dynamic adverse selection problem: the principal must minimize his costs subject to delivering a given amount of continuation utility to *every* type.

Here's how Fernandes and Phelan's insight works in this context. Suppose the principal wants to induce an unemployed agent to choose p_* in every period. Given an incentive-compatible contract (c^E, c^U), we can define $V(S)/(1-\beta)$ to be the ex-ante utility of the agent if he begins life with S units of savings (as opposed to with zero) and then chooses an optimal effort and savings strategy in response to the contract. Thus, V is a *value function*. Define *DOM* to be the set of all such value functions (as we vary the incentive-compatible contract (c^E, c^U)). Further, given a value function V in *DOM*, let $\Pi(V)$ be the minimal cost to the principal of all incentive-compatible contracts that generate the value function V.

Then, the function $\Pi : DOM \to R_+$ satisfies the following functional equation (FE):

$$\Pi(V) = \min_{c^E, c^U, W} p_* c^E/(1-\beta) + (1-p_*)\left\{c^U + \beta \Pi(W)\right\}$$

s.t.

$$(0, p_*) \in \arg \max_{\substack{c^U \geqslant S' \geqslant 0 \\ 1 \geqslant p \geqslant 0}} pu\left(c^E\right) - v(p)(1-\beta) + (1-p)\left\{u\left(c^U - S'\right)(1-\beta) + \beta W(S')\right\},$$

$$V(S) = \max_{\substack{S\beta^{-1} + c^U \geqslant S' \geqslant 0 \\ 1 \geqslant p \geqslant 0}} \left\{pu\left(c^E + S(1-\beta)/\beta\right) - v(p)(1-\beta)\right\} + (1-p)\left\{u\left(c^U + S/\beta - S'\right)(1-\beta) + \beta W(S')\right\},$$

$$c^E, c^U \geqslant 0,\ W \in DOM.$$

At a given point in time, the principal seeks to minimize the expected value of his discounted costs, given that he wishes to induce an agent with no assets to choose effort p_* and to choose not to save. The possibility of hidden savings means that, in order to make sure the contract is in fact incentive-compatible, the agent needs to know how much utility he will get from choosing values of savings other than the principal's preferred level of savings (zero). Hence, the principal needs to satisfy a promise-keeping constraint that applies to all values of S, not just $S = 0$, and needs to pick a continuation value function W, not just a continuation utility.

We now have a recursive approach to *UIP*. Let Π be the solution to (FE). Then, the principal's first step is to solve the minimization problem:

$$\min_{V \in DOM} \Pi(V)$$
$$\text{s.t.} \quad V(0) \geqslant u^*.$$

He obtains a solution V_0 to this minimization problem. Next, the principal solves the minimization problem in (FE) with V_0 substituted in for V, and obtains a solution (c_1^E, c_1^U) and a continuation value function V_1. He again solves the minimization problem in (FE), now with V_1 substituted in for V. This will deliver a (c_2^E, c_2^U), as well a continuation value function V_2. The principal can continue in this recursive fashion; the resulting $(c_t^E, c_t^U)_{t=1}^{\infty}$ solves *UIP*.

Note that because of hidden savings, the relevant state variable is now a *function*, not a number as when only effort is hidden. This is inevitable, because we have to keep track of continuation utility for all types—that is, for all savings levels. As well, we have to use a generalization of Abreu et al. (1990) to iterate on (infinite-dimensional) sets of functions until we find *DOM*. These infinities pose significant computational difficulties. Hence, at this point in time, it is not known how to implement Fernandes and Phelan's recursive approach in practice when the agent has a continuum of possible savings levels.[5]

3.2. The first-order approach

Much of the analysis of moral hazard problems uses the *first-order approach*. To see how this approach works, it's useful to look at a two-period version of the unemployment insurance problem posed in the previous section. I set the discount rate equal to zero, and assume that the agent has preferences of the form

$$\ln(c_1) + \ln(c_2) - v(p_2)$$

and a technology of the form

$$y = \begin{cases} E & \text{with probability } p_2, \\ U & \text{with probability } 1 - p_2. \end{cases}$$

The agent can secretly save at a zero rate of return. The principal cannot observe the agent's choice of storage level or the agent's choice of p_2; the principal can condition the agent's second period consumption on the realization of y.

The principal's problem is to (weakly) implement a choice $p_2^* \in (0,1)$ at minimal expected cost, given that the agent must receive at least reservation utility u_*. Mathematically, the principal's problem is:

$$\min_{c_1, c_E, c_U \geqslant 0} c_1 + p_2^* c_E + \left(1 - p_2^*\right) c_U$$

s.t.

$$\left(S, p_2^*\right) \in \max_{\substack{1 \geqslant p_2 \geqslant 0 \\ c_1 \geqslant S \geqslant 0}} \ln(c_1 - S) + p_2 \ln(c_E + S) + (1 - p_2)\ln(c_U + S) - v(p_2),$$

$$\ln(c_1) + p_2^* \ln(c_E) + \left(1 - p_2^*\right)\ln(c_U) - v\left(p_2^*\right) \geqslant u_*.$$

It is simple to show that given a solution to this problem (c_1, c_E, c_U), then $(c_1 - S, c_E + S, c_U + S)$ is also a solution which leads the agent not to store. Hence, the principal's minimal costs are not increased by considering the following problem with a smaller constraint set:

$$\min_{c_1, c_E, c_U \geqslant 0} c_1 + p_2^* c_E + \left(1 - p_2^*\right) c_U$$

s.t.

$$\left(0, p_2^*\right) \in \max_{\substack{1 \geqslant p_2 \geqslant 0 \\ c_1 \geqslant S \geqslant 0}} \ln(c_1 - S) + p_2 \ln(c_E + S) + (1 - p_2)\ln(c_U + S) - v(p_2),$$

$$\ln(c_1) + p_2^* \ln(c_E) + \left(1 - p_2^*\right)\ln(c_U) - v\left(p_2^*\right) \geqslant u_*.$$

[5] However, if the agent had only a finite number of possible savings levels, we might be able to use this approach to some effect (à la Doepke and Townsend, 2003).

I call this problem $P1$.

A difficulty with this problem is that there is no obvious way to attack it using standard Lagrangian methods. The first-order approach gets around this difficulty by replacing the agent's decision problem with its first-order necessary conditions. This creates the following problem $P2$:

$$\min_{c_1,c_E,c_U \geqslant 0} c_1 + p_2^* c_E + \left(1-p_2^*\right)c_U$$
$$\text{s.t.} \quad \ln(c_E) - \ln(c_U) = v'\left(p_2^*\right),$$
$$1/c_1 \geqslant p_2^*/c_E + \left(1-p_2^*\right)/c_U,$$
$$\ln(c_1) + p_2^* \ln(c_E) + \left(1-p_2^*\right)\ln(c_U) - v\left(p_2^*\right) \geqslant u_*.$$

This problem has two advantages relative to $P1$. The first is obvious: the constraint set is such that the problem is easily amenable to Lagrangian methods. The second advantage is more subtle. In the previous subsection, we saw that the recursive formulation of $P1$ is difficult to implement computationally. Werning (2002) considers multiperiod versions of the problem $P2$. He shows that, in each period, the principal chooses current consumption and next period's continuation utility subject to the incentive constraints on effort, subject to delivering a pre-specified amount of continuation utility and subject to not exceeding a pre-specified upper bound on marginal utility of consumption. The last constraint guarantees that the principal is satisfying the agent's intertemporal Euler equation at each point in time. Thus, multiperiod versions of $P2$ are recursive in two state variables: continuation utility and an upper bound on continuation marginal utility. This is much simpler than multi-period versions of $P1$ (like our original problem *UIP*), where we have to keep track of an infinite-dimensional state variable (and solve as well for the domain of that state variable).

So, it seems like a good idea to attack $P2$ instead of $P1$. Unfortunately, solving $P2$ may not be the same as solving $P1$. The problem is that the agent's objective function is not globally concave in savings and effort. It follows that the first order conditions of the agent's decision problem are only necessary: the constraint set to $P2$ is in general larger than the constraint set to $P1$. This creates the possibility that the solution to $P2$ may not be incentive-compatible. I now show that this possibility is in fact realized if v has sufficiently low curvature.

To do so, I first solve $P2$. In this two-period context, the solution is simple: at an optimum, the two weak inequalities must hold with equality. If the last constraint is an inequality, simply lower c_1: this lowers the principal's objective without violating the other two constraints. If the second constraint is an inequality, raise c_1 by εc_1, lower c_E by εc_E and lower c_U by εc_U. This keeps the agent's ex-ante utility the same, and does not affect the agent's effort decision. The principal's objective is lowered because $c_1 < p_2^* c_E + (1-p_2^*)c_U$ (by Jensen's inequality).

Thus, the solution to $P2$ is the unique triple (c_1^*, c_E^*, c_U^*) that satisfies all constraints with equality. However, (c_1^*, c_E^*, c_U^*) is not in the constraint set of $P1$. Here's why. Given (c_1^*, c_E^*, c_U^*), the agent's objective is supposedly maximized at $S = 0$ and $p_2 = p_2^*$. By construction, the agent's first order conditions are satisfied (with equality). But look at the Hessian of his objective:

$$\begin{matrix} -1/(c_1^*)^2 - p_2^*/(c_E^*)^2 - (1-p_2^*)/(c_U^*)^2 & (1/c_E^* - 1/c_U^*) \\ (1/c_E^* - 1/c_U^*) & -v''(p_2^*) \end{matrix}$$

A necessary condition for $S = 0$ and $p_2 = p_2^*$ to solve the agent's problem is that this Hessian be negative semi-definite. It is true that the diagonal elements are non-positive. But the determinant of the Hessian is negative if $v''(p_2^*)$ is sufficiently small, and so, even though the agent's first order conditions are satisfied at $S = 0$ and $p_2 = p_2^*$, he can experience a *second-order* gain by increasing S above 0 and lowering p_2 below p_2^*.

Thus, even in this simple example, the first-order approach is invalid if v has sufficiently low curvature. This possibility is generated by the fact that the agent's objective function is not guaranteed to be non-concave as a function of p_2 and S. The same kind of reasoning can be applied in the infinite-horizon setting of Section 2 to show that we cannot always use the first-order approach.[6]

There is no set of known conditions in the infinite horizon problem *UIP* that are sufficient to guarantee that the first-order approach is valid with hidden savings. Abraham and Pavoni (2003) point out, though, that it is possible to verify whether a particular solution to the first-order approach problem is actually a solution to the true problem. They use a two-step numerical procedure in their analysis of optimal unemployment insurance with hidden borrowing and lending. First, they solve the first-order approach problem (the infinite-horizon analog of $P2$). Second, they verify whether the solution is incentive-compatible, by checking whether the agent finds it optimal to choose p^* when confronted with the solution contract. They conclude that for all of their parameterizations, the solution to the first-order approach problem is in fact the solution to the true problem.

Werning (2002) also attacks *UIP* by solving the first-order approach problem. For some specifications of u and v, he shows numerically and analytically that in the solution to this problem, the difference between c_{Ut} and c_{Et} is falling over time. (He interprets this falling differential as implying that unemployment *benefits* should be *increasing* in the duration of unemployment.) His paper does not have the kind of explicit verification step contained in Abraham and Pavoni (2003). Hence, his paper contains no information about whether his characterization of the solution to the first-order approach problem carries over to the true problem *UIP* or not.

4. Solving the insurance problem in the linear disutility case

We now return to the problem *UIP*: what is the principal's preferred contract among all those incentive-compatible contracts that provide the agent with ex-ante utility no less than u^*? We have seen in the previous section that there are no generally valid approaches that

[6] Even without hidden savings, it is possible that the first-order approach is invalid. The basic problem, again, is that the agent's problem may not be globally concave in effort. However, there are known sufficient conditions that preclude this possibility and $P1$ satisfies those sufficient conditions. See Rogerson (1985b) for a full discussion.

are currently computationally feasible to solving contracting problems with hidden effort and hidden savings. In this section, I specialize the problem by assuming that

$$v(p_2) = \alpha p_2, \quad \alpha > 0.$$

Under this assumption, the first-order approach is definitely invalid, because the Hessian of the agent's objective is guaranteed not to be negative definite. Nonetheless, I can provide a complete analytical characterization of the optimal contract.

4.1. A relaxed problem

We begin by first constructing a superset for the set of incentive-compatible contracts. Any incentive-compatible contract must satisfy

$$(1-\beta)\sum_{s=0}^{\infty}\beta^s u\left(c_{t+s}^U\right) = u\left(c_t^E\right) - \alpha(1-\beta) \quad \text{for all } t. \tag{R1}$$

This restriction derives from the linearity of the agent's problem. In particular, in period t, the agent's problem is linear in p_{t+s} for $s \geqslant 0$. Hence, if he chooses $p_t > 0$ in every period, he must be indifferent among all possible p sequences, including setting $p_{t+s} = 0$ for all s, and setting $p_t = 1$.

The restriction (R1) is implied by efforts being hidden. In addition, hidden savings implies that any incentive-compatible contract must also satisfy

$$c_t^U \leqslant c_{t+1}^U \quad \text{for all } t. \tag{R2}$$

Suppose $c_{t+1}^U < c_t^U$. Then, an unemployed agent in period t prefers to set ($S_t > 0$, $p_{t+1} = 0$) to setting ($S_t = 0$, $p_{t+1} = 0$). But (R1) implies that the agent is indifferent between setting ($S_t = 0$, $p_{t+1} = 0$) and ($S_t = 0$, $p_{t+1} = p_{t+1}^*$). Hence, if $c_{t+1}^U < c_t^U$, it is not optimal for an unemployed agent in period t to set $S_t = 0$ and $p_{t+1} = p_{t+1}^*$. This is a contradiction of incentive-compatibility.

Note that (R2) implies that a contract may satisfy both (R1) and

$$u'\left(c_t^U\right) \geqslant p_{t+1}^* u'\left(c_{t+1}^E\right) + \left(1 - p_{t+1}^*\right)u'\left(c_{t+1}^U\right) \quad \text{for all } t, \tag{R3}$$

and still not be incentive-compatible. (R1) and (R3) are the first-order necessary conditions that are implied by the optimality of effort strategy p^* and zero savings. But, just as in the discussion of the first-order approach in the previous section, they do not take into account the second-order consequences of simultaneous changes in savings and effort.

Thus, the set of contracts that satisfy (R1) and (R2) are a superset of the incentive-compatible contracts. We now pose a *relaxed problem*: among the contracts that satisfy (R1) and (R2), and provide the agent with at least u_* in ex-ante utility, which ones does the principal prefer?

4.2. Solving the relaxed problem

To solve the relaxed problem, we begin with two straightforward observations. First, in any solution to the relaxed problem, the ex-ante utility constraint must hold with equality.

If it does not, we can lower $u(c_1^U)$ by ε and $u(c_1^E)$ by $\varepsilon(1-\beta)$. This change improves the principal's objective without violating any of the constraints for ε small. Second, note that the constraints (R1) and (R2) together imply that in any solution, $c_t^U < c_t^E$ for all t.

The next step is the key one: in any solution to the relaxed problem,

$$c_t^U = c_{t+1}^U.$$

Suppose not, and $c_t^U < c_{t+1}^U$. Then, we can construct a new contract by using a perturbation similar to that in Rogerson (1985a): raising $u(c_t^U)$ by ε, lowering $u(c_{t+1}^U)$ by $\varepsilon\beta^{-1}$ and lowering $u(c_{t+1}^E)$ by $\varepsilon\beta^{-1}(1-\beta)$. This new contract satisfies (R1) for any ε, and satisfies (R2) as long as ε is sufficiently small. The new contract's change in cost relative to the old one is given by

$$\begin{aligned}&\varepsilon/u'(c_t^U) - \varepsilon p_t^*/u'(c_{t+1}^U) - \varepsilon(1-p_{t+1}^*)/u'(c_{t+1}^E)\\ &\quad\leqslant \varepsilon/u'(c_t^U) - \varepsilon/u'(c_{t+1}^U) \quad \left(\text{because } u''<0 \text{ and } c_{t+1}^E > c_{t+1}^U\right)\\ &\quad< 0,\end{aligned}$$

and so the old contract was not optimal.

Hence, in any contract that solves the relaxed problem, $c_t^U = \bar{c}^U$ for all t. From (R1), we know that

$$c_t^E = \bar{c}^E = u^{-1}\left(u\left(\bar{c}^U\right) + \alpha(1-\beta)\right)$$

for all t. We can then find the unique solution to the relaxed problem by substituting into the ex-ante utility constraint to find

$$\bar{c}^U = u^{-1}\left[u^*(1-\beta)\right].$$

4.3. *The optimal contract*

We have characterized the unique solution to the relaxed problem. To verify that it in fact solves the original problem, we need to show that this solution is in fact incentive-compatible. But note that for any p,

$$u'\left(\bar{c}^U\right) \geqslant pu'\left(\bar{c}^E\right) + (1-p)u'\left(\bar{c}^U\right).$$

Hence, no matter what p sequence that he chooses, an unemployed agent never wants to save. As well, the agent is indifferent between all levels of p in each period. It follows that the contract is indeed incentive-compatible, and must be in fact the principal's preferred incentive-compatible contract.

4.4. *Discussion*

It is useful to contrast this contract with the optimal contract when the agent cannot secretly save. When savings are observable, it is optimal in this setting for the principal to leave unemployed agents *savings-constrained*, so that

$$u'\left(c_t^U\right) < p_{t+1}u'\left(c_{t+1}^E\right) + (1-p_{t+1})u'\left(c_{t+1}^U\right).$$

Intuitively, the optimal way to provide incentives in period $(t+1)$ is to punish the agent so severely when he is unemployed that he would like to save from period t to period $(t+1)$.

Once the agent can save secretly, it is no longer possible to punish the agent so severely. The key principle underlying the optimal contract is that it is designed to punish the agent as much as is possible ex-post, given the agent's ability to undermine such punishments using secret savings. One might think that this principle means that the optimal contract would adjust to secret savings by making the above inequality an equality. Indeed, had we incorrectly used the first-order approach to "solve" *UIP*, the "solution" would in fact have had this property.

The problem with this thinking is that even if a contract satisfies the intertemporal Euler equation $u'(c_t^U) = p_{t+1}u'(c_{t+1}^E) + (1-p_{t+1})u'(c_{t+1}^U)$, the agent can still undermine the punishment inherent in the contract by saving secretly. In particular, suppose the intertemporal Euler equation holds but $c_t^U > c_{t+1}^U$. Then, the agent will find it optimal to save $(c_t^U - c_{t+1}^U)/2$ from period t to period $(t+1)$ and then set $p_{t+1} = 0$ in period $(t+1)$. In other words, the possibility of a *joint* deviation of saving and shirking imposes the even tighter intertemporal restriction of $c_t^U \leqslant c_{t+1}^U$ on the optimal contract. Given this restriction, the optimal contract imposes the most severe punishment on the agent in period $(t+1)$—and this implies that c_t^U equals c_{t+1}^U in the optimal contract.

The structure of the optimal contract implies that for all t:

$$u'\big(c_t^U\big) > p_{t+1}u'\big(c_{t+1}^E\big) + (1-p_{t+1})u'\big(c_{t+1}^U\big),$$

so that for all t the agent is *borrowing-constrained*. Earlier, I restricted attention to contracts which induce zero savings on the part of the agent. This raises the question of whether there are other optimal contracts which induce the same consumption allocation for the agent, but a positive amount of private savings in at least some period. But it is optimal for the agent to be borrowing-constrained at every date in the optimal contract; hence, private savings must be zero at every date.

It is useful to note as well that the optimal contract in this setting is renegotiation-proof: it is Pareto optimal at the beginning of each period. (Of course, it is not Pareto optimal after the agent has exerted effort within a period, but before the realization of his employment status.) In contrast, Chiappori et al. (1994) find that the ex-ante optimal contract is not renegotiation-proof when they consider a principal-agent problem in which the agent has only two possible effort choices and can secretly borrow and lend.

5. Conclusions

This paper considers the optimal provision of unemployment insurance for an agent who can secretly exert effort to find a job and who can secretly save. The paper argues that it is not practical to compute an approximate solution to the contracting problem using currently available recursive methods. As well, the first-order approach is not generally valid: the complementary nature of shirking and saving makes the agent's problem non-concave.

Despite these difficulties, it is possible to completely and analytically characterize the optimal contract when the agent's disutility of effort is a linear function of his probability of finding a job. The paper uses this characterization to show that the nature of optimal unemployment insurance is considerably changed if the agent can engage in secret saving. In particular, the agent's compensation when he is unemployed or when he gets a job is independent of his history, instead of depending in complicated ways on the duration of unemployment. As well, rather than being savings-constrained, the agent faces binding borrowing constraints at each date.

It is natural to ask whether these findings are robust to introducing small amounts of curvature in v. I suspect that the exact history independence result will collapse—although my guess is that even in those cases, there will not be much loss in welfare in restricting the contract to be history independent. As well, I suspect too that the optimal contract will continue to leave the agent borrowing-constrained (which also means that the first-order approach will not work). The challenge that remains is to develop robust and practical numerical methods to assess these, and other, conjectures. The continuous-time approach of Williams (2003) may be a promising step in this direction.

Acknowledgments

I thank Harold Cole, Russell Cooper, Barbara McCutcheon, Chris Phelan, Ivan Werning, the referee, and especially the editor (Richard Rogerson) for their comments. I also thank the participants at the many seminars at which I presented the paper. I acknowledge the support of NSF SES 0076315. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

References

Abdulkadiroglu, A., Kuruscu, B., Sahin, A., 2002. Unemployment insurance and the role of self-insurance. Review of Economic Dynamics 5, 681–703.

Abraham, A., Pavoni, N., 2003. Efficient allocations with moral hazard and hidden borrowing and lending. Working paper. University College London.

Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. Econometrica 58, 1041–1064.

Chiappori, P.-A., Macho, I., Rey, P., Salanie, B., 1994. Repeated moral hazard: The role of memory, commitment, and the access to credit markets. European Economic Review 38, 1527–1553.

Cole, H.L., Kocherlakota, N., 2001. Efficient allocations with hidden income and hidden storage. Review of Economic Studies 68, 523–542.

Doepke, M., Townsend, R., 2003. Dynamic mechanism design with hidden income and hidden actions. Working paper. Univ. of California—Los Angeles.

Fernandes, A., Phelan, C., 2000. A recursive formulation for repeated agency with history dependence. Journal of Economic Theory 99, 223–247.

Hopenhayn, H., Nicolini, J.-P., 1997. Optimal unemployment insurance. Journal of Political Economy 105, 412–438.

Rogerson, W., 1985a. Repeated moral hazard. Econometrica 53, 69–76.

Rogerson, W., 1985b. The first-order approach to principal–agent problems. Econometrica 53, 1357–1367.

Shimer, R., Werning, I., 2003. Optimal unemployment insurance with sequential search. Working paper. Univ. of Chicago.

Spear, S., Srivastava, S., 1987. Repeated moral hazard with discounting. Review of Economic Studies 54, 599–617.

Werning, I., 2002. Optimal unemployment insurance with hidden savings. Working paper. Univ. of Chicago.

Williams, N., 2003. On dynamic principal–agent problems in continuous time. Working paper. Princeton University.