Duality in Contracting

Peter Bardsley
The University of Melbourne

June 2011

Abstract

In a linear contracting environment the Fenchel transform provides a complete duality between the contract and the information rent. Through an appropriate generalised convexity this can be extended to provide a complete duality in the supermodular quasilinear contracting environment that covers the majority of applications. Along the way, a supermodular envelope theorem is proved, somewhat different in nature to the Milgrom Segal result.

Key Words: mechanism design, contract theory, duality, Fenchel transform, abstract convexity.

JEL Codes: D82, D86

Contents

1 Introduction

In the linear agency problem, when the valuation is linear in type, the agent is confronted with a decision problem

$$\max_{x \in X} \theta x - \tau(x).$$

Here θ is the agent's type, $x \in X$ is the allocation chosen by the agent, and $\tau(x)$ is the contract or tariff designed by the principal. This problem is identical in structure to the standard producer problem in price theory

$$\max_{x \in X} px - c(x)$$

where p is the price, $x \in X$ is the output quantity chosen by the agent, and $c(x)$ is the cost function.

The duality framework of price theory embeds the producer's problem as one of a pair of dual problems. The primal problem, described by the cost function, lives on the commodity space and the dual problem, described by the profit function, lives on the price space. These two structures are interlinked by the envelope theorem and its variants. This framework clarifies the producer problem and underlies many useful technical results, in particular the regularity and convexity of the profit function. The duality transform of price theory is a particular example of the Fenchel transform, which is studied systematically in convex analysis (Rockafellar 1970).

Since the linear agency problem and the producer problem are isomorphic, there is a completely parallel duality framework for the agency problem, with the type θ playing the role of the price p. The primal problem is described by the contract function, living on the allocation space, and the dual problem is described by the information rent, living on the type space. While this observation is completely elementary, it is already useful and not always fully evident in standard treatments of the theory. It is exploited, for example, by Krishna and Maenner (2001) in auction theory.

The contribution of this paper is to show that this duality can be extended, virtually in its totality, to the quasilinear framework that is standard in most applications of contract theory.[1] The extension requires a non-trivial but quite intuitive extension of the concept of convexity, which is of interest in itself. The construction is interesting because we get a full duality if and only if the agent's valuation is, after perhaps relabeling some types, supermodular. This observation provides some explanation of why supermodularity or singlke crossing so often leads to tractable models. There is a connection between this incentive condition and the geometrical structures that are exploited in convex analysis.

Just as in the classical case, generalised convexity leads to strong regularity properties for envelope functions. The arguments are very geometric, and are modelled on those of classical convex analysis. We prove a non-convex separation theorem, and display a full duality between the contract and the information rent. In particular, not only the information rent but also the contract function can be represented as an integral over marginal valuations. Along the way we prove an envelope theorem that is related to but different from the envelope theorem of Milgrom and Segal (2002). The envelope theorem proved here is

[1] We assume throughout that the type space is 1 dimensional.

adapted to the quasilinear environment, where it is easier to apply and produces stronger results. The Milgrom Segal Theorem applies in a less restrictive environment, but requires the checking of a side condition.

The results in this paper are related to those of Rochet (1987), who characterised implementability in a quasilinear environment, to Krishna and Maenner (2001), Border (1991), Border (2007) and Vohra (2011) who exploit linear duality and convex geometry in auction theory, and to Milgrom and Segal (2002). The literature on nonsmooth analysis, generalised convexity and the generalised Fenchel transform is comprehensively addressed in Rockafellar and Wets (1998).

The structure of the paper is as follows. The general framework is introduced in Section 1.1. Section **??** sets out the structure of the Fenchel duality that arises in the linear agency case (that is, the agent's valuation is linear; the contract is of course not restricted to be linear). The discussion is set out in such a way that the arguments generalize in as straight forward a way as possible to the general case. In section 3 the necessary machinery is developed to support an appropriate non-convex duality and a generalised Fenchel transform. Along the way, links are made to the convexity literature and an envelope theorem is set out. Duality results for the quasilinear case are set out in Section 4. In Section 5 these results are applied to study the question of when a general nonlinear contract can be approximated by a family of simple contracts. *Which sections are easy to read. Remark: no assumption of compacteness.*

1.1 The Framework

We consider the standard quasi-linear adverse selection contracting problem where an agent of type $\theta \in \Theta$ chooses $\xi(\theta) \subset X$ to solve

$$\rho(\theta) = \sup_{x \in X} v(\theta, x) - \tau(x) \tag{1a}$$

$$= \sup_{(x,t) \in \Gamma} v(\theta, x) - t \tag{1b}$$

$$\xi(\theta) = \operatorname*{argmax}_{x \in X} v(\theta, x) - \tau(x) \tag{1c}$$

$$= \operatorname*{argmax}_{(x,t) \in \Gamma} v(\theta, x) - t. \tag{1d}$$

Here X and Θ are intervals in $\mathbb{R}$, $\Gamma = \operatorname{epi} \tau = \{(x,t) : t \geq \tau(x)\}$ is the epigraph of the contract function $\tau : X \to \bar{\mathbb{R}}$, and $\bar{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$ is the extended real line, with ∞ being an element greater than any real number[2]. As is standard in convex analysis, allowing extended real values is a convenient way of restricting the domain of τ. Points where $\tau(x) = \infty$ are infeasible, since the transfer $\tau(x)$ is infinite. It will be assumed that τ is proper, that is, $\operatorname{dom}(\tau) = \{x : \tau(x) < \infty\} \neq \emptyset$ in order to exclude uninteresting vacuous cases. Assumptions on the agent's valuation function $v(x, \theta)$ will be discussed below.

A standard interpretation of this problem is that $x \in \xi(\theta)$ is a quantity sold by a principal to an agent of type θ in return for a payment $\tau(x)$. Equivalently, Γ is a menu of contracts offered by the principal. Under this interpretation the value function $\rho(\theta)$ is the information rent accruing to the agent of type θ. There are of course many other interpretations of this canonical model.

[2] Note that there is no assumption that either X or Θ be compact.

If the agent's valuation is linear, that is $v(x, \theta) = \theta x$, then Problem (1) is a standard construct in convex analysis, the Fenchel transform, and the properties of this transform provide a great deal of information about the contracting problem. The main contribution of this paper is to show that a nonconvex duality can be used to define a generalised Fenchel transform which throws almost as much light on the quasi-linear case.

The key idea of convex duality is that a convex object can be approximated by the affine half spaces that support it. Broadly speaking, a set is amenable to such approximation if it has no concave sections — hollows or dents that cannot be penetrated by a supporting affine hyperplane. Nonconvex duality extends this idea by using a family of primitive objects, for example conical or parabolic sets, that can penetrate hollows and approximate a larger family of objects. The main innovation in this paper is to construct these primitive approximating objects not in an ad hoc way but from the problem at hand: from the agent's indifference curves in contract space. It will transpire that a satisfactory nonconvex duality can be defined whenever the valuation $v(x, \theta)$ is supermodular.

The focus in this paper is on the agent's problem. Relatively little will be said about the principal's problem or about the structure of optimal contracts. The approach is relatively direct and elementary, in so far as there will be no use of control theory. In particular there is no need for any a priori assumption that the agent's choice $x(\theta)$ is an absolutely continuous function or for delicate arguments in control theory, as is required in most standard control theoretic approaches to the agency problem (see, for example, Guesnerie and Laffont (1984), Fudenberg and Tirole (1991), Jullien (2000)). There will be no discussion of participation constraints or individual rationality, as the agent's participation decision can be separated from the agent's action decision conditional on participation, which is what we study here. The participation decision is straight forward, as it is monotonic in type, and the ideas presented here throw no new light on this.

2 Linear Agency

In this section we study Problem (1) under the assumption that the agent's valuation $v(x, \theta) = \theta x$ is linear in both the quantity x and the agent's type θ, so the information rent $\rho(\theta)$ is the Fenchel transform of the contract $\tau(x)$ (Rockafellar 1970). When we refer to this as a *linear agency problem* we mean that the agent's valuation θx is linear; the contract $\tau(x)$ will of course not in general be linear. It is useful to explore the linear case in some detail as a first step because the structure of the results can be seen with a minimum of technical machinery. This structure, and most of the proofs, will carry across virtually unchanged to the general case, once some machinery is developed.

Fenchel duality is most familiar to economists as the duality between the cost function $c(x)$, on commodity space, and the profit function $\pi(p) = \sup_x px - c(x)$, on price space. Under standard assumptions $c(x)$ and $\pi(p)$ are Fenchel duals, and Fenchel duality underlies the interplay between economic analysis in commodity space and in price space, in particular the well known applications of the envelope theorem. But this framework applies equally well to the linear agency problem where the primal function $\tau(x)$ lives on the choice space X and the dual function $\rho(\theta)$ lives on the type space Θ. In this case the dual variable is the agent's type θ, which plays a role exactly analogous to the dual variable p in price theory, and all the standard duality results of price theory carry over. For standard references on Fenchel duality see Rockafellar (1970) or Rockafellar and Wets (1998, Section 11.L).

There are two ways to approach the concept of convexity: internal and external. The internal route characterizes convexity of a set by the property that if the set contains two points then it must contain the interval between them. The external route approaches convexity by requiring that a set be the intersection of the affine half spaces containing it. Of course this leads to a slightly different concept: in finite dimensions a set is the intersection of affine half spaces if and only if it is both closed and convex. In particular, if a proper function is bounded below[3] then it is the envelope of its affine supports if and only if it is convex and lower semicontinuous. We will call such functions *envelope functions*. The interplay between these two convexity concepts, internal and external, generates many of the important results in convexity and will be fundamental to this paper as well.

It is useful to recall some basic definitions and results from convex analysis, as a basis for generalizations that will be presented later. A proper function $\tau(x)$ is convex if its epigraph $\operatorname{epi}\tau$ is a convex set. It is lower semicontinuous if $\operatorname{epi}\tau$ is a closed set. An affine function $\phi(z) = \alpha z + \beta$ is an affine minorant of τ if its graph lies below $\operatorname{epi}\tau$; that is, $\phi(z) \leq \tau(z)$ for all z. An affine support at x is an affine minorant ϕ whose graph touches $\operatorname{epi}\tau$ at x; that is, $\phi(x) = \tau(x)$. The slope α of an affine support at x is a subgradient at x. The set of all subgradients at x is written $\partial\tau(x)$. The subgradient $\partial\tau$ is a generalisation of the derivative that allows us to handle both smooth points and kinks in the convex function τ in a consistent manner. If a convex function τ is differentiable at x then the subgradient is a singleton set, containing just the derivative $\tau'(x)$, but in general the subgradient is a set valued concept.[4]

Given a proper convex function $\tau(x)$, its Fenchel transform or conjugate is

$$\tau^*(\theta) = \sup_{x \in X} \theta x - \tau(x)$$
$$\xi(\theta) = \operatorname*{argmax}_{x \in X} \theta x - \tau(x).$$

The conjugate τ^* has attractive properties. For a start, τ^* is by construction the envelope of its affine minorants, so it is an envelope function. It is convex and lower semicontinuous. Furthermore it is locally Lipshitz continuous on its domain, absolutely continuous, differentiable almost everywhere, and it is the integral of its derivative. The allocation correspondence ξ is monotone: if $\theta' \geq \theta$ and $x' \in \xi(\theta')$, $x \in \xi(\theta)$ then $x' \geq x$. If the function τ that we start with is an envelope function then the correspondence ξ is in fact maximal monotone. Maximal monotonicity means that ξ cannot be expanded as a correspondence without violating monotonicity. Maximal monotone correspondences on $\mathbb{R}$ have a simple characterization that we will use repeatedly. Let $x(\theta) \in \xi(\theta)$ be a selection from the correspondence. Then $x(\theta)$ is a monotonic function so it is continuous except at a countable number of points where it has a jump discontinuity. The maximal monotone correspondence ξ can be reconstructed from the

[3] It is sufficient that it be bounded below by some affine function.

[4] This is illustrated by the absolute value function $\tau(x) = |x|$, for which $\partial\tau(x) = \{1\}$ if $x > 0$, $\partial\tau(x) = \{-1\}$ if $x < 0$, but $\partial\tau(x) = [-1, 1]$ if $x = 0$.

graph of the function $x(\theta)$ by "filling in the jumps". The correspondence ξ, considered as a set valued function, is given by the formula $\xi(\theta) = [x_-(\theta), x_+(\theta)]$, where $x_-(\theta)$ and $x_+(\theta)$ are the limits of $x(\theta)$ from the left and the right respectively.

There is a close correspondence between $\tau(x)$ and its biconjugate $\tilde{\tau}(x) = \tau^{**}(x)$, constructed by applying the Fenchel transform twice. If τ is an envelope function then $\tilde{\tau} = \tau$. $\tilde{\tau}$ is is the supremum of the affine minorants of τ, and it is the greatest envelope function that lies below τ. It can be constructed from τ by filling in the gaps: adding in any limit points to make it lower semicontinuous and then filling in the non-convex parts of the epigraph. See Figure 1 below.

2.1 Duality

Fenchel conjugation leads to a complete duality theory for the linear agency problem. As mentioned above, the agent's problem is isomorphic to the standard producer's problem but with type θ playing the role of price p. The results are cleanest if the contract function τ is an envelope function. That is, τ is convex and lower semicontinuous, and hence the envelope of its affine minorants. We treat this case first.

As above, write $\rho(\theta) = \tau^*(\theta) = \max_{x\in X} \theta x - \tau(x)$, $\xi(\theta) = \operatorname{argmax}_{x\in X} \theta x - \tau(x)$ and $\eta(x) = \operatorname{argmax}_{\theta\in\Theta} \theta x - \rho(\theta)$. The interpretation is that $\tau(x)$ is the contract function, $\rho(\theta)$ is the information rent function, $\xi(\theta)$ is the set of choices over which an agent of type θ can optimally mix, and $\eta(x)$ is the set of types who can optimally choose x. Let $x(\theta) \in \xi(\theta)$ and $\theta(x) \in \eta(x)$ be selections from the respective correspondences.

The duality structure can be summarised as follows (see Rockafellar (1970), Rockafellar and Wets (1998, Chapters 11, 12)).

1. Both τ and ρ are envelope functions; in particular, they are both absolutely continuous and differentiable almost everywhere. The correspondences ξ and η are maximal monotone;

2. The functions τ and ρ are Fenchel conjugates: $\rho = \tau^*$ and $\tau = \rho^*$. In particular, $\tau = \tau^{**} = \tilde{\tau}$ and $\rho = \rho^{**} = \tilde{\rho}$;

3. $\partial\rho = \xi$ and $\partial\tau = \eta$; in particular, at points of differentiability $\rho'(\theta) = x(\theta)$ and $\tau'(x) = \theta(x)$;

4. The Fenchel inequality holds: $\tau(x) + \rho(\theta) \geq \theta x$, with equality iff $\theta \in \partial\tau(x) = \eta(x)$ iff $x \in \partial\rho(\theta) = \xi(\theta)$

5. ξ and η are inverse correspondences: $x \in \xi(\theta)$ if and only if $\theta \in \eta(x)$;

6. We have integral representations

$$\rho(\theta) = \rho(\theta_0) + \int_{\theta_0}^{\theta} x(\theta)\, d\theta \tag{2}$$

$$\tau(x) = \tau(x_0) + \int_{x_0}^{x} \theta(x)\, dx. \tag{3}$$

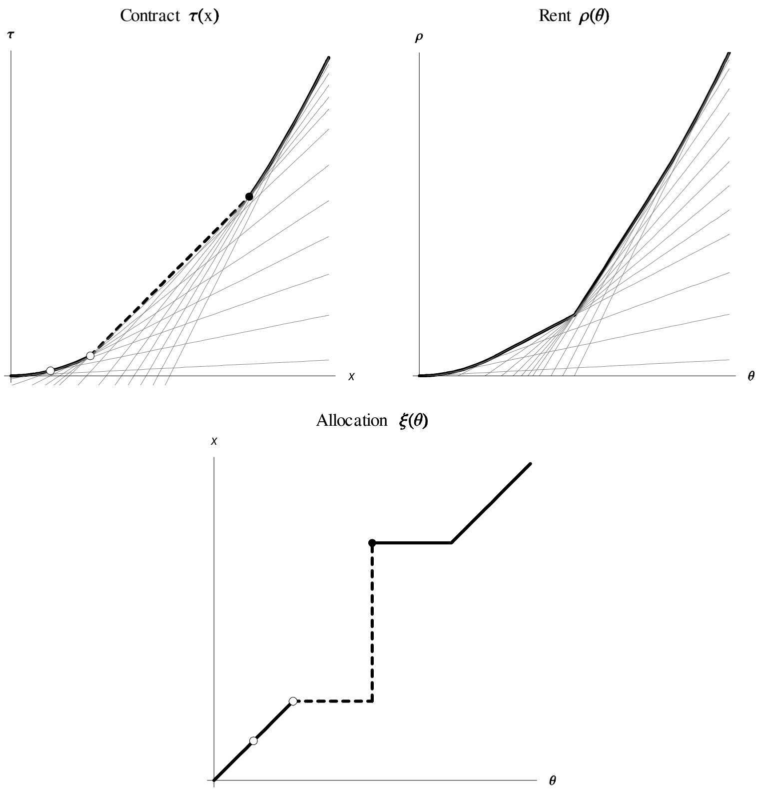

Figure 1: Linear duality

The the integral envelope representation 2 of the information rent is standard. The entirely symmetric relationship for the contract function is striking, but often unnoticed. It is worthwhile to take note of what underlies these representations, since these will be generalised below. There are two steps. First of all, the envelope relationship (3 above) establishes that the subgradient of the value function is the subgradient of the valuation $v\left(x\,,\theta\right)$; in the linear case this is just the allocation correspondence ξ. The second step uses the regularity properties of envelope functions (1 above) to rewrite this as an integral representation.

In the case that τ is not an envelope function τ, ρ, ξ, and η can be constructed in two steps. It is convenient to impose the convention, which will be in force for the rest of the paper, that $\tau\left(x\right)=\infty$ if $\tau\left(x\right)>\tilde{\tau}\left(x\right)$. This will not change the decision of any agent. In the first step, since $\tilde{\tau}$ is an envelope function, the relationship between $\tilde{\tau}$, ρ, $\tilde{\xi}$, and $\tilde{\eta}$ is as described above. In particular, $\tilde{\tau}$ and ρ have envelope representations as in the equations above. In the second step we construct τ by restricting $\tilde{\tau}$ to $\operatorname{dom}\tau$ and construct ξ by restricting $\tilde{\xi}$ such that $\operatorname{im}\xi=\operatorname{dom}\tau$ (that is, $\xi=\left\{(\theta,x):x\in\tilde{\xi}\left(\theta\right)\cap\operatorname{dom}\tau\right\}$). For an illustration, refer to Figure 1. This shows the typical relationship between τ, ρ, ξ and their biconjugates $\tilde{\tau}$, $\tilde{\rho}$, $\tilde{\xi}$, which are in duality as described above.

Construction of $\tilde{\tau}$ by closure and linear interpolation corresponds to max mon by horizontal and vertical segments. ???

2.2 When will the principal choose an envelope contract?

Is it reasonable to focus on envelope contracts? Is there any loss of generality in such an assumption? In price theory, it is generally argued that the predictive content of the theory is unchanged if a cost function $c(x)$ is replaced with its lower semicontinuous convexification $\tilde{c}(x)$. That is to say, there is no significant loss of generality in assuming that $c(x)$ is an envelope function. The argument is that points on the cost function where $c(x) \neq \tilde{c}(x)$ are economically irrelevant, since are inaccessible to any hyperplane and they will never be chosen at any price. For an exposition of this point of view, set out with respect to consumer theory rather than producer theory, see Jehle and Reny (2000, Chapter 2.1.2). This line of argument suggests that without significant loss of generality we may assume that the principal will offer an envelope contract $\tau(x)$.

Here we diverge from the main line of the paper in order to explore this issue briefly. Before doing so let us first note in a practical vein that, since $\rho = \tau^* = \tilde{\tau}^*$, one cannot distinguish between τ and $\tilde{\tau}$ through the information rent ρ. So in any model where the optimal contract can be computed by characterizing ρ by a control theoretic calculation, and appealing to a revenue equivalence argument, one can assume without loss of generality that the principal chooses an envelope contract $\tau = \tilde{\tau}$. This will certainly be so under a monotone likelihood assumption.

Let us now consider the issue a little more generally.[5] In the interest of simplicity we assume, for the remainder of this section, that the type space is non-atomic. We also assume, as is standard in mechanism design, that we are considering weak implementability: if there are multiple equilibria, one can focus on the equilibrium in which the agent chooses the contract more favourable to the principal. For a discussion, see (Jackson 2001). Let τ be an optimal contract, according to the objectives of the principal, and let ξ be the allocation correspondence implemented by τ. Since ξ is a correspondence, the behaviour of the agent is not entirely determined by τ. At mixing points, where there is a vertical jump in the correspondence, there is a range of optimal actions, and we must specify how the agent mixes. But this can be ignored since there can be at most a countable number of jumps in a monotonic correspondence, and the type space is non-atomic, so the set of mixing types is of measure zero.

Let $\tilde{\tau}$ be the biconjugate of τ, and $\tilde{\xi}$ the allocation correspondence implemented by $\tilde{\tau}$. If the principal is to choose a non-envelope contract, then this means that $\tau \neq \tilde{\tau}$. Since $\tau \subset \tilde{\tau}$, and $\xi \subset \tilde{\xi}$, then this means that the principal will offer a contract with gaps in it. Assume, for example, that $\tilde{\tau}$ is smooth and strictly concave. Then one could delete all points (θ, t) with θ rational. But this will not change the pay-off to the principal since the rationals have measure zero. The assertion is that, provided the principal's payoff is continuous in the allocation x, it can never be strictly optimal for the principal to exclude types in this way.

To argue more formally, assume that it is optimal for the principal to offer a non-envelope contract τ. We compare this with the envelope contract $\tilde{\tau}$. Then (under the normalisation convention introduced above that $\tau(x) = \infty$ if $\tau(x) > \tilde{\tau}(x)$,) and identifying the functions set theoretically with their graphs) we have $\tau \subset \tilde{\tau}$. So it must be the case that it is optimal for the principal delete points

[5]These arguments apply mutatis mutandis to the quasilinear case once the appropriate machinery is developed below.

from the menu of contracts by deleting points from the graph of $\tilde{\tau}$. We will show that such deletions must occur at points where there is a discontinuity in the principal's payoff, or must give rise to a non-monotonic participation set (that is, there are types θ_1, θ_2 who will participate and accept contracts, but types θ with $\theta_1 < \theta < \theta_2$ who will not), and that these can be excluded under reasonable assumptions. Let $\tau_E \subset \tau$ and $\tilde{\tau}_E \subset \tilde{\tau}$ be the sets of exposed points (these are points that do not lie in the interior portion of any affine segment of the graph of the function). Then it can be shown the τ_E is dense in $\tilde{\tau}_E$.[6] In particular, any isolated exposed points in $\tilde{\tau}$ must also lie in τ.

Proposition 1 *Assume that the distribution of types is nonatomic and that the principal maximizes the expected value of a continuous function $\phi(x)$ of the decision x (for example, $\phi(x) = \tau(x) - c(x)$, where $c(x)$ is a continuous function), subject to incentive compatibility constraints. Then we may assume without loss of generality that the principal offers an envelope contract.*

Proof. Note first that $\phi(x) = \tau(x) - c(x)$ is continuous, as asserted in the statement of the proposition, since $\tau = \tilde{\tau}$ on $\operatorname{dom}\tau$ and $\tilde{\tau}$ is continuous on its domain. We can assume that $\phi(x)$ is non-negative on $\operatorname{dom}\tau$, since the principal could delete any points where $\phi(x) < 0$. This could only improve the principal's payoff.

To modify $\tilde{\tau}$ into τ the principal must delete points from $\operatorname{dom}\tilde{\tau}\diagdown\operatorname{dom}\tau$. We consider this in two steps, deleting first any non-exposed points in $\operatorname{dom}\tilde{\tau}\diagdown\operatorname{dom}\tau$ and then any exposed points. In each case we check that the expected return to the principal cannot increase as a result of these deletions.

If x is a non-exposed point then it lies in an affine segment of the graph of $\tilde{\tau}$. Let θ be the slope of this segment. Then x is optimal for type θ, and x lies in a non-degenerate interval $\xi(\theta)$ over which θ mixes. This corresponds to a vertical step in the allocation correspondence ξ. If x is deleted then type θ may well change their behaviour, choosing a different point in $\xi(\theta)$. But there can be at most a countable number of vertical steps in ξ, and a countable number of types whose behaviour might change. Since the type space is non-atomic, this is a set of measure zero.

The behaviour of an agent θ at an exposed point x is quite different. By definition, x is the unique optimal choice of this agent. If x is deleted from the menu then the agent's optimal choice set will be empty. If x' were a candidate optimum then x' yields a strictly lower payoff that x. But, since τ_E is dense in $\tilde{\tau}_E$, there are points in $\operatorname{dom}\tau$ yielding payoffs arbitrariliy close to that of x. So x' cannot be an optimum. So when exposed points are deleted, agents who were choosing these points drop out and no longer participate. But the principal's payoff is continuous and non-negative on $\operatorname{dom}\tau$, so $\phi(x) \geq 0$ at any exposed point in $\operatorname{dom}\tilde{\tau}\diagdown\operatorname{dom}\tau$. Deleting these points cannot improve the expected payoff to the principal. ■

The assumption about the principal's objective is that she cares about the transacted quantity and the price, but not otherwise about the identity of the agent. This assumption might fail if, for example, the principal were altruistic or if higher types were more expensive to serve. But even in this case one must remark that any such non-envelope contract would have the implausible feature that a positive measure of agents have been excluded by the principal's deleting

[6] By the Minkowski-Krein-Milman theorem (Rockafellar 1970).

exposed points of the contract space. These agents would reject the contract not because they have a better outside option but because they cannot make up their mind. There are contract points arbitrarily close to the deleted point, but no actual optimal response.

3 Abstract Convexity

The basic idea that drives convex duality is to approximate a set (or a function) by a family of primitive objects: affine hyperplanes. Broadly speaking, a set is amenable to such approximation if it has no concave sections — hollows or dents that cannot be penetrated by a supporting affine hyperplane. Nonconvex duality extends this idea by using a family of primitive objects, for example conical or parabolic sets, that can penetrate hollows and approximate a larger family of objects. See for example Rockafellar and Wets (1998, Chapter 11.K). This is the approach that will be adopted here in constructing a duality for contracting problems. The primitive approximating objects will not, however be arbitrary. They will be adapted to the problem at hand, and will be constructed from the payoff functions for the agents in contract space. Rather than begin with sets, it is more convenient to start with an abstract convexity concept for functions. Convexity for sets will then follow by requiring that their characteristic function be convex.

We now leave the linear case and return to the quasilinear framework of Problem 1. Let Φ be the set of functions $f : \mathbb{R} \to \mathbb{R}$ of the form $f(z) = v(z, \theta) - \alpha$, where $\theta \in \Theta$, and $\alpha \in \mathbb{R}$. Following Rockafellar and Wets (1998) (see also Pallaschke and Rolewicz (1997) and Rubinov (2000)), regard Φ as a set of elementary functions which will be used to define an abstract convexity class. If Φ were the set of affine functions then this will lead to the class of convex functions of classical convex analysis (Rockafellar 1970), but we will work with an abstract convexity class that is more closely adapted to Problem 1.[7]

The existing literature moves directly from Φ to the class of *externally* Φ-convex functions, the envelopes of elementary functions in Φ, and does not address the issue of an appropriate definition of *internal* convexity or the relationship between the two ideas (Pallaschke and Rolewicz 1997)(Rubinov 2000). We have sufficient structure to address this issue properly, and there are significant payoffs from doing so. The basic construct is to define a Φ-interval in $\mathbb{R} \times \mathbb{R}$. This is a set of the form $\{(x, v(x, \theta) - \alpha) : x \in [x_0, x_1]\}$, where $\theta, \alpha \in \mathbb{R}$ and $[x_0, x_1]$ is an interval in $\mathbb{R}$. A Φ-interval in $\mathbb{R} \times \mathbb{R}$ is thus the graph of an elementary Φ-function $v(x, \theta) - \alpha$ restricted to an interval $[x_0, x_1]$. If we note that the equation for the indifference curve for an agent of type θ is $t = v(x, \theta) - \alpha$, then it is apparent that a Φ-interval is just a connected segment of an agent's indifference curve in contract space. A Φ-interval joining two points plays the same role in Φ-convexity as rectilinear intervals play in classical convexity. A set $S \subset \mathbb{R} \times \mathbb{R}$ is *internally Φ-convex* if, for any $(x, t), (x', t') \in S$, and any Φ-interval I with endpoints (x, t) and (x', t'), $I \subset S$.

As in standard convex analysis, there are two ways to define convexity of a function: internal and external. The function $\tau(x)$ is *internally Φ-convex* if the epigraph $\operatorname{epi} \tau = \{(x, t) : t \geq \tau(x)\}$ is an internally Φ-convex subset of $\mathbb{R} \times \mathbb{R}$.

[7] The value of $v(x, \theta)$ outside $X \times \Theta$ is immaterial. It may be convenient for it to take the value ∞ outside this domian. The main thing is that it is finite valued on $X \times \Theta$.

It is *externally* Φ*-convex* if it is the upper envelope of elementary functions in Φ. That is, in the classical case, if its epigraph is the intersection of the affine half spaces containing it. These two concepts are of course not the same.

The existing literature on abstract convexity focuses solely on outer convexity; the internal concept appears not to have been studied. Rather awkwardly, and in conflict with classical convex analysis, in this literature a function $\tau(x)$ is defined to be Φ-convex if it is the upper envelope of elementary functions in Φ.[8] We will attempt to avoid this confusion by calling outer convex functions envelope functions. If we occasionally use the term Φ-convex without qualification it will always mean internally Φ-convex, as in standard convex analysis.

If this concept is to be useful, the minimal requirement is that the elementary functions Φ be indeed (internally) convex. We will say that $v(x,\theta)$ has *monotone differences* if, for all $\theta, \theta' \in \Theta$, the function $v(\cdot,\theta') - v(\cdot,\theta)$ is monotonic. That is either, for all $x' > x$, $v(x',\theta') - v(x,\theta') \geq v(x',\theta) - v(x,\theta)$ or, for all $x' > x$, $v(x',\theta') - v(x,\theta') \leq v(x',\theta) - v(x,\theta)$. This is implied if $v(x,\theta)$ is supermodular or if $v(x,\theta)$ is submodular, or if there is a mixture of cases. It is equivalent to the requirement that, after an appropriate relabelling of types, $v(x,\theta)$ is supermodular.

Lemma 1 *The elementary functions* Φ *are internally* Φ*-convex if and only if* $v(x,\theta)$ *has monotone differences.*

Proof. The following are equivalent

1. $v(x,\theta)$ does not have monotonic differences
2. There exist θ, $\theta' \in \Theta$ and $x < x' < x'' \in X$ such that $h(x) = v(x,\theta') - v(x,\theta)$ is strictly increasing over one of the intervals $\{x,x'\}, \{x',x''\}$ and strictly decreasing over the other
3. There exist θ, $\theta' \in \Theta$ and $x < x' < x'' \in X$ such that $h(x) = v(x,\theta') - v(x,\theta)$ is strictly increasing over the interval $\{x,x'\}$ and strictly decreasing over $\{x',x''\}$ (just switch θ, θ' if necessary)
4. There exist θ, $\theta' \in \Theta$, $x < x' < x'' \in X$ and a constant c such that $h(x) + c > 0$, $h(x') + c < 0, h(x'') + c > 0$
5. There exist elementary functions $f(x) = v(x,\theta')$ and $g(x) = v(x,\theta) + c$ such that $f(x) > g(x)$, $f(x') < g(x')$, $f(x'') > g(x'')$
6. There exists an elementary function $g(x)$ that is not internally convex.

■

This Lemma creates a useful link between the incentive-related single crossing properties that are pervasive in contract theory and the geometric insights associated with convexity.

The Fenchel transform generalizes in a straight forward way to this abstract convexity setting (see Rockafellar and Wets (1998, Section 11.L)) .

[8] Rockafellar and Wets sensibly avoid this terminology.

Let

$$
\begin{aligned}
\tau^{*}(\theta) &= \sup_{x\in X} v(x,\theta) - \tau(x) \\
\xi(\theta) &= \operatorname*{argmax}_{x\in X} v(x,\theta) - \tau(x) \\
\tilde{\tau}(x) &= \tau^{**}(x) = \sup_{\theta\in\nleq} v(x,\theta) - \tau^{*}(\theta)
\end{aligned}
$$

be the generalised Fenchel Φ-conjugate and bi-conjugate of τ . We note that in defining τ^{*} we could equivalently have written $\tau^{*}(\theta) = \sup_{(x,t)\in\operatorname{epi}\tau} v(x,\theta) - t$, so τ^{*} depends only on $\operatorname{epi}\tau$ and the class Φ.The following result is routine, but we include a proof for completeness.

Lemma 2 (abstract Fenchel transform) *Let $\tau^{*}(\theta)$ and $\tilde{\tau}(x)$ be defined by equation 3. Then*

1. *τ^{*} is the upper envelope of the family Φ of functions $v(x,\theta) - t$ parametrised by points $(x,t) \in \operatorname{epi}\tau$*
2. *$(\theta, r) \in \operatorname{epi}\tau^{*}$ parametrizes the set Φ of functions $f(x) = v(x,\theta) - r$ that are dominated by $\tau(x)$*
3. *$\tilde{\tau}$ is the greatest Φ-envelope function that lies below Γ*
4. *the Fenchel inequality holds: $\tau^{*}(\theta) + \tau(x) \leq v(x,\theta)$, with equality if and only if $x \in \xi(\theta)$*
5. *$\tau(x) = \tilde{\tau}(x)$ if and only if $x \in \operatorname{im}\xi = \cup_{\theta\in\Theta}\xi(\theta)$.*

Proof. We have already noted that $\tau^{*}(\theta) = \sup_{(x,t)\in\Gamma} v(x,\theta) - t = \sup_{(x,t)\in\Gamma} f_{x,t}(\theta)$. This establishes (1).Now $(\theta,\rho) \in \operatorname{epi}\tau^{*}$ iff, for all x, $\rho \geq v(x,\theta) - \tau(x)$ iff, for all x, $\tau(x) \geq f_{\theta,\rho}(x)$. This establishes (2). Then by (1), $\tilde{\tau}$ is the upper envelope of the Φ functions parametrised by points in $\operatorname{epi}\tau^{*}$. But by (2) these are just the Φ functions that lie below Γ. The Fenchel inequality is immediate from this. ■

As in the classical case, the Fenchel transform sets up a one to one correspondence between Φ-envelope functions and their conjugates.

3.1 A characterization of Φ-envelope functions

The key step in exploiting the power of the classical Fenchel transform is to characterise the functions that can arise as the envelope of a family of affine functions. The main result in this section will be an analogous characterization of the Φ-envelope functions under a strict supermodularity[9] assumption.

Theorem 1 *Assume that $v(x,\theta)$ is continuous and strictly supermodular, and that $\tau(x)$ is proper, internally Φ-convex and lower semicontinuous. Then the correspondence $\xi(\theta) = \operatorname{argmax}_{x\in X} v(x,\theta) - \tau(x)$ is maximal monotone.*

[9] If we allow nonstrict supermodularity the results become much messier as pooling can arise not only because of kinks in the contract but because the agent's marginal utility is not responsive to type in some regions. Since this behaviour is driven by the way that the agents' preferences are parametrised, and this can be arbitrarily complicated, we avoid these cases.

Proof. We first check that the correspondence $\theta \mapsto \xi(\theta)$ is monotonic. Let $x \in \xi(\theta), x' \in \xi(\theta')$, with $\theta' < \theta$. We must check that $x' \leq x$. Let $t = \tau(x) - v(x, \theta)$, and $t' = \tau(x') - v(x', \theta')$. By incentive compatibility $v(x, \theta) - t \geq v(x', \theta) - t'$ and $v(x', \theta') - t' \geq v(x, \theta') - t$. Adding these inequalities, $v(x', \theta') - v(x, \theta) - v(x', \theta') + v(x, \theta) \geq 0$. But this contradicts the supermodularity assumption unless $x' \leq x$.

Consider first the case where X is compact, and note that in this case the correspondence ξ is closed, non-empty compact valued, and upper hemicontinuous. This is, in essence, the Berge maximum theorem (Berge 1963), but since this is usually stated under an assumption of continuity, not semi-continuity, we outline the argument. It is immediate that ξ is non-empty compact valued since X is compact and τ is lower semicontinuous. Assume that ξ is not closed. Then there exist $x_i \to x$, $\theta_i \to \theta$, $x_i \in \xi(\theta_i)$, but $x \notin \xi(\theta)$. That is, there exists z such that $v(x, \theta) - \tau(x) < v(x, \theta) - \tau(x)$. Now, by the lower semicontinuity of τ and the continuity of v we have

$$
\begin{aligned}
\limsup v(x_i, \theta_i) - \tau(x_i) &\leq v(x, \theta) - \tau(x) \\
&< v(z, \theta) - \tau(z) \\
&= \lim v(z, \theta_i) - \tau(z)
\end{aligned}
$$

so eventually there exists i such that $v(x_i, \theta_i) - \tau(x_i) < v(z, \theta_i) - \tau(z)$, which contradicts the assumption that $x_i \in \xi(\theta_i)$. So the correspondence ξ is closed. It then follows by standard results in the cited references that ξ is upper hemicontinuous.

Since for all θ it is the case that $\xi(\theta) \neq \emptyset$, to show that ξ is maximal monotone it is sufficient to show that the image $\Xi = \operatorname{im} \xi$ of ξ is connected.[10] Let $a < b$ be in the range Ξ, and let x be such that $a < x < b$ but $x \notin \Xi$. Let $x_- = \sup\{z \in \Xi : z < x\}$, and let $x_+ = \inf\{z \in \Xi : z > x\}$. Now there exist x_i, θ_i such that $x_i \in \xi(\theta_i)$ and $x_i \downarrow x_+$, and hence θ_i is decreasing. Moving if necessary to a subsequence, we may assume that $\theta_i \downarrow \theta_+$ for some θ_+. By upper hemicontinuity, $x_+ \in \xi(\theta_+)$. Similarly, θ_- is the greatest θ such that $x_- \in \xi(\theta)$. If $\theta_- < \theta < \theta_+$ then by monotonicity $\xi(\theta) \subset [x_-, x_+]$ but it cannot take on any interior value. Thus $\xi(\theta) \subset \{x_-, x_+\}$. Arguing again by monotonicity, there exists a $\tilde{\theta}$ between θ_- and θ_+ such that $\xi(\theta) = \{x_-\}$ for $\theta < \tilde{\theta}$, and $\xi(\theta) = \{x_+\}$ for $\tilde{\theta} < \theta$. But this would contradict the minimality of θ_+ or the maximality of θ_-. Thus $\theta_- = \theta_+$, and $\xi(\theta_-) = \xi(\theta_+) = \{x_-, x_+\}$. So we have shown that there exists a θ such that $\xi(\theta) = \{x_-, x_+\}$. Since both x_- and x_+ are optimal for this type we must have $v(x_-, \theta) - \tau(x_-) = v(x_+, \theta) - \tau(x_+)$. Thus we have constructed a Φ-interval such that both endpoints lie in Γ but the interior points do not lie in Γ. But this would contradict the Φ-convexity of Γ, so it must be that the image of ξ is a connected interval.

Now consider the case where X is not necessarily compact. If ξ is not maximal monotone then there exists a point $(\theta, x) \in \Theta \times X$ such that $(\theta, x) \notin \xi$ yet (θ, x) is comparable to ξ in the product order on $\Theta \times X$; this means that for any $(\theta', x') \in \xi$ either $(\theta', x') \leq (\theta, x)$ or $(\theta', x') \geq (\theta, x)$. For then (θ, x) could be added to the correspondence ξ without destroying monotonicity. Let ϕ be the elementary function of type θ passing through $(x, \tau(x))$. If we can show

[10] This follows, for example, from the construction of 1 dimensional maximal monotone correspondences by filling in the gaps in the graph of a monotone function, as outlined previously.

that ϕ supports τ at x then we will have established by contradiction that ξ is maximal monotone.

Consider first the case where X has a maximal element, and x is this maximal element. We show first that τ is supported at x by an elementary function $\tilde{\phi}$ whose type we will denote $\tilde{\theta}$. Let $I = [x', x]$ be a non-degenerate closed interval in X with with upper bound x. We consider the problem restricted from X to I. Since I is compact we know by the paragraph above that τ is supported on I by an elementary function $\tilde{\phi}(x)$, whose type will be $\tilde{\theta}$. We need to check that $\tilde{\phi}$ supports τ on the whole of X, not just on I. Assume not. Then $\tilde{\phi}(x) = \tau(x)$, $\tilde{\phi}(z) \leq \tau(z)$ for $z \in I$, but $\tilde{\phi}(x'') > \tau(x'')$ for some $x'' \in X \diagdown I$. Consider Φ-interval formed by the graph of $\tilde{\phi}$ restricted to the interval $[x'', x]$. We have just shown that the endpoints of the interval lie in $\operatorname{epi} \tau$, so by internally Φ-convexity the whole Φ-interval lies in $\operatorname{epi} \tau$. In particular, $\tilde{\phi}(z) \geq \tau(z)$ for $z \in I$. But we have already shown the opposite inequality, so we conclude that $\tilde{\phi}(z) = \tau(z)$ for $z \in I$. We now consider the larger interval $J = [x'', x]$, and a point $\hat{x}$ in the interior of I. Appealing once again to the result that has been proven on compact intervals, there exists an elementary function $\hat{\phi}$ of type $\hat{\theta}$ that supports τ at $\hat{x}$ on J. But τ is equal to the elementary function $\tilde{\phi}$ on a neighbourhood of $\hat{x}$, so by strict single crossing $\hat{\phi} = \tilde{\phi}$ and $\hat{\theta} = \tilde{\theta}$. In particular, this means that $\tilde{\phi}$ must support τ on the whole of $J = [x'', x]$. But this contradicts the assumption that $\tilde{\phi}(x'') > \tau(x'')$. So we have proven that there is an elementary function of type $\tilde{\theta}$ that supports τ at x. That is, $\left(\tilde{\theta}, x\right) \in \xi$. Since by assumption $(\theta, x) \geq \xi$ it follows that $\theta \geq \tilde{\theta}$, so ϕ is steeper than $\tilde{\phi}$, and $\tilde{\phi}(x) = \phi(x)$. Thus ϕ supports τ at x. The case where X has a minimal element is similar.

We now turn to the general case. If $(\theta, x) \leq \xi$ or $(\theta, x) \geq \xi$ then by truncating X at x we are back to the case where the result has been proven. So we can decompose $\xi = \xi_1 \cup \xi_2$ into two disjoint, nonempty pieces such that $\xi_1 \leq (\theta, x) \leq \xi_2$. We split X into two corresponding pieces $X_1 = \{z \in X : z \leq x\}$ and $X_2 = \{z \in X : z \geq x\}$ as well. Arguing as above, there is an elementary function ϕ_1 of type θ_1 supporting τ at x on X_1, and an elementary function ϕ_2 of type θ_2 supporting τ at x on X_2. By monotonicity $\theta_1 \leq \theta_2$, so in fact ϕ_1 and ϕ_2 support τ on the whole of X. Thus $(\theta_1, x), (\theta_2, x) \in \xi_1, \xi_2$ respectively and $(\theta_1, x) \leq (\theta, x) \leq (\theta_2, x)$, so ϕ is sandwiched between ϕ_1 and ϕ_2. Thus ϕ supports τ at x on X. This contradiction establishes the result. ■

The following condition will be required in characterizing envelope functions.

Definition 1 (growth condition) *A function $f(x)$ satisfies the growth condition if, for all $x < x'$ such that at least one of $f(x)$, $f(x')$ is finite*

$$\inf_{\theta} (v(x', \theta) - v(x, \theta)) \leq f(x') - f(x) \leq \sup_{\theta} (v(x', \theta) - v(x, \theta))$$

Note that, under supermodularity, the elementary functions in Φ satisfy 1, and that this condition is preserved when passing to upper envelopes, as are both internal and external Φ-convexity.

Theorem 2 *Assume that $v(x, \theta)$ is continuous and strictly supermodular. The following are equivalent*

1. *$\tau(x)$ is a Φ-envelope function. That is, it is externally Φ-convex;*

2. *$\tau(x)$ is proper, lower semicontinuous, internally Φ-convex and satisfies the growth condition 1;*

3. *$\tau(x)$ is proper, internally Φ-convex and satisfies the growth condition 1, and the correspondence ξ is maximal monotone.*

Proof. Condition (1) implies (2) because the elementary functions have these properties, and they are preserved in passing to the envelope by taking the intersection of epigraphs. Condition (2) implies (3) by Theorem 1. It remains to be shown that condition (3) implies (1).

Let $\tilde{\tau}(x) = \tau^{**}(x)$ be the biconjugate of $\tau(x)$. We must show that $\tau = \tilde{\tau}$. From Lemma 2, $\tau(x) = \tilde{\tau}(x)$ if and only if $x \in \operatorname{im}\xi$, so we must show that $\operatorname{im}\xi = X$. Since ξ is maximal monotone $\operatorname{im}\xi$ is a connected interval, so if there exists a point such that $\tilde{\tau}(x) \neq \tau(x)$ then it must lie either above of below the interval $\operatorname{im}\xi$.

We consider the first alternative, and assume that there exists $x' \in X$ such that $x' > \operatorname{im}\xi$ and $\tilde{\tau}(x') < \tau(x')$. Choose t' such that $\tilde{\tau}(x') < t' < \tau(x')$ and consider the contract (x', t'). No agent will accept this contract since any elementary function supporting the contract must pass through or below $(x', \tilde{\tau}(x'))$. Thus, for all $\theta \in \Theta$ and $x \in \xi(\theta)$, $v(x,\theta) - \tau(x) > v(x',\theta) - t'$. This implies that $\tau(x) + v(x',\theta) - v(x,\theta)$ is uniformly bounded above by t', and this implies, by the growth condition, that $\tau(x') < \infty$. The growth condition applies here since $x \in \operatorname{im}\xi$, so $\tau(x)$ is finite). So we can write $t' = \tau(x') - \delta$. The incentive compatibility condition can then be written

$$\begin{aligned}
\tau(x') - \tau(x) &> \delta + v(x',\theta) - v(x,\theta) \\
&\geq \delta + \sup_{\theta} v(x',\theta) - v(x,\theta) \\
&> \sup_{\theta} v(x',\theta) - v(x,\theta)
\end{aligned}$$

which contradicts the growth condition. A similar argument applies if $x > \operatorname{im}\xi$ and $\tau(x) < \infty$. ■

The classical Fenchel duality theorem is equivalent to a separating hyperplane theorem of convex analysis (Rockafellar 1970). In a similar manner, Theorem 2 can be reformulated as a non-convex separation theorem.

Corollary 1 *Assume that $\tau(z)$ is proper, lower semicontinuous, internally Φ-convex and satisfies the growth condition 1, and let $t \leq \tau(x) < \infty$. Then there exists an elementary function $\phi \in \Phi$ such that $t \leq \phi(x)$and, for all z, $\phi(z) \leq \tau(z)$.*

3.2 The structure of envelope functions.

In the classical convex case, the envelope function $\tilde{\tau}$, the envelope of the affine functions supporting τ, is constructed by closing and convexifying the epigraph of τ. This construction is described above in Section 2.1. Theorem 2 shows that a completely analogous construction applies to the Φ-envelope. This construction is valuable because it makes very explicit what is lost or gained in passing from a function to its biconjugate. It makes clear that the arguments set out in Section 2.2 that in most contexts we can assume without loss of generality that the principal will offer an envelope contract apply essentially without change in the general case. The following geometrical Lemma, which generalizes a simple classical property of convex functions, will be used.

Lemma 3 (chord lemma) *Assume that $v(x,\theta)$ is continuous and strictly supermodular, and that $\tau(x)$ is internally Φ-convex and satisfies the growth condition 1. Consider points $x' \in \xi(\theta')$, $x'' \in \xi(\theta'')$ with $x' < x''$.*

1. *There exists a Φ-interval I, the graph of an elementary function $f(z)$ restricted to $[x', x'']$, joining $(x', \tau(x'))$ and $(x'', \tau(x''))$ in $\operatorname{epi}\tau$.*

2. *$f(z) \geq \tau(z)$ if $x' \leq z \leq x''$, and $f(z) \leq \tau(z)$ if $z \leq x'$ or $z \geq x''$*

Proof. To check the existence of such an interval, consider the elementary function $\phi(z,\theta) = \tau(x') + v(z,\theta) - v(x',\theta)$. For any θ we have $\phi(x',\theta) = \tau(x')$. The quantity $\phi(x'',\theta)$ is continuous in θ. By incentive compatibility $\phi(x'',\theta') \geq \tau(x'')$ and $\phi(x'',\theta'') \leq \tau(x'')$, so for some intermediate θ with $\theta' \leq \theta \leq \theta''$ we must have $\phi(x'',\theta) = \tau(x'')$. Choosing this θ, we set $f(z) = \phi(z,\theta)$. This proves 1 and the first part of 2.

To prove the second part of 2, consider the case $z > x''$. If $z \in \xi(\theta''')$ for some θ''' then we argue as follows. First of all, we extend the interval I constructed above by enlarging its domain of definition to $[x', z]$. By part 1, there exists a Φ-interval J joining $(x', \tau(x'))$ and $(z, \tau(z))$ in $\operatorname{epi}\tau$. The intervals I and J cross at $(x', \tau(x'))$. J is above I at x'' since J lies in $\operatorname{epi}\tau$ over the whole interval $[x', z]$ and I crosses the graph of τ at x''. If $f(z) > \tau(z)$ then J would be below I at z, contradicting single crossing.

If $v(z,\theta)$ satisfied the strict condition $\min_\theta (v(z,\theta) - v(x',\theta)) \leq \tau(z) - \tau(x') \leq \max_\theta (v(z,\theta) - v(x',\theta))$ — this is just the growth condition over the interval $[x', z]$ with inf and sup replaced with min and max — then the argument is similar. We construct the interval J just as in part 1, and then argue as before. In the general case we can modify this construction slightly, replacing $(z, \tau(z))$ by $(z, \tau(z) + \varepsilon)$ for some small ε, and the argument still goes through. ■

The construction of $\tilde{\tau}$ from τ can now be described in two steps that are exactly analogous to the classical construction described in Section 2.1.

1. We add limit points so that τ becomes lower semicontinuous and ξ has a closed graph. At each limit point x of $\operatorname{dom}\tau$ we modify τ by setting $\tau(x) = \tilde{\tau}(x)$ and modify ξ by setting $\xi^{-1}(x) = \tilde{\xi}^{-1}(x)$.

2. We then fill gaps in $\operatorname{dom}\tau$ so that τ becomes Φ-convex and ξ becomes maximal monotone. Let $J = (a, b)$ be a maximal open connected interval in $\operatorname{dom}\tilde{\tau} \setminus \operatorname{dom}\tau$.If $a, b \in \operatorname{dom}\tau$ then the elementary function through $(a, \tau(a))$ and $(b, \tau(b))$ constructed in the Chord Lemma supports $\operatorname{epi}\tau$ at these points. We modify τ by replacing it with this Φ-convex interpolation between these points. Let θ be the type of the agent attached to this interpolation. We modify ξ by replacing $\xi(\theta)$ with $\tilde{\xi}(\theta)$. If $b \notin \operatorname{dom}\tilde{\tau}$, so $\tilde{\tau}(b) = \infty$, let $b' = \sup\{z : \tilde{\tau}(z) < \infty\}$. Then $\tilde{\tau}(z) = \infty$ for $b' < z \leq b$. If $a \neq b'$ then let $b'' \in (a, b')$, and let $f(z)$ be the elementary function of type θ defining the chord through $(a, \tau(a))$ and $(b'', \tilde{\tau}(b''))$. Then, by single crossing, $f(z)$ and θ do not depend on b''. Thus we set $\tau(z)$ equal to the elementary function $f(z)$ on (a, b'), and $\tau(z) = \infty$ on (b', b). $\tau(b')$ is determined by continuity. We modify ξ by replacing $\xi(\theta)$ with the vertical section $\tilde{\xi}(\theta)$. The construction is similar if $\tilde{\tau}(a) = \infty$.

3.3 Regularity properties, and an envelope theorem

We can now use Φ-convexity arguments to establish some important regularity properties of envelope functions. These generalize standard properties of convex functions (see, for example, Rockafellar (1970, 13.3.3)).

Lemma 4 *Let $\tau(x)$ be a Φ-envelope function. If $v(x,\theta)$ is locally Lipshitz in x for each θ then $\tau(x)$ is locally Lipshitz , and hence continuous and absolutely continuous, at interior points of its domain.*

Proof. We note first that, by Lemma 3, for any points $x' \in \xi(\theta')$, $x'' \in \xi(\theta'')$ with $x' < x''$ there exists a Φ-interval I, the graph of an elementary function $f(z)$ restricted to $[x', x'']$, joining $(x', \tau(x'))$ and $(x'', \tau(x''))$ in $\operatorname{epi}\tau$.

Now let $x'' \in \xi(\theta'')$ be any point in $\operatorname{dom}\tau$. We will show that τ is locally Lipshitz on $[\underline{x}, x'']$. Let $x' < x''$, with $x' \in \xi(\theta')$. And consider the interval I described by the elementary function $f(z)$, as constructed in the previous paragraph. We will now extend this interval, by extending the domain of the elementary function to the whole of $[\underline{x}, x'']$. By Φ-convexity, $f(z) \geq \tau(z)$ on $[x', x'']$. Furthermore, $f(z)$ is associated with a type $\theta > \theta'$ so to the left of x' it lies below the elementary function supporting $\tau(z)$ at x'. Thus $f(z) \leq \tau(z)$ on $[\underline{x}, x']$. So, by monotonicity, $\tau(z)$ is sandwiched between $f(z)$ and the constant function $\tau(x')$ on $[\underline{x}, x'']$. As these are locally Lipschitz, $\tau(z)$ is locally Lipshitz at x'. ■

Theorem 3 (Supermodular Envelope Theorem) *Let*

$$\rho(\theta) = \sup_{x \in X} v(x,\theta) - \tau(x)$$
$$\xi(\theta) = \operatorname*{argmax}_{x \in X} v(x,\theta) - \tau(x) .$$

Assume that $v(x,\theta)$ is continuous, differentiable in θ, and strictly supermodular. Then $\rho(\theta)$ is differentiable almost everywhere, $\rho'(\theta) = v_2(x(\theta), \theta)$ almost everywhere, and

$$\rho(\theta) = \rho(\theta_0) + \int_{\theta_0}^{\theta} v_\theta(x(\theta), \theta)\, d\theta$$

where $x(\theta)$ is a selection from $\xi(\theta)$.

Proof. Since ρ is absolutely continuous it is differentiable almost everywhere and $\rho(\theta) = \rho(\theta_0) + \int_{\theta_0}^{\theta} \rho'(\theta)\, d\theta$. Let θ be a point of differentiability of ρ. Since ρ is an envelope function, it is supported at θ by an elementary function $\phi(\theta) = v(x,\theta) - t$ for some $x \in \xi(\theta)$. Since ϕ supports ρ at θ, and both functions are differentiable at this point, it follows that $\rho'(\theta) = \phi'(\theta) = v_\theta(x,\theta)$. ■

This result is related to, but different from, the Envelope Theorem of Milgrom and Segal (2002) who study the parametric optimisation problem

$$\rho(\theta) = \sup_{x \in X} f(x,\theta) .$$

Milgrom and Segal assume that $f(z,\theta)$ is absolutely continuous in θ, hence almost everywhere differentiable, and require the side condition that the derivative $f_\theta(z,\theta)$ be uniformly integrable, something that is not necessarily easy to

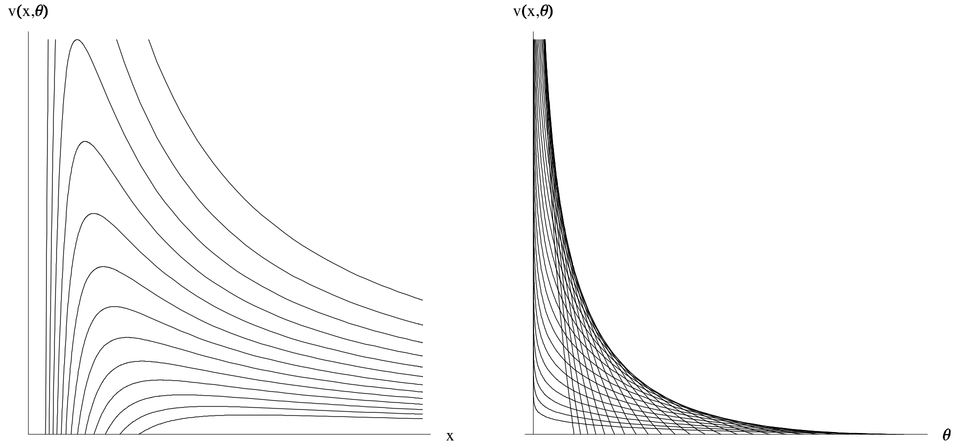

Figure 2: Envelope Example: $v(x,\theta) = -\frac{\log\theta}{x} - \frac{1}{x^2}$.

check if the type space or the action space is not compact.[11] They show that the value function $\rho(\theta)$ is absolutely continuous and derive an envelope relationship between $\rho(\theta)$ and the parametric family $f(z,\theta)$.

In the quasilinear environment (which of course is a more restrictive environment than that studied by Milgrom and Segal) things are somewhat simpler. Instead of their side condition it is necessary only to check single crossing, a condition that is natural in applications. We also get a stronger result: $\rho(\theta)$ is not just absolutley continuous but locally Lipshitz. Consider the example

$$\rho(\theta) = \sup_x -\frac{\log\theta}{x} - \frac{1}{x^2}$$

with $X = \Theta = (0,1]$, which satisfies the conditions of Theorem 3 (see Figure 2). In this case the cross derivative is $f_\theta(z,\theta) = -\frac{1}{\theta x}$, which is not uniformly integrable unless both the type space and the action space are restricted to compact intervals away from 0.

4 Quasilinear Agency

We now return to the quasilinear contracting Problem 1. As discussed above in 2.2, in the vast majority of applications it is without loss of generality to assume that the principal offers an envelope contract. The following theorem collects together results that have already been proved.

Theorem 4 *Assume that $v(x,\theta)$ is continuously differentiable and strictly su-*

[11] Milgrom and Segal argue in a footnote to their paper that the uniform integrability condition will hold in a quasilinear environment provided that the type space and the action space are compact intervals.

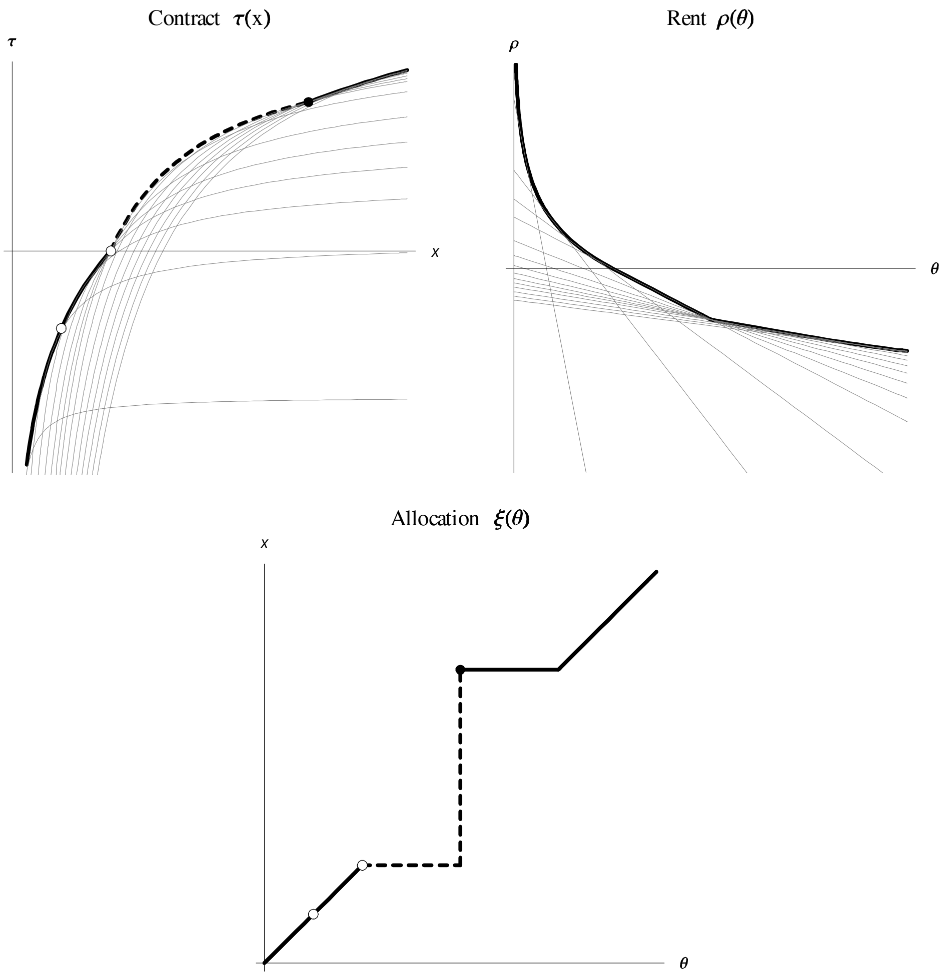

Figure 3: Nonlinear duality

permodular, and that τ is an envelope function. Let

$$\rho(\theta) = \sup_{x \in X} v(\theta, x) - \tau(x)$$
$$\xi(\theta) = \operatorname*{argmax}_{x \in X} v(\theta, x) - \tau(x)$$
$$\eta(x) = \operatorname*{argmax}_{\theta \in \Theta} v(\theta, x) - \rho(\theta).$$

Let $x(\theta) \in \xi(\theta)$ and $\theta(x) \in \eta(x)$ be selections from the respective correspondences. Then

1. *Both τ and ρ are Φ-envelope functions; in particular, they are both absolutely continuous and differentiable almost everywhere. At points of differentiability $\rho'(\theta) = v_2(x(\theta), \theta)$ and $\tau'(x) = v_1(x, \theta(x))$.*

2. *The correspondences ξ and η are maximal monotone.*

3. *The functions τ and ρ are generalised Fenchel conjugates: $\rho = \tau^*$ and $\tau = \rho^*$. In particular, $\tau = \tau^{**} = \tilde{\tau}$ and $\rho = \rho^{**} = \tilde{\rho}$.*

4. *The generalised Fenchel inequality holds: $\tau(x) + \rho(\theta) \geq \theta x$, with equality iff $\theta \in \eta(x)$ iff $x \in \xi(\theta)$.*

5. *ξ and η are inverse correspondences: $x \in \xi(\theta)$ if and only if $\theta \in \eta(x)$;*

6. *We have integral representations*

$$\rho(\theta) = \rho(\theta_0) + \int_{\theta_0}^{\theta} v_2(x(\theta), \theta)\, d\theta$$
$$\tau(x) = \tau(x_0) + \int_{x_0}^{x} v_1(x, \theta(x))\, dx.$$

As above, the interpretation is that $\tau(x)$ is the contract function, $\rho(\theta)$ is the information rent function, $\xi(\theta)$ is the set of choices over which an agent of type θ can optimally mix, and $\eta(x)$ is the set of types who can optimally choose x. The Lipshitz continuity condition is certainly satisfied if $v(x, \theta)$ is continuous and piecewise differentiable.

If the contract $\tau(x)$ is not an envelope function, then we can compare it with its biconjugate $\tilde{\tau}(x)$, which is an envelope function. $\tilde{\tau}(x)$ is constructed from $\tau(x)$ by Φ-convexifying and taking the closure of $\operatorname{epi}\tau$. Φ-convexification means adding to $\operatorname{epi}\tau$ any points in an interval in an agent's indifference curve if the endpoints of that interval lie in $\operatorname{epi}\tau$ — this generalizes the usual notion of convexity. Both τ and $\tilde{\tau}$ give rise to the same information rent: $\tau^* = \tilde{\tau}^* = \rho$, and the duality between $\tilde{\tau}$ and ρ is as described in Theorem 4 above. The contract τ can be constructed by deleting points from $\tilde{\tau}$ in such a way that $\tilde{\tau}$ remains the closed convexification of τ.

All in all, the Fenchel duality properties that are immediate in the linear case carry over with virtually no change to the quasilinear case. The main difference is that in the linear case the correspondences ξ and η have a geometrical interpretation as the subgradient correspondences of ρ and τ respectively: $\xi = \partial\rho$ and $\eta = \partial\tau$. For non-convex functions the subgradient is no longer an appropriate concept to describe the locally convex behavior of functions. There are several non-smooth generalizations of the subgradient appropriate to locally Lipshitz functions, for example the proximal and Clarke subgradients ((Clarke 1990), (Clarke, Ledyaev, Stern, and Wolenski 1998), (Rockafellar and Wets 1998)). It seems likely that a geometric interpretation of ξ and η can be formulated using these concepts. A second difference is that in the linear case the value function $v(x, \theta)$ is, by definition, linear and hence smooth. Since the techniques used in this paper apply to a wide range of non-smooth contexts, it seems likely that the assumption that the value function be differentiable could be weakened to allow, for example, kinks in the valuation.

5 Application: Menus of Simple Contracts

It is common in applications to assume that the principal offers a menu of linear contracts, but this is a very restrictive assumption (Rogerson 1987). Consider for example a linear agency model with a smooth, strictly convex contract

$$\rho(\theta) = \sup_{x \in X} v(\theta, x) - \tau(x)$$

and a linear contract $l(x)$ that supports $\tau(x)$ at x_0. That is, $l(x_0) = \tau(x_0)$ and $l'(x_0) = \tau'(x_0) = \theta_0$. Then every point on the graph of $l(x)$, apart from $(x_0, l(x_0))$ lies strictly outside the contract set $\operatorname{epi}\tau$. If this contract is offered as part of a menu then any agent of type $\theta > \theta_0$ will misrepresent their type as θ_0 and select $l(x)$.

Let us define a simple contract as one that is continuous and piecewise linear with a single kink: $\tau(x) = \begin{cases} \tau_0 + \alpha(x - x_0), & x \leq x_0 \\ \tau_0 + \beta(x - x_0), & x \geq x_0. \end{cases}$ That is to say, τ is characterized by a target x_0, a linear penalty for falling short of x_0, and a linear reward for exceeding x_0.

Theorem 5 *Assume that $v(x,\theta)$ is continuous, continuously differentiable in x, and strictly supermodular. Let ξ be an allocation that is supported by an incentive compatible contract τ. Assume that the agent's information rent $\rho(\theta)$ is non-negative, and that the total value of the contract $v(x,\theta)$ is bounded. Then ξ is supported by a menu of simple contracts.*

Proof. Without loss of generality one can assume that τ is an envelope contract, since any simple contract that supports $\tilde{\tau}$ will support τ. By the Fenchel inequality $\tau(x) + \rho(\theta) \leq v(x,\theta)$ so $\tau(x)$ is bounded above. But τ is convex and lower semicontinuous, so $\operatorname{dom}\tau = [a,b]$ is a closed interval. Let $x \in [a,b]$. It is sufficient to show that τ is supported from above by a globally Lipshitz function at x. Consider types $\theta_a \in \eta(a)$ and $\theta_b \in \eta(b)$. By the proof of Lemma 4 there exist elementary functions $\phi_1(z)$ and $\phi_2(z)$ such that $\phi_1(x) = \phi_2(x) = \tau(x)$, $\phi_1(z) \geq \tau(z)$ for $a \leq z \leq x$, and $\phi_2(z) \geq \tau(z)$ for $x \leq z \leq b$. That is, τ is supported from above at x by $\phi = \max[\phi_1, \phi_2]$. But ϕ_1 and ϕ_2 are Lipshitz since they are continuously differentiable on $[a,b]$, and the max of two Lipshitz functions is Lipshitz. ■

It is worth noting that this result is different in nature to properties established in, for example, the elementary part of Theorem 2. Envelope functions are by definition constructed by approximation from below, and properties like lower semicontinuity and Φ-convexity that are stable under approximation from below are more or less automatic. Theorem 5, which relates to approximation from above, is more delicate.

The hypotheses of Theorem 5 require that the information rent be bounded below (that is to say, the agent has an outside option) and that the total value created for the agent by the transaction be bounded uniformly in type. Since the choice space X is compact, the force of this assumption is not that the value be bounded, but that the bound be independent of type.[12]

6 Conclusion

If the agent's valuation is linear, then the standard adverse selection contracting problem is isomorphic to the producer's problem in elementary price theory,

[12]If the choice space were unbounded then the appropriate condition would be that the average value $\frac{v(x,\theta)}{x}$ be bounded uniformly in type as $x \to \infty$.

with price replaced by the agent's type. Standard duality results of price theory, which are most conveniently expressed using the Fenchel transform, apply more or less directly. In particular, not only the information rent but also the contract can be expressed as an integral of marginal valuations. If the environment is quasilinear, and the agent's valuation is strictly supermodular, then an abstract convexity and an abstract Fenchel transform can be developed which allow these results to be generalised almost in their entirety to the quasilinear contracting problem. The supermodular quasilinear contracting problem is one of the canonical models of microeconomics, with applications ranging from optimal regulation to auction theory.

This framework can be used to establish an envelope theorem and a complete duality theory between the information rent and the contract. Supermodularity plays a key role in two respects. Firstly, it creates a link between incentive theory (the Spence-Mirrlees single crossing property) and an abstract convexity that is rich enough to capture both inner and outer concepts of convexity and the relations between them. This structure allows standard techniques of convex optimisation to be applied in this generalised setting. Secondly, it implies that the allocation correspondence is maximal monotone. Maximal monotone correspondences have strong properties that drive the results.

The framework developed in this paper has potential for applications that go beyond those presented here. In this paper we focus on the duality structure inherent in the agent's decision problem, and draw out some implications for the structure of implementable contracts. There is no discussion of optimal contracts. It seems likely that there are useful things to be said in this direction, building in particular on implications of maximal monotonicity that go beyond the techniques used here. The other interesting direction is multidimensional type spaces. While the details of particular arguments in this paper draw heavily on the one dimensionality assumption, the general framework of Fenchel duality on which we draw is inherently multidimensional. It is natural to conjecture that there may be useful generalisations to this setting.

References

Berge, C. (1963): *Topological Spaces*. Oliver and Boyd, London.

Border, K. (1991): "Implementation of Reduced Form Auctions: a Geometric Approach," *Econometrica*, 59, pp. 1175–1187.

——— (2007): "Reduced Form Auctions Revisited," *Economic Theory*, 31, pp. 167–181.

Clarke, F. H. (1990): *Optimization and nonsmooth analysis*, vol. 5 of *Classics in Applied Mathematics*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, second edn.

Clarke, F. H., Y. S. Ledyaev, R. J. Stern, and P. R. Wolenski (1998): *Nonsmooth analysis and control theory*, vol. 178 of *Graduate Texts in Mathematics*. Springer-Verlag, New York.

Fudenberg, D., and J. Tirole (1991): "Game theory. 1991," .

Guesnerie, R., and J. Laffont (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(3), 329–369.

Jackson, M. O. (2001): "A Crash Course in Implementation Theory.," *Social Choice and Welfare*, 18(4), 655 – 708.

Jehle, G. A., and P. J. Reny (2000): *Advanced Microeconomic Theory.* Addison Wesley, second edn.

Jullien, B. (2000): "Participation Constraints in Adverse Selection Models* 1," *Journal of Economic Theory*, 93(1), 1–47.

Krishna, V., and E. Maenner (2001): "Convex Potentials with an Application to Mechanism Design," *Econometrica*, 69(4), pp. 1113–1119.

Milgrom, P., and I. Segal (2002): "Envelope theorems for arbitrary choice sets," *Econometrica*, 70(2), 583–601.

Pallaschke, D., and S. Rolewicz (1997): *Foundations of mathematical optimization*, vol. 388 of *Mathematics and its Applications*. Kluwer Academic Publishers Group, Dordrecht, Convex analysis without linearity.

Rochet, J. C. (1987): "A Necessary and Sufficient Condition for Rationalizability in a Quasilinear Context," *Econometrica*, 16, pp. 191–200.

Rockafellar, R. T. (1970): *Convex analysis*, Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J.

Rockafellar, R. T., and R. J.-B. Wets (1998): *Variational analysis*, vol. 317 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin.

Rogerson, W. (1987): "On the Optimality of Menus of Linear Contracts," *Northwesten Discussion Paper*, 714R.

Rubinov, A. (2000): *Abstract convexity and global optimization*, vol. 44 of *Nonconvex Optimization and its Applications*. Kluwer Academic Publishers, Dordrecht.

Vohra, R. (2011): *Mechanism Design a Linear Programming Approach.* Cambridge University Press, New Yoprk.