Economic Theory **18**, 207–235 (2001)

Economic Theory

Dynamic labor contracts with temporary layoffs and permanent separations

Christopher Sleet[1] **and Şevin Yeltekin**[2]

[1] Department of Economics, University of Iowa, Iowa City IA 52242, USA
(e-mail: christopher-sleet@uiowa.edu)
[2] Kellogg Graduate School of Management, Northwestern University, Evanston, IL 60208, USA
(e-mail: s-yeltekin@kellogg.nwu.edu)

Received: February 28, 2000; revised version: January 21, 2001

Summary. We study the implications of optimal dynamic contracts in private information environments for fluctuations in effort and employment across time and productivity states. To this end, we incorporate temporary layoffs and permanent separations as well as on-the-job effort variations into a dynamic model of moral hazard. We consider two different "commitment" environments. In a "full commitment" environment, although the firm can temporarily lay a worker off, neither party can dissolve the contractual relationship once it has been initiated. On the other hand, in a "limited commitment" environment, both parties can dissolve the relationship at the beginning of any period in order to pursue an outside option.

We use our model to study the implications of optimal contracts for incentives, employment histories, layoffs and separations across full information, full commitment and limited commitment settings. We compute solutions to the relevant principal-agent problems, endogenously determining the set of states in which separations occur and the domain of the firm's value function, as well as the value function itself.

Keywords and Phrases: Dynamic contracts, Layoffs, Separations.

JEL Classification Numbers: C63, D82, J23, J33.

Correspondence to: Ş. Yeltekin

1 Introduction

Over the last ten years a number of papers have analyzed the nature of optimal contracts between a principal and an agent in dynamic private information environments[1]. Many of these earlier papers have been concerned with the asymptotic distribution of consumption and utility. In contrast, the aim of this paper is to shift attention to the implications of these contracts for fluctuations in effort and employment across time and productivity states. To this end our model, unlike its predecessors in the literature, incorporates temporary layoffs and permanent separations as well as on-the-job effort variations.

To make matters concrete, consider the following setup. A firm and a worker match and enter into a "contractual relationship"[2] . The worker operates a stochastic technology for the firm that converts worker effort into output. Although the worker's output is observed by the firm, the worker's effort is not (unless the worker is laid off, in which case the firm knows that the worker's effort is zero). The firm's per period profits, the difference between its revenues and the amount of consumption that it delivers to the worker, is perturbed by a persistent and publicly observable revenue shock.[3] Assume that the firm and the worker live forever, and potentially, could match forever. One can then conceive of two different "commitment" environments. In the first, although the firm can temporarily lay a worker off, neither party can dissolve the contractual relationship once it has been initiated. Hence the label "full commitment" for such environments. In the second type of environment commitment is more limited. While both parties can commit not to dissolve a relationship during a period, they are permitted to separate at the beginning of each period in order to pursue an outside option. Furthermore the firm can not commit to making payments to a worker after separation has occurred. Such settings are referred to as "limited commitment" environments.[4] To be credible, a contract in such an environment must deliver continuation payments to both parties that weakly exceed their outside options after each history. In particular, an optimal limited commitment contract must contain rules for optimal separations and be consistent with the fact that each party receives its outside option after a separation.

We use these setups to investigate the extent to which private information and limited commitment constrain firm-worker risk sharing and the optimal provision of incentives. To do this we compute solutions to the relevant principal-agent problems. In the limited commitment case this requires endogenously determining the set of states in which separations occur (the separation region) and the domain

[1] The following is an indicative, rather than exhaustive list: Atkeson and Lucas (1992), Phelan (1995), Phelan and Townsend (1991), Spear and Srivastava (1987), Thomas and Worrall (1990).

[2] For analysis of dynamic firm-multiple worker problems see Yeltekin (1999).

[3] This revenue shock may be interpreted in several ways. It may be regarded as an observable productivity shock, independent of the activities of workers. Alternatively, it may be regarded as a price shock.

[4] One might call an environment in which even one period contracts are unenforceable an environment without commitment. Kocherlakota (1996) considers a dynamic, *full information* environment without commitment.

of the firm's value function, as well as the value function itself. We compare the implications of optimal contracts for incentives, employment histories, layoffs and separations across full information, full commitment and limited commitment settings.

1.1 Related literature

This subsection briefly comments on related literature. Green (1987) and Spear and Srivastava (1987) introduced a recursive formulation for dynamic contracts that uses a utility promise from a principal to an agent as a state variable. This formulation is now standard and is used here. The computational method is similar in flavor to that contained in Phelan and Townsend (1991) who convexify a discretized version of the agency problem. They solve this problem using a value iteration and a linear programming algorithm. We employ a number of techniques to accelerate the computations. In particular, we utilize modified policy iteration and the Dantzig-Wolfe decomposition described by Prescott (1998). We also extend the Phelan-Townsend approach in order to consider the limited commitment problem described above.

Phelan (1995) considers a limited commitment problem with private information. There are three key differences between his formulation and ours. Firstly, Phelan does not include a publicly observable revenue shock. Hence the principal and the agent either sign a contract or they do not. Once they have signed they do not separate, although the need to provide continuation payoffs in excess of outside options constrains the set of continuation contracts. Thus, Phelan's setup, unlike ours, can not be used to analyze efficient separation decisions. Secondly, Phelan's model is of the hidden information variety while ours is of the hidden action variety. The former is easier to analyze theoretically, and Phelan provides a useful collection of analytical results. Thirdly, Phelan goes beyond the analysis in this paper by endogenizing the outside option value available to the worker. To conclude this section, we repeat that much of the earlier literature (including Phelan and Phelan and Townsend) has been chiefly concerned with the implications of dynamic contractual models for wealth and utility distributions. On the other hand we are mainly interested in their implications for effort, layoffs and employee-firm separations.

2 Full commitment problems

To begin with the physical environment is described. This is similar to that used by Phelan and Townsend (1991). Two changes are made: we explicitly incorporate a publicly observable zero effort "layoff" state and we add a publicly observed revenue shock.

2.1 The physical environment

In the initial period a firm hires a worker to operate a stochastic technology and a contract is signed between the two parties. Workers and firms are infinitely lived and once they have entered into a contractual arrangement they are committed to remain in it. The technology stochastically converts worker effort into output. In each period the worker randomizes over effort levels, hence generating a probability distribution over output. The firm rewards the worker with a lottery over current and future consumption and future effort levels.

To facilitate the computations a discretized model is considered[5]. Suppose then that the set of feasible outputs is given by a finite set $Q = \{q_i\}_{i=1}^{n_q}$ with $q_{i+1} > q_i$, $q_1 = 0$ and $q_{n_q} = \overline{q}$ and that the set of feasible worker efforts is also a finite set $A = \{a_i\}_{i=1}^{n_a}$ with $a_{i+1} > a_i$, $a_1 = \underline{a} = 0$ and $a_n = \overline{a}$. Also let $A_e = \{a_i\}_{i=2}^{n_a}$, i.e. A_e does not include the first (layoff) effort level. Let C be a finite, ordered set of worker consumptions with n_c elements, minimal element $\underline{c}$ and maximal element $\overline{c}$. Similarly, let Θ denote a finite, ordered set of "revenue shocks" that perturb the revenue of the firm. Define $\Omega \equiv A \times Q \times C \times \Theta$ and let ω denote an element of Ω. Elements ω^t of Ω^t will be referred to as t-period "histories" or t-period "careers". The set of probability distributions over $\Omega^\infty \equiv \Omega \times \Omega \times \cdots$ will be denoted $P(\Omega^\infty)$. Note that Ω^1 is to be interpreted as the set of period 0 efforts, outputs, consumptions and period 1 revenue shocks. Similarly, Ω^t contains the $t+1$-th period shock. None of the Ω^t sets include the zeroth period shock which is fixed exogenously. Finally, let $\Phi^t \equiv \Omega^t \times A \times Q \times C$, $t > 0$ and $\Phi^0 \equiv A \times Q \times C$.

Suppose that the firm's technology is described by a matrix P of dimension $n_a \times n_q$. The (i, j)-th element of the matrix, $P_{i,j}$, gives the probability of output q_j occurring given that the worker has exerted effort level a_i. P satisfies two assumptions. Firstly, it is monotone: if $i > j$ then the probability distribution described by the i-th row of the matrix first order stochastically dominates that described by the j-th row. Also $P_{1,1} = 1$, if the worker exerts no effort then no output will be produced, and $P_{i,1} = 0$ $i > 1$, if the worker exerts a positive amount of effort some output will be produced. If the firm instructs the worker to exert no effort, then the worker will be said to be laid off.

The worker's preferences over effort and consumption are described by a function $u : A \times C \to \Re$. Hence, u may be regarded as a vector of size $n_a \times n_c$. The elements of u are assumed to be consistent with strict concavity of the worker's preferences over consumption and effort. Additionally, they are assumed to be strictly decreasing in effort and strictly increasing in consumption. Firms provide workers with lotteries over infinite sequences ("careers") of efforts, consumptions, outputs and shocks. Such lotteries are denoted $\pi^\infty \in P(\Omega^\infty)$. [6] The preferences of workers over lotteries are described by the functional U, where:

$$\text{Workers:} \quad U(\pi^\infty) = E_{\pi^\infty} \sum_{t=0}^{\infty} \beta^t u(a_t, c_t).$$

[5] Sleet and Yeltekin (2000) discuss the ability of such discretized models to approximate continuous models.

[6] Where Ω^∞, along with other sets, is endowed with the trivial discrete σ-algebra.

Here $\beta \in (0,1)$ is the agent's subjective discount rate and E_{π^∞} denotes an expectation under the probability measure π^∞. Workers can not be compelled to work. Consequently, the lowest expected utility that a worker can receive is $u(\underline{c}, \underline{a})/(1-\beta)$. Clearly, the highest expected utility that a worker can receive is $u(\overline{c}, \underline{a})/(1-\beta)$. Denote these expected utility bounds by $\underline{w}$ and $\overline{w}$ respectively. Any utility level w between $\underline{w}$ and $\overline{w}$ can be achieved by randomizing between these two bounds, thus the set of possible ex-ante expected utilities that a worker can obtain is the interval $[\underline{w}, \overline{w}]$. Let W be a finite subset of this interval with smallest element $\underline{w}$ and largest element $\overline{w}$.

The firm is assumed to be risk neutral and, in any period, it receives the residual after the worker has been paid: $\theta q - c$. Here $\theta \in \Theta$ is a publicly observable revenue shock that perturbs the firm's payoff. It is assumed to evolve stochastically according to a Markov chain with transition matrix G. The firm is assumed to discount expected future profits at the same rate as the worker and values lotteries over infinite realizations of outputs, consumptions and revenue shocks according to:

$$\text{Firm:}\quad E_{\pi^\infty} \sum_{t=0}^{\infty} \beta^t (\theta_t q_t - c_t).$$

Two different *information* environments will be considered. In the first the worker's effort choice is publicly observed. In the second it is the private information of the worker. Our interest is primarily in the second information environment. The first serves as a benchmark. In both environments, everything apart from effort is common knowledge [7]. In particular, it is assumed throughout that the firm observes the worker's consumption c.

A full information dynamic contract is an element of $P(\Omega^\infty)$ that delivers some initial prescribed (lifetime) utility promise ($w \in W$) to the worker and respects the exogenously given kernel P and the transition matrix G. More formally, if $w \in W$ is the prescribed utility promise and θ is the initial (period 0) revenue shock then $\pi^{\infty,w,\theta} \in P(\Omega^\infty)$ is a full information contract delivering w and consistent with $\theta_0 = \theta$ if:

$$w = E_{\pi^{\infty,w,\theta}}\left[\sum_{t=0}^{\infty} \beta^t u(c_t, a_t)\right] \tag{1}$$

$$\text{For all } \omega^t \in \Omega^t, q_t \in Q, a_t \in A, t$$
$$\sum_{\omega^{t-1}\times a_t \times q_t \times C \times \Theta} \pi^{\infty,w,\theta}(\omega^t) = \sum_{\omega^{t-1}\times a_t \times Q \times C \times \Theta} P(q_t|a_t)\pi^{\infty,w,\theta}(\omega^t), \tag{2}$$

$$\text{For all }\ \phi^t \in \Phi^t, \theta_t, \theta_{t+1} \in \Theta, t;\ \text{with } \theta_0 = \theta$$
$$\sum_{\phi^t \times \theta_t \times A \times Q \times C \times \theta_{t+1}} \pi^{\infty,w,\theta}(\omega^{t+1}) = \sum_{\phi^t \times \theta_t \times A \times Q \times C \times \Theta} G(\theta_{t+1}|\theta_t)\pi^{\infty,w,\theta}(\omega^{t+1}). \tag{3}$$

[7] An interesting extension is to the case in which the firm has some private information in addition to the worker. This double sided moral hazard problem is briefly discussed in the conclusion.

Here $E_{\pi^{\infty,w,\theta}}$ denotes the expectation under the probability measure $\pi^{\infty,w,\theta}$. Notice that the last constraint must hold for the exogenously given initial shock, $\theta_0 = \theta$.

A private information contract must, additionally, respect an "incentive compatibility" condition[8]. This requires that it is in the worker's interest to implement the conditional distributions over effort implied by the contract. Notationally, let $\delta = \{\delta^t\}_{t=0}^{\infty}$ and $\delta_t : \Omega^{t-1} \times A \to A$. Each δ_t is required to satisfy: $\delta_t(\omega^{t-1}, a_1) = a_1$, $\forall \omega^t \in \Omega^t$, and $\delta_t : \Omega^t \times A_e \to A_e$. δ represents a worker deviation strategy. For each history and each recommended effort in A_e, δ gives a possibly alternative action choice in A_e. No deviations are permitted if the firm has recommended the layoff effort state (recall that this effort level is observable). Let $\pi_\delta^{\infty,w,\theta}$ denote a probability measure whose conditional distributions over output have been "modified" by the deviation strategy, i.e.

$$\pi_\delta^{\infty,w,\theta}(\omega^t) \equiv \prod_{i=1}^{t} \frac{P(q_i|\delta_i(\omega^{i-1}, a_i))}{P(q_i|a_i)} \pi^{\infty,w,\theta}(\omega^t) \, . \tag{4}$$

The worker's payoff from adopting a deviation strategy δ, given a contract $\pi^{\infty,w,\theta}$, can then be calculated as follows:

$$U^\delta(\pi_\delta^{\infty,w,\theta}) = E_{\pi_\delta^{\infty,w,\theta}} \sum_{t=0}^{\infty} \beta^t u(c_t, \delta_t(\omega^{t-1}, a_t)). \tag{5}$$

We say that $\pi^{\infty,w,\theta}$ is incentive compatible if:

$$U(\pi^{\infty,w,\theta}) \geq \sup_\delta U^\delta(\pi_\delta^{\infty,w,\theta}) \, . \tag{6}$$

In the sequel, if $\pi^{\infty,w,\theta}$ is a contract (either full or private information) then the conditional probability distribution $\pi^{\infty,w,\theta}(\cdot|\omega^t) \in P(\Omega^\infty)$ will be referred to as the continuation of $\pi^{\infty,w,\theta}$ after ω^t, or just "the continuation contract" when the context is clear.

In the absence of the persistent revenue shocks, one possible dynamic contract is that which repeats the solution to a simpler static problem (in the appropriate information environment) and delivers an expected utility to the worker of $w(1-\beta)$ in each period. However, incentive compatibility requires a spread in the worker's utility conditional on output to induce effort levels above the lowest. Concavity of the worker's utility function ensures that the least costly way for the firm to implement this spread is to adjust both the current and the future utilities of the worker in response to a particular q realization (i.e. both rewards and punishments are smoothed across dates). This introduces additional dynamics into a dynamic private information contract beyond those induced by the persistent revenue shock process. More precisely, even if the θ-revenue shocks are i.i.d., i.e. the rows of G are the same, $\pi^{\infty,w,\theta}(\cdot|\omega^t)$ need not equal $\pi^{\infty,w,\theta}(\cdot|\omega^{t-1})$.

[8] The following formulation follows Phelan and Townsend (1991).

In general, a dynamic contract will condition lotteries over efforts and consumptions at each date on the entire publicly observable history of the relationship between the firm and the worker, i.e. if $\omega^t \neq \omega'^t$ then, typically, $\pi^{\infty,w,\theta}(\cdot|\omega^t) \neq \pi^{\infty,w,\theta}(\cdot|\omega'^t)$. Consequently, dynamic contracts are potentially complicated. A number of papers beginning with Green (1987) and Spear and Srivastava (1987) have shown that dynamic agency problems can be reformulated recursively, with a lifetime utility promise from the principal to the agent acting as a state variable. This means that optimal contracts have a simple Markov representation with the utility promise augmenting the set of "natural" physical state variables.[9] This sort of formulation is now standard in the literature and we proceed in an analogous fashion. Specifically, a recursive formulation is used in which a utility promise w and the observable revenue shock, θ, act as state variables.

2.1.1 The recursive formulation

The details of the recursive formulation are now sketched.

The timing of moves within a period is as follows. The firm and the worker enter the period knowing the current state - i.e. the value of the revenue shock and the worker's current life time utility promise. Each worker then randomizes over efforts. Given an effort realization $a_i \in A$, an output q is realized, according to $P_{i,\cdot}$. The firm then randomizes over worker consumption levels, c, and continuation utility promises, w'. The firm may be thought of as selecting, subject to constraints, a stochastic kernel that gives the probability distribution over current choices and the next period's shock-promise states conditional on each current shock-promise state. To this end let π be a matrix with $n_s \equiv n_w \times n_\theta$ rows and $n_v \equiv n_a \times n_q \times n_c \times n_w \times n_\theta$ columns. Suppose that each row defines a probability distribution, i.e. the elements are non-negative and sum to one. Call the set of such matrices Π. Define $\mathscr{T}_w = \{1, 2, \cdots, n_w\}$ and $\mathscr{T}_\theta = \{1, 2, \cdots, n_\theta\}$. Let $\pi^{i,j}$, $(i,j) \in \mathscr{T}_w \times \mathscr{T}_\theta$ denote the $(i-1) \times n_\theta + j$-th row of $\pi \in \Pi$. We interpret $\pi^{i,j}$ as the probability distribution over current effort, output, consumption and the future promise and revenue shock states given that the current promise is w_i and the current revenue shock is θ_j. π will be called a full information Markov contract if it is consistent with the keeping of promises and the exogenously determined technology P and shock transition G. π will be called a private information Markov contract if it is, in addition, consistent with the provision of appropriate incentives to workers.

More precisely, consider the following constraints:

- The "promise keeping" constraint guarantees that π respects the inherited worker promise. For all $(i,j) \in \mathscr{T}_w \times \mathscr{T}_\theta$,

$$w_i = \sum_{A\times Q\times C\times W\times \Theta} \{u(a,c) + \beta w'\}\ \pi^{i,j}(a,q,c,w',\theta'). \tag{7}$$

[9] This justifies optimizing over sets of Markov contracts to find the optimal contract below.

– Consistency with P. For all $(i,j) \in \mathscr{I}_w \times \mathscr{I}_\theta$ and $(\overline{a},\overline{q}) \in A \times Q$,

$$\sum_{C \times W \times \Theta} \pi^{i,j}(\overline{a},\overline{q},c,w',\theta') = P(\overline{q} \mid \overline{a}) \sum_{Q \times C \times W \times \Theta} \pi^{i,j}(\overline{a},q,c,w',\theta'). \quad (8)$$

– Consistency with G. For all $(i,j) \in \mathscr{I}_w \times \mathscr{I}_\theta$,

$$G(\theta'|\theta_j) = \sum_{A \times Q \times C \times W} \pi^{i,j}(a,q,c,w',\theta'). \quad (9)$$

– Incentive compatibility. In private information environments, contracts must be incentive compatible. More specifically, for each recommended effort, $a \in A_e$, and any possible deviation, $\hat{a} \in A_e$, current utility plus the discounted future utility of taking a should be greater than that of deviating to $\hat{a}$[10]. Thus, for all effort pairs $(a,\hat{a}) \in A_e \times A_e$,

$$\sum_{Q \times C \times W \times \Theta} \{u(a,c) + \beta w'\}\ \pi^{i,j}(a,q,c,w',\theta')$$
$$\geq \sum_{Q \times C \times W \times \Theta} \{u(\hat{a},c) + \beta w'\}\ \frac{P(q \mid \hat{a})}{P(q \mid a)} \pi^{i,j}(a,q,c,w',\theta'). \quad (10)$$

Define $\mathscr{D}_F(w_i,\theta_j) = \{\lambda \in \Delta^{n_v - 1} | \lambda \text{ satisfies } 7-9 \text{ at } (w_i,\theta_j)\}$ and $\mathscr{P}_F = \{\pi \in \Pi | \forall (i,j) \in \mathscr{I}_w \times \mathscr{I}_\theta, \pi^{i,j} \in \mathscr{D}_F(w_i,\theta_j)\}$. In the previous definition Δ^k denotes the k-dimensional simplex. Similarly define $\mathscr{D}_P(w_i,\theta_j) = \{\lambda \in \Delta^{n_v - 1} | \lambda \text{ satisfies } 7-10 \text{ at } (w_i,\theta_j)\}$ and $\mathscr{P}_P = \{\pi \in \Pi | \forall (i,j) \in \mathscr{I}_w \times \mathscr{I}_\theta, \pi^{i,j} \in \mathscr{D}_P(w_i,\theta_j)\}$.

Definition 1 *A full information Markov contract is an element of $\mathscr{P}_F$. A private information contract is an element of $\mathscr{P}_P$.*

An *optimal private information contract* is an element π_* of $\mathscr{P}_P$ such that for each $(i,j) \in \mathscr{I}_w \times \mathscr{I}_\theta$ $\pi_*^{i,j}$ solves the following programming problem:

$$S_P(w_i,\theta_j) = \max_{\pi^{i,j} \in \mathscr{D}_P(w_i,\theta_j)} \sum_{A \times Q \times C \times W \times \Theta} \{(\theta q - c) + \beta S_P(w',\theta')\} \pi^{i,j}(a,q,c,w',\theta') \quad (11)$$

Similarly an *optimal full information contract* is an element π_* of $\mathscr{P}_F$ such that for each $(i,j) \in \mathscr{I}_w \times \mathscr{I}_\theta$ $\pi_*^{i,j}$ solves:

$$S_F(w_i,\theta_j) = \max_{\pi^{i,j} \in \mathscr{D}_F(w_i,\theta_j)} \sum_{A \times Q \times C \times W \times \Theta} \{(\theta q - c) + \beta S_F(w',\theta')\} \pi^{i,j}(a,q,c,w',\theta') \quad (12)$$

[10] Note a_1 can be distinguished from the other actions and therefore it is not necessary to impose incentive compatibility with respect to this action.

3 The limited commitment environment

We now turn to the limited commitment environment. In this environment, a worker can costlessly quit the employment relationship and receive some (possibly market determined) outside utility option, $w_a \in W$. Similarly the firm can costlessly quit the relationship and receive an outside profit option, s_a. Any contract, in addition to specifying a probability distribution over effort, consumption, etc., will also specify a probability distribution over separation dates. To formalize this idea let υ^T denote a T period career, where $\upsilon^T = \{\{a_t, q_t, c_t, \theta_{t+1}\}_{t=0}^{T-1}, T\}$ and let Υ denote the set of possible careers. Here T is a random variable (taking values in $\overline{\Re}_+$) that gives the date at which the worker and firm separate. Let $\psi^{w,\theta}$ denote a probability distribution over the set of possible careers that delivers an initial utility promise w and is consistent with the initial shock θ. Formally, $\psi^{w,\theta}$ satisfies

$$w = \tilde{U}(\psi^{w,\theta}) \equiv E_{\psi^{w,\theta}}\left[\sum_{t=0}^{T-1}\beta^t u(c_t, a_t) + \beta^T w_a\right] \tag{13}$$

and the analogues of (2) and (3):

$$\text{For all } \omega^t \in \Omega^t, q_t \in Q, a_t \in A, t$$
$$\sum_{\omega^{t-1}\times a_t\times q_t\times C\times\Theta}\psi^{w,\theta}(\omega^t) = \sum_{\omega^{t-1}\times a_t\times Q\times C\times\Theta}P(q_t|a_t)\psi^{w,\theta}(\omega^t), \tag{14}$$

$$\text{For all } \phi^t \in \Phi^t, \theta_t, \theta_{t+1} \in \Theta, t;\ \text{with } \theta_0 = \theta$$
$$\sum_{\phi^t\times\theta_t\times A\times Q\times C\times\theta_{t+1}}\psi^{w,\theta}(\omega^{t+1}) = \sum_{\phi^t\times\theta_t\times A\times Q\times C\times\Theta}G(\theta_{t+1}|\theta_t)\psi^{w,\theta}(\omega^{t+1}) \tag{15}$$

In the sequel, we restrict attention to private information limited contracts, which we refer to simply as "limited commitment contracts" to economize on labeling. Suppose then that $\psi^{w,\theta}$ satisfies the incentive compatibility conditions:

$$\tilde{U}(\psi^{w,\theta}) \geq \sup_{\delta}\tilde{U}^{\delta}(\psi^{w,\theta}_{\delta}) \tag{16}$$

where δ is a deviation strategy as before and $\psi^{w,\theta}_{\delta}$ and $\tilde{U}^{\delta}$ are defined analogously to $\pi^{w,\theta}_{\delta}$ and U^{δ} in the obvious way. In order for $\psi^{w,\infty}$ to be a limited commitment contract it must satisfy two additional "limited commitment" or credibility constraints. To begin with, the availability of the worker's outside option means that the worker must receive no less than w_a from the continuation of $\psi^{w,\theta}$ after each ω^t, $t < T$. Formally,

$$E_{\psi^{w,\theta}(\cdot|\omega^t)}\left[\sum_{s=t}^{T-1}\beta^{s-t}u(c_s, a_s) + \beta^T w_a\right] \geq w_a \quad \forall\omega^t \in \Omega^t, t. \tag{17}$$

Furthermore, the firm can sever the employment relationship in order to take an outside profit option, s_a. Hence, it must receive more than s_a from the continuation of $\psi^{w,\theta}$ after each ω_t, $t < T$, i.e.

$$E_{\psi^{w,\theta}(\cdot|\omega^t)}\left[\sum_{s=t}^{T-1}\beta^{s-t}(\theta_s q_s - c_s) + \beta^T s_a\right] \geq s_a \quad \forall \omega^t \in \Omega^t, t. \tag{18}$$

We summarize the above discussion with the following definition:

Definition 2 *Suppose that $\psi^{w,\theta}$ is a probability distribution over the set of careers Υ. Then $\psi^{w,\theta}$ is a limited commitment contract that delivers expected utility w to the worker and is consistent with the shock θ if it satisfies 13-18.*

Define the correspondence $W' : \Theta \rightrightarrows W$ as follows: $W'(\theta) = \{w \in W | \exists$ a limited commitment contract $\psi^{w,\theta}$ that delivers w and is consistent with $\theta\}$. In other words, $W'(\theta)$ gives the set of worker payoffs that can be delivered to a worker by some limited commitment contract in state θ.[11] Note that there may exist some θ such that $W'(\theta) = \emptyset$. These states are ones in which there exists no limited commitment contract that gives both the worker and the firm more than their outside option. We will call these states separation states. The following proposition characterizes the structure of the correspondence W'.

Proposition 1 *W' has the following structure. There is some $j \in \{1, \cdots, n_\theta\} \cup \{\infty\}$ and a pair $\underline{w} \in W^{n_\theta}$ and $\overline{w} \in W^{n_\theta}$, with $\overline{w} > \underline{w}$ such that*

$$W'(\theta_k) = \begin{cases} \emptyset & \textit{if } k < j, \\ \{\underline{w}_k, \cdots, \overline{w}_k\} & \textit{otherwise.} \end{cases}$$

Proof. Sleet and Yeltekin (2000).

In other words, there is a limited commitment contract available that can deliver w and is consistent with θ provided $\theta = \theta_k \geq \theta_j$ and the contract promises the worker $w \in \{\underline{w}_k, \cdots, \overline{w}_k\}$. If, however, $\theta < \theta_j$, then the worker and firm must separate since there is no $w \geq w_a$ such that the firm's payoff from continuing with the contract weakly exceeds s_a.

The optimal limited commitment contract also has a Markov representation (see Sleet and Yeltekin, 2000). Hence, we proceed, as in previous sections, to define a class of Markov contracts. We then state a dynamic programming problem whose solution yields the optimal limited commitment contract. To begin with fix $\pi \in \mathscr{P}_P$. Let $\psi \equiv \pi(\text{Graph}(W'))$ be the matrix comprised of those rows of π that are associated with states in Graph(W'), i.e. if ψ^m denotes the m-th row of ψ then there is a pair $(i(m), k(m))$ such that $(w_{i(m)}, \theta_{k(m)}) \in$ Graph (W') and $\psi^m = \pi^{i(m),k(m)}$. Next define, for $(w_i, \theta_k) \in$ Graph (W'), $\mathscr{E}(w_i, \theta_k)$ to equal the set of probability distributions λ in $\mathscr{D}_P(w_i, \theta_k)$ such that:

– If $\theta' \in \{\underline{\theta}, \cdots, \theta_{j-1}\}$ then all mass is placed on w_a:

[11] Abreu, Pearce and Stacchetti (1990) develop recursive methods for analyzing sequential equilibrium payoff sets in repeated games. Similarly, one could model the interactions between the firm and the worker as a dynamic game. Associated with this game is a subgame perfect equilibrium payoff correspondence, W^c, that maps θ states to sets of equilibrium firm-worker payoffs. W' can then be thought of as the correspondence obtained by projecting the images of W^c onto the worker's utility axis.

$$\sum_{A\times Q\times C\times\{w_a\}\times\{\theta'\}} \lambda(a,q,c,w',\theta') = \sum_{A\times Q\times C\times W\times\{\theta'\}} \lambda(a,q,c,w',\theta'), \quad (19)$$

– otherwise, all mass is placed on $W'(\theta')$.

for each $\theta' \in \Theta$

$$\sum_{A\times Q\times C\times[\underline{w}(\theta'),\bar{w}(\theta')]\times\{\theta'\}} \lambda(a,q,c,w',\theta') = \sum_{A\times Q\times C\times W\times\{\theta'\}} \lambda(a,q,c,w',\theta')$$

We define ψ to be a Markov limited commitment contract if for each row, ψ^m, of ψ,

$$\psi^m \in \mathscr{E}(w_{i(m)}, \theta_{k(m)}) \, . \quad (20)$$

The firm's objective is to maximize its surplus. This maximization can be broken into two steps. Firstly, given its utility promise to its worker, w_i, and the realization of the revenue shock, θ_k, the firm decides whether or not to sever the relationship and take its outside option value:

$$S^*(w_i, \theta_k) = \max(s_a, S(w_i, \theta_k)) \quad (21)$$

where $S(w_i, \theta_k)$ is the value to the firm of continuing with the relationship. S solves the following Bellman like equation, if $(w_i, \theta_k) \in \mathrm{Graph}(W')$,

$$S(w_i, \theta_k) = \max_{\lambda\in\mathscr{E}(w_i,\theta_k)} \sum_{A\times Q\times C\times W\times\Theta} [\theta q - c + \beta S^*(w', \theta')]\lambda(a,q,c,w',\theta') \quad (22)$$

otherwise, $S(w_i, \theta_k) = -\infty$.

Given j and the credible utility sets, $\{\underline{w}_k, \bar{w}_k\}_{k=j}^{n_\theta}$, the above dynamic programming problems can be solved by the same methods used in the full commitment case. Determining the bounds on utility promises and j , however, involves solving for both S and W' simultaneously. In Sleet and Yeltekin (2000) we consider a more general problem, in which the sets A, Q, C and W are not restricted to be finite. We argue that for this problem the value function S and the correspondence W' are fixed points of a monotone operator. This justifies the application of the following algorithm that exploits this underlying monotonicity.

The Algorithm

Step 0: **Initialize**.

Set $\hat{\theta}_0 = \underline{\theta}$.

Set $\underline{w}_0(\theta) = \underline{w}$, $\forall\theta \in \Theta$.

Set $\overline{w}_0(\theta) = \overline{w}$, $\forall\theta \in \Theta$.

These can be used to define an initial domain W'_0 and an initial constraint correspondence $\mathscr{E}_0$ in the obvious way.

Set $S^*_0(\theta, w)$ to some large initial value (e.g. $\theta\overline{q}/(1-\beta)$.)

Set $n = 0$.

Step 1: **Optimization**

Solve:

$$S_n(w,\theta) = \max_{\pi \in \mathscr{E}_n(w,\theta)} \sum_{A\times Q\times C\times W\times\Theta} [\theta q - c + \beta S_n^*(w',\theta')]\pi(a,q,c,w',\theta') \tag{23}$$

Step 2: **Updating**.

2.1. Update $\hat{\theta}_{n+1}$:

Set $\hat{\theta}_{n+1} = \min\{\theta : S_n(w,\theta) \geq s_a \text{ some } w \in W_n'(\theta)\}$

2.2. Update utility bounds:

For $\theta \geq \hat{\theta}_{n+1}$, set $\underline{w}_{n+1}(\theta) = \min\{w : S_n(w,\theta) \geq s_a\}$

And

For $\theta \geq \hat{\theta}_{n+1}$, set $\overline{w}_{n+1}(\theta) = \max\{w : S_n(w,\theta) \geq s_a\}$

Hence update W_{n+1}' and $\mathscr{E}_{n+1}$.

2.3. Update value function:

$$\begin{aligned} S_{n+1}^*(w,\theta) \quad = \quad & s_a \text{ if } w < \underline{w}_{n+1}(\theta) \text{ or } w > \bar{w}_{n+1}(\theta) \text{ or } \theta < \hat{\theta}_{n+1} \\ & S_{n+1}(w,\theta) \text{ otherwise.} \end{aligned} \tag{24}$$

Step 3: **Check for Convergence**

Check if $W_{n+1}' = W_n'$ and $\sup_{(w,\theta)\in \mathrm{Gr}(W_{n+1}')} |S_{n+1}^*(w,\theta) - S_n^*(w,\theta)| < \epsilon$, for some tolerance criterion ϵ. If so STOP, otherwise set $n = n + 1$ and go back to STEP 1.

3.1 Full commitment versus limited commitment

The recursive problems in the full and limited commitment cases differ in their determination of the domain of the firm's value function. Under full commitment this domain is fixed exogenously. Moreover, in this case the sets of value promises available to the firm are chosen independently of the observable productivity (θ) shock. In contrast, under limited commitment the domain of the value function is chosen to be consistent with a pair of outside options available to the firm and worker. Since these are chosen independently of the θ, the set of value promises available to the firm depends upon the θ shock.

4 Parameterizations and computational issues

4.1 Worker's utility

Functional forms and parameterizations are selected that are standard in the macroeconomics literature. This facilitates comparisons with results obtained in that literature. Specifically, the worker's utility function is assumed to be of the following form:

$$u(a,c) = \frac{\left(c^{\alpha}(1-a)^{1-\alpha}\right)^{1-\sigma} - 1}{1-\sigma},$$

where α is the share parameter in the composite consumption-leisure good consumed by the worker. Our base line choice for σ is 1, or, equivalently:

$$u(a,c) = \alpha \ln c + (1-\alpha)\ln(1-a)$$

Our baseline choice for α is 0.36. We undertake some sensitivity analysis around these choices. In particular, we consider the alternative values for σ of 0.5 and 1.5.

4.2 Full information grids

The set of feasible efforts A is set to be $\{0.0, 0.25, \cdots, 0.35\}$ with up to 12 elements between 0.25 and 0.35. 0 corresponds to a "layoff". Recall that the firm can always detect 0 effort. The positive values correspond to non-layoff levels of effort, these are chosen to be relatively close together. The structure of this effort grid can be justified by supposing that a fixed quantity of effort must be expended before any production occurs (in, for example, commuting). Various micro studies (e.g. Ghez and Becker, 1975; Juster and Stafford, 1991) have found that households allocate about one third of their time to market activities. If effort is roughly proportional to hours worked then it might be reasonable to suppose that worker's allocate about one third of their effort endowment to production. We do suppose this. We also suppose that the agent can undertake small on the job variations in effort that can not be directly observed. This variation is described by the set $\{0.25, \cdots, 0.35\}$.

The output grid Q is set to $\{0, 1, 2\}$. 0 is the layoff level of output. The other output levels occur if some effort is expended. The θ grid is varied across applications. The baseline grid is set to $\{0.9, 1.0, 1.1\}$. The "large" shock grid is set to $\{0.5, 1.0, 1.5\}$. These grids and the stochastic kernels described below are selected to ensure that for at least some utility promises the full commitment private information contract in the baseline case selects "interior" effort levels between 0.25 and 0.35 with positive probability (i.e. they are selected so that the optimal contract delivers some on the job effort variation across utility promise levels).

Atkeson and Lucas (1995) and Phelan (1995) have shown analytically that the asymptotic properties of optimal dynamic contracts in hidden information environments depend on the imposed (or derived) bounds of the feasible promise set. This model is a hidden action one where analytical results are more scarce. Suppose that the worker's utility is additively separable in consumption and effort ($u(c,a) = v(c) - h(a)$) and that $1/v'(c)$ is convex. Then, absent grid constraints, it may be shown that worker consumption is a non-negative supermartingale (see Rogerson, 1985). Hence, by Doob's theorem, worker consumption almost surely converges. It is also possible to show that there exists some $\hat{w} \in (-\infty, \infty]$ such that for all $w \geq \hat{w}$, the firm sets worker effort to its lowest level and completely smoothes the worker's consumption. We conjecture that the convergence of worker consumption implies either that the worker's utility promise converges to $\hat{w}$ or consumption converges to zero and the promise converges to a small number. If this conjecture is correct, it suggests that the upper and lower bounds of the worker's promise grid will constrain the asymptotic behavior of the contract. We set these grid bounds consistently with our consumption and effort grids. Thus, $w_1 = u(c_1, a_1)/(1-\beta)$. and $w_{n_w} = u(c_{n_c}, a_1)/(1-\beta)$. In the calculations below, we set the w and c grids to contain between 50 and 100 points.

4.3 Stochastic kernels

The technology relating efforts to the probability of each output, namely $P(q|a)$, is set according to the following rules. $P(0|0)$ is set equal to 1. For $a > 0$, if there are n_a actions, then for the n-th action ($n > 1$) $P(q_1|a_n) = 0$, $P(q_2|a_n) = 2/3 - 1/3 \times (n-2)/(n_a - 2)$ and $P(q_3|a_n) = 1/3 + 1/3 \times (n-2)/(n_a - 2)$. The baseline Markov chain for Θ is constructed to be monotone and to give an autocorrelation coefficient of 0.8.

4.4 Outside utility options

The limited commitment model introduces two additional parameters - the outside utility options of the worker and the firm - w_a and s_a. The baseline values are taken to be $u(c_1, a_1)/(1-\beta)$ and 2 respectively. However, we undertake some sensitivity analysis around these values.

5 Numerical approach

All of the coding was undertaken in FORTRAN. Heavy use was made of the IMSL linear programming routine DDLPRS. This routine implements a revised simplex algorithm.

In addition to discretization and the revised simplex algorithm our algorithms incorporate modified policy iteration and the Dantzig-Wolfe decomposition. The former is a well known acceleration technique for dynamic programming problems (see Judd, 1998). The latter is a well known method in the operations research community for solving linear programming problems with a so called block angular structure. Prescott (1998) recognized that contracting models had such a structure and was the first to suggest applying this approach to these models. His paper contains full details of the method. However, we sketch the key issues below in the context of a full commitment problem. Roughly speaking, the idea is to split the firm's choice problem in any state into a master problem and a collection of sub-problems. Each of the latter has as its objective a weighted sum of the firm's and the worker's expected payoffs. These weights are determined by the dual variables of the master problem. Each sub-problem is parameterized in a particular agent action. The i-th sub-problem maximizes its objective subject to the constraint that the i-th action is implemented in an incentive compatible way. The solution to this sub-problem generates a pair of expected payoffs to the firm and the worker. Similarly, the set of sub-problems generates a set of payoff pairs. The elements of this set are then checked to see if they satisfy an optimality criterion. This criterion also depends on the dual variables obtained from the master problem. If all elements of the set satisfy the optimality criterion the algorithm terminates. If not those that fail are added to the existing set of payoff pairs obtained in previous iterations of the algorithm. The master program then solves for the optimal lottery over these payoff pairs consistent with promise keeping. The dual variables from this problem are then used to form the objective functions and the optimality constraints in the next round of sub-problems. During each iteration of the Dantzig-Wolfe algorithm a collection of payoff sets are built.

Convergence may be further accelerated by a number of straightforward, practical steps. The Dantzig-Wolfe algorithm builds up sets of payoff pairs for each optimization. We found it useful to save these sets at each optimization and reuse them in subsequent optimizations within the policy iteration step. Additionally, we relaxed the Dantzig-Wolfe optimality condition in the early stages of the policy iteration so that in these stages "rough" solutions are found. We then progressively tightened this criterion in later iterations. Once converged value functions have been calculated they may, of course, be used as initial conditions in other related policy iterations. In particular, the value functions calculated as solutions to full commitment problems may be used as starting points in the calculation of limited commitment ones.

Table 1 displays the run times for a single dynamic optimization (i.e. for a specific θ and utility promise w) with and without the Dantzig-Wolfe decomposition, for different n_w numbers. Note that for $n_w = 50$ and $n_\theta = 3$, one iteration in the dynamic programming problem consists of 150 of these optimizations. The following run times were computed for $n_q = 3$, $n_a = 4$ and $n_\theta = 3$. These numbers suggest that for $n_w = 50$, using the Dantzig-Wolfe decomposition saves about 10 minutes per iteration.

Table 1. Run times

	$n_w = 15$	$n_w = 30$	$n_w = 50$
With Dantzig-Wolfe	0.1s	0.7s	2.3s
Without Dantzig-Wolfe	0.8s	3.0s	6.6s

6 Results for the full commitment case

This section provides some initial characterization of full commitment contracts.

6.1 Effort functions and temporary layoffs

We refer to the expected effort of a worker conditioned on the current θ shock and utility promise as an "effort policy function". Figure 1 illustrates the effort policy

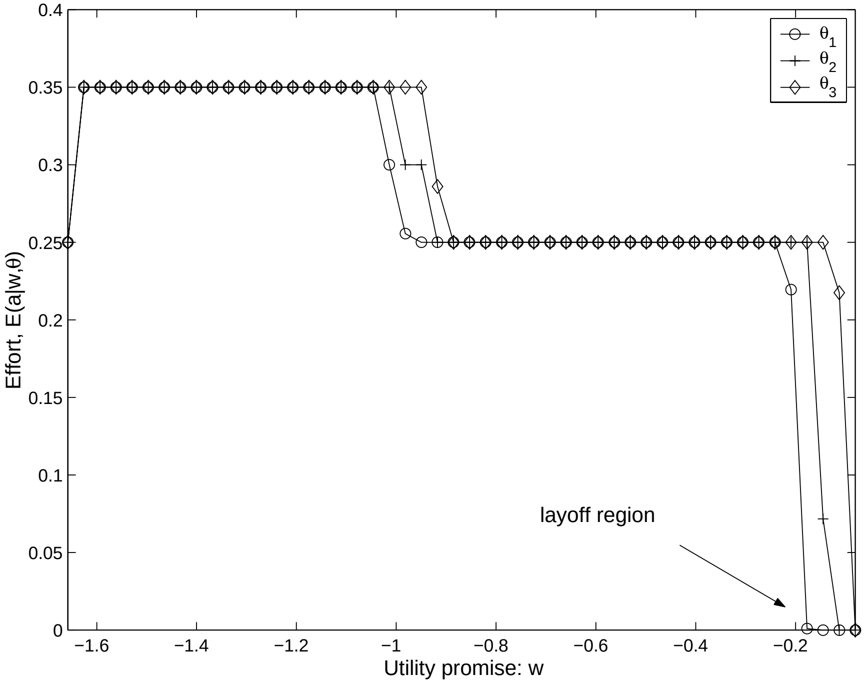

Figure 1. Full commitment: effort, benchmark case

function implied by the benchmark full commitment contract. In the figure, each curve shows conditional expected effort as a function of the utility promise for a fixed θ. Separate curves are drawn for different θ values.

For each θ the effort policy functions have a similar structure. They rise sharply from 0.25 to the maximal effort level, 0.35. They are then stable at 0.35 until some critical utility promise is reached, at which point they come down to 0.25. They are then stable at 0.25 until a second critical utility promise level is reached, at which point they come down to zero. The initial low effort level is a consequence of the construction of the grids. Effort levels above 0.25 require a spread of current consumptions and future utilities to render them incentive compatible. This in turn requires that higher outputs are rewarded with consumptions and utility promises above the lowest values in the corresponding grids. But then by choosing an effort level of 0.25 the worker can ensure himself a utility in excess of $\underline{w}$. Incentive compatibility implies he must obtain at least this from a higher effort level which violates promise keeping. As the utility promise rises above $\underline{w}$ and the constraints associated with the grid boundaries are relaxed, the effort function jumps. Further increases in w raise the cost of inducing higher effort[12]. Consequently, the effort level falls, firstly to 0.25 and

[12] There are two aspects to this cost. One is the resource cost of providing the utility necessary to compensate the worker for the extra effort. The other is the cost of compensating the worker for the variation in his consumption stream necessary to render the extra effort incentive compatible.

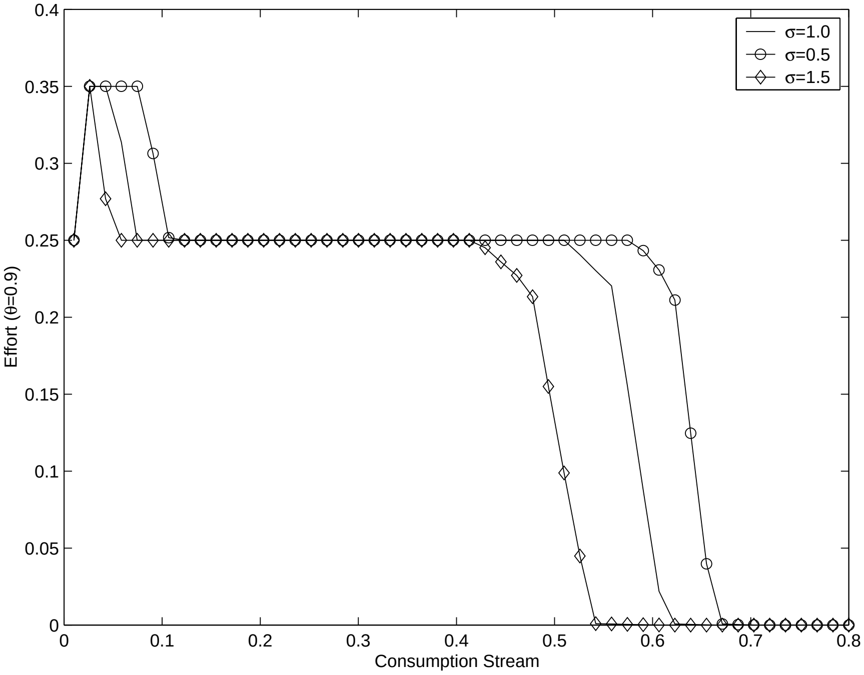

Figure 2. Effort, $\theta = 0.9$

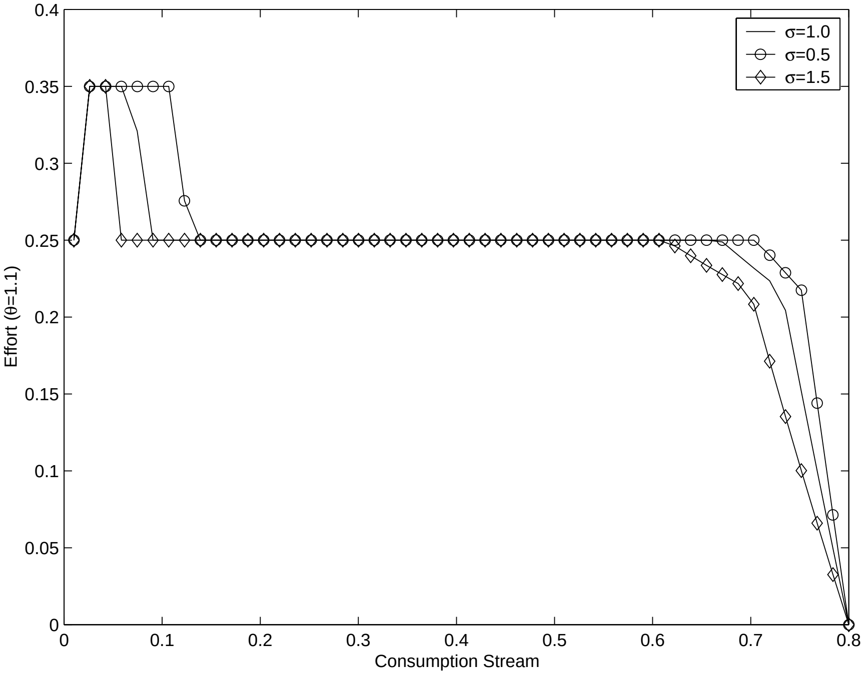

Figure 3. Effort, $\theta = 1.1$

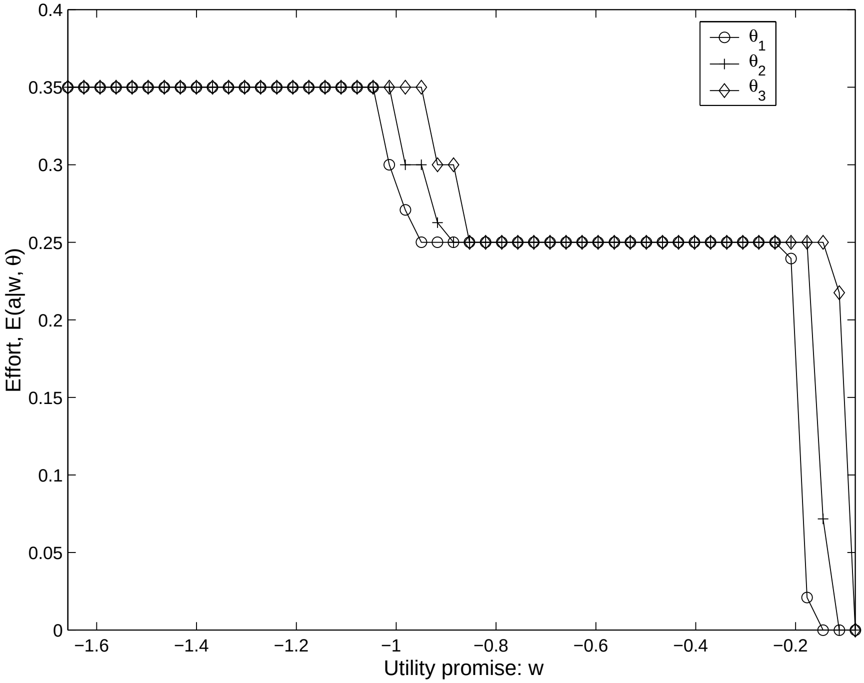

Figure 4. Full info.: effort, benchmark

ultimately to 0. For a particular θ, the region of the utility promise space in which effort is zero represents that θ's layoff region. Unsurprisingly, the critical promise levels described above tend to rise as θ rises. In this sense higher θ shocks induce higher effort levels.

Figures 2 and 3 compare effort functions as the utility parameter σ is altered. In order to make the comparisons we "renormalize" the domains of the effort functions. To do this we calculate the value to the worker of receiving various fixed consumption streams (and zero effort) under different σ values. We then calculate the expected effort at these values for $\theta = 0.9$ (Figure 2) and $\theta = 1.1$ (Figure 3) under the optimal contracts. The figures indicate that higher effort levels are associated with lower values for σ. Variations in utility are necessary to render higher effort levels incentive compatible. As σ (the worker's coefficient of relative risk aversion) falls, these variations become cheaper to implement and effort levels tend to rise.

We conclude this subsection with a comparison of the effort levels under full commitment and full information[13]. Figure 4 shows the full information effort function under the benchmark parameterization. A comparison with Figure 1 shows that full information effort policy functions lie slightly above the full commitment policy functions. This is most marked for very low utility promises and for promises in the region in which the policy functions adjust from 0.35

[13] i.e. contracts in which the worker's effort is observed.

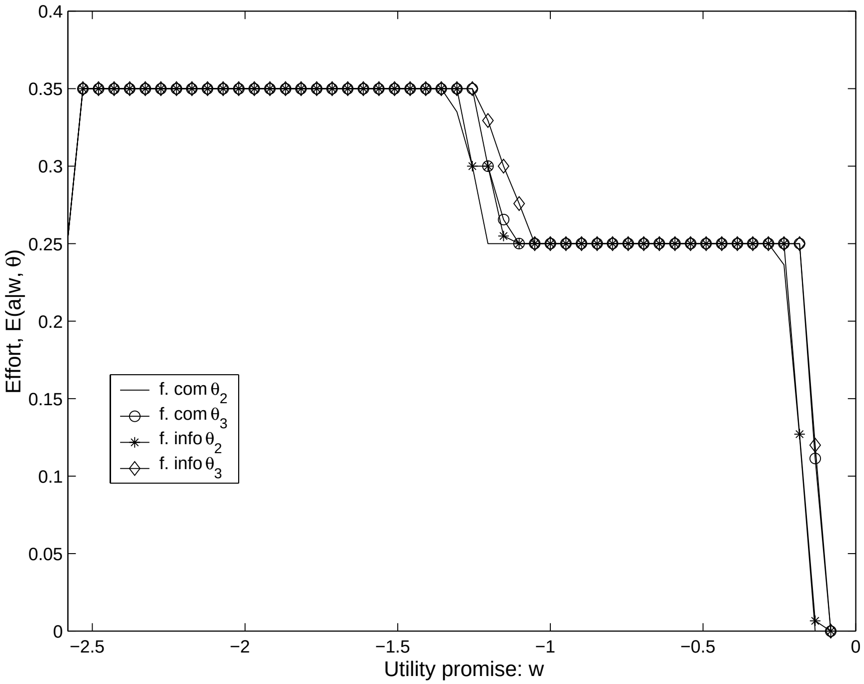

Figure 5. Full info, commit.: effort, benchmark

down to 0.25. Figure 5 incorporates full commitment *and* full information effort functions for θ values θ_2 and θ_3 under a different parameterization. In this case σ is raised to 1.5. Now the difference between the effort policy functions is more marked (compare the functions for $\theta = \theta_3$ and promises between -1 and -1.4). The reason is simply, that the higher worker risk aversion raises the cost of providing the utility spreads necessary for the incentive compatibility of higher effort levels. The firm responds by economizing on effort in the full commitment case.

6.2 Consumption and utility promises

Figures 6 and 7 show the consumption and continuation utility promise functions implied by the benchmark full commitment contract. In this case the θ shock is held constant, and the different curves show the expected consumption and utility levels conditional on the utility promise and the realized output. Again the current utility promise is placed on the horizontal axis.

Variations in ex post utility are necessary to render higher effort levels incentive compatible. These are achieved by varying current consumption and the continuation utility promise contingent on the output realization (raising it for the good output realization q_3, and cutting it for the low realization, q_2). As is well known the use of future utility promises to smooth rewards and punishments

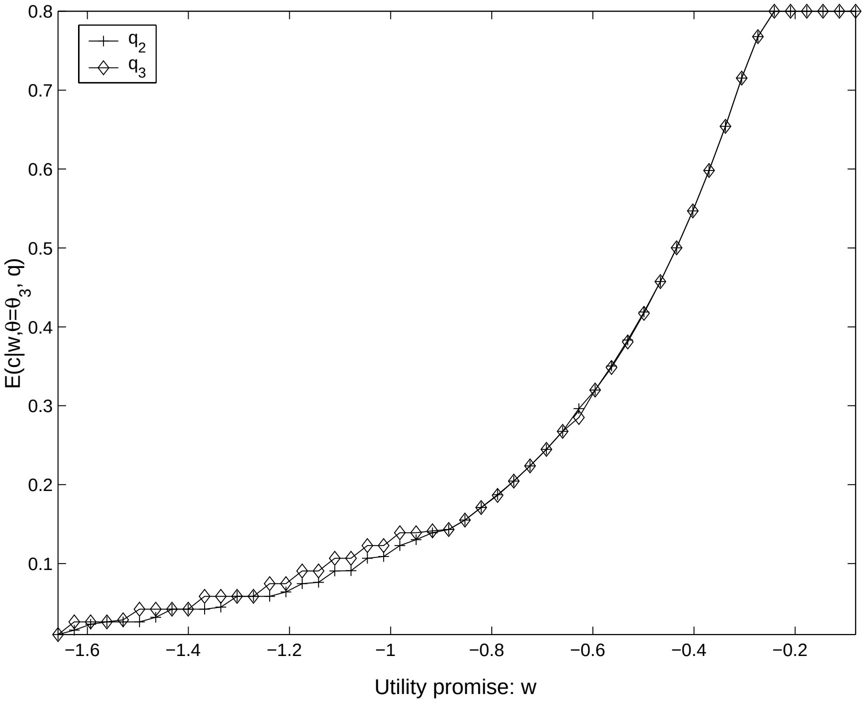

Figure 6. Consumption, benchmark

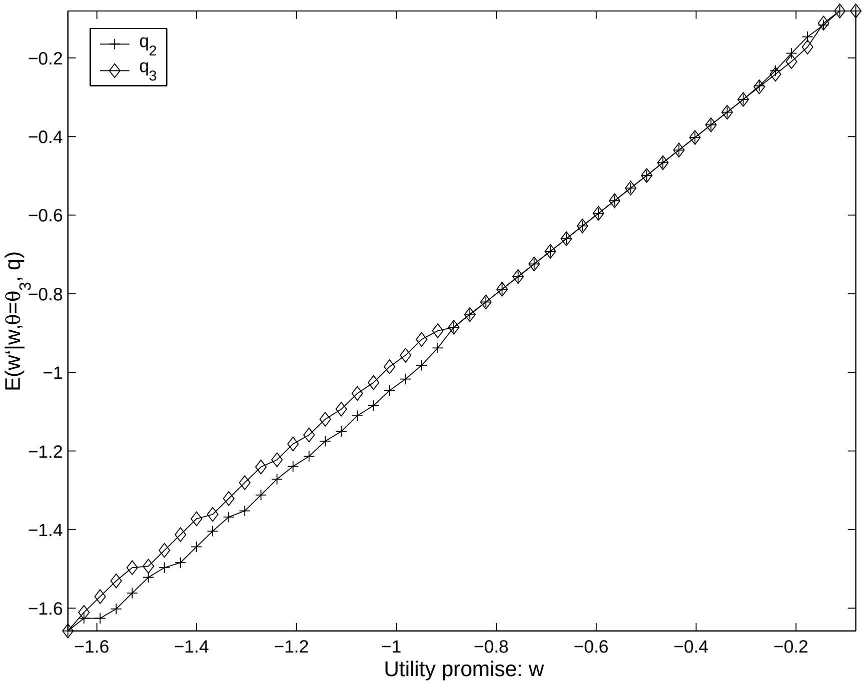

Figure 7. Utility promise, benchmark

enables the firm to motivate the worker in an efficient way. Recall that the layoff (zero effort) state is publicly observable. Hence defections from 0 effort or from the 0.25 effort level down to 0 effort are also observable. If (as the calculations confirm) upward incentive compatibility constraints do not bind on the optimal contract then no variation in utility is necessary to ensure the incentive compatibility of the 0.25 (or 0) effort levels. Consequently, as Figure 7 illustrates, for utility promises sufficiently high, effort is set to 0.25 and both consumption and the continuation promise are held constant across output realizations. Although, Figure 7 is drawn only for $\theta = \theta_3$, the policy functions for the other values of θ are similar. Thus, there is a critical utility promise above which the worker obtains complete consumption insurance and effort is set to 0.25. It follows that with the indivisibility in the provision of effort that we have incorporated into the model the endogenous utility promise dynamics can not drive the worker from a non-layoff to a layoff state. It is the incentive compatibility constraint that induces these endogenous dynamics.[14] This is relevant in those regions of the utility promise domain for which the effort function exceeds 0.25. The region of the domain in which the effort function takes the value of 0.25 in all θ states effectively separates the “endogenous dynamics region” of the state space from the “layoff region”. With this parameterization layoffs only occur if the initial utility promise is high enough. This result need not occur if the effort indivisibility was reduced or the worst shock was decreased in value.

The utility promise and effort functions shed light on the stochastic process induced by the optimal contract. Roughly speaking high θ shocks tend to induce higher effort (and hence output) in the current period. They also tend to raise the probability of a high θ shock in the next period, and, hence, the probability of a high effort in that period. On the other hand, increases in current output lead to increases in the utility promise and this tends to depress future effort.

6.3 Careers

We conclude this section by showing some sample “career paths” induced by the optimal contract. The first sample (Figure 8) is a “long career” of 900 periods. It converges to a utility promise level close to the left hand edge of the 0.25 effort region. From period 578 onwards effort is kept at 0.25 and the worker obtains complete insurance. This convergence is consistent with the conjecture of Section 4. Interestingly, however, it is not to immiseration as in Thomas and Worrall (1990).[15]

Figures 9 and 10 show shorter careers. In the first the worker is successful. Starting at a fairly low initial utility level, he provides high effort (periods 1-30, and, intermittently, from 30 to 40). He produces sufficient high output levels that

[14] More precisely, by endogenous dynamics we mean variations in continuation contracts that are induced by the need to provide incentives rather than by exogenously determined changes in the physical environment.

[15] Thomas and Worrall provide a theoretical result proving the convergence of the worker’s utility to “immiseration” ($-\infty$) in a hidden information model. Phelan (1996) links this result to specific properties of the worker’s utility function.

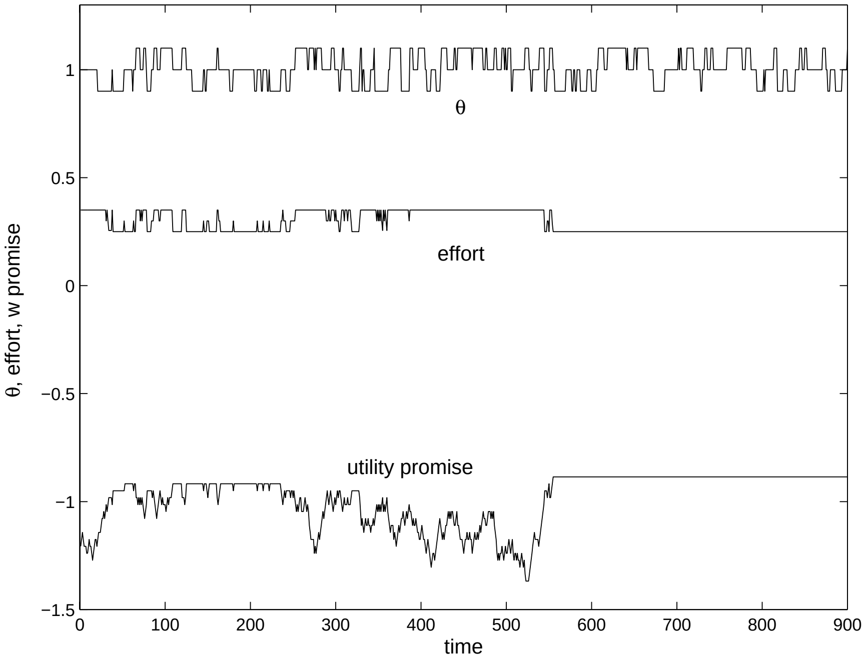

Figure 8. A long career

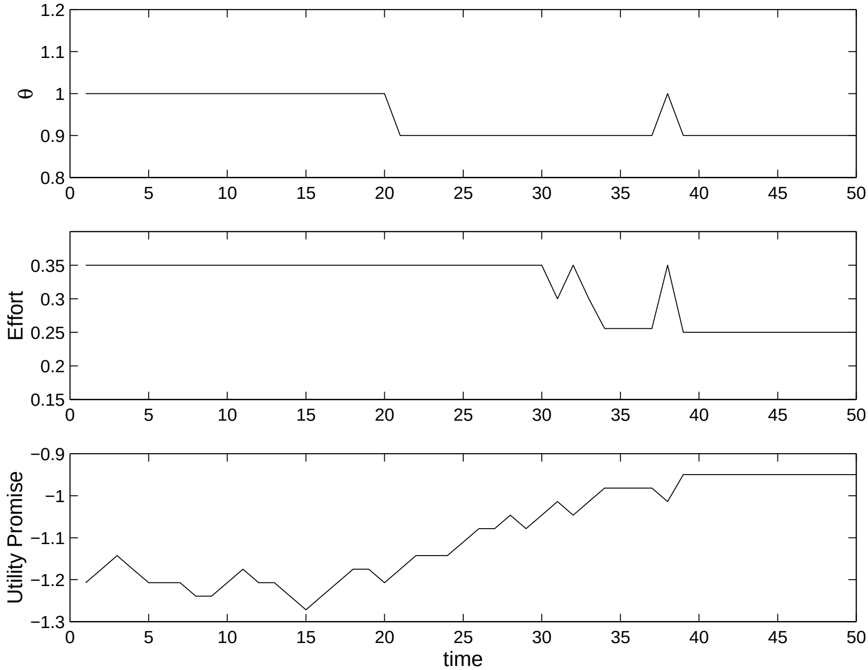

Figure 9. A successful career

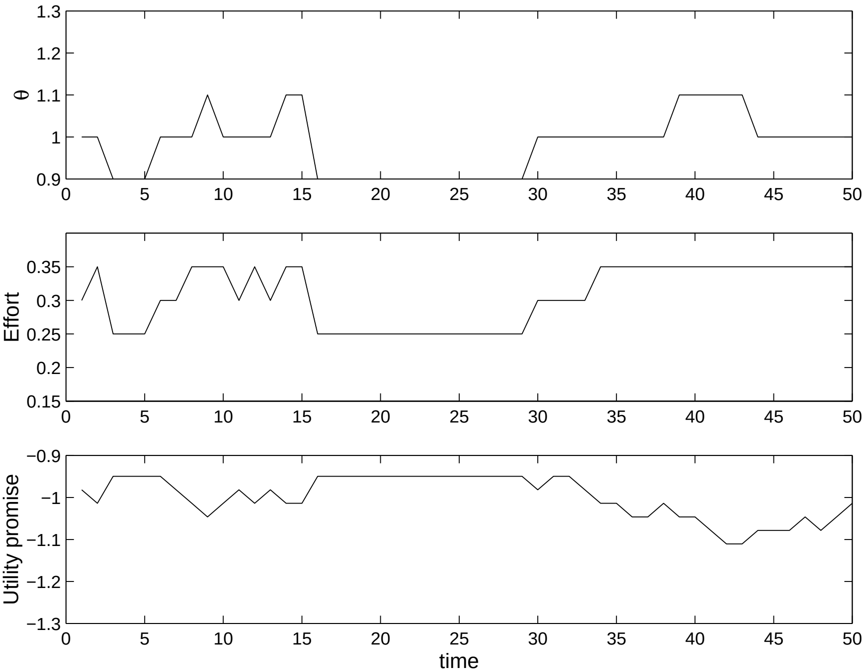

Figure 10. A mixed fortunes career

his promise drifts upwards (periods 15 to 40) and eventually his utility promise converges. In the second figure, the fortunes of the agent are more mixed. The realization of a series of higher θ shocks from period 30 causes the firm to request higher effort from the worker. The worker is unlucky, however, producing several low outputs. This causes his utility promise to drift downwards.

7 Results for the limited commitment case

7.1 Value functions

We begin by showing the value functions for the "large shock" full commitment and limited commitment cases. In the latter case the outside options of the firm and worker are set to 1.75 and $\underline{w}$ respectively. In Figure 11, the solid curves illustrate the value functions for the full commitment case. The lines with $*$'s in them are for the limited commitment case. For each type of contract there are separate lines for the different θ shocks. In both cases the value functions are increasing in θ. They are initially increasing and then decreasing in the utility promise. The limited commitment value function is only defined on the credible utility promise set (i.e. the graph of the correspondence W'). For θ_3 and θ_2, $W'(\theta) \subset [\underline{w}, \overline{w}]$. However, $W'(\theta_1) = \emptyset$ and θ_1 is a separation state. Hence, no limited commitment value function is drawn for θ_1. As the figure shows, even for those (w, θ)'s for which both the limited and full commitment value functions are

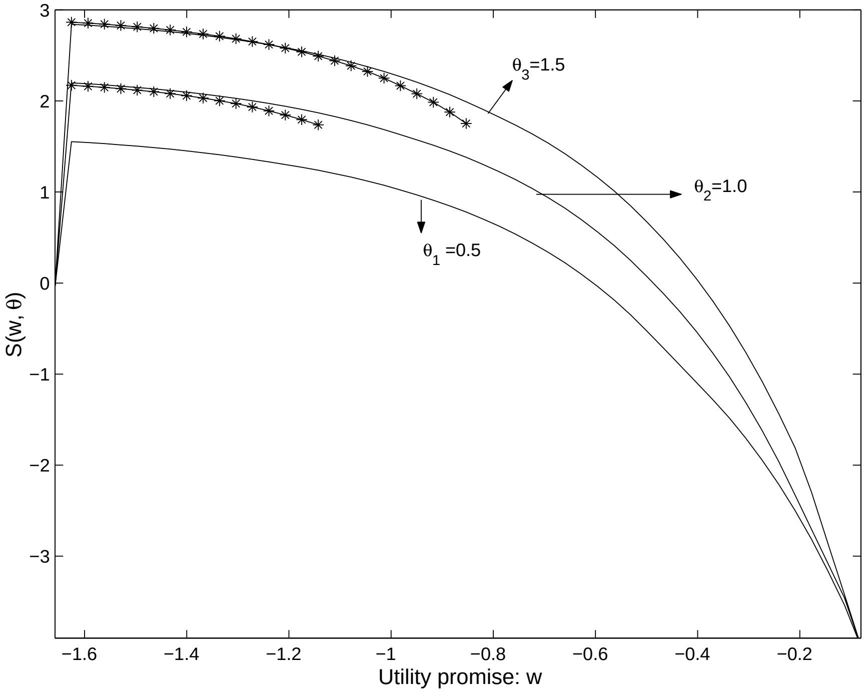

Figure 11. Value functions

defined, the former lies below the latter. This is especially noticeable for (w, θ) pairs such that w is close to the maximum value of $W'(\theta)$ - in these regions the firm is particularly constrained in its ability to use continuation utility promises.

7.2 Effort functions

Figures 12 and 13 show effort functions for full and limited commitment contracts on different θ sets. Figure 12 shows them on the benchmark θ set, $\{0.9, 1.0, 1.1\}$. In this case, for those (w, θ)'s for which the limited commitment value function is defined, the full commitment and limited commitment effort functions coincide. They equal 0.35. The credibility constraint on the firm has the effect of ruling out those high (and very low utility promise states) associated with lower effort levels. There are no separation states in this case and the calculations suggest that the Markov chain induced on the finite (w, θ)-grid is irreducible, and, hence, ergodic.

Figure 13 repeats the exercise, but for the "large" θ shock set, $\{0.5, 1, 1.5\}$. Now, under limited commitment, θ_1 is a separation state. Clearly, it is absorbing. Also at some (w, θ) pairs in GraphW', effort is lower under the limited commitment contract than under the full commitment one. For example, check the full commitment θ_3 curve (solid line) against the limited commitment θ_3 curve (line with diamonds). These reduced effort levels stem from the fact that if a firm wishes to motivate a worker by making a high expected future lifetime utility

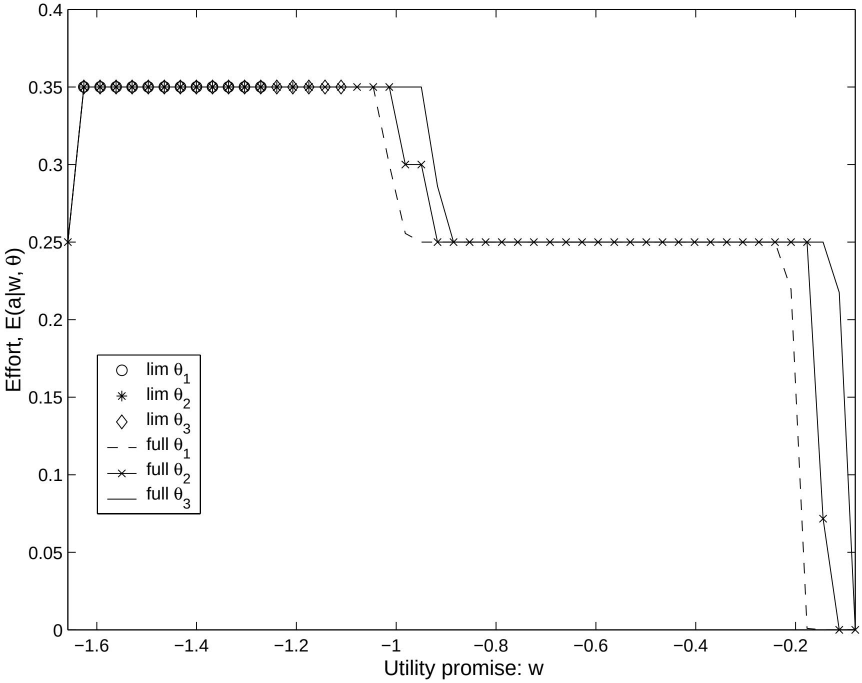

Figure 12. Effort functions: Benchmark shock set

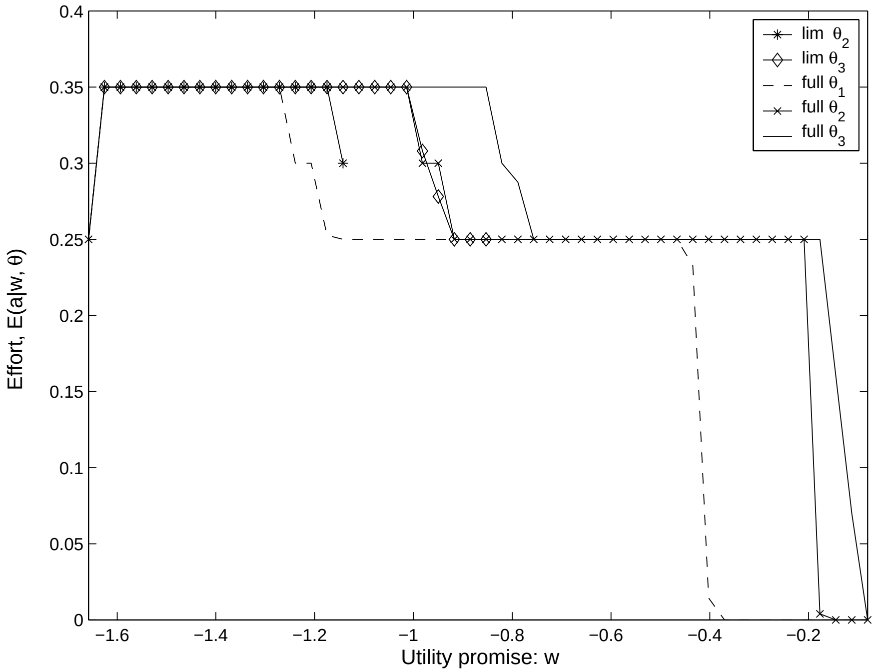

Figure 13. Effort functions: Big shock set

promise it must do so by concentrating high future rewards in high θ states. The effect will be to raise the cost of motivating workers and hence to either reduce average worker effort or to raise worker consumption volatility or both.

7.3 Consumption and utility promise functions

Figures 14 and 15 show consumption and utility promise policy functions under the large θ shocks assumption. These policy functions are shown for $\theta = \theta_3$ and are drawn on the endogenously determined utility promise domain. Comparison with the corresponding full commitment case reveals that the consumption functions are tilted upwards; the utility promise functions are tilted downwards slightly. The separation state, and the firm's inability to commit to payments after a separation, constrains firm-worker risk sharing. It also means that there is a positive probability of the worker receiving a low outside option utility. In order to satisfy the promise keeping constraint, utility (and consumption in particular) must be concentrated on the higher (non-separation) θ states. This results in the upward tilting of the consumption policy function.

7.4 Careers

We conclude by showing a figure illustrating a "limited commitment" career under the large shocks parameterization. In this case a series of good θ shocks

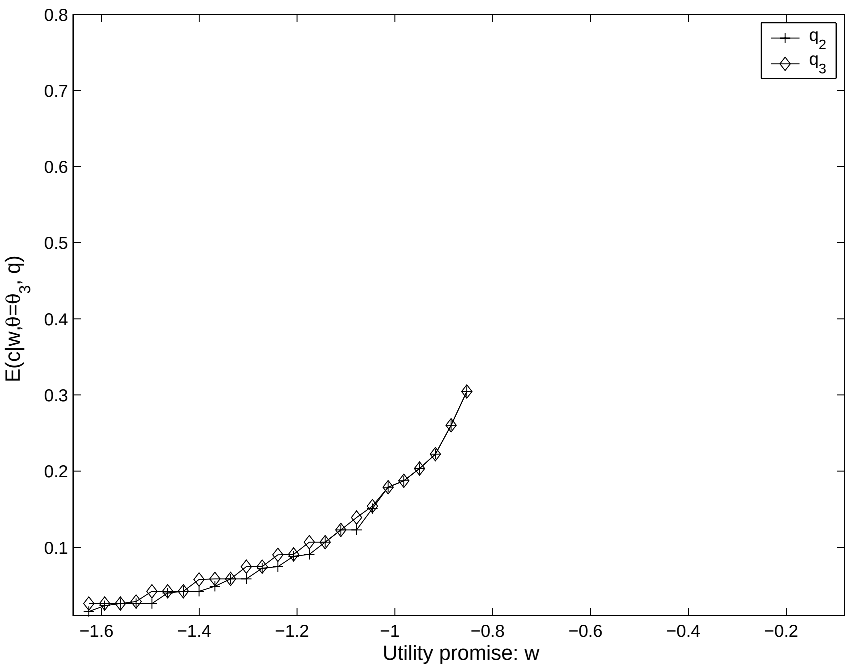

Figure 14. Lim. Commit. consumption function

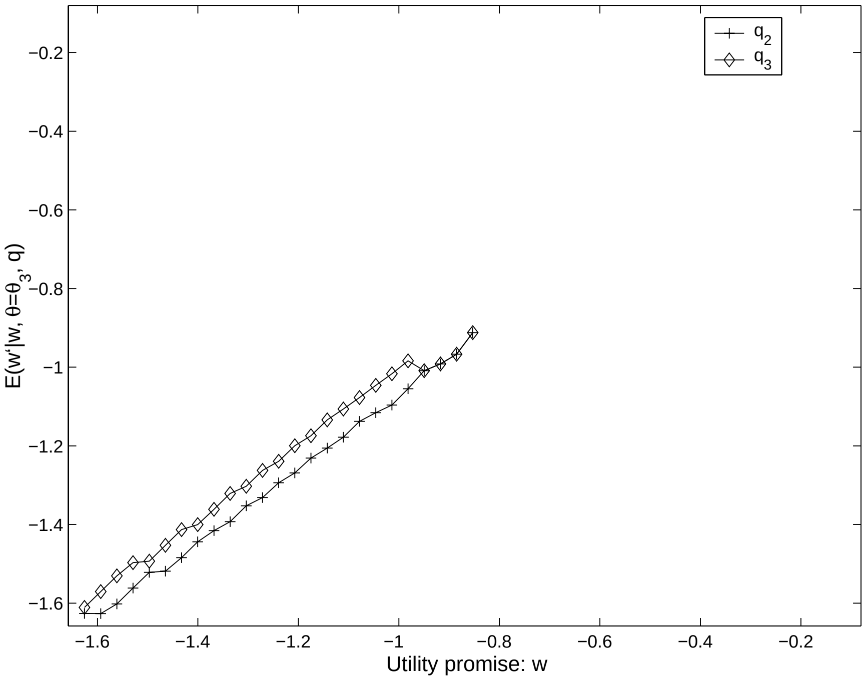

Figure 15. Lim. Commit. promise function

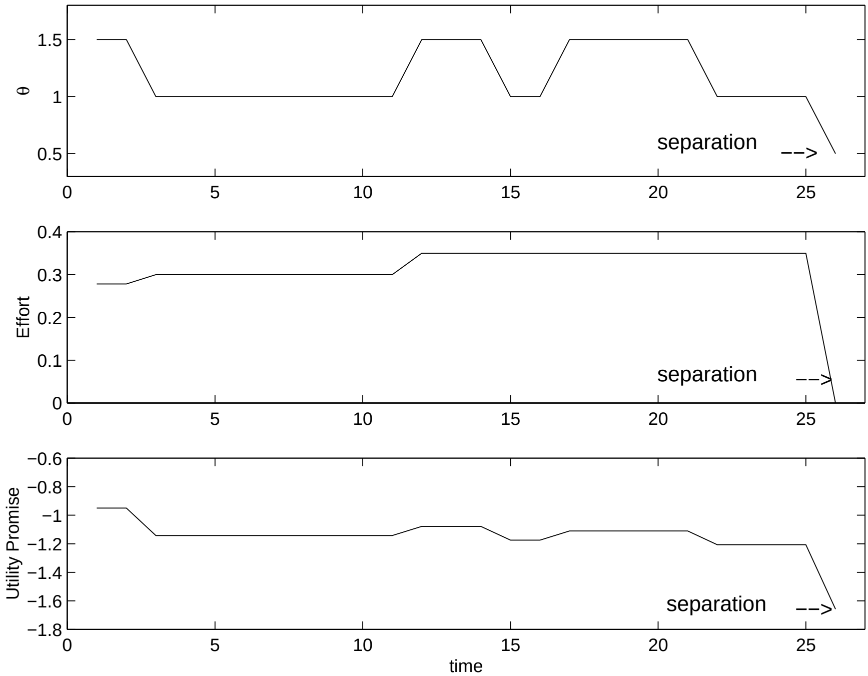

Figure 16. Lim. Commit. career

occur in the second half of the career. These are associated with an increase in effort, but only modest variations in the worker's expected utility. Eventually, a bad θ shock occurs, the firm and the worker separate and the worker's expected utility dips to his outside option.

8 Conclusion

This paper analyzes the fluctuations of effort and employment across dates and across productivity states in a model of moral hazard. Two types of environments are considered. In the first contracts are permanent and layoffs only temporary. In the second both firms and workers have access to an outside utility option. When the contractual relationship can not guarantee a payoff to both parties in excess of these outside options a separation occurs. In general contractual relationships are finitely lived and their durations are determined endogenously. A recursive formulation of the firm's problem in each of these environments is provided. For each environment the temporary layoff and permanent separation regions of the state space are computed. In the full commitment case the level of effort obtained by the contract is below that obtained by the full information case. Incentive compatibility constraints reduce the amount of consumption smoothing relative to this case. We find simulated examples in which the worker's career converges to a low effort-high consumption-complete consumption smoothing outcome. Limited commitment considerations further constrain the firm - effort levels and the amount of consumption smoothing are further reduced. Examples are found in which separation states occur. Restricted in its ability to provide credible dynamic incentives, a firm relies more heavily on current consumption to compensate and motivate workers in a limited commitment environment.

The techniques applied in this paper can be used to address issues of job security, contract duration and the optimal provision of incentives in a variety of related settings. Indeed, several important extensions suggest themselves. Firstly, one could incorporate shocks that are privately observed by the firm. This would convert the current problem into one of "double sided moral hazard". Static labor contracts under private firm information were extensively studied in the early 1980's. It would be interesting to see how far the insights of that literature survive the incorporation of hidden worker effort and the extension to a dynamic contract. Such a problem could be solved by introducing an additional set of firm based incentive compatibility constraints into the optimizations undertaken in this paper. Secondly, the current analysis could, in principle, be extended to a full general equilibrium one. This would entail the endogenous determination of the outside utility options $\{w_a, s_a\}$ and, perhaps, the interest rate paid by firms. One approach to this problem is that of Phelan (1995) who assumes that firms face an exogenous cost, c, of separating (and recontracting). Phelan exogenously sets the firm's interest factor to β and s_a to 0. In the absence of any observable shocks, w_a is endogenously determined by:

$$s_a = 0 = S(w_a) - c \tag{25}$$

where S is the firm's value function. Costly separation could also be obtained by embedding the current contracting problem into a search model with some rule for bargaining over the joint surplus after a match. Alternatively, w_a could be obtained by modeling a worker home production opportunity.[16] We leave these topics for further research.

References

Abreu, D., Pearce, D., Stacchetti, E.: Towards a theory of repeated games with imperfect monitoring. Econometrica **58**, 1041–1063 (1990)

Atkeson, A., Lucas, R.: On efficient distributions with private information. Review of Economic Studies **59**, 427–453 (1992)

Ghez, G., Becke, G.: The allocation of time and goods over the life cycle. New York: Columbia University Press 1975

Green, E.: Lending and the smoothing of uninsurable income. In: Prescott, E., Wallace, N. (eds.) Contractual arrangements for international trade. Minneapolis, MN: University of Minnesota Press 1987

Judd, K.: Numerical methods in economics. Cambridge, MA: MIT Press 1998

Juster, F., Stafford, F.: The allocation of time: empirical findings, behavior models and problems of measurement. Journal of Economic Literature **29**, 471–522 (1991)

Kocherlakota, N.: Efficient bilateral risk sharing without commitment. Review of Economic Studies **63**, 595–609 (1996)

Phelan, C.: Repeated moral hazard and one-sided commitment. Journal of Economic Theory **66**, 488–506 (1995)

Phelan, C., Townsend, R.: Computing multi-period information constrained optima. Review of Economic Studies **58**, 853–881 (1991)

Prescott, E.S.: Computing moral hazard using the Dantzig-Wolfe decomposition. Working Paper 98-06, Federal Reserve Bank of Richmond (1998)

Rogerson, W.: Repeated moral hazard. Econometrica **53**, 69–76 (1985)

Spear, S., Srivastava, S.: On repeated moral hazard with discounting. Review of Economic Studies **54**, 599–618 (1987)

Sleet, C., Yeltekin, Ş.: Dynamic labor contracts with temporary layoffs and permanent separations. Mimeo, Northwestern University and University of Texas-Austin (2000)

Jonathan, T., Worrall, T.: Income fluctuations and asymmetric information. Journal of Economic Theory **51**, 367–390 (1990)

Yeltekin, Ş.: Dynamic principal-multiple agent contracts. Manuscript, KGSM, Northwestern University (1999)

[16] In general the extension to an equilibrium model will require searching for variables (prices or outside utility options) that, in some sense, clear markets. The large linear programs used to solve for the dynamic contracts are costly to compute and any additional "equilibrium loops" necessitated by this search could lead to programs with long running times. Such research would, therefore, need either to find simpler formulations of the problem or additional acceleration techniques for the computations. In the first regard, see Phelan (1995) who assumes hidden information and exponential utility.