Contract Design for Biodiversity Procurement

Peter Bardsley

And

Ingrid Burfurd

A Contributed Paper to the Australian Agricultural & Resource Economics Society's Annual Conference,
Cairns, February 11-13, 2009.

Contract Design for Biodiversity Procurement

Peter Bardsley[1] | Ingrid Burfurd

University of Melbourne | University of Melbourne

Revision 1, August 2008

[1]Corresponding Author: Peter Bardsley, Department of Economics, The University of Melbourne, Victoria 3010 Australia. Email p.bardsley@unimelb.edu.au Fax +61 3 8344 6899 Phone +61 3 8344 6931.

Abstract

Market based instruments are proving increasingly effective in biodiversity procurement and in regulatory schemes to preserve biodiversity. The design of these policy instruments brings together issues in auction design, contract theory, biology, and monitoring technology. Using a mixed adverse selection, moral hazard model, we show that optimal contract design may differ significantly between procurement and regulatory policy environments.

Keywords biodiversity, procurement, adverse selection, moral hazard, contract theory

JEL Codes D82, D86, Q57

The preservation of biodiversity is an objective shared by governments around the world. Biodiversity assets are often located on private land, but provide environmental services with a strong public good element. Private landholders may not have the incentive to manage biodiversity in a socially optimal manner, and so governments use a range of policies to influence landholders' behaviour. These policies include regulation to prohibit actions that endanger or deplete environmental assets, and incentive programs to increase investment in nature conservation.

Regulation, Pigouvian taxes and simple incentive programs – such as untargeted subsidies – have in many cases proven to be relatively inefficient and ineffective, especially for non-point-source problems, and a range of market based schemes have been introduced with the aim of improving economic efficiency. Examples include tradeable emission permit schemes for Sulphur, Nitrogen and Carbon, biodiversity procurement auctions, and environmental offset markets such as the US EPA's wetlands offset program or the Victorian Bush Broker Scheme [1,5,33,35,37]. For a recent review of procurement in the context of biodiversity economics see [15].

There are three broad classifications of procurement markets for environmental public goods: public procurement of environmental services (where the government is the purchaser), private procurement of environmental services (where a private developer is purchasing an offset to satisfy regulatory requirements), and markets designed to match private buyers and sellers of offsets.

One of the salient features of both private and public procurement programs for non-point-source outcomes is the delegation of conservation activities to private landholders. In these policy institutions the principal (whether the government or a developer) faces asymmetric information problems. Adverse selection is a problem because the principal cannot gauge the key characteristics of suppliers[1] – for example, the principal cannot gather detailed information on suppliers' environmental production functions, or a supplier's opportunity costs. Moral hazard is a problem because procurement of environmental services requires the supplier to undertake actions and to protect a site through time, and it is difficult for the principal to observe and monitor compliance with a contract. Because environmental procurement is subject to exogenous risks (for example, drought, disease and predators), it is critical that contract design accounts for risk as well as for these asymmetrical information problems. Contract design (and therefore the incentive to comply with a contract) is fundamental to the environmental integrity of all these policy institutions.

Contract design for environmental services has been considered by a number of authors. Antle et al [2] discuss contracts to procure carbon sequestration services,

[1] The principal also has private information, for example scientific information and knowledge of environmental stocks so we are, at least formally, dealing with an informed principal problem [20]. However Maskin and Tirole show that in a quasilinear private values environment, which we have here, this reduces to the standard principal agent problem. The principal's information affects her valuation of the purchase, but creates no additional strategic effects to be addressed through contract design.

and Feng et al [8] consider the design of conservation incentive schemes, but neither of these consider issues of asymmetric information or incentive compatibility. Following the foundational contribution of Baron and Myerson [4] on regulation under asymmetric information, there has been a steadily growing literature on optimal environmental policy under asymmetric information. This literature reflects the importance of heterogeneity and highly disaggregated private information in the environmental sector, and the fundamental constraints that these impose on policy. Spulber [32] considers optimal quotas and taxes for a polluting industry under asymmetric information. Smith [31] applies a model of this type to design a least cost mechanism to retire agricultural land for environmental purposes. Wu et al [40] study the optimal design of an input tax or subsidy in a model of joint production of agricultural production and a pollutant. Moxey, White, Ozanne et al [22, 25, 26, 39] study optimal regulation of the use of an environmentally harmful input, allowing for both adverse selection and moral hazard. Crepin [6] models wetland procurement, and Sheriff [30] studies optimal regulation in the presence of political constraints.

There are common features in most of these models that make them inappropriate to model biodiversity procurement. While adverse selection is taken into account, these are all essentially single agent models that assume universal participation. When participation is voluntary, as is the case for environmental markets, the principal can maximise efficiency by contracting with low-cost agents. An appropriate mechanism must therefore inducing truth-telling and allow selective contracting with agents. This leads naturally to models that incorporate an optimal auction framework [12]. More significant is the fact that these models do not fully address moral hazard. The actions that are procured (for example, not collecting fallen firewood from a forest) are intrinsically difficult to observe, and are related by complicated and indirect means to desired outcomes (for example the presence and prevalence of certain species). Risk and uncertainty, both in the environmental production function and in the monitoring and measurement technology, are also important. Ozanne and White [26] put forward the model that is of most relevance here. They allow for hidden action by the agent, which is addressed by a random monitoring mechanism. While this is reasonable for actions that are readily verifiable (for example acreage set aside), it appears less relevant to biodiversity procurement where production is stochastic and actions are intrinsically unverifiable, where a more conventional principal-agent approach, linking rewards to outcomes rather than verifiable inputs, would seem to be more reasonable. In the Ozanne and White model all agents face a linear incentive contract with the same slope; in the more general approach as considered here they would face a non-linear contract (which may be implemented as a type-revealing menu of linear incentive contracts). We will be particularly concerned with how such a menu of contracts might vary in different policy environments. Ferarro et al [9, 10] also emphasise the need to deal with both adverse selection and moral hazard and the importance of risk. They do so however in only general terms, without getting into the specifics of a model that can give practical guidance in program design.

We thus need a mixed model, allowing for voluntary participation, multiple agents, risky production, adverse selection and moral hazard. In this paper we apply a version of the Laffont and Tirole [16] procurement model, modified to allow for risk aversion. This model is very tractable, and sufficiently flexible to address the issues that we are interested in. We examine issues that arise in the design of contracts under different assumptions about the policy institution. The main question that we ask is whether one can design a common type of contract that can be used across a range of institutions (public procurement, private procurement, procurement by a market maker in an offset market), or whether contract design must be sensitive to the institutional framework.

In order to clarify the nature of the economic problems we first summarise what a typical biodiversity procurement and preservation scheme looks like in practice. We then formally address the contract design issue, and finally we discuss some practical implications of the model in a real world design context.

1 Biodiversity procurement

In order to establish a concrete context we will outline a typical biodiversity procurement program in Victoria, Australia. We will identitify informational constraints and will clarify which constraints must be resolved through contract design, and which can be addressed elsewhere in the institutional design problem.

Since European settlement in the late eighteenth century there has been substantial biodiversity loss across the Australian continent. Some of Australia's rarest and most important temperate ecosystems – for example, extensive and diverse grasslands – have been reduced to remnant habitat fragments on private land [2]. Government has responded by taking increased responsibility for biodiversity management. Over the past decade Australian State and Federal Governments have allocated well over $300 million to biodiversity and habitat-related initiatives, mostly relating to the management of assets on private land.

Traditionally government has sought to preserve and rehabilitate biodiversity through untargeted conservation subsidies for volunteer programs and through regulatory prohibitions on any further damage to biodiversity assets. Both of these policies face incentive problems: regulation creates an incentive for landholders to conceal environmentally valuable biodiversity located on their land [14, 21], and although voluntary conservation is extremely valuable, incentives for participants

[2] Since European settlement, 90 per cent of the native vegetation in the eastern temperate zone has been removed [3].

to create value over the longer term are weak or absent. In addition, while voluntary measures are well meaning, these interventions may not always be supported by good science or directed to the highest value activities [29]. In the last decade government has moved toward a market based approach to biodiversity rehabilitation that is based on scientific modelling and a formal application of economic design principles.

Biodiversity management presents difficult public policy problems; asymmetric information is the source of many of these difficulties. Information about the location and status of biodiversity assets, landholders' ability to produce conservation gains, and landholders' opportunity cost of actions, is highly disaggregated and subject to adverse selection. The adverse selection problem is compounded by the complexity and heterogeneity of biodiversity resources. Moral hazard is also a problem because it is difficult and costly to observe and monitor private landholders' compliance with procurement contracts. The moral hazard problem is exacerbated by the highly stochastic nature of environmental production.

Three strategies have proven useful in addressing informational problems. First, a direct attempt has been made to reduce the informational disadvantage as much as possible through the systematic use of technology and science. Field ecologists have developed efficient environmental classification and scoring methods [28], which allow the classification and valuation of environmental assets with relatively high accuracy and low transaction costs. This science also underpins advice to landholders regarding the potentially valuable actions that the government may be willing to purchase. However, while improved science allows the government to identify and value the characteristics of remnant vegetation patches, conservation production technology and inputs are the private information of agents. For example, the conservation potential of a vegetation patch is influenced by the unique biophysical characteristics of the site and historical land management practices. Landholders have also observed their site over many seasons, and therefore have information on patterns of species presence and the factors that influence this [13]. This information is private, as are opportunity costs, giving rise to an adverse selection problem. Environmental service contracts are therefore procured through auctions. Landholders nominate the conservation actions that they will undertake, with bids ranked according to the estimated environmental benefits that they generate and the price of the contract. Bids are then accepted or rejected competitively. Finally, moral hazard is addressed by writing incentive contracts contingent on observable actions and outcomes.

This approach has been found to improve the efficiency of government spending [29, 34, 36, 38]. Because of the emphasis on incentive-compatibility, the interests of landholders and government are brought much more closely into alignment, and there has been wide-spread acceptance of this approach across a range of stakeholder groups [7]. Importantly, landholders have an incentive to reveal private information, and a number of previously unknown valuable biodiversity assets have been identified and rehabilitated.

It is natural to extend a competitive approach to the procurement of environmental offsets. Under a typical "no net loss" regulatory framework, developers are able to destroy existing habitat assets provided that the loss of habitat is offset by the rehabilitation and permanent protection of a commensurate asset. Offsets are underpinned by a stringent regulatory framework, and procurement incurs high transaction costs. Developers must identify an environmentally acceptable offset site, negotiate with potential suppliers on a bilateral basis, design and enter into a contract with suppliers and meet the Government's compliance requirements. Given the Government's investment in the infrastructure to support efficient environmental procurement, there is an obvious role for government as a facilitator of the market for third-party offsets. A government-facilitated matching market not only reduces transaction costs but improves the environmental integrity of the offset policy. Such a market has recently been announced in Victoria [23].

There are thus three distinct frameworks for the purchase of environmental assets. The first is the government's procurement by auction of biodiversity as a public good. The second is the private purchase of environmental off-sets by developers. The third is the facilitation of trade in the off-set market through government[3] intermediation in the procurement of environmental assets on behalf of developers. In each case, the question arises of how to write an incentive compatible and efficient contract with the seller of the environmental service.

2 The model

We will use a version[4] of Laffont and Tirole's basic procurement model [16], which we modify to allow for risk aversion[5]. This model allows for voluntary participation, adverse selection and moral hazard, with the advantage that the interaction between contract design and auction design is particularly transparent [17]. We will follow the notation of their monograph [18], which differs in some respects from the original paper.

By varying the principal's objective function, we can examine, in a unified manner, the three different policy environments mentioned above. The key parameters will be θ, the weight placed by the principal on the welfare of the agent, and λ, the deadweight cost of raising taxation revenue. For public procurement $\theta = 1$ (we assume that the government is concerned with maximising aggregate welfare and is unconcerned about the distribution of information rents) and $\lambda > 0$; for

[3] In principle one could imagine a private intermediary playing this role; this seems unlikely due to economies of scale and government's existing role in guaranteeing the integrity of traded off-sets.

[4] We reinterpret effort as output enhancing rather than cost reducing.

[5] See [16, p. 663] for some general remarks about risk aversion in their model.

private procurement $\theta = \lambda = 0$, since the principal maximises his own profit; in a government intermediated off-set market $\theta = 1$ but $\lambda = 0$ since the government is concerned with efficiency, and revenue is not raised through taxation. Since we are concerned with optimal contract design we will first consider an environment where a single principal contracts with a single agent. We will discuss below how this case can be modified if there are many agents, and the right to participate is auctioned.

A participating agent can produce environmental output $x + \varepsilon = \beta + e + \varepsilon$, where β is the agent's type, $e \geq 0$ is effort exerted by the agent, and ε is a random shock with mean zero and variance σ^2 reflecting the fact that production is risky (it may also incorporate risk introduced by errors in measuring output). The type β is a parameter in the agent's conservation production function. Although the government can estimate and value the quantity of remnant habitat (measured in quality adjusted units), agents hold private information on the unique mix of technology and inputs that will influence the rehabilitation and protection of this vegetation (such as a site's biophysical characteristics and resource management history). Through the exertion of effort e the agent can improve the quality of remnant vegetation. Both e and β are private information, known only to the agent, but total output x can be measured (possibly with some error incorporated in ε) by the principal. We assume that the type β is distributed, independently and identically, in the interval $\left[\underline{\beta}, \overline{\beta}\right]$ with distribution function $F(\beta)$ and density function $f(\beta)$ and that the inverse hazard rate $h(\beta) = \frac{1-F(\beta)}{f(\beta)}$ is non-increasing. We assume that the principal is risk neutral.

The agent's utility $v(t,e) = g(t) - \psi(e)$ is separable in the income transfer t that will be made by the principal, and in effort e, and is strictly concave (risk averse) in t. For any random transfer z, let $c(z)$ be the cost of risk bearing (the insurance premium that would make the agent indifferent between z and its expected value $\bar{z}$, net of the insurance premium) defined by $Eg(z) = g(\bar{z} - c(z))$. By the separability assumption, we can re-label indifference curves to linearise $g(z)$ (see, for example, [19, p 45]). We can thus assume that the agent's utility[6], given a random transfer z and an effort level e, is of the form

$$\tilde{v}(z,e) = \bar{z} - c(z) - \psi(e).$$

Thus, given this normalisation, the agent's utility depends on the expected transfer, subject to adjustments for a risk premium and for the disutility of effort. Following Laffont and Tirole, we will assume[7] that disutility of effort is increasing and convex in effort: $\psi(0) = 0$, $\psi(e) \geq 0$, $\psi'(e) \geq 0$, and $\psi''(e) > 0$. We will also

[6] This may no longer be an expected utility functional of z. We note that this type of normalisation is conventional in finance, where a mean-variance specification of preferences is standard. Such preferences are of course not of expected utility form, except in exceptional cases.

[7] The plausibility of these assumptions, which of course depend upon our normalisation linearising $f(t)$, is ultimately an empirical matter.

assume, as do they, that $\psi'''(e) \geq 0$.

If the variance is small then the risk premium takes the form

$$c(z) = \frac{1}{2}\eta \operatorname{Var}(z)$$

where η is an appropriate coefficient of risk aversion [24, pp. 69-80]. This relationship holds exactly if z is normally distributed and the underlying utility function $v(t, e)$ displays constant absolute risk aversion with coefficient of risk aversion η; the reader may wish to keep this case in mind[8]. In general, it will hold approximately if the variance is small; the appropriate coefficient of risk aversion will then depend upon wealth and background risk (for a recent discussion in the context of a mean variance specification see [11]). In the interest of tractability we will assume that risk preferences are of this form, and we will ignore wealth effects, due to the contract, on risk preferences. That is, we assume that the impact of the contract on the agent's risk and wealth position is not sufficient to significantly affect their risk preferences, and we will treat η as a constant. This seems reasonable provided that the agent's risk and return exposure through the contract is not dominant in the agent's portfolio of investments and activities. This seems a reasonable assumption in the policy context with which we are concerned.

[8]It also holds exactly if the underlying utility function is quadratic.

The principal offers a menu of linear contracts[9]

$$T\left(x+\varepsilon,\hat{\beta}\right)=a\left(\hat{\beta}\right)+b\left(\hat{\beta}\right)(x+\varepsilon)$$

contingent on announced type $\hat{\beta}$ and linear in observed output $x+\varepsilon$. If the agent announces type $\hat{\beta}$ and chooses effort e consistent with expected output $x=e+\beta$, then they will receive a transfer z with mean $a\left(\hat{\beta}\right)+b\left(\hat{\beta}\right)(e+\beta)$ and variance $b\left(\hat{\beta}\right)^2\sigma^2$, and achieve utility

$$\begin{aligned}
U\left(\hat{\beta},e,\beta\right) &= \tilde{v}\left(a\left(\hat{\beta}\right)+b\left(\hat{\beta}\right)(e+\beta+\varepsilon),e\right)\\
&= a\left(\hat{\beta}\right)+b\left(\hat{\beta}\right)(e+\beta)-\frac{1}{2}\eta b\left(\hat{\beta}\right)^2\sigma^2-\psi(e)\\
&= \left(a\left(\hat{\beta}\right)-\frac{1}{2}\eta b\left(\hat{\beta}\right)^2\sigma^2\right)+b\left(\hat{\beta}\right)(e+\beta)-\psi(e)\\
&= \tilde{T}\left(x,\hat{\beta}\right)-\psi(e)
\end{aligned}$$

where

$$\tilde{T}\left(x,\hat{\beta}\right)=\left(a\left(\hat{\beta}\right)-\frac{1}{2}\eta b\left(\hat{\beta}\right)^2\sigma^2\right)+b\left(\hat{\beta}\right)x$$

is the implied risk adjusted menu of contracts in the agent's decision space. $\tilde{T}\left(x,\hat{\beta}\right)$ is linear in expected outcome (or equivalently effort e). We note that the slopes of the contracts $T\left(x+\varepsilon,\hat{\beta}\right)$ and $\tilde{T}\left(x,\hat{\beta}\right)$ are the same, but that the intercept

[9] In contradistinction to [16], we restrict ourselves to linear contracts in order to accommodate risk aversion.

of $\tilde{T}\left(x,\hat{\beta}\right)$ is adjusted by an amount that depends on both $a\left(\hat{\beta}\right)$ and $b\left(\hat{\beta}\right)$ to accommodate the cost of risk bearing. We notice[10] that, irrespective of $\hat{\beta}$, the optimal effort level e does not depend on the variance σ^2.

We write

$$u(\beta) = \max_{\hat{\beta},e} U\left(\hat{\beta}, e, \beta\right)$$

for the information rent earned by an agent of type β. We write $e(\beta)$ for the optimal effort function implied by the contract and $x(\beta) = e(\beta) + \beta$ for the implied expected output. We will write

$$\begin{aligned} t(\beta) &= \mathrm{E}\,T(x(\beta)+\varepsilon, \beta) \\ &= T(x(\beta), \beta) \\ &= a(\beta) + b(\beta)\,x(\beta) \\ \tilde{t}(\beta) &= \tilde{T}(x(\beta), \beta) \\ &= a(\beta) + b(\beta)\,x(\beta) - \frac{1}{2}\eta b(\beta)^2 \sigma^2 \\ &= t(\beta) - \frac{1}{2}\eta b(\beta)^2 \sigma^2 \end{aligned}$$

for the expected transfer and the risk adjusted expected transfer respectively. To reduce notation we will when convenient drop the argument β and write x, t, $\tilde{t}$,

[10] This is because of the additive specification of the error ε and our assumption that, for wealth variations implied by the contract, η is constant.

e, u, a, b, f, F, h, instead of $x(\beta)$, $t(\beta)$, $\tilde{t}(\beta)$, $e(\beta)$, $u(\beta)$, $a(\beta)$, $b(\beta)$, $f(\beta)$, $F(\beta)$, $h(\beta)$, and we will denote differentiation with respect to β by a dot. If the principal offers an incentive compatible individually rational contract, then by standard arguments we have

$$\begin{aligned} u &= \tilde{t} - \psi(e) \\ x &= e + \beta \\ \tilde{t} &= t - \frac{1}{2}\eta\sigma^2\psi'(e)^2 \\ \dot{u} &= \psi'(e) \\ b &= \psi'(e) \\ \dot{\tilde{t}} &= \psi'(e)\,\dot{x} \\ \dot{x} &\geq 0 \\ u(\underline{\beta}) &= 0 \end{aligned}$$

These are, respectively, the definitions of u, x and $\tilde{t}$, the envelope condition, the first and second order conditions, incentive compatibility, and individual rationality. The agent's utility depends on x and $\tilde{t}$, and as β varies the contract $\left(x(\beta), \tilde{t}(\beta)\right)$ traces out a locus, the contract curve, in agent's $\left(x, \tilde{t}\right)$ space. The slope of the contract curve is $\frac{d\tilde{t}}{dx} = \frac{\dot{\tilde{t}}}{\dot{x}} = \psi'(e)$. This contract curve is the envelope of the menu of linear contracts $\tilde{T}\left(x, \hat{\beta}\right)$.

The principal's payoff depends on x and t; she chooses a contract (x, t) (and hence implicitly e and u), subject to incentive compatibility and individual rationality constraints, to maximise the objective

$$\max_{\{x(\cdot),u(\cdot)\}} \int_{\underline{\beta}}^{\overline{\beta}} \{x - (1+\lambda)\,t + \theta u\}\, dF$$

where λ is the cost of raising the revenue t, and $\theta \in \left[\underline{\beta}, \overline{\beta}\right]$ is the weight placed by the principal on the agent's utility. We shall interpret the value of these parameters in the policy environments of interest below. After an integration by parts, making use of the envelope condition and the individual rationality constraint, the principal's problem becomes

$$\max_{\{e(\cdot)\}} \int_{\underline{\beta}}^{\overline{\beta}} \left\{ e + \beta - (1+\lambda)\left(\psi\,(e) + \frac{\eta\sigma^2}{2}\psi'\,(e)^2\right) - (1+\lambda-\theta)\,\psi'\,(e)\,h \right\} dF. \quad (1)$$

We note the term containing η in the expression for the virtual surplus; it is through this term that risk aversion enters into the contracting problem.

By a standard argument[11] the integrand is concave in e, and the optimal effort

[11] See for example Chapter 6 of [12].

is determined at interior points by the principal's first order condition

$$1 = (1+\lambda)\,\psi'(e)\left(1+\eta\sigma^2\psi'(e)\,\psi''(e)\right) + (1+\lambda-\theta)\,h\psi''(e)\,.$$

It can then be shown that the contract curve is convex in $\left(x,\tilde{t}\right)$ space, confirming that the contract may be implemented by a menu of linear contracts, tangent to the contract curve.

A linear contract is of the form $T(x) = a + bx$, promising a fixed up-front payment a, and a conditional payment bx that depends upon the output achieved. Since typically $0 \leq b \leq 1$, this takes the form of a surplus sharing rule. For example, if $a = 3$ and $b = .5$, then the contract would specify an up front payment of 3 plus 50% of the value of the output generated. By offering a menu of such contracts $T(x,\beta) = a(\beta) + b(\beta)\,x$, one for each type β, the agent can choose either a soft contract (with a low b and a weak incentive to maximise output) or a tough contract (with a high b and a strong incentive to maximise output). The tough contract will be more attractive to the low cost agents. By the revelation principal, if the menu of contracts is designed to be incentive compatible then each agent will choose the contract appropriate to their type. By choosing this contract, they are in effect induced to reveal their private information, and optimal effort is induced.

In general, to write a linear contract it is sufficient to specify the slope b and a fixed point on the line, which will be some base combination of output and payment. The fixed component may be interpreted as a sign-up fee paid to the agent, an entry fee paid by the agent, an output target, or some combination of these. For example if $a = -1$ and $b = .5$, then the contract $T(x) = -1 + .5x$ can be interpreted to mean that there is a fixed payment (in this case an entry fee of -1), plus 50% of the value of the output generated. But this can also be written $T(x) = .5(x-1)$, and interpreted to say that the agent gets nothing until they reach an output target (of 1 in this case), and then they get 50% of the value of the output that they produce beyond this target.

We now consider the case of multiple agents. Under risk neutrality contract design is particularly simple [17]. The optimal mechanism can be implemented by conducting a preliminary auction of the right to participate, and then offering winners exactly the menu of contracts derived above for the single agent case. If an agent wins the right to participate, paying p, and then selects the contract $a + bx$, then the total fixed payment will be $p + a$. The effect of the auction is to contract the type space of the agents participating in the contract to $\left[\beta_0, \overline{\beta}\right]$, where β_0 is the type of the highest non-participating agent. Thus competition between agents in the preliminary auction reduces, but does not entirely eliminate, uncertainty about the agents' types at the contracting stage. This residual uncertainty is managed through the contract design. An important implication is that contract design is invariant to the number of participants in the auction.

Under risk aversion the analysis is not straight forward, except in one case that we will focus on. If the cost of effort function $\psi(e)$ is quadratic[12] (this includes the case of constant marginal effort) and we write

$$\phi(e) = \psi(e) + \frac{\eta\sigma^2}{2}\psi'(e)^2$$

for the total cost including risk bearing, then the Hamiltonian (1) can be written

$$e + \beta - (1+\lambda)\phi(e) - \frac{(1+\lambda-\theta)}{(1+\delta\eta\sigma^2)}\phi'(e)h$$

where $\delta = \psi''(e)$, which is a constant. The problem is then isomorphic to the risk neutral case, but with the hazard rate adjusted by a constant factor, becoming $\frac{h}{(1+\delta\eta\sigma^2)}$. Laffont and Tirole's proofs then go through exactly as before, so we get the basic separability result. Thus in this quadratic cost case the optimal mechanism can be implemented by a preliminary auction of the right to participate, and then by allowing participants to choose a contract from the menu derived above exactly as in the single agent case. In particular, the design of the menu of contracts does not depend on the number of potential participants in the mechanism.

[12] This is clearly a restrictive assumption, but not unreasonable as a first approximation.

3 An example

We illustrate with a numerical example. Let us assume that β is distributed uniformly on the interval $\left[\underline{\beta}, \overline{\beta}\right]$, so $h(\beta) = \left(\beta - \underline{\beta}\right) / \left(\overline{\beta} - \underline{\beta}\right)$, and that $\psi(e) = \frac{e^2}{2}$. For simplicity, we will assume that $\left[\underline{\beta}, \overline{\beta}\right] = [0,1]$. In this case the principal's first order condition is

$$1 = (1 + \lambda - \theta)(1 - \beta) + (1 + \lambda)\left(1 + \eta\sigma^2\right) e$$

so the contractual level of effort is

$$e = \frac{1 + (1 + \lambda - \theta)(\beta - 1)}{(1 + \lambda)(1 + \eta\sigma^2)}.$$

The principal's virtual surplus (the expression in the integrand (1)) is

$$k(\beta) = \frac{(1 + \lambda - \theta)^2 \beta^2 + 2\left(1 + \theta + (1 + \lambda)\eta\sigma^2 - (\theta - \lambda)^2\right)\beta + \frac{1}{2}(\theta - \lambda)^2}{(1 + \lambda)(1 + \eta\sigma^2)}$$

which is convex, and non-negative at $\beta = 0$. A simple calculation shows that $k(\beta)$ is minimised at $\beta_0 = -\frac{\left(1-\lambda^2\right)+\left(\theta-\theta^2\right)+2\theta\lambda+\eta\sigma^2(1+\lambda)}{(1+\lambda-\theta)^2}$ which is negative provided that $0 \leq \theta \leq 1$ and $0 \leq \lambda \leq 1$. Thus $k(\beta)$ is non negative over the whole interval $[0,1]$, and u can be calculated by integrating $\dot{u} = \psi'(e)$ from zero, yielding

$$u = \frac{(1+\lambda-\theta)\,\beta^2 + 2\,(\theta-\lambda)\,\beta}{2\,(1+\lambda)\,(1+\eta\sigma^2)}.$$

In general, the virtual surplus will not be non-negative over the whole interval. There will be some cut-off type β_0 such that $k(\beta_0) = 0$. The principal will wish to exclude all types below this, and u will be determined by integrating from β_0. The transfers t and $\tilde{t}$ are readily calculated from e and u.

3.1 Policy settings

We show in Figures 1 and 2 the behaviour of variables of interest for typical parameters, using the example calculated above. In these Figures we have allowed the type space to be more general than the $[0, 1]$ used for simplicity above. In fact we have set $\left[\underline{\beta}, \overline{\beta}\right] = [-1, 0.5]$, to illustrate the scenario where many landholders will own a significantly degraded asset which requires a threshold level of effort before it can be considered an environmental off-set, while some may hold assets of considerable value. These Figures also illustrate that there is a critical cut-off type, and types below this would not participate.

We consider, as above, three policy settings. The first is government procurement. The government is efficiency maximising, putting equal weight on the utility of the principal (the tax payer) and the agent, so $\theta = 1$, but raises revenue through distortionary taxation, so $\lambda > 0$. In this example, we have put $\lambda = .2$. This policy setting is shown with a dashed line. The second setting is private procurement of environmental offsets. In this case, we have assumed that the principal (the developer) is a monopsonistic[13] purchaser, facing no funding distortion, and placing no weight on the utility of the agent, so we set $\lambda = \theta = 0$. This case is shown with a solid line. In the third case the government, acting as an intermediary in the off-set market, procures off-sets and sells them on to developers. We assume that the government is benevolent, placing equal weight on the utility of both parties (so we implicitly assume that the government acts to correct the monopsony distortion, maximising total social surplus rather than just the monopsonist's share). Since the off-set is paid for by the developer, there is no funding distortion, and we set $\theta = 1$, $\lambda = 0$. This case is shown with a dotted line. Figure 1 shows effort, output, transfer and the agent's information rent by type. Figure 2 shows the optimal contract. This may be represented as a nonlinear contract in either observable $(x + \varepsilon, t)$ output-transfer space, or in the agent's (e, u) effort-utility decision space.

We also show, in the final panel of Figure 2, the menu of linear contracts that implements the optimal contract. As noted in Section 2, the fixed component of a linear contract may be interpreted as a sign-up fee paid to the agent, an entry fee paid by the agent, an output target, or some combination of these. An output

[13] This means that they can unilaterally design the contract, as is usually assumed in principal-agent problems.

target may be a more reasonable way to write a realistic contract (rather than to expect the payment of up front entry fees), and in the third panel of Figure 2 we have chosen to present the menu of contracts in this form, as a locus of pairs (x_0, b) where x_0 is an output target and b is an output value share.

It is clear from these Figures that the optimal contract is quite different in each policy setting. The offset market contract is particularly simple and it is easy to see why this should be so, especially if agents are risk neutral. In this case, the government wants to maximise total surplus and does not care how it is allocated. In particular, it is happy for the agent to retain all the surplus (this surplus may be extracted separately, through a preliminary auction, as discussed above). The natural solution is to "sell the project to the agent", allowing the agent to retain all the surplus at the margin, and hence to internalise all externalities. The principal thus offers a simple pooling contract[14] inducing the first best effort level. If the agents are risk averse then it is no longer optimal for them to bear all the risk, and the slope b will be less than 1. But risk aversion is unrelated to type under our assumptions, so it is not useful to screen agents by exposure to risk, and a simple pooling contract is still optimal. The same contact $T(x) = a + bx$, with $b < 1$, is offered to everyone. After the right to participate is auctioned the actual payment made to a participating agent will be $p+a+bx$, where p is the participation payment

[14]That is to say, the slope of the contract is the same for everyone; the intercept, which is implicitly determined at the auction stage, may vary between types.

determined competitively through the auction. The payment p will reflect both the type β and risk aversion; p may vary according to the bidder or be uniform across bidders, depending upon the format of the auction (first price or second price).

In contrast, the private purchaser has a strong incentive to minimise transfers to the agent ($\theta = 0$), and implements a diverse menu that separates types strongly, inducing a wide variation in effort levels. The government procurement contract is in some ways intermediate between the two. The government has an incentive to minimise transfers to the agent because of the distortionary cost of raising revenue ($\lambda = .2$). It also implements a screening contract, but one that screens less aggressively than that of the private developer.

The contracts, as actually implemented, look a little different (the details here are sensitive to the parameters chosen). When procuring for the off-set market, the government offers a simple take it or leave it contract (not a menu), characterised by an up front payment and a relatively high performance incentive. The developer offers a wide range of linear contracts, characterised by an output target and a range of performance incentives, some of them quite weak. When procuring to augment the stock of public goods, government offers a menu of contracts that is intermediate between the two. With the parameters chosen here, the high power contracts are characterised by an output target, while the low power contracts are characterised by an up front payment.

4 Discussion

Market based instruments for biodiversity procurement, either to enhance the stock of assets or to facilitate trade in off-set markets, are increasingly important. Standard procurement models can be adapted to provide a workable framework for contract design in such markets. What do we learn from this modelling exercise?

First of all, contract design cannot be treated as a discrete problem without considering the policy environment. Contracting frameworks that make sense in one environment cannot necessarily be applied without thought to another, despite the apparent similarity of the problems.

In the case of biodiversity regulation and procurement, we find that contract design is simplest for the case of off-set markets. In this case there is no need to screen agents through the contracting framework, and a simple off-the-shelf linear contract will do to job. This simple contract will have two components: a fixed component a (which may be either positive or negative) and a variable component b. Participation is then determined by an auction, in which landholders compete to bid down the size of the payment p that they would be paid in order to sign a contract. The net immediate payment, on signing of the contract, is $a + p$. The variable payment, received at the end of the contract, depends on the outcome. b

is the proportion of the value created that goes to the landholder.

In practice, the fixed component may take various forms. As noted above, it may be any combination of payment and output target. The contract may also specify some payments against the verifiable provision of various inputs (for example, erecting fences) in the earliest phase of the contract. In our framework such payments would be part of the fixed component a of the contract. There is clearly considerable freedom in putting together the fixed component bundle a. If the bundle a were varied we would expect the equilibrium bid p to adjust so as to keep $a + p$ constant.

The variable component of the contract, in the case of an off-set market, is likely to be quite steep. If agents are risk neutral, then at the margin they should be paid the full social value of the increment in environmental services that they provide beyond a contractually agreed threshold. If agents are risk averse, then they must be compensated to some degree for bearing risk: we would expect this cost to be built into their initial bid. If agents are risk averse it is optimal to moderate the variable component of the contract to introduce some risk sharing. The slope of the contract will depend on the amount of risk (both production risk and measurement risk) and on the degree of risk aversion; the higher these are, the higher is the cost of risk bearing, and the flatter (softer) will be the optimal contract. Naturally, any actions that can be taken to reduce exposure to risk will lead to more favourable (and steeper) contracts from the principal's point of view. For example, where it is efficient, some level of insurance could be built in (for example, output targets might be conditional on drought conditions, bushfires, or other events beyond the landowner's control). The reduction of measurement and monitoring risk through appropriate science and technology will also lead to more favourable contracts from the principal's point of view.

Optimal contracts for procurement markets will differ from those in an off-set market. In this case it is optimal to screen bidders, inducing them to reveal information about their private characteristics, and reducing the information rents that are paid to them. The incentive to screen is stronger for private procurement than for public procurement, and the private procurer will offer a relatively wider range of contracts, including some that are quite aggressive and will appeal only to high types, and some of lower power. Such screening reduces information rents and the payment to agents, but potentially induces lower effort levels. The public procurer will screen less aggressively, offering a menu of contracts intermediate in structure between the private procurer and the off-set market contract.

Contract design is thus sensitive to the institutional framework. In contrast, at least within the theoretical framework[15] used here, it is not sensitive to the number of bidders. Contract design need not differ according to the intensity of

[15] Most importantly, the separability of effort in the utility function and the quadratic effort function.

competition, which is harnessed by auctioning the right to participate. If there are many potential suppliers, then they will bid down the information rents in the auction. If the asset has no good substitutes, and there are few potential suppliers then rents will be higher, and the designer will be relying more on the screening properties of the contract to minimise rents.

The model that we have used here, based on the Laffont and Tirole [16] procurement model, leads to useful insights into the nature of the contract design problem for biodiversity procurement. In particular, it suggests that contracts designed for off-set matching markets should be simpler in structure, and provide harder incentives, than those designed for procurement. It would be useful to know how robust these conclusions are if we relax the assumptions embedded in this specification, though it is unlikely that closed form solutions can be found, and numerical simulation would be required. Empirical evidence on the validity of these assumptions would be valuable, particularly regarding key inputs, including the degree of risk aversion, the distribution of types in the population, and the shape of the effort curve (and more generally the nature of the biodiversity production function) and the monitoring technology. With respect to the distribution of types, there is a growing body of data from environmental procurement auctions which may be amenable to econometric investigation [27]. Getting information on the effort function and the production function for environmental goods may require a different approach. Biological and biophysical modelling and simulation is widely used to model agricultural systems, and may be able to be adapted to model the production of habitat and biodiversity assets.

As is usual in the application of theory to policy, the main contribution of the framework is probably its value in organising concepts and ideas. However some basis for empirical judgement exists, and any contract design put forward in practice implicitly takes a position on these empirical questions. The model also provides a framework to test hypotheses and improve parameter estimates. Direct experimentation and research, in the form of carefully designed field pilots, would allow economists to begin the process of testing the hypotheses underlying the standard contracting framework. Field pilots, simultaneously delivering environmental gains and research findings, may also be useful in building confidence in policy makers and facilitating implementation.

References

[1] ANAO. *Performance Information: Commonwealth Financial Assistance under the National Heritage Trust*, volume 43. Australian National Audit Office, Canberra, 2001.

[2] John Antle, Susan Capalbo, Siân Mooney, Edward Elliott, and Keith Paustian. Spatial heterogeneity, contract design, and the efficiency of carbon sequestration policies for agriculture. *Journal of Environmental Economics and Management*, 46(2):231–50, 2003.

[3] ANZECC. *National Framework for the Management and Monitoring of Australias Native Vegetation.* Australian and New Zealand Environmental Conservation Council, Canberra, 2000.

[4] D. P. Baron and R. B. Myerson. Regulating a monopolist with unknown costs. *Econometrica*, 50(4):911–930, 1982.

[5] BEST. *Compensating for Wetland Losses Under the Clean Water Act.* Board on Environmental Studies and Toxicology, National Academy Press, Washington, 2001.

[6] Anne-Sophie Crepin. Incentives for wetland creation. *Journal of Environmental Economics and Management*, 50(3):598–616, 2005.

[7] DSE. *Bush Tender - a Landholder Perspective.* Department of Sustainability and Environment, Melbourne, 2006.

[8] Hongli Feng, Lyubov A. Kurkalova, Catherine L. Kling, and Philip W. Gassman. Environmental conservation in agriculture: Land retirement vs.

changing practices on working land. *Journal of Environmental Economics and Management*, 52(2):600–614, 2006.

[9] Paul J. Ferraro. Asymmetric information and contract design for payments for environmental services. *Ecological Economics*, forthcoming.

[10] Paul J. Ferraro and R. David Simpson. Cost-effective conservation: A review of what works to preserve biodiversity. In Wallace E. Oates, editor, *The RFF Reader in Environmental and Resource Policy*, pages 163–168. Resources for the Future, Washington, DC, second edition edition, 2006.

[11] Guenter Franke, Richard C. Stapleton, and Marti G. Subrahmanyam. Background risk and the demand for state-contingent claims. *Economic Theory*, 23(2):321–335, 02 2004.

[12] Drew Fudenberg and Jean Tirole. *Game Theory*. Cambridge University Press, Cambridge, Mass. and London, 1991.

[13] Russell Gorddard, Stuart Whitten, and Andrew Reeson. When should biodiversity tenders contract on outcomes? *Paper presented at 52nd Annual Conference of the Australian Agricultural and Resource Economics Society*, 2008.

[14] Robert Innes, Stephen Polasky, and John Tschirhart. Takings, compensation and endangered species protection on private land. *Journal of Economic Perspectives*, 12(3):35–52, 1998.

[15] Andreas Kontoleon, Unai Pascual, and Timothy Swanson. *Biodiversity Economics: Principles, Methods and Applications*. Cambridge University Press, 2007.

[16] Jean-Jacques Laffont and Jean Tirole. Using cost observation to regulate firms. *Journal of Political Economy*, 94(3):614–641, 1986.

[17] Jean-Jacques Laffont and Jean Tirole. Auctioning incentive contracts. *Journal of Political Economy*, 95(5):921–937, 10 1987.

[18] Jean-Jacques Laffont and Jean Tirole. *A Theory of Incentives in Procurement and Regulation*. Cambridge and London:; MIT Press, 1993.

[19] Andreu Mas-Colell, Michael Dennis Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[20] Eric Maskin and Jean Tirole. The principal-agent relationship with an informed principal: The case of private values. *Econometrica*, 58(2):379–409, 1990.

[21] John Merrifield. A market approach to conserving biodiversity. *Ecological Economics*, 16:217–226, 1996.

[22] Andrew Moxey, Ben White, and Adam Ozanne. Efficient contract design for agri-environment policy. *Journal of Agricultural Economics*, 50:187–202, 1999.

[23] Veronika Nemes, Charles R. Plott, and Gary Stoneham. Electronic bushbroker exchange: Designing a combinatorial double auction for native vegetation offsets. *SSRN eLibrary*, 2008.

[24] David M. G. Newbery and Joseph E. Stiglitz. *The Theory of Commodity Price Stabilization : A Study in the Economics of Risk*. Clarendon Press ; Oxford University Press, Oxford; New York, 1981.

[25] Adam Ozanne, Tim Hogan, and David Colman. Moral hazard, risk aversion and compliance monitoring in agri-environmental policy. *European Review of Agricultural Economics*, 28:329–347, 2001.

[26] Adam Ozanne and Ben White. Equivalence of input quotas and input charges under asymmetric information in agri-environmental schemes. *Journal of Agricultural Economics*, 58:260–268, 2007.

[27] Harry J. Paarsch and Han Hong. *An Introduction to the Structural Econometrics of Auction Data*. MIT Press, Cambridge and London, 2006.

[28] David Parkes, Graeme Newell, and David Cheal. Assessing the quality of native vegetation: The habitat hectares approach. *Ecological Management and Restoration*, 4:29–38, 2003.

[29] PC. *Impacts of Native Vegetation and Biodiversity Regulations*, volume Report No 29. Productivity Commission, Melbourne, 2004.

[30] Glenn Sheriff. Optimal environmental regulation of politically influential sectors with asymmetric information. *Journal of Environmental Economics and Management*, 55(1):72–89, 2008.

[31] Rodney B. W. Smith. The conservation reserve program as a least-cost land retirement mechanism. *American Journal of Agricultural Economics*, 77(1):93–105, 1995.

[32] Daniel F. Spulber. Optimal environmental regulation under asymmetric information. *Journal of Public Economics*, 35:163–181, 1988.

[33] Robert N. Stavins. What can we learn from the grand policy experiment? lessons from so2 allowance trading. *Journal of Economic Perspectives*, 12(3):69–88, 1998.

[34] G. Stoneham, V. Chaudhri, A. Ha, and L. Strappazzon. Auctions for conservation contracts: An empirical examination of victoria's bushtender trial.

Australian Journal of Agricultural and Resource Economics, 47(4):477–500, 12 2003.

[35] G. Stoneham, V. Chaudhri, A. Ha, and L. Strappazzon. Auctions for conservation contracts: An empirical examination of victoria's bushtender trial. In Andreas Kontoleon, Unai Pascual, and Timothy Swanson, editors, *Biodiversity Economics: Principles, Methods and Applications*, pages 389–416. Cambridge University Press, 2007.

[36] Gary Stoneham. Creating markets for environmental goods and services: A mechanism design approach. *SSRN eLibrary*, 2008.

[37] T. Tietenberg. The tradable permits approach to protecting the commons: What have we learned? *The Drama of the Commons*, pages 197–232, 2002.

[38] B. White and M. Burton. Measuring the cost effectiveness of conservation auctions relative to alternate policy mechanisms. *University of Western Australia School of Agriculture and Resource Economics Working Paper*, 2005.

[39] Ben White. Designing voluntary agri-environment policy with hidden information and hidden action: A note. *Journal of Agricultural Economics*, 53:353–360, 2002.

[40] JunJie Wu and Bruce A. Babcock. Contract design for the purchase of environmental goods from agriculture. *American Journal of Agricultural Economics*, 78:935–945, 1996.

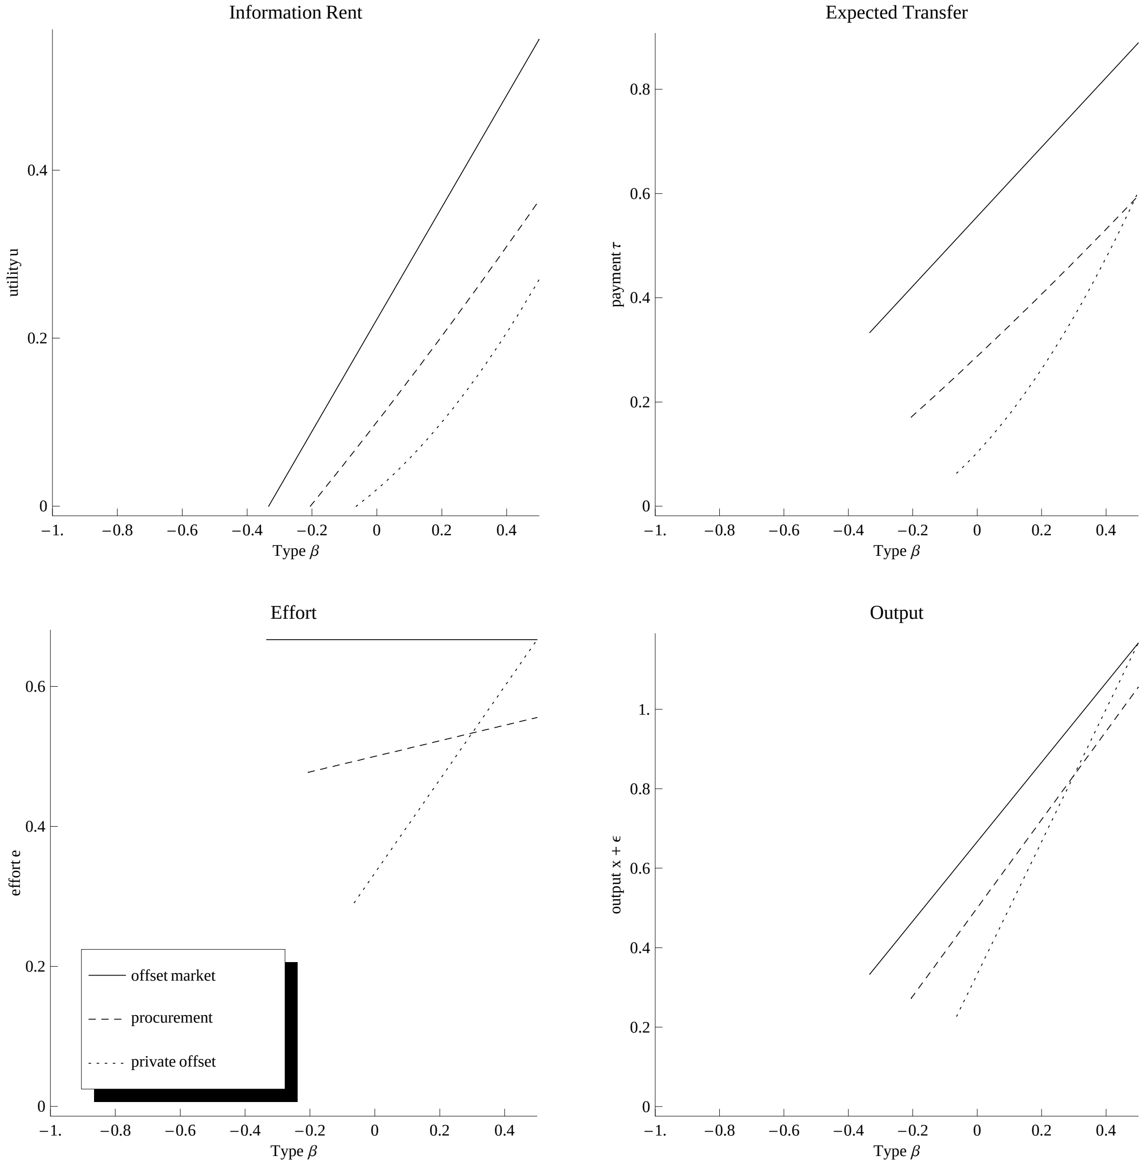

Figure 1: Effort, outcomes, and transfers to type β.

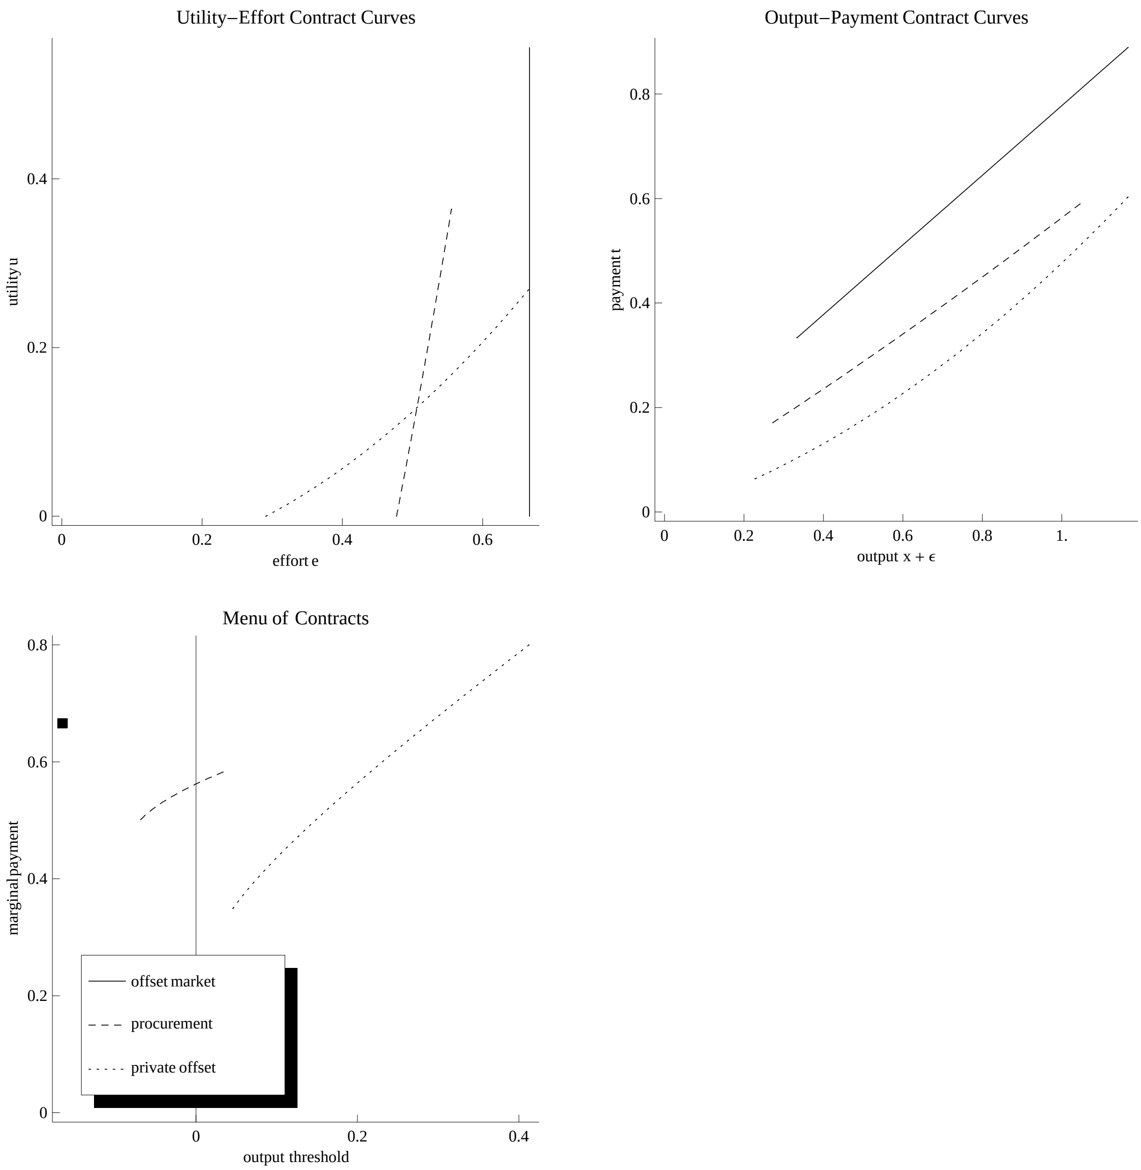

Figure 2: Optimal Contracts