Risk & Sustainable Management Group

Risk & Uncertainty Program Working Paper: 6/R04

Research supported by an Australian Research Council Federation Fellowship
http://www.arc.gov.au/grant_programs/discovery_federation.htm

Fixed wages and bonuses in agency contracts: the case of a continuous state space

Maria Racionero
Australian National University

and

John Quiggin
Australian Research Council Federation Fellow, University of Queensland

Schools of Economics and Political Science
University of Queensland
Brisbane, 4072
rsmg@uq.edu.au
http://www.uq.edu.au/economics/rsmg

THE UNIVERSITY OF QUEENSLAND AUSTRALIA

Fixed wages and bonuses in agency contracts: the case of a continuous state space

Maria RACIONERO* and John QUIGGIN†
Risk and Sustainable Management Group
Risk and Uncertainty Working Paper 6/R04

22 June 2004

Abstract

In this paper, we extend the state-contingent production approach to principal-agent problems to the case where the state space is an atomless continuum. The approach is modelled on the treatment of optimal tax problems. The central observation is that, under reasonable conditions, the optimal contract may involve a fixed wage with a bonus for above-normal performance. This is analogous to the phenomenon of 'bunching' at the bottom in the optimal tax literature.

Keywords: Moral hazard; State-contingent production; Principal-agent problem

JEL Classification: D8, D2

*School of Economics, Australian National University, ACT 0200, Australia
(e-mail: maria.racionero@anu.edu.au)
†School of Economics, University of Queensland, QLD 4072, Australia

Thanks to seminar participants at the ANU, Universitat Autonoma de Barcelona, Universitat de Girona, University of Melbourne and University of New South Wales.

1 Introduction

The agency problem (Spence and Zeckhauser 1971; Ross 1973) has played a central role in the development of the economic analysis of problems involving asymmetric information. Analysis of this problem has produced important insights, but it has also been plagued by technical difficulties and the derivation of solutions that appeared unrealistic in economic terms.

Analysis was initially undertaken in the standard 'state space' framework in which random variables are represented by a mapping from a space of states of nature to a space of outcomes. However, Mirrlees (1974) discovered that, if uncertainty is represented by a stochastic production function not exhibiting weak disposability of output and the number of states of nature is finite, the principal can always achieve the first-best outcome. This is done simply by specifying an arbitrarily large penalty if output falls below the lowest level consistent with the first-best effort, and offering the fixed payment from the first-best contract otherwise.

Mirrlees (1974) proposed a reformulation of the problem, replacing the stochastic-production function with a parametrized distribution formulation, in which the scalar effort variable indexes a family of cumulative distribution functions over a discrete output space. This reformulation permitted the consideration of nontrivial problems, but gave rise to a range of technical difficulties. First, the parametrized distribution formulation is inconsistent with the standard microeconomic framework in which outputs are continuous quantities. Second, some aspects of the solution seem implausible in economic terms. In particular, in the absence of restrictive conditions, it is difficult to preclude the possibility that the optimal contracts will make payments a decreasing function of output over some range [1] (see Grossman and Hart (1983) for discussion on non-monotonicity). The simplest way to overcome these difficulties has been to assume that output takes only two possible values, or at most a finite number, but this assumption is obviously restrictive.

Quiggin and Chambers (1998) argued that the difficulties of the state-space approach were not inherent in the representation, but reflected the unrealistic properties of the technology implied by the use of a stochastic production function representation with a scalar input. With such a technology, observation of output in any state of nature allows a principal to infer the entire vector of state-contingent outputs and therefore to achieve the first-best. This analysis was extended further by Chambers and Quiggin (2004) who derived conditions under which the optimal contract would involve a fixed wage and a non-stochastic output.

Quiggin and Chambers examined a state-contingent representation of the agency problem, deriving a closed form solution for the case of two states of nature. They observed, but did not discuss in detail, the similarities between the state-contingent representation of the agency problem and the optimal tax problem considered by Mirrlees (1971), and subsequent writers, including Lollivier and Rochet (1983) and Weymark (1986ab) for the quasi-linear utility speci-

[1] One popular alternative has been the model of Holmstrom and Milgrom (1987), which gives rise to simple optimal payment rules at least as a limiting case.

fication.

The two-state representation analyzed by Quiggin and Chambers (2004) yields a number of useful insights, and is consistent with the standard microeconomic framework. However, like the two-outcome representation commonly used in applications of the parametrized distribution function approach, it is unduly restrictive. In particular it does not permit the modelling of many frequently observed wage contracts, such as that of a fixed basic wage with output-contingent bonuses.

In this paper, we extend the analysis of Quiggin and Chambers (1998) to the case where the state space is an atomless continuum. The approach is modelled on the treatment of optimal tax problems by Lollivier and Rochet (1983) and Boadway et al. (2000). The responses of agents to incentive structures may usefully be compared to the response of taxpayers to an income tax schedule. The states of nature in the representation of the principal-agent problem proposed by Quiggin and Chambers correspond naturally to the individual taxpayers in the optimal tax problem. However, whereas the social planner in the optimal tax problem deals with a set (possibly a continuum) of agents, each choosing the individually optimal level of effort and consumption, the problem considered here involves a single agent choosing an optimal level of effort and a state-contingent output plan. The state-contingent representation allows for a more flexible response, since effort may affect output differently in different states of nature. But the technological interpretation also raises some technical difficulties regarding sets of measure zero, which are addressed through the use of the Frechet derivative.

The central observation is that, under reasonable conditions, the optimal solution involves the specification of a fixed output level and fixed payment to the agent over an interval at the lower end of the state space. Thus the optimal contract may be regarded as a fixed wage with a bonus for above-normal performance. This is analogous to the phenomenon of 'bunching' commonly found in the optimal tax literature.

2 States, production and preferences

We consider a state-space Ω and an associated probability measure μ. For simplicity, we focus on the case of a single state-contingent output. In the discrete case, considered by Chambers and Quiggin (2000), Ω is a finite set $\{1...S\}$ with elements denoted s, and production involves the transformation of nonstochastic (*ex ante*) input vectors $\mathbf{x} \in \Re_+^N$ into state-contingent output vectors $\mathbf{z} \in Z \subseteq \Re_+^\Omega$. Thus, z_s is the realized output contingent on the occurrence of state s. General features of the production problem have been considered by Chambers and Quiggin (2000).

As noted by Chambers and Quiggin (2000) in the case of a discrete state space, the distinction between nonstochastic inputs $\mathbf{x}$ and stochastic outputs $\mathbf{z}$ is most appropriately interpreted in terms of an *ex ante* technology. Inputs are committed before uncertainty about the state of nature is resolved. Thus, the input vector is known with certainty, and is hence appropriately viewed as an element of the finite-dimensional space $\Re_+^{\mathbf{N}}$.

The agent maximizes an objective functional of the general form

$$V(\mathbf{y},\mathbf{x}) = \int u(y(\omega))d\mu(\omega) - g(\mathbf{x}), \tag{1}$$

increasing in the first argument and decreasing in the second, where $\mathbf{y}$ is an outcome space consisting of state-contingent income vectors, that is mappings $\mathbf{y} : \Omega \rightarrow \Re$, and $g : \Re_+^N \rightarrow \Re$ is an effort-cost function. Given the *ex ante* interpretation, an objective of this general form is implied by the assumption that preferences are additively separable over time. We assume that u is concave and twice continuously differentiable.

For the continuous case, we will assume, without significant loss of generality, that Ω is an interval. To simplify comparisons with the optimal tax literature we will make the specific assumption that Ω is an interval of the form $[\underline{\omega}, \overline{\omega}]$ where $\overline{\omega} > \underline{\omega} > 0$.[2] State-contingent outputs are given by measurable mappings $\mathbf{z} : \Omega \rightarrow \Re_+$, where output in state $\omega \in \Omega$ is denoted $z(\omega)$. We will denote the zero mapping by $\mathbf{0}$ and the constant mapping yielding output 1 for all ω by $\mathbf{1}$. Thus a uniform change in state contingent output may be represented by a shift from $\mathbf{z}$ to $\mathbf{z} + \boldsymbol{\delta}\mathbf{1}$ for $\delta \in \Re$.

We will denote the L_2 norm by $\|\mathbf{z}\|$ and will consider $\Re_+^{\Omega}$ as an inner-product space, with the topology of weak convergence. That is, letting $\left\{\mathbf{z}^k\right\}$ denote a sequence in $\Re_+^{\Omega}$, $\mathbf{z}^k \rightarrow \mathbf{z}$ if for all $\mathbf{z}'$, $\left\langle \mathbf{z}', \mathbf{z}^k \right\rangle \rightarrow \langle \mathbf{z}', \mathbf{z} \rangle$.[3] Note that this definition includes all the usual cases of convergence, including that of a continuous distribution collapsing to a point mass, and of pointwise convergence in the space of probability distributions over a (fixed) finite set of outcomes.

3 Input sets and the cost functional

For both discrete and continuous cases, the technology may be characterized by a technology set

$$T = \{(\mathbf{x},\mathbf{z}) \in \Re_+^{\mathbf{N}} \times Z : (\mathbf{x},\mathbf{z}) \text{ is feasible}\},$$

or, equivalently, by input sets

$$X(\mathbf{z}) = \{\mathbf{x} \in \Re_+^{\mathbf{N}} : (\mathbf{x},\mathbf{z}) \in T\}.$$

The elements of $X(\mathbf{z})$ are the input vectors consistent with the state-contingent output schedule $\mathbf{z} : \Omega \rightarrow \Re$. Most of the properties of $X(\mathbf{z})$ stated by Chambers and Quiggin (2000) for the discrete case depend only on ordering and convexity properties and therefore apply directly

[2] In the optimal tax literature, ω is a parameter corresponding to the wage. For the quasilinear utility specification (Lollivier and Rochet 1983; Weymark 1986ab; Boadway et al. 2000), ω is strictly positive. In the analysis of uncertainty, it is conventional to represent the state space by the unit interval $[0,1]$. Results presented here can be translated back to this conventional representation with an affine transform.

[3] This could prove particularly convenient in finance applications since $\Re_+^{\Omega}$ is self-dual and, the space of state-contingent price vectors is also $\Re_+^{\Omega}$.

for arbitrary state spaces Ω. The only technical difficulties arise with concepts of closure. In addition to the standard properties of convexity and weak disposability, we assume that X is a closed correspondence or, more precisely,

$$\left\|\mathbf{z}^k - \mathbf{z}\right\| \to 0,\ \mathbf{x}^k \in X\left(\mathbf{z}^k\right) \to \mathbf{x} \Rightarrow \mathbf{x} \in X\left(\mathbf{z}\right). \tag{2}$$

As noted above, closure is defined with respect to the topology of weak convergence.

In any analysis of mappings from an atomless measure space, it is necessary to consider the problems raised by sets of zero measure. In the continuous framework a particular state has zero measure. We adopt the standard notation $\mathbf{z} = \mathbf{z}'$ *ae* (almost everywhere) if

$$\mu\left\{\omega : z\left(\omega\right) \neq z'\left(\omega\right)\right\} = 0,$$

that is, if $\mathbf{z}$ and $\mathbf{z}'$ differ on a set of measure zero. Note that, given (2), we have:

Lemma 1 *If* $\mathbf{z} = \mathbf{z}'$ *ae*, $X(\mathbf{z}) = X(\mathbf{z}')$.

That is, if $\mathbf{z}$ and $\mathbf{z}'$ differ on a set of measure zero, the input set consistent with the production of $\mathbf{z}$ is equal to the input set consistent with the production of $\mathbf{z}'$.

3.1 The cost functional

Given preferences of the general separable form (1), technology may be summarized by a cost functional, $c : Z \to \Re_+$,

$$c(\mathbf{z}) = \inf\{g(\mathbf{x}) : (\mathbf{x}, \mathbf{z}) \in T\} = \inf\{g(\mathbf{x}) : \mathbf{x} \in X\left(\mathbf{z}\right)\}.$$

That is, for measurable $\mathbf{z} \in Z \subseteq \Re_+^\Omega$, $c(\mathbf{z})$ is the minimum cost such that $(\mathbf{x}, \mathbf{z})$ is feasible.

For the case where Z is a finite-dimensional vector space $\Re_+^S$, Chambers and Quiggin (2000) analyze the properties of $c(\mathbf{z})$ under a range of conditions. The central focus of this paper will be the derivation of tools for the analysis of the cost functional in the case where Z is of the form $\Re_+^\Omega$, for a general space Ω, with particular emphasis on the case $\Omega = [\underline{\omega}, \overline{\omega}]$. The derivation of the properties of convexity, monotonicity and continuity is analogous to that for the discrete case. Given a cost functional c, the objective function (1) may be restated as

$$V\left(\mathbf{y}, \mathbf{z}\right) = \int u(y\left(\omega\right))d\mu\left(\omega\right) - c(\mathbf{z}). \tag{3}$$

3.2 Marginal cost and the Frechet derivative

As will be shown below, the case of state-allocable linear costs yields a tractable closed form representation for the solution to the agency problem. More generally, if the cost function is differentiable, it may be approximated locally by a linear functional, with coefficients interpretable as marginal costs. In the discrete case, the marginal cost is represented by the derivatives of the cost functional c. For the general case $Z = \Re_{+}^{\Omega}$, the appropriate generalization of the derivative, $c_{\omega}(\mathbf{z})$, is given by consideration of the Frechet derivative, which encompasses general measure spaces Ω, including the unit interval with Lebesgue measure and other probability measures as well as discrete state spaces $\Omega = \{1, ...S\}$ with associated probability measures.

Suppose that $c : \mathbf{Z} \rightarrow \Re$ is Frechet differentiable, and denote its derivative by $c'(\bullet; \mathbf{z})$. c' is a linear mapping from $\mathbf{\Delta}$, the space of differences in Z, to $\Re$. Hence, by the Riesz representation theorem, there exists a measure $\mu_c(\bullet; \mathbf{z})$ (absolute continuous with respect to Lebesgue measure) such that, for any $\delta \in \Delta$,

$$c'(\delta; \mathbf{z}) = \int_{\Omega} \delta d\mu_c(\omega; \mathbf{z})$$

and

$$c(\mathbf{z} + \boldsymbol{\delta}) - c(\mathbf{z}) = c'(\delta; \mathbf{z}) + o\,\|\delta\|.$$

Now we denote the marginal cost of output in state ω by

$$c_{\omega}(\mathbf{z}) = \mu_c(\omega; \mathbf{z}).$$

Observe that, since, for any given ω, c_{ω} induces a mapping from Z to $\Re$, it is possible to define higher derivatives, cross derivatives and so on. In particular, following the approach of Ryder and Heal (1973),[4] examining preferences with respect to consumption streams over time, it is possible to use the derivative concept to determine whether outputs in different states of nature are substitutes or complements and, more generally, whether the cost functional is submodular or supermodular in $\mathbf{z}$.

The cost functional c may be approximated in a neighborhood of any $\mathbf{z}^0$ by

$$c(\mathbf{z}) = c\left(\mathbf{z}^0\right) + \int \mu_c(\omega; \mathbf{z}^o)\left(z(\omega) - z^0(\omega)\right) d\omega.$$

In the special case of linear costs, the approximation is exact and we have

$$c(\mathbf{z}) = \int \mu_c(\omega; \mathbf{0})\, z(\omega)\, d\omega = \int \mu_c(\omega)\, z(\omega)\, d\omega,$$

noting that μ_c is independent of $\mathbf{z}$. Hence, the state-specific cost functions are given by

$$c_{\omega}(z) = \mu_c(\omega)\, z.$$

More generally, and treating $c(\mathbf{z}^0) - \int \mu_c(\omega; \mathbf{z}^o)\, z^0(\omega)$ as a fixed cost c^0, we may regard the affine cost functional[5]

$$\hat{c}(\mathbf{z}) = c^0 + \int \mu_c(\omega; \mathbf{z}^o)\, z(\omega)\, d\omega$$

as a local linear approximation to c in a neighborhood of $\mathbf{z}^0$.

[4] Ryder and Heal (1973) do not consider the Frechet derivative explicitly. Rather they define a closely related derivative measure which they refer to as the Volterra derivative, following Volterra (1930).

[5] Note that the 'local cost function' does not satisfy the property $c(\mathbf{0}) = 0$ unless $c^0 = 0$. This does not raise any serious difficulties.

4 The principal-agent problem

We now consider a principal, contracting with the agent to produce output z in return for payment y. We assume that the principal cannot observe ω, the state of the world, directly. Hence the principal must offer a payment schedule of the form $y(z)$, $y : \Re_+ \rightarrow \Re$, excluding dependence of the payment y on the state of the world ω. We will assume that the payment schedule must be piecewise continuous. We do not require that the payment $y(z)$ be non-negative, thereby allowing for the possibility of sanctions such as dismissal with loss of accrued benefits.

Given such a payment schedule, the agent chooses an output vector

$$\mathbf{z} \in \arg\max \left\{ \int u(y(z(\omega)))d\mu(\omega) - c(\mathbf{z}) \right\}. \tag{4}$$

The principal's profit is given by

$$\pi(\omega) = z(\omega) - y(z(\omega)).$$

We assume that competition among risk-neutral principals drives expected profits to zero. We therefore consider the problem

$$\max_{\mathbf{y},\mathbf{z}} \int \{u(y(z(\omega))) - c(\mathbf{z})\} \, d\mu(\omega) \tag{5}$$

subject to (4) and the zero-expected-profit constraint:

$$\int \pi(\omega) \, d\mu(\omega) \geq 0. \tag{6}$$

This formulation is dual to that of Quiggin and Chambers (1998), where the objective is to maximize profit subject to a participation constraint. Following the arguments of Grossman and Hart (1983), Quiggin and Chambers (1998) show that, in a constrained-optimal solution, the participation constraint must bind. Given this result, it is straightforward to show that, in the dual problem considered here, the zero-profit constraint must bind. We record this observation as a Lemma.

Lemma 2 *In any solution to (5), subject to (4) and (6), the zero-profit constraint (6) is binding.*

4.1 The agent's optimization problem

In order to characterize the agent's optimal behavior, we use some simplifying assumptions that do not involve any significant loss of generality.

We will assume that the states have been ordered from worst to best, in the sense that, for any given distribution of output, the least cost $\mathbf{z}$ yielding that distribution is monotonic, that is, $\omega \leq \omega' \Rightarrow z(\omega) \leq z(\omega')$. More precisely, let

$$F(z;\mathbf{z}) = \mu\{\omega : z(\omega) \leq z\}.$$

Then for any monotonic $\mathbf{z}$ and any $\mathbf{z}'$ such that

$$F(z;\mathbf{z}) = F(z;\mathbf{z}')\,\forall z,$$

we have $c(\mathbf{z}) \leq c(\mathbf{z}')$.

Under this assumption, the agent will never benefit from the choice of a non-monotonic $\mathbf{z}$ in the problems under consideration. Similarly, without loss of generality, we can assume that $\mathbf{z}$, considered as a function of ω, is lower semi-continuous. Now if $\mathbf{z}$ and $\mathbf{z}'$ are monotonic and lower semi-continuous, and $\mathbf{z}' \neq \mathbf{z}$, then it cannot be true that $\mathbf{z}' = \mathbf{z}$ *a.e.* By confining attention to the set $Z^* \subseteq Z$ of monotonic, lower semi-continuous output vectors $\mathbf{z}$, we therefore avoid complications associated with sets of zero measure without any associated loss of generality.

We can also characterize the nature of the agent's optimal output:

Lemma 3 *The state ordering assumption implies, for any pricing scheme of the form $y = y(z)$, that the optimal response $z(\omega)$ is nondecreasing in ω.*

In view of Lemma 3, we will assume without loss of generality that $\mathbf{z}$ is upper semicontinuous in ω.

Lemma 3 also holds in the Grossman–Hart moral hazard model, where output is a stochastic function of scalar input. We now consider a crucial result which holds generally in the state-contingent model with strictly monotone costs, but only under special and restrictive conditions in the Grossman–Hart model.

Lemma 4 *Let $\mathbf{z}$, satisfy (4) for the piecewise continuous payment schedule y. Then for any interval $[\omega^0, \omega^1]$ for which z is a continuous function of ω, y is a monotone increasing function on $[z(\omega^0), z(\omega^1)]$.*

Proof: Suppose not. Then there exists some interval $[z^2, z^3]$ with $z(\omega^0) \leq z^2 < z^3 \leq z(\omega^1)$ on which y is nonincreasing. Consider the output vector z^3 such that

$$\widetilde{\mathbf{z}}(\omega) = \begin{cases} z(\omega) & z(\omega) \leq z^2 \\ z^2 & z^2 \leq z(\omega) \leq z^3 \\ z(\omega) & z(\omega) > z^3. \end{cases}$$

By strict monotonicity of c, $c(\widetilde{\mathbf{z}}) < c(\mathbf{z})$, and, since y is nonincreasing on $[z^0, z^1]$,

$$\int u(y(z(\omega)))d\mu(\omega) \leq \int u(y(\tilde{z}(\omega)))d\mu(\omega),$$

so that the choice of $\mathbf{z}$ violates (4).

From Lemmas 3 and 4, we can assume, without loss of generality, that y is a monotone increasing function of z. This in turn implies that it is never rational for the agent to choose an output $\mathbf{z}$ if there exists $\widetilde{\mathbf{z}} \geq \mathbf{z}$ with $c(\mathbf{z}) = c(\widetilde{\mathbf{z}})$. Thus we can confine attention to the subset of Z^* where $c(.)$ is strictly monotone increasing.

We next consider conditions under which z will be continuously differentiable. Since any monotone increasing y can be approximated arbitrarily closely by a differentiable function, there is no loss of generality in assuming that y is a differentiable function of z. Now all that is required is:

Assumption 1: For all z continuous in ω, $\mu_c(\omega; \mathbf{z})$ is continuously differentiable in ω.

Lemma 5 *Under Assumption 1, and with y a differentiable function of z, any optimal z is differentiable as a function of ω.*

We can now show that for any desired output $\mathbf{z}$, the incentive compatibility requirement precisely specifies the feasible incentive structure $y(z)$ on any interval $[\omega_0, \omega_1]$ for which z is continuously differentiable. We begin by observing that the truth-telling condition (4) requires that, for all $\widetilde{\mathbf{z}}$ in a neighborhood of the agent's optimal choice of $\mathbf{z}$,

$$\int u(y(z(\omega))) - \mu_c(\omega; \mathbf{z}) z(\omega) \, d\mu(\omega) \geq \int u(y(\widetilde{z}(\omega))) - \mu_c(\omega; \mathbf{z}) z(\omega) \, d\mu(\omega), \tag{7}$$

which holds if and only if, for all ω and all $\widetilde{\mathbf{z}}$ in a neighborhood of the agent's optimal choice of $\mathbf{z}$,

$$u(y(z(\omega))) - \mu_c(\omega; \mathbf{z}) z(\omega) \geq u(y(\tilde{z}(\omega))) - \mu_c(\omega; \mathbf{z}) \tilde{z}(\omega). \tag{8}$$

Letting $z(\omega)$ be the state-contingent output plan desired by the principal, assume that $y(z)$ is arbitrarily negative for z outside the image set $\{z(\omega) : \omega \in \Omega\}$ so that any candidate $\widetilde{\mathbf{z}}$ must also lie in this range. That is, for any $\widetilde{\mathbf{z}}$ and ω there exists $\tilde{\omega} \in \Omega$ such that $\tilde{z}(\omega) = z(\tilde{\omega})$.

Now we can immediately derive:

Lemma 6 *Let z be continuous. Then (7) holds if and only if, for all ω, $\tilde{\omega}$,*

$$u(y(z(\omega))) - \mu_c(\omega; \mathbf{z}) z(\omega) \geq u(y(z(\tilde{\omega}))) - \mu_c(\omega; \mathbf{z}) z(\tilde{\omega}). \tag{9}$$

Taking limits as $\tilde{\omega}$ approaches ω from above and below, we obtain our main result:

Theorem 1 *If Assumption 1 is satisfied, the incentive structure y must satisfy*

$$\left(u'(y(z(\omega)))y'(z(\omega)) - \mu_c(\omega; \mathbf{z})\right) z'(\omega) = 0 \tag{10}$$

with complementary slackness.

We will refer to condition (10) as the first-order incentive compatibility condition (FOIC), and to the requirement that z be a monotone increasing function of ω, derived in Lemma 3, as the second-order incentive compatibility condition (SOIC).

This condition requires either $z'(\omega) = 0$ or

$$u'(y(z(\omega)))y'(z(\omega)) - \mu_c(\omega; \mathbf{z}),$$

which may be rewritten as

$$u'(y(z(\omega)))y'(z(\omega)) = \mu_c(\omega; \mathbf{z}). \tag{11}$$

That is, the marginal cost of increasing output in a neighborhood of ω, given by $\mu_c(\omega; \mathbf{z})$, must be equal to the marginal benefit, given by $u'(y(z(\omega)))y'(z(\omega))$.

This condition may usefully be compared to that derived by Quiggin and Chambers (1998) for the case of two states of nature, where $z_1 > z_2$. Quiggin and Chambers show that, under plausible conditions, truthtelling will imply the binding condition

$$\pi_2(u(y_2) - u(y_1)) = c(z_1, z_2) - c(z_1, z_1). \tag{12}$$

For linear costs, that is, $c(\mathbf{z}) = \pi_1 c_1 z_1 + \pi_2 c_2 z_2$, equation (12) becomes

$$u(y_2) - u(y_1) = c_2(z_2 - z_1),$$

or, for some $z^* \in [z_1, z_2]$,

$$u'(y(z^*))y'(z^*) = c_2,$$

which is identical to (11).

In the optimal tax literature, situations where $y'(\omega) = z'(\omega) = 0$ are referred to as 'bunching'. In the context of the principal-agent problem, the natural interpretation is that of a fixed wage contract, requiring the agent to produce a given output z for all ω in some interval, and not providing sufficient incentive to induce higher outputs for more favorable states within that interval. In general, bunching can occur over any set of cost levels. However, following the arguments of Boadway et al. (2000) and Guesnerie and Laffont (1984), it can be shown that, under plausible conditions, violations will be confined to the bottom end of the distribution, that is, to high cost levels.

There are two possible interpretations for a solution of this kind. One is that of a fixed wage, with bonuses paid for higher outputs in some subset of favorable states. The second arises when, for low values of ω, $z(\omega) = 0$. That is, in unfavorable states, no production takes place. The solution may then involve a discontinuous jump to a strictly positive level of output z_{min}. In effect, under conditions where the incentive structure would lead the worker to produce less than z_{min}, a separation (fire or quit) takes place. This may involve no payment, a positive separation payment or a negative payment, representing dismissal with loss of accrued benefits.

4.2 An explicit solution for the case of linear costs

A useful special case is that of linear state-allocable costs. That is, for each $\omega \in [\underline{\omega}, \overline{\omega}]$, there exists $c(\omega)$ such that, for all $\mathbf{z}$,

$$c(\mathbf{z}) = \int_{\underline{\omega}}^{\overline{\omega}} c(\omega) z(\omega) d\mu(\omega).$$

We will assume, without loss of generality, that $c(\omega)$ is a non-increasing function of ω, so that the states are ordered from worst to best. In particular, for the derivation of a closed form solution, allowing comparative statics, it is useful to adopt the normalization

$$c(\omega) = \frac{c}{\omega}, \tag{13}$$

where c is a cost parameter. This normalization can be used to represent an arbitrary $c(\omega)$ on Ω by an appropriate choice of the density function $f(\omega)$. Following Boadway et al. (2000), we may now derive the explicit solution for the case of linear state-allocable costs. The main interest is in the reinterpretation of terms from the optimal tax setting to that of the principal–agent problem.

Note that the incentive schedule $y(z)$ can be represented by a set of payment–output bundles $(y(\omega), z(\omega))$, one for each $\omega \in \Omega$, designed to induce the agent to produce the appropriate output at each state. With

$$U(\omega) = \omega u(y(\omega)) - cz(\omega),$$

the maximization problem is:

$$\begin{array}{ll} \max\limits_{y(\omega),z(\omega)} & \int_{\underline{\omega}}^{\overline{\omega}} \frac{U(\omega)}{\omega} f(\omega)\, d\omega \\ \text{s.t.} & \int_{\underline{\omega}}^{\overline{\omega}} [z(\omega) - y(\omega)] f(\omega)\, d\omega \geq 0 \\ & \dot{U}(\omega) = u(y(\omega)) \\ & \dot{y}(\omega) \geq 0. \end{array}$$

Following Boadway et al. (2000), $x(\omega) \equiv \dot{y}(\omega) = y'(z(\omega))\, \dot{z}(\omega)$ is a natural choice for the control variable in this dynamic optimization problem. The Hamiltonian function is:

$$\begin{aligned} H(\omega) &= \frac{U(\omega)}{\omega} f(\omega) + \lambda [\omega u(y(\omega)) - U(\omega) - cy(\omega)] f(\omega) + \\ &\quad + v(\omega) u(y(\omega)) + \mu(\omega) x(\omega) + \kappa(\omega) x(\omega), \end{aligned}$$

where λ is the shadow price of the zero expected profits constraint, $v(\omega)$ is the co-state variable associated with $\dot{U}(\omega) = u(y(\omega))$, $\mu(\omega)$ is the co-state variable associated with $x(\omega) \equiv \dot{y}(\omega)$ and $\kappa(\omega)$ is the shadow-price of the non-negativity constraint on $x(\omega)$. We can now derive the solution

$$v(\omega) = \int_{\underline{\omega}}^{\omega} \left(\lambda - \frac{c}{\omega}\right) f(m)\, dm = \lambda F(\omega) - G(\omega),$$

where

$$G(\omega) = \int_{\underline{\omega}}^{\omega} \frac{c}{m} f(m)\, dm \tag{14}$$

and

$$\lambda = G(\overline{\omega}). \tag{15}$$

The cumulative values $F(\omega)$ and $G(\omega)$ characterize the production technology. $F(\omega)$ is simply the probability that the unit cost level is greater than c/ω and $f(\omega)$ is the associated density. $G(\omega)$ is the expected value of c/m over the interval $\underline{\omega} \leq m \leq \omega$, multiplied by the probability of states in that interval: $E[c/m|\, m \leq \omega]\, F(\omega)$. The shadow price of the expected profits condition, denoted by λ, is the expected value of c/m over the entire distribution, $E[c/m]$ and, accordingly, depends only upon the technology.

Let $\pi(z) = z - y(z)$ denote the principal's unit profit for output level z. For the case when z is strictly increasing in ω (that is, the SOIC condition is satisfied and the first order approach is valid) we obtain, following Boadway et al. (2000),

$$\pi'(z(\omega)) = 1 - y'(z) = 1 - \frac{c}{\omega u'(y(\omega))} = \frac{\left[\dfrac{G(\omega)}{G(\overline{\omega})} - F(\omega)\right]}{\omega f(\omega)}.$$

It is worth noting that $\pi'(z(\omega)) > 0$, except at the end points where $\pi'(z(\omega)) = 0$. The optimal contract effectively makes the agent the residual claimant in those extreme states. This is consistent with the result that Quiggin and Chambers (1998) obtain for the best state.

However, as noted above, there is no guarantee that the SOIC condition will always be satisfied. When it is violated, the optimal contract involves a fixed wage requiring the agent to produce a given output z for all ω in some interval. Violations of the SOIC can occur at any level of cost and depend only on the technology. However, under plausible conditions, these can be confined to high cost level states. There will be violations at the bottom if $\underline{\omega}$ is small enough (in particular, $\underline{\omega} \leq 1/2G(\overline{\omega})$ for any density $f(.)$). If there is bunching at the bottom, $\pi'(z(\omega))$ is positive at the end of the bunching interval.

5 Concluding comments

In this paper, we have extended the analysis of Quiggin and Chambers (1998) to the case where the state space is an atomless continuum. The crucial analytical innovation is the definition of a Frechet differentiable generalization of the state-contingent cost function for a discrete case space. In particular, the Frechet derivative may be interpreted as specifying state-contingent marginal costs, even in the case of an atomless state space. This innovation allows the use of tools developed for the analysis of optimal tax problems with a continuum of abilities and quasilinear preferences (Lollivier and Rochet 1983; Boadway et al. 2000), with the set of individual ability levels being reinterpreted as a set of states of nature with different cost levels.

The central result is that, under reasonable conditions, the optimal solution involves the specification of a fixed output level and fixed payment to the agent over an interval at the lower end of the state space. Thus, the optimal contract may be regarded as a fixed wage with a bonus for above-normal performance. This is analogous to the phenomenon of 'bunching' at the bottom commonly found in the optimal tax literature.

It is natural to consider whether it is possible to translate results in the opposite direction. In particular, Holmstrom and Milgrom (1987) have shown that, when the state space is sufficiently complex (in their model, it arises from a Brownian motion), the optimal solution for the principal involves the setting of an affine payment schedule. In view of the great interest that has been shown in the derivation of conditions under which an optimal tax scale will be linear or affine, it is of interest to consider whether there are representations of the optimal income tax problem analogous to the principal-agent problems considered by Holmstrom and Milgrom. This issue will be addressed in future research.

References

[1] Boadway, R., Cuff, K. and Marchand, M. (2000). Optimal income taxation with quasi-linear preferences revisited, *Journal of Public Economic Theory* 2(4), 435–460.

[2] Chambers, R. G. and Quiggin, J. (2000). *Uncertainty, Production, Choice, and Agency : The State-Contingent Approach,* Cambridge University Press, New York.

[3] Chambers, R. G. and Quiggin, J. (2004), 'Incentives and standards in agency contracts", *Journal of Public Economic Theory* ,forthcoming .

[4] Grossman, S. and Hart, O. (1983). An analysis of the principal-agent problem, *Econometrica* 51, 7–45.

[5] Guesnerie, R. and Laffont, J-J. (1984). A complete solution to a class of principal–agent problems with an application to the control of a self-managed firm, *Journal of Public Economics* 25, 329–369.

[6] Holmstrom, B. and Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives, *Econometrica* 55, 303–328.

[7] Lollivier, S. and Rochet, J-C. (1983). Bunching and second-order conditions: a note on optimal tax theory, *Journal of Economic Theory* 31, 392–400.

[8] Mirrlees, J.A. (1971). An exploration in the theory of optimum income taxation, *Review of Economic Studie*s 38, 175–208.

[9] Mirrlees, J. (1974). Notes on welfare economics, information, and uncertainty, in *Essays on Economic Behaviour Under Uncertainty* M. Balch, D. McFadden and Shih-Yen Wu (eds) North-Holland, Amsterdam.

[10] Quiggin, J. and Chambers, R. (1998). A state-contingent production approach to principal-agent problems with an application to point-source pollution control, *Journal of Public Economics* 70, 441–472.

[11] Ross, S.A. (1973). The economic theory of the agency: the principal's problem, *American Economic Review* 63(2), 134–139.

[12] Ryder, H. and Heal, G. (1973). Optimal growth with intertemporally dependent preferences, *Review of Economic Studies* 40(1), 1–31.

[13] Spence, M. and Zeckhauser, R. (1971). Insurance, information, and individual action, *American Economic Review* 61(2), 380–387.

[14] Volterra, V. (1930), *Theory of Functionals and of Integral and Integro-Differential Equations*, republished by Dover, New York, 1959.

[15] Weymark, J. (1986a). A reduced-form optimal nonlinear income tax problem, *Journal of Public Economics* 30, 199–217.

[16] Weymark, J. (1986b). Bunching properties of optimal nonlinear income taxes, *Social Choice and Welfare* 3, 213–232.