Dynamic Implementation*

David Rahman
University of Minnesota

Preliminary and Incomplete

January 16, 2009

Abstract

Consider a strategic environment subject to moral hazard and adverse selection across multiple stages, with rich communication protocols. In this paper, we prove that for any allocation, there exist linear transfers to make it incentive compatible if and only if every undetectable deviation from honesty and obedience is unprofitable when the transfers equal zero, where 'undetectable' means that the distributions of actual and reported types coincide. The set of transfers that implement a given implementable allocation is also characterized. These results extend Rochet's (1987) characterization of implementability to a dynamic context. The paper also characterizes optimal allocations, profit-maximizing mechanisms, virtual implementation, implementation subject to dynamic budget balance, and dynamic revenue equivalence.

JEL Classification: D82, D83, D86.
Keywords: dynamic mechanism design, multistage games with communication, duality, mediated transfers, private monitoring.

*I owe many thanks to Narayana Kocherlakota, Chris Phelan and Itai Sher for helpful discussions.

Contents

1 Introduction

Providing dynamic incentives is a crucial topic in economics, with applications ranging from income taxation and the provision of public goods to auctions and partnerships. Recently, several papers have emerged (Athey and Segal, 2007; Bergemann and Välimäki, 2007; Gershkov and Moldovanu, 2007; Pai and Vohra, 2008; Pavan et al., 2008) with the common goal of understanding aspects of dynamic mechanism design. The present paper contributes to this literature by characterizing implementable allocations in a dynamic quasi-linear environment with both moral hazard and adverse selection, in the spirit of Myerson's (1986) multistage games with communication.

To intuitively describe this characterization, Theorem 1, fix an arbitrary allocation. Call a deviation any dynamic strategy that possibly involves different behavior from that stipulated by the allocation. A deviation is called undetectable if the distribution over actual types equals the distribution of reported types given the deviation. Theorem 1 says that a given allocation is implementable, i.e., there is a schedule of linear transfers that makes it incentive compatible, if and only if every undetectable deviation is unprofitable when the transfers are restricted to equal zero. Therefore, in order to verify whether or not an allocation is implementable, we do not need to search for a transfer scheme such that incentive compatibility will be satisfied with respect to every deviation. Instead, we only need to check whether or not incentive compatibility holds with respect to the subset of undetectable deviations after fixing transfers to equal zero. If (and only if) this latter condition holds, then there exists a possibly non-zero scheme that implements the allocation in question. When one exists, Theorem 2 characterizes all such transfer schemes.

This characterization extends Rochet's Theorem to a dynamic context, thus delivering a dynamic version of cyclic monotonicity. It also generalizes results by Cremer and McLean (1988). In a static context without moral hazard, they found a condition that guarantees implementability of every allocation, regardless of preferences. In the present framework, their condition means that there are no undetectable deviations, so every allocation is implementable by Theorem 1. Extending Cremer and McLean's logic to a dynamic context requires dealing with the possibility that types depend on past allocations, an issue that clearly does not arise in a static environment. Hence, whether or not a deviation is detectable can crucially depend on the allocation in question. Naturally, such a problem suggests considering virtual implementation as a solution concept. These issues are discussed at length in Section 4.

The main results of this paper are based on the static model of Rahman (2008c). Regarding recent literature, the work of Athey and Segal (2007) and Bergemann and Välimäki (2007) fits naturally into our model. Assuming independent private values, they extend the VCG mechanism to implement a dynamically efficient allocation. By Theorem 1, it is easy to see that an efficient allocation is implementable. Indeed, if values are private then an agent's undetectable deviation cannot affect anyone else's expected utility. Therefore, it will be profitable if and only if it is welfare-improving, contradicting efficiency.[1] As for other relevant papers, Gershkov and Moldovanu (2007) characterize implementable (deterministic, Markovian) allocations in a restricted setting, also with independent private values. Pai and Vohra (2008) study revenue-maximizing dynamic auctions in a restricted environment, and relate weak—rather than cyclic—monotonicity to implementability, even though types are multi-dimensional. Finally, Pavan et al. (2008) study dynamic mechanism design but fail to provide simultaneously necessary and sufficient conditions for implementation. In addition, their results rely on revenue equivalence; the results in this paper do not. Corollary 1 characterizes dynamic revenue equivalence à la Heydenreich et al. (2008).

Theorem 1 is extended below in several ways, some technical (infinite type spaces, infinite horizon) and some with more economic content (virtual implementation, budget balance, limited liability, individual rationality, and optimal allocations). However, Theorem 1 is viewed here as a "straw man" to which these results add "meat."

Specifically, adding infinite type spaces and an infinite horizon does not change the main message of Theorem 1, save technicalities. Heuristically, these guarantee that approximately undetectable deviations are approximately unprofitable. Virtual implementation adds a nuance to Theorem 1, but the main message prevails that verifying (virtual) implementability boils down to checking for detectable deviations. Budget balance shifts the notion of detectability to minimally distinguishing between those who deviated and those who did not. This case is explored in some depth, by characterizing allocations that are implementable subject to history-dependent budgets. This result leads to a "golden rule" for budgeting incentives, i.e., balancing the decision of saving money to provide future incentives with spending it now to encourage current effort. Limited liability and individual rationality also introduce manageable complications; they are useful because they allow a characterization of profit-maximizing mechanisms. Finally, the study of optimal mechanisms leads to a Bellman equation that characterizes optimality in a general dynamic context.

[1] This proof also holds for allocations that maximize any non-negatively-weighted sum of utilities.

2 Model

Let $I = \{1, \ldots, n\}$ be a finite set of agents who interact over finitely many stages, $T \in \mathbb{N}$. At every stage $t \in \{1, \ldots, T\}$, agent i is asked by some mediator to make a choice $c_{it} \in C_{it}$ and decides $d_{it} \in C_{it}$ which in principle could be different from c_{it}, then observes some signal privately, $s_{it} \in S_{it}$, and finally submits a report $r_{it} \in S_{it}$ to the mediator which could differ from s_{it}. Let $C_{it} \times S_{it}$ be the (nonempty) finite[2] set of all signal-choice pairs, with typical elements (c_{it}, r_{it}), (d_{it}, s_{it}), etc. For now, assume that everyone is honest and obedient, so $d_{it} = c_{it}$ and $r_{it} = s_{it}$ for all (i, t).

An *information structure* is any sequence of maps $\Pr = \{\Pr_t : C^t \times S^{t-1} \to \Delta(S_t)\}$ indexed by t.[3] At every stage t, $\Pr_t$ describes the probability that a given profile of signals is observed by the agents conditional on previous choices and observations. We will usually just write $\Pr(r_t|c^t, r^{t-1})$ and drop the time subscript on Pr. An *allocation* is a sequence $\mu = \{\mu_t : C^{t-1} \times S^{t-1} \to \Delta(C_t)\}$ describing the probability with which the mediator makes recommendations to players conditional on previous recommendations and reports. Every pair $(\Pr, \mu)$ induces a stochastic process over partial histories. For any partial history (c^t, r^t), denote its induced probability by

$$\Pr(c^t, r^t|\mu) \;=\; \prod_{\tau=1}^{t} \mu_\tau(c_\tau|c^{\tau-1}, r^{\tau-1}) \Pr(r_\tau|c^\tau, r^{\tau-1}). \tag{2.1}$$

An *incentive scheme* is any sequence of functions $\{\zeta_t : I \times C^t \times S^t \to \mathbb{R}\}$ indexed by t, where $\zeta_{it}(c^t, r^t)$ denotes the money paid by agent i in stage t if the mediator's partial history of recommendations and reports is (c^t, r^t).[4] A *mechanism* is any pair (μ, ζ) where μ is an allocation and ζ is an incentive scheme.

Utility functions are defined as follows. Let $v_{it}(c^t, r^t) \in \mathbb{R}$ be a "utility flow" accrued by agent i after partial history (c^t, r^t). The expected utility to agent i from a given mechanism (μ, ζ) equals

$$U_i(\mu, \zeta) \;=\; \sum_{(t, c^t, r^t)} \Pr(c^t, r^t|\mu)[v_{it}(c^t, r^t) - \zeta_{it}(c^t, r^t)].$$

Let us now define incentive compatibility, which requires more notation. Fix an agent i, assume everyone else is honest and obedient, and consider i's incentive to deviate.

[2]All the finiteness assumptions above can be relaxed—see Section 6 below.

[3]For any family $\{X_{it}\}$ of sets indexed by i and t, let $X_{i0} = \{\emptyset\}$, $X_t = \prod_i X_{it}$, $X_{-it} = \prod_{j \neq i} X_{jt}$, $X_i^t = \prod_{\tau=1}^t X_{i\tau}$ and $X^t = \prod_i X_i^t$. Thus, S^t is the space of t-stage histories of signal profiles.

[4]We could just as easily restrict payments to occur at the end, but this formulation allows for richer restrictions on budgets, such as dynamic budget constraints—see Section 5 below.

A partial history for the mediator is a tuple (c^t, r^t) of recommendations made to agents, c^t, and reports made by agents, r^t. Agent i's partial history is a tuple $(c_i^t, d_i^t, s_i^t, r_i^t)$ of recommendations, actual decisions, observed signals and reports. Let

$$\Pr(c^t, s_i^t, r_{-i}^t | d_i^t, r_i^t, \mu) \;=\; \prod_{\tau=1}^{t} \mu_\tau(c_\tau | c^{\tau-1}, r^{\tau-1}) \Pr(s_{i\tau}, r_{-i\tau} | d_i^\tau, c_{-i}^\tau, s_i^{\tau-1}, r_{-i}^{\tau-1}) \quad (2.2)$$

be the probability that recommendations c^t and signals (s_i^t, r_{-i}^t) are observed if agent i makes decisions d_i^t, submits reports r_i^t, and everyone else is honest and obedient.

A (behavior) *strategy* for agent i is a sequence $\sigma_i = \{(\delta_{it}, \rho_{it})\}$ indexed by t such that $\delta_{it} : C_i^t \times C_i^{t-1} \times S_i^{t-1} \times S_i^{t-1} \to \Delta(C_{it})$ and $\rho_{it} : C_i^t \times C_i^t \times S_i^t \times S_i^{t-1} \to \Delta(S_{it})$. Intuitively, a strategy consists of two parts: a plan to make a decision, δ_{it}, and a plan to submit a report, ρ_{it}. These plans are allowed to depend on all available information, namely the history of previous reports and decisions as well as the history of previous signals observed and recommendations received. An important example of a strategy is honesty and obedience: $\delta_{it}(d_{it} | r_i^t, c_i^t, s_i^t, d_i^{t-1}) = 1$ if and only if $d_{it} = c_{it}$ and $\rho_{it}(r_{it} | r_i^{t-1}, c_i^{t-1}, s_i^t, d_i^{t-1}) = 1$ if and only if $r_{it} = s_{it}$. Let θ_i denote this strategy. A strategy σ_i is called a *deviation plan* (or simply *deviation*) if $\sigma_i \neq \theta_i$.

For any given strategy σ_i, define the following conditional probability:

$$\sigma_{it}(d_i^t, r_i^t | c_i^t, s_i^t) \;=\; \prod_{\tau=1}^{t} \delta_{i\tau}(d_{i\tau} | c_i^\tau, d_i^{\tau-1}, s_i^{\tau-1}, r_i^{\tau-1}) \rho_{i\tau}(r_{i\tau} | c_i^\tau, d_i^\tau, s_i^\tau, r_i^{\tau-1}). \quad (2.3)$$

The sequence $\{\sigma_{it}\}$ captures all relevant information about σ_i. Each σ_{it} is clearly a conditional probability. It describes the probability that i makes decisions d_i^t and submits reports r_i^t conditional on having observed recommendations c_i^t and signals s_i^t.

Given a mechanism (μ, ζ), denote the expected utility from a strategy σ_i by

$$U_i(\sigma_i | \mu, \zeta) \;=\; \sum_{(t, c^t, d_i^t, s_i^t, r^t)} \sigma_{it}(d_i^t, r_i^t | c_i^t, s_i^t) \Pr(c^t, s_i^t, r_{-i}^t | d_i^t, r_i^t, \mu)[v_{it}(d_i^t, c_{-i}^t, s_i^t, r_{-i}^t) - \zeta_{it}(c^t, r^t)].$$

Definition 1. Given any mechanism (μ, ζ), a strategy σ_i is (μ, ζ)-*unprofitable* if

$$U_i(\sigma_i | \mu, \zeta) \;\leq\; U_i(\mu, \zeta).$$

Call (μ, ζ) *incentive compatible* if every deviation is (μ, ζ)-unprofitable. An allocation μ is *implementable* if (μ, ζ) is incentive compatible for some incentive scheme ζ.

For simplicity, Definition 1 ignores incentives "off the equilibrium path" and asks for $(\theta_1, \ldots, \theta_n)$ to be a Nash equilibrium given (μ, ζ). Alternatively, we could have made the generic assumption that Pr has full support for this restriction to incur no loss.

3 Results

We will now present our two main results. The first characterizes implementable allocations just in terms of the information structure and preferences. The second constructs an incentive scheme that implements any implementable allocation.

Definition 2. A deviation σ_i is called supp μ-*undetectable*[5] if

$$\Pr(c^t, r^t|\mu) \;=\; \sum_{(d_i^t, s_i^t)} \sigma_{it}(d_i^t, r_i^t|c_i^t, s_i^t)\Pr(c^t, s_i^t, r_{-i}^t|d_i^t, r_i^t, \mu) \qquad \forall (t, c^t, r^t). \tag{3.1}$$

Intuitively, σ_i is supp μ-undetectable if the probability distribution over the mediator's partial histories generated by σ_{it} coincides with what it would have been had the agent chosen to always behave honestly and obediently, i.e., had he followed θ_i instead of σ_i. Indeed, the left-hand side of (3.1) is the probability that (c^t, r^t) is the profile of observed recommendations and reports given that everyone is honest and obedient, whereas the right-hand side is the probability that c^t is recommended and r^t is reported if agent i employs the deviation σ_i. To illustrate, if $t = 1$ and $|C_1| = 1$ (so there is only adverse selection in the first stage) then (3.1) yields

$$\Pr(r^1) \;=\; \sum_{s_i^1} \sigma_{i1}(r_i^1|s_i^1)\Pr(s_i^1, r_{-i}^1) \qquad \forall r^1 \in S^1.$$

Let $\Pr(r^t|c^t) = \prod_{\tau=1}^{t} \Pr(r^\tau|c^\tau, r^{\tau-1})$ for all (t, c^t, r^t). By substituting (2.1) and (2.2) into (3.1), it is easy to see that (3.1) simplifies to

$$\Pr(r^t|c^t) \;=\; \sum_{(d_i^t, s_i^t)} \sigma_{it}(d_i^t, r_i^t|c_i^t, s_i^t)\Pr(s_i^t, r_{-i}^t|d_i^t, c_{-i}^t) \qquad \forall (t, c^t, r^t) \text{ s.t. } \Pr(c^t, r^t|\mu) > 0.$$

This condition only depends on μ via supp μ, hence the term "supp μ-undetectable."

Finally, it is worth emphasizing that detectability is only defined statistically, i.e., the probability distribution over outcomes given a strategy differs from that induced by honest and obedient behavior, rather than necessarily the outcomes themselves. Section 4 compares detectability with the literature on mechanism design.

Theorem 1. *An allocation μ is implementable if and only if every* supp *μ-undetectable deviation is $(\mu, \mathbf{0})$-unprofitable ($\mathbf{0}$ is the incentive scheme that always equals zero).*

[5] By definition, supp $\mu = \{(c^t, r^t) : \Pr(c^t, r^t|\mu) > 0\}$ is the set of partial histories with positive probability under μ.

Theorem 1 is the main result of this paper. Intuitively, it says that implementability is equivalent to honesty and obedience being optimal in a hypothetical problem with fewer strategies than the original problem and no transfers. Hence, to check for implementability, instead of verifying that for some incentive scheme every deviation is unprofitable, it is sufficient (and necessary) to just check that, for the incentive scheme that is identically equal to zero, every undetectable deviation is unprofitable.

Our next result is to derive, for any implementable allocation, the set of incentive schemes that implement it. We begin with preliminary definitions and notation.

For any convex function $f : \mathbb{R}^m \to \mathbb{R}$, the *subdifferential* of f at $x \in \mathbb{R}^m$ equals the set $\partial f(x) = \{p \in \mathbb{R}^m : p \cdot (y - x) \leq f(y) - f(x) \ \forall y \in \mathbb{R}^m\}$. Let $\mathscr{D}_i$ be the set of vectors $\lambda_i \geq 0$ that are proportional to a strategy σ_i, i.e., there exists $q \in \mathbb{R}_+$ such that $\lambda_i = q\sigma_i$. For any partial history (c^t, r^t), let us write

$$\lambda_{it} \cdot \Delta \Pr(c^t, r^t|\mu) \ = \sum_{(d_i^t, s_i^t)} \lambda_{it}(d_i^t, r_i^t|c_i^t, s_i^t)[\Pr(c^t, s_i^t, r_{-i}^t|d_i^t, r_i^t, \mu) - \Pr(c^t, r^t|\mu)]$$

for (an amount proportional to) the difference between the probability of actual types and reported types. Consider the following convex function:

$$F_i(\mathbf{z}^{\pm}|\mu) = \max_{\lambda_i \in \mathscr{D}_i} \{U_i(\lambda_i|\mu, \mathbf{0}) : -\mathbf{z}_t^-(c^t, r^t) \leq \lambda_{it} \cdot \Delta \Pr(c^t, r^t|\mu) \leq \mathbf{z}^+(c^t, r^t) \ \forall (t, c^t, r^t)\}.$$

$F_i(\mathbf{z}^{\pm}|\mu)$ is proportional to the maximum expected utility with respect to strategies for which the change in the probability of reported types from the strategy relative to honesty and obedience is bounded $\mathbf{z}^{\pm}$. Clearly, by revealed preference F_i is a convex function of $\mathbf{z}^{\pm}$, and by Theorem 1, if μ is implementable then $F_i(\mathbf{0}|\mu) = 0$.

Theorem 2. *Suppose that μ is an implementable allocation. A given incentive scheme $\zeta = (\zeta_1, \ldots, \zeta_n)$ implements μ if and only if for every agent i, there exists $\zeta_i^{\pm} \in \partial F_i(\mathbf{0}|\mu)$ such that $\zeta_i = \zeta_i^+ - \zeta_i^-$.*

By Theorem 2, any incentive scheme that implements a given allocation is a subgradient of some suitably chosen function. Revenue equivalence is now an easy corollary. Intuitively, it obtains if a related function has a unique subgradient.

Corollary 1. *An allocation μ exhibits dynamic revenue equivalence (i.e., any two schemes that implement μ differ by a constant) if and only if the function*

$$G_i(\mathbf{z}^{\pm}) = \max_{\lambda_i \in \mathscr{D}_{i,\varepsilon}} \{U_i(\lambda_i|\mu, \mathbf{0}) : -\mathbf{z}_t^-(c^t, r^t) \leq \lambda_{it} \cdot \Delta \Pr(c^t, r^t|\mu) + \varepsilon \mathbf{1}\{(\hat{c}^{\hat{t}}, \hat{r}^{\hat{t}})\} \leq \mathbf{z}^+(c^t, r^t)\}$$

is differentiable at $\mathbf{0}$ for each agent i, where $(\hat{c}^{\hat{t}}, \hat{r}^{\hat{t}})$ is some fixed partial history.[6]

[6]This revenue equivalence is *not* just in expectation, unlike Heydenreich et al. (2008).

4 Discussion

In this section, we discuss the relationship between Theorem 1 and static mechanism design, corollaries to dynamic mechanism design, and virtual implementation. We will specify the model to involve just adverse selection. Following most of the literature, and to draw better comparisons with it, the results focus on deterministic allocations, although relaxing them to allow for random ones is a trivial exercise.

4.1 Relation to Static Mechanism Design

Consider a two-stage problem with an agent (agent 1) and a principal (agent 2), where the agent has private information in the first stage and the principal commits to a contingent choice in the second stage: $n = T = 2$, $|C_{11}| = |C_{12}| = |C_{21}| = 1$, $|S_{21}| = |S_{22}| = |S_{12}| = 1$ and v_2 is a constant function.[7] Without loss of generality, let us drop all subscripts. Fix a deterministic allocation $\mathbf{x} : S \to C$. By Theorem 1, $\mathbf{x}$ is implementable if and only if every supp $\mathbf{x}$-undetectable deviation is $(\mathbf{x}, \mathbf{0})$-unprofitable, i.e., a given deviation σ satisfies $\Pr(r) = \sum_s \sigma(r|s) \Pr(s)$ for every r only if $\sum_{(r,s)} \sigma(r|s) \Pr(s) v(\mathbf{x}(r), s) \leq \sum_s \Pr(s) v(\mathbf{x}(s), s)$.

Rochet's Theorem states that an allocation $\mathbf{x}$ is implementable if and only if it is cyclically monotone.[8] Therefore, Theorem 1 generalizes cyclic monotonicity to a dynamic context.[9] The key difference is that the set of undetectable deviations (which in the static context would correspond to cycles) now depends on the allocation.

Cremer and McLean's (1988) Theorem is also a special case of Theorem 1. They find a condition that characterizes implementability of every allocation in a static model. Indeed, with all the assumptions above except that now there are several agents, a deviation plan is supp $\mathbf{x}$-undetectable if $\Pr(r) = \sum_{s_i} \sigma_i(r_i|s_i) \Pr(s_i, r_{-i})$ for every agent i and signal profile r. Once again, this condition does not depend on $\mathbf{x}$, so we may talk about detectability without reference to an allocation. Cremer and McLean's condition may be interpreted as saying that there are no undetectable deviations. Therefore, every undetectable deviation is vacuously unprofitable.

[7] Making v_2 constant implies that the principal choices are not subject to incentive constraints.

[8] An allocation $\mathbf{x}$ is *cyclically monotone* if $\sum_{k=1}^{m} [v(\mathbf{x}(s_k), s_{k+1}) - v(\mathbf{x}(s_k), s_k)] \leq 0$ for every finite cycle $(s_1, \ldots, s_m, s_{m+1})$ such that $s_{m+1} = s_1$.

[9] See Rahman (2008a) for a detailed comparison of Theorem 1 with Rochet's Theorem, as well as a direct proof that cyclic monotonicity is equivalent to every deviation being $(\mathbf{x}, \mathbf{0})$-unprofitable.

How does Theorem 1 add to these static results? In this dynamic context, there are two ways in which agents' types may be correlated: across agents and over time. Cremer and McLean (1988) showed how to implement allocations when types are correlated across agents in a static environment. On the other hand, a setting with perfect serial correlation over time is just a static problem, since each agent knows all of his future types by learning his current type, and may be treated by applying Rochet's Theorem to each agent. Hence, Theorem 1 adds value by accounting for noisy serial correlation. In addition, Theorem 1 accommodates the possibility that future types are affected not just by past types, but also by past decisions.

4.2 Dynamic Adverse Selection

Consider a principal-agent problem with dynamic adverse selection, i.e., $n = 2$, $|C_{1t}| = |C_{21}| = 1$, $|S_{2t}| = 1$ and v_{2t} is a constant function for every t, so we may drop subscripts denoting individuals. Agent 1 is the "agent" with private information over time and agent 2 is the "principal" who can commit to taking actions contingent on the agent's reports. A deterministic allocation $\mathbf{x}$ is any sequence $\{\mathbf{x}_t : S^t \to C^{t+1}\}$ of maps. As a matter of notation, let $\mathbf{x}^t(r^t) = (\mathbf{x}_1(r^1), \ldots, \mathbf{x}_t(r^t))$.

Corollary 2. *Fix any principal-agent problem with dynamic adverse selection. A deterministic allocation* $\mathbf{x}$ *is implementable if and only if for any deviation* σ*,*

$$\Pr(r^t|\mathbf{x}^t(r^t)) \;=\; \sum_{s^t} \sigma_t(r^t|s^t) \Pr(s^t|\mathbf{x}^t(r^t)) \qquad \forall (t, r^t)$$

implies that

$$\sum_{(t,s^t,r^t)} \sigma_t(r^t|s^t) \Pr(s^t|\mathbf{x}^t(r^t)) v_t(\mathbf{x}^t(r^t), s^t) \;\le\; \sum_{(t,r^t)} \Pr(r^t|\mathbf{x}^t(r^t)) v_t(\mathbf{x}^t(r^t), r^t),$$

where we use the notation $\Pr(s^t|\mathbf{x}^t(r^t)) = \prod_{\tau=1}^{t} \Pr(s_\tau|\mathbf{x}^{\tau-1}(r^{\tau-1}), s^{\tau-1})$.

Corollary 2 is the dynamic generalization of Rochet's Theorem to environments with pure adverse selection.

Next, let us extend Corollary 2 to include several agents. Label the principal as player 0 and suppose that all others are agents. A deterministic allocation is still a sequence of maps $\mathbf{x} = \{\mathbf{x}_t : S^t \to C_{0t+1}\}$.

Corollary 3. *Fix a dynamic adverse selection problem with several agents. A deterministic allocation* $\mathbf{x}$ *is implementable if and only if for every* i *and* σ_i*,*

$$\Pr(r^t|\mathbf{x}^t(r^t)) \;=\; \sum_{s_i^t} \sigma_t(r_i^t|s_i^t) \Pr(s_i^t, r_{-i}^t|\mathbf{x}^t(r^t)) \qquad \forall (t, r^t)$$

(*i.e.,* σ_i *is* supp $\mathbf{x}$*-undetectable*) *implies that*

$$\sum_{(t,s_i^t,r^t)} \sigma_{it}(r_i^t|s_i^t) \Pr(s_i^t, r_{-i}^t|\mathbf{x}^t(r^t)) v_{it}(s_i^t, r_{-i}^t, \mathbf{x}^t(r^t)) \;\leq\; \sum_{(t,r^t)} \Pr(r^t|\mathbf{x}^t(r^t)) v_{it}(r^t, \mathbf{x}^t(r^t)).$$

This result extends Cremer and McLean's Theorem to a dynamic environment. The key difference in a dynamic setting is that "full surplus extraction" does not follow immediately from Corollary 3. Indeed, Cremer and McLean's argument may be paraphrased as follows. Suppose that $T = 1$. For any allocation $\mathbf{x}$, a deviation plan σ_i is supp $\mathbf{x}$-undetectable if $\Pr(r^1) = \sum_{s_i^1} \sigma_{i1}(r_i^1|s_i^1) \Pr(s_i^1, r_{-i}^1)$ for every r^1. Crucially, notice that this condition does not depend on the allocation at all. Therefore, a deviation plan is supp $\mathbf{x}$-undetectable if and only if it is supp $\mathbf{x}'$-undetectable for any two allocations $\mathbf{x}$ and $\mathbf{x}'$. So it is meaningful to describe deviation plans as simply being undetectable without reference to an allocation. Cremer and McLean (1988) show that every deviation is detectable if and only if every allocation is implementable, regardless of agents' utility functions. In particular, a surplus-extracting allocation is always implementable in this case.

Such logic no longer extends to a dynamic environment because the relevant notion of dynamic detectability now does depend on the allocation, and therefore, one cannot apply the logic of Cremer and McLean (1988) unless one makes the further restriction that the probability of types does not depend on the principal's choices. This argument is summarized below.

Corollary 4. *Fix a dynamic adverse selection problem, and suppose that the distribution over signals doesn't depend on the principal's choices, i.e.,* $\Pr_t : S^t \to \Delta(S_{t+1})$ *for every* t*. Every allocation is implementable regardless of agents' utility functions if and only if every deviation is detectable, i.e., for every* i *and* σ_i*,*

$$\Pr(r^t) \;=\; \sum_{s_i^t} \sigma_{it}(r_i^t|s_i^t) \Pr(s_i^t, r_{-i}^t) \qquad \forall (t, r^t)$$

implies that $\sigma_{it}(r_i^t|s_i^t) = 1$ *if* $r_i^t = s_i^t$ *and* 0 *otherwise, where we are using the notation* $\Pr(r^t) = \prod_{\tau=1}^t \Pr(r_\tau|r^{\tau-1})$ *for all* (t, r^t).

We can also characterize implementable allocations under the assumption that types do not depend on the principal's choices. An allocation $\mathbf{x}$ is implementable if and only if every undetectable deviation is $\mathbf{x}$-unprofitable, where the notion of detectability is independent of the allocation. Clearly, the notion of profitability unavoidably isn't.

We end by extending Cremer and McLean's logic when types depend on the principal's choices. Given a subset B of partial histories, a B-*deviation* is any deviation σ_i that is dishonest or disobedient with positive probability at some partial history in B.

Theorem 3. *An allocation $\mathbf{x}$ is implementable regardless of agents' utility functions if and only if each* supp $\mathbf{x}$*-deviation is* supp $\mathbf{x}$*-detectable, i.e., not* supp $\mathbf{x}$*-undetectable.*

4.3 Virtual Implementation

The previous subsection suggested that deviations may be detectable with respect to some allocations but not others. This begs the following question: can a lottery over allocations increase the set of detectable deviations? Applying a key result in Rahman (2008c), we now consider this possibility with virtual implementation.

Definition 3. An allocation μ is *virtually implementable* if there exists a sequence $\{\mu^m\}$ of implementable allocations such that $\mu^m \to \mu$. Given a partial history (c^t, r^t), say that σ_i is $\{(c^t, r^t)\}$-*undetectable* if

$$\Pr(r^t|c^t) \;=\; \sum_{(d_i^t, s_i^t)} \sigma_{it}(d_i^t, r_i^t|c_i^t, s_i^t)\,\Pr(s_i^t, r_{-i}^t|d_i^t, c_{-i}^t).$$

Call σ_i *undetectable* if it is $\{(c^t, r^t)\}$-undetectable at every partial history (c^t, r^t). Otherwise, call σ_i *detectable.*

Theorem 4. *An allocation μ is virtually implementable regardless of agents' utility functions if and only if every* supp μ*-deviation is detectable.*

The crux of this theorem is that behavior outside of supp μ may be required to detect a supp μ-deviation. However, deviations from this detecting behavior need not be detectable, and thus Theorem 4 strictly generalizes Theorem 3.

We end by remarking that it is also possible to provide necessary and sufficient conditions for virtual implementation given a fixed profile of utility functions, based on a result in Rahman (2008c, Theorem 3)—the details are available on request.

5 Applications

5.1 Dynamically Optimal Mechanisms

5.1.1 A Characterization of Dynamically Optimal Allocations

Consider the problem of finding an optimal allocation subject to being implementable. We begin by providing sufficient conditions for an implementable allocation to maximize the value of a given function, followed by necessary and sufficient ones. The sufficient conditions have the advantage of being relatively easier to verify.

Let $f = \{f_t : C^t \times S^t \to \mathbb{R}\}$ be a sequence of functions indexed by t, and consider the following optimization problem.

$$\sup_{\mu} \sum_{(t,c^t,r^t)} f_t(c^t, r^t) \Pr(c^t, r^t|\mu) \quad \text{s.t. } \mu \text{ is an implementable allocation.} \tag{5.1}$$

A difficulty with this optimization is that the set of implementable allocations need not be closed, so the sup above may not be attained. For an instance of this difficulty, see Rahman (2008c, Example 1). However, the sup will be attained by an allocation if it satisfies certain properties implied by duality. If μ^* is an optimal solution then call it *f-optimal*, and let $F(\mu^*) = \sum_{(t,c^t,r^t)} f_t(c^t, r^t) \Pr(c^t, r^t|\mu^*)$.

It will be useful to introduce additional notation. Given $(i, t, c^t, r^t, d_i^t, s_i^t)$, we will denote by $w_{it}(d_i^t, c_{-i}^t, s_i^t, r_{-i}^t|c_i^t, r_i^t) = v_{it}(d_i^t, c_{-i}^t, s_i^t, r_{-i}^t)\mathrm{Lr}(d_i^t, s_i^t|c^t, r^t)$ the likelihood-weighted utility to any agent i from (d_i^t, s_i^t) relative to (c^t, r^t), where

$$\mathrm{Lr}(d_i^t, s_i^t|c^t, r^t) = \begin{cases} \Pr(s_i^t, r_{-i}^t|d_i^t, c_{-i}^t)/\Pr(r^t|c^t) & \text{if } \Pr(r^t|c^t) > 0 \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

is a kind of *likelihood ratio* between $(d_i^t, c_{-i}^t, s_i^t, r_{-i}^t)$ and (c^t, r^t).

Let $\mathscr{D}$ be the set of all vectors $\lambda \geq 0$ that are proportional to a deviation profile, i.e., there exists a number $q \in \mathbb{R}_+$ and a deviation profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ such that $\lambda = q\sigma$. Let $\mathscr{U} \subset \mathscr{D}$ be the subset of vectors that are proportional to an undetectable deviation profile, i.e., a profile σ of deviations such that each σ_i is $\{(c^t, r^t)\}$-undetectable (see Definition 3) for all (i, t, c^t, r^t). Given $\lambda \in \mathscr{D}$, denote by $\lambda_t \cdot \Delta w_t$ be the function defined pointwise by

$$\lambda_t \cdot \Delta w_t(c^t, r^t) = \sum_{(i,d_i^t,s_i^t)} \lambda_{it}(d_i^t, r_i^t|c_i^t, s_i^t)[w_{it}(d_i^t, c_{-i}^t, s_i^t, r_{-i}^t|c_i^t, r_i^t) - w_{it}(c^t, r^t|c_i^t, r_i^t)].$$

Intuitively, $\lambda_t \cdot \Delta w_t(c^t, r^t)$ is proportional to the sum across agents of the change in likelihood-weighted utility from unilaterally deviating according to the deviation profile to which λ is proportional.

Theorem 5. *For any $\lambda \in \mathscr{U}$, let $\{J_t\}$ be any solution to the following Bellman equation, defined pointwise for all (t, c^{t-1}, r^{t-1}) by $J_{T+1}(c^T, r^T) = 0$ and*

$$J_t(c^{t-1}, r^{t-1}) \;=\; \max_{c_t} \sum_{r_t} [J_{t+1}(c^t, r^t) + f_t(c^t, r^t) - \lambda_t \cdot \Delta w_t(c^t, r^t)] \Pr(r_t | c^t, r^{t-1}).$$

An implementable allocation μ^* is f-optimal if

$$F(\mu^*) \;=\; \min_{\lambda \in \mathscr{U}} J_1.$$

For an example where this Bellman equation is unable to characterize optimal implementable allocations (because the condition above is not necessary), see Rahman (2008c, Example 4). When Theorem 5 does not apply, either because there is no optimal implementable allocation or the dynamic programming problem above fails to characterize optimality, we must resort to a different approach. Write

$$\left\|\lambda_t \cdot \Delta \mathrm{Lr}(c^t, r^t)\right\| \;= \sum_{i \in I} \Big| \sum_{(d_i^t, s_i^t)} \lambda_{it}(d_i^t, r_i^t | c_i^t, s_i^t)[\mathrm{Lr}(d_i^t, s_i^t | c^t, r^t) - \mathrm{Lr}(c_i^t, r_i^t | c^t, c^t)] \Big|.$$

Theorem 6. *For any $\lambda \in \mathscr{D}$ and $z \in \mathbb{R}_+$, let $\{J_{zt}\}$ be any solution to the following Bellman equation, defined pointwise for all (t, c^{t-1}, r^{t-1}) by $J_{zT+1}(c^T, r^T) = 0$ and*

$$\begin{aligned} J_{zt}(c^{t-1}, r^{t-1}) \;=\; \max_{c_t} \sum_{r_t} [&J_{zt+1}(c^t, r^t) + f_t(c^t, r^t) - \lambda_t \cdot \Delta w_t(c^t, r^t) \\ &+ z \left\|\lambda_t \cdot \Delta \mathrm{Lr}(c^t, r^t)\right\|] \Pr(r_t | c^t, r^{t-1}). \end{aligned}$$

A virtually implementable allocation μ^* is f-optimal if and only if

$$F(\mu^*) \;=\; \sup_{z \geq 0} \min_{\lambda \in \mathscr{D}} J_{z1}.$$

If an f-optimal implementable allocation exists then this family of Bellman equations indexed by z still characterizes the value of optimal implementable allocations. Hence, by virtue of automatically being virtually implementable, if μ^* is implementable then it is f-optimal if and only if $F(\mu^*) = \sup_{z \geq 0} \min_{\lambda \in \mathscr{D}} J_{z1}$.

5.1.2 Optimal Mechanisms and Revenue-Maximizing Auctions

We now consider the problem of finding an optimal mechanism subject to individual rationality. Assume that each agent has an outside option—that they may take at any time—to permanently exit a given mechanism. Let $\overline{v} = \{\overline{v}_{it} : C^t \times S^t \to \mathbb{R}\}$ be a sequence of contingent utility flows from each agent's outside option that determine its expected value. A mechanism (μ, ζ) is *individually rational with respect to* $\overline{v}$ if

$$\sum_{(c^\tau, r^\tau) \geq (c_i^t, r_i^t)} \Pr(c^\tau, r^\tau|\mu)[v_{i\tau}(c^\tau, r^\tau) - \zeta_{i\tau}(c^\tau, r^\tau) - \overline{v}_{i\tau}(c^\tau, r^\tau)] \ \geq \ 0 \qquad \forall (i, t, c_i^t, r_i^t).$$

This inequality says that the expected net present value of continuing in the mechanism is nonnegative for every agent after any partial history with positive probability.

Consider the following optimization problem:

$$\sup_{(\mu,\zeta)} \sum_{(t,c^t,r^t)} [\sum_{i \in I} \zeta_{it}(c^t, r^t) - g_t(c^t, r^t)] \Pr(c^t, r^t|\mu) \ \text{ s.t.} \tag{5.2}$$

(μ, ζ) is an incentive compatible mechanism, individually rational with respect to $\overline{v}$.

Let $G(\mu, \zeta) = \sum_{(t,c^t,r^t)}[\sum_i \zeta_{it}(c^t, r^t) - g_t(c^t, r^t)] \Pr(c^t, r^t|\mu)$ be the value of the above objective at any (μ, ζ). A mechanism (μ^*, ζ^*) is called $(g, \overline{v})$-*optimal* if it solves problem (5.2) above. Next, extending Theorem 5, we provide sufficient conditions for a mechanism to be $(g, \overline{v})$-optimal. We leave deriving a general characterization of optimality (i.e., necessary and sufficient conditions) to the reader on the grounds that it follows the same lines as Theorem 6 above together with Theorem 7 below.

We will need additional notation. Let $\Delta v_t(c^t, r^t|\overline{v}) = \sum_i v_{it}(c^t, r^t) - \overline{v}_{it}(c^t, r^t)$ be the sum across agents of the difference between their utility flow, $v_{it}(c^t, r^t)$, and the flow value of their outside option, $\overline{v}_{it}(c^t, r^t)$, at a given partial history (c^t, r^t). Let us write $\overline{w}_{it}(d_i^t, c_{-i}^t, s_i^t, r_{-i}^t|c_i^t, r_i^t) = \overline{v}_{it}(d_i^t, c_{-i}^t, s_i^t, r_{-i}^t)\mathrm{Lr}(d_i^t, s_i^t|c^t, r^t)$ and

$$\begin{aligned}\lambda_t \cdot \Delta\widehat{w}_t(c^t, r^t) \ &= \sum_{(i,d_i^t,s_i^t)} \lambda_{it}(d_i^t, r_i^t|c_i^t, s_i^t)[(w_{it}(d_i^t, c_{-i}^t, s_i^t, r_{-i}^t|c_i^t, r_i^t) - \overline{w}_{it}(c^t, r^t|c_i^t, r_i^t)) \\ &\qquad -(v_{it}(c^t, r^t) - \overline{v}_{it}(c^t, r^t))\mathrm{Lr}(d_i^t, s_i^t|c^t, r^t)]\end{aligned}$$

for the sum of differences in the value of each agent's inside option when playing the deviation to which λ is proportional versus behaving honest and obediently. Finally, let $\mathscr{R}$ be the set of $\lambda \geq 0$ proportional to a deviation profile and satisfying

$$\sum_{(d_i^t,s_i^t)} \lambda_{it}(d_i^t, r_i^t|c_i^t, s_i^t)[\mathrm{Lr}(d_i^t, s_i^t|c^t, r^t) - \mathrm{Lr}(c_i^t, r_i^t|c^t, c^t)] \ = \ \gamma_{it}(c_i^t, r_i^t) - 1 \qquad \forall (i, t, c^t, r^t),$$

where $\gamma_i \geq 0$ for every agent i.

Theorem 7. *For any $\lambda \in \mathscr{R}$, let $\{J_t\}$ be any solution to the following Bellman equation, defined pointwise for all (t, c^{t-1}, r^{t-1}) by $J_{T+1}(c^T, r^T) = 0$ and*

$$J_t(c^{t-1}, r^{t-1}) = \max_{c_t} \sum_{r_t} \Pr(r_t|c^t, r^{t-1})[J_{t+1}(c^t, r^t) - g_t(c^t, r^t) \\ + \Delta v_t(c^t, r^t|\overline{v}) - \lambda_t \cdot \Delta\widehat{w}_t(c^t, r^t)].$$

An implementable allocation μ^* is $(g, \overline{v})$-optimal if

$$G(\mu^*) = \min_{\lambda \in \mathscr{R}} J_1.$$

An immediate implication of Theorem 7 is the following sufficient condition for revenue-maximizing auctions, which naturally generalizes Myerson's (1981) results.

Consider the following *optimal dynamic auction* problem: a dynamic adverse selection environment with several agents (i.e., $|C_{it}| = 1$ for all t and $i \neq 0$, see Section 4.2) and $g \equiv 0$, so the principal's objective in (5.2) is to maximize revenue. The principal's choices are whom to allocate an object at every stage: $C_{0t} = X = \{0, 1, \ldots, n\}$, where 0 stands for nobody getting the object. Denote by $x^t = (x_1, \ldots, x_t) \in X^t$ the history of choices by the principal. Suppose that the value of everyone's outside option satisfies $\overline{v}_{it} \equiv 0$. These restrictions simplify Theorem 7 as follows.

Corollary 5. *For any $\lambda \in \mathscr{R}$, let $\{J_t\}$ be any solution to the following Bellman equation, defined pointwise for all (t, c^{t-1}, r^{t-1}) by $J_{T+1}(c^T, r^T) = 0$ and*

$$J_t(x^{t-1}, r^{t-1}) = \max_{x_t \in X} \sum_{r_t} \Pr(r_t|x^t, r^{t-1})[J_{t+1}(x^t, r^t) + \Delta v_t(x^t, r^t|\mathbf{0}) - \lambda_t \cdot \Delta\widehat{w}_t(c^t, r^t)].$$

A given implementable allocation is part of an optimal dynamic auction if the revenue it generates equals $\min_{\lambda \in \mathscr{R}} J_1$.

The key difference between this problem and its static version is that past decisions influence current ones. Thus, if $v_{it}(x^t, r^t) = v_{i\hat{t}(t)}(r_{i\hat{t}(t)})$ and $\hat{t}(t) = \min\{\tau \leq t : x_\tau = i\}$ (if $\{\tau \leq t : x_\tau = i\} = \emptyset$ then $v_{it} = 0$) then the problem becomes maximizing virtual welfare at every stage, dropping agents after each stage for a subproblem with fewer people to whom the good ought to be allocated.

— *To be completed.* —

5.2 Dynamic Budget Constraints

5.2.1 Budget Balance

Let us begin by imposing dynamic budget balance on transfers to obtain a similar characterization to that of Theorem 1. Using similar techniques to those in Rahman and Obara (2008), we will characterize budget balanced implementation as follows.

Definition 4. Given a mechanism (μ, ζ), say that ζ exhibits *budget balance* if

$$\sum_{i \in I} \zeta_{it}(c^t, r^t) \;=\; 0 \qquad \forall (t, c^t, r^t).$$

A deviation profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is supp μ-*unattributable* if

$$\sum_{(d_i^t, s_i^t)} \sigma_{it}(d_i^t, r_i^t | c_i^t, s_i^t) \Pr(c^t, s_i^t, r_{-i}^t | d_i^t, r_i^t, \mu) \;=\; \sum_{(d_j^t, s_j^t)} \sigma_{jt}(d_j^t, r_j^t | c_j^t, s_j^t) \Pr(c^t, s_j^t, r_{-j}^t | d_j^t, r_j^t, \mu)$$

for every pair (i, j) of agents and every (t, c^t, r^t). Finally, a deviation *profile* σ is called (μ, ζ)-*unprofitable* if the sum of payoffs across agents from each unilateral deviation plan σ_i is not positive, i.e.,

$$\sum_{i \in I} U_i(\sigma_i | \mu, \zeta) \;\leq\; \sum_{i \in I} U_i(\mu, \zeta).$$

Attribution is simply a weak requirement for distinguishing agents with respect to their behavior. A deviation profile σ is unattributable if the same probability distribution over reports is generated after a unilateral deviation in σ, regardless of the identity of the unilateral deviator. Therefore, it is not only impossible to identify a deviator, but also it is impossible to identify an obedient agent. Intuitively, a lack of attribution stifles budget balanced implementation because in order to provide budget-balanced incentives some agents must be rewarded while others are being punished. Intuitively, if those who ought to be rewarded cannot be distinguished from those who ought to be punished then budget-balanced incentives must fail. It turns out that attribution is the weakest such distinguishability condition that guarantees budget-balanced implementation, as the next result shows.

Theorem 8. *An allocation μ is implementable with budget-balanced linear transfers if and only if every* supp *μ-unattributable deviation profile is $(\mu, \mathbf{0})$-unprofitable.*

Now consider the problem of dynamic, possibly history-contingent budget constraints. For instance, the group's available budget may depend on the history of output.

To model this situation, assume that there is a zeroth agent who takes no actions but observes signals that help determine the budget. This zeroth agent is indifferent over everything and cannot be used as a budget-breaker, i.e., no payments can be made to him, and without loss always tells the truth. We assume that the budget does not depend on other agents' reports. This assumption is without loss of generality. Indeed, one might argue that at the heart of the problem of budget levels that are private information is that agents might misreport the amount of budget available in order to perhaps keep some of it. In this case, submitting the report that the budget is low would affect an agent's utility as much as the amount that the agent chose to keep. However, in this model we assume that reports are costless. On the other hand, this behavior can be modeled as an action, which is allowed to affect budgets.

A *budget* is any sequence $\{B_t : C^t \times S_0^t \to \mathbb{R}\}$ of maps indexed by t, where $B_t(c^t, r_0^t)$ is interpreted as the amount of budget available to provide incentives after every agent has chosen c^t and the zeroth agent's signal is r_0^t. Given a mechanism (μ, ζ), say that ζ *attains the budget* B if

$$\sum_{i \in I} \zeta_{it}(c^t, r^t) \;=\; B_t(c^t, r_0^t) \qquad \forall (t, c^t, r^t).$$

An allocation μ is *implementable with budget* B if there exists ζ that attains the budget B and with which (μ, ζ) is incentive compatible.

Finally, to characterize implementation with a budget, we need one more definition. Rewrite $U_i(\sigma_i|\mu, \zeta, v)$ instead of $U_i(\sigma_i|\mu, \zeta)$ to denote the dependence of U_i on the utility profile v. Define $\hat{v}_{it}^B(c^t, r^t) = v_{it}(c^t, r^t) - \frac{1}{n} B_t(c^t, r_0^t)$ for each (i, t, c^t, r^t). A deviation profile σ is called (μ, ζ, B)-*unprofitable* if

$$\sum_{i \in I} U_i(\sigma_i|\mu, \zeta, \hat{v}^B) \;\leq\; \sum_{i \in I} U_i(\mu, \zeta, \hat{v}^B).$$

Theorem 9. *An allocation μ is implementable with budget B if and only if every* supp *μ-unattributable deviation profile is $(\mu, \mathbf{0}, B)$-unprofitable.*

Intuitively, Theorem 9 describes the kind of budget that is preferable for dynamic incentive provision. For instance, it precisely formalizes the statement that providing incentives is easier when the group's budget is more likely to diminish after an individually desirable unilateral deviation. Indeed, given two budgets, B and C, if B is more likely to diminish after an unattributable deviation profile than C but otherwise they are the same, then the left-hand side of the inequality defining $(\mu, \mathbf{0}, B)$-unprofitability will be lower with B than with C, making it easier for the deviation profile to be $(\mu, \mathbf{0}, B)$-unprofitable than $(\mu, \mathbf{0}, C)$-unprofitable.

5.2.2 Budgeting Incentives

Consider the following profit maximization problem of *budget incentives*, where a principal dynamically allocates amounts of money in order to provide incentives to agents with limited liability. Specifically, the principal's problem is the following:

$$\min_{(\mu,\zeta)} \sum_{(t,c^t,r^t)} \zeta_{0t}(c^t,r^t)\Pr(c^t,r^t|\mu) \quad \text{s.t.} \tag{5.3}$$

(μ,ζ) is an incentive compatible mechanism, $\zeta \leq 0$, and

$$\sum_{\tau=1}^{t}[R_\tau(c^\tau,r^\tau) + \sum_{i\in I\cup\{0\}} \zeta_{i\tau}(c^\tau,r^\tau)] \;\geq\; 0 \qquad \forall(t,c^t,r^t).$$

The constraint $\zeta \leq 0$ ensures that agents enjoy limited liability, so that they never have to pay any money to the principal. We denote by $\zeta_0 \leq 0$ the amount of money paid by the principal. Hence, for simplicity, we assume that the principal wants to maximize the present value of expected money holdings and that he is not able to borrow, only save. The last family of constraints describes the principal's budget allocation problem. At every stage t, the principal obtains—perhaps as revenue from outside of this model—the amount $R_t(c^t,r^t) \geq 0$ of money. This amount, together with previous stages' amounts of unspent money, is available to the principal for either consumption, investment in current incentives, or investment in future incentives.

The decision problem faced by the principal involves not only how much to reward workers for generating revenue, but also when to do so. The dual of this problem reveals some insight into how much surplus that the principal may extract. The following notation will be useful. For any $\lambda \in \mathscr{D}$ and any partial history (c^t,r^t), write $\lambda_{it}\cdot\Delta\Pr(c^t,r^t|\mu) = \sum_{(d_i^t,s_i^t)}\lambda_{it}(d_i^t,r_i^t|c_i^t,s_i^t)[\Pr(c^t,s_i^t,r_{-i}^t|d_i^t,r_i^t,\mu) - \Pr(c^t,r^t|\mu)]$ for (an amount proportional to) the change in probability that the mediator observes (c^t,r^t) under the allocation μ and the deviation σ_i relative to honesty and obedience.

Theorem 10. *For any $\lambda \in \mathscr{D}$, let $\{J_t\}$ be any solution to the following Bellman equation, defined pointwise for all (t,c^{t-1},r^{t-1}) by $J_{T+1}(c^T,r^T) = 0$ and*

$$J_t(c^{t-1},r^{t-1}) \;=\; \max_{c_t}\sum_{r_t}\Pr(r_t|c^t,r^{t-1})[J_{t+1}(c^t,r^t) + \sum_{\tau=1}^{t}R_\tau(c^\tau,r^\tau) - \lambda_t\cdot\Delta\widehat{w}_t(c^t,r^t)].$$

The principal's maximum revenue from problem (5.3) above solves

$$\min_{\lambda\in\mathscr{D}}\{J_1 : \lambda_t\cdot\Delta\Pr(c^t,r^t|\mu) \leq 1 \;\forall(t,c^t,r^t)\}.$$

It would be interesting to explore this problem further in the future.

6 Extensions

Now we extend the model in three ways. Firstly, we consider infinitely many types. Secondly, we solve the model for $T = \infty$. Finally, we allow for a form of risk aversion.

6.1 Infinitely Many Types

The model can easily be extended to include infinitely many types, using results from Rahman (2008b, Theorem 4) and Rahman (2008a, Theorem 1) for static mechanisms.

The main complication introduced by having infinitely many types is continuity. Intuitively, a sequence of detectable deviations may be "asymptotically undetectable." For instance, suppose that the set of types is the interval $[0, 1]$ and that the actual type is 0. Conceivably, reporting $1/m$ may be detectable for every m yet the relative change in probabilities diminishes faster than the change in utilities, yielding a sequence of deviations that is "asymptotically undetectable" but "asymptotically profitable."

— *To be completed.* —

6.2 Infinite Horizon

Let us extend the model to the case where $T = \infty$. Basically, all the previous results continue to hold verbatim in this case after making some standard assumptions, by applying an infinite-dimensional version of the proof of Theorem 1. Below, we will make these assumptions and explicitly prove the extension of Theorem 1. Proofs of extensions of the other results are left to the reader and available on request.

— *To be completed.* —

6.3 Risk Aversion

Here, we will describe a situation where individuals have separable—but no longer necessarily linear—utility over transfers, using a standard trick due to Mirrlees.

— *To be completed.* —

7 Conclusion

This paper contributed to the recent literature on dynamic mechanism design by characterizing dynamic implementation in a multistage environment with communication and linear transfers by asking that every undetectable deviation be unprofitable.

Virtual dynamic implementation was also characterized, broadly also in terms of detectability. Furthermore, optimal dynamic mechanisms were characterized in terms of a Bellman equation. Although this equation is not recursive, any hope of finding a tractable recursive structure must be abandoned under dynamic private monitoring, like here. Conversely, imposing tractable recursiveness would incur a cost in terms of obtaining suboptimal mechanisms. It would be interesting in the future to understand how constrained-optimal mechanisms would perform relative to fully optimal ones.

A Proofs

Theorem 1. The proof proceeds in three steps. Firstly, Step 1 describes implementability of an allocation μ as a system of (finitely many) linear inequalities. Step 2 applies the Theorem of the Alternative to find an equivalent dual system of linear inequalities that characterizes existence of a solution to the original system. Finally, Step 3 shows that this alternative system is equivalent to every supp μ-undetectable deviation being $(\mu, \mathbf{0})$-unprofitable.

– *Step 1.* We begin by providing an equivalent description of incentive compatibility for a given mechanism (μ, ζ) in two parts. The first part defines the gains from one-step deviations after any partial history. The second part aggregates these deviation gains to impose dynamic incentive compatibility, which is intuitively expressed as requiring that after any partial history, as long as agents have behaved honestly and obediently hitherto, they must remain willing to behave honestly and obediently henceforth.

We need one additional piece of notation. For $t \leq \tau$, let $x^\tau[y^t] = (y_1, \ldots, y_t, x_{t+1}, \ldots, x_\tau)$. For instance, $c^\tau[d_i^t] = (d_{i1}, \ldots, d_{it}, c_{it+1}, \ldots, c_{i\tau}, c_{-i}^\tau)$.

We begin by defining two kinds of deviation gains. The first kind ensures obedience and the second honesty. For any agent i, stage t, and partial history $(c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1})$, write

$$V_{it}(c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1}) = \sum_{(c^\tau, r^\tau) \geq (c_i^t, r_i^{t-1})} \Pr(c^\tau, r^\tau[s_i^{t-1}] | c_i^\tau[d_i^t], r_i^\tau, \mu)[v_{i\tau}(c^\tau[d_i^t], r^\tau[s_i^{t-1}]) - \zeta_{i\tau}(c^\tau, r^\tau)]$$

$$- \Pr(c^\tau, r^\tau[s_i^{t-1}] | c_i^\tau[d_i^{t-1}], r_i^\tau, \mu)[v_{i\tau}(c^\tau[d_i^{t-1}], r^\tau[s_i^{t-1}]) - \zeta_{i\tau}(c^\tau, r^\tau)].$$

The quantity $V_{it}(c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1})$ denotes the change in utility from a one-step deviation at a partial history where the mediator recommended to and was told by agent i the profile (c_i^t, r_i^{t-1}), i actually played and observed (d_i^{t-1}, s_i^{t-1}), and considers a deviation in stage t from c_{it} to d_{it}. For any agent i, stage t, and partial history $(c_i^t, d_i^t, s_i^t, r_i^t)$, write

$$W_{it}(c_i^t, d_i^t, s_i^t, r_i^t) \;=\; \sum_{(c^\tau, r^\tau) \geq (c_i^t, r_i^t)} \Pr(c^\tau, r^\tau[s_i^t] | c_i^\tau[d_i^t], r_i^\tau, \mu)[v_{i\tau}(c^\tau[d_i^t], r^\tau[s_i^t]) - \zeta_{i\tau}(c^\tau, r^\tau)]$$

$$- \Pr(c^\tau, r^\tau[s_i^t] | c_i^\tau[d_i^t], r_i^\tau[s_i^t[r_i^{t-1}]], \mu)[v_{i\tau}(c^\tau[d_i^t], r^\tau[s_i^t]) - \zeta_{i\tau}(c^\tau, r^\tau[s_i^t[r_i^{t-1}]])].$$

The quantity $W_{it}(c_i^t, d_i^t, s_i^t, r_i^t)$ describes the gain from a one-step deviation at a partial history $(c_i^t, d_i^t, s_i^t, r_i^{t-1})$, where agent i lies by reporting r_{it} instead of s_{it}.

After either kind of one-step deviation, either the disobedience defining V or the dishonesty defining W, agent i is assumed to be subsequently honest and obedient. Now, we can describe the gains from any dynamic deviation as the aggregate value from one-step deviation gains, which leads to the following equivalent description of incentive compatibility.

A mechanism (μ, ζ) is *incentive compatible* if and only if

$$\sum_{c_{i1}} V_{i1}(c_{i1}, \delta_{i1}(c_{i1})) + \sum_{s_{i1}} W_{i1}(c_{i1}, \delta_{i1}(c_{i1}), s_{i1}, \rho_{i1}(s_{i1})) + \cdots +$$
$$\sum_{c_{iT}} V_{iT}(c_i^T, \delta_{i1}(c_{i1}), \ldots, \delta_{iT}(c_{iT}), s_i^{T-1}, \rho_{i1}(s_{i1}), \ldots, \rho_{iT-1}(s_{iT-1})) +$$
$$\sum_{s_{iT}} W_{iT}(c_i^T, \delta_{i1}(c_{i1}), \ldots, \delta_{iT}(c_{iT}), s_i^T, \rho_{i1}(s_{i1}), \ldots, \rho_{iT}(s_{iT})) \;\leq\; 0$$

for every agent i and every tuple (δ_i, ρ_i) such that $\delta_{it} : C_{it} \to C_{it}$ and $\rho_{it} : S_{it} \to S_{it}$.

It is easy but tedious to verify that this condition is equivalent to that in Definition 1. The reader is spared the details, which are available on request. To see why this equivalence holds, notice that the left-hand side of the subtraction in any $V_{it}(c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1})$ cancels out with the sum of right-hand sides in $W_{it}(c_i^t, d_i^t, s_i^t, r_i^t)$ with respect to s_{it}. Therefore, the incentive compatibility constraints above are obtained by constructing the telescoping series derived from V_i and W_i with respect to any deviation (δ_i, ρ_i) from honesty and obedience. Finally, by linearity, defining incentive compatibility with respect to all pure deviations (δ_i, ρ_i) as above is equivalent to defining it with respect to all deviations, as in Definition 1.

This equivalent description of incentive compatibility yields the following linear system of equations and inequalities. The *primal* system consists of (i) the family of equations defining $V_{it}(c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1})$, indexed by $(i, t, c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1})$, (ii) the family of equations defining $W_{it}(c_i^t, d_i^t, s_i^t, r_i^t)$, indexed by $(i, t, c_i^t, d_i^t, s_i^t, r_i^t)$, and (iii) the family of inequalities deswcribing incentive compatibility, indexed by (i, β_i, ρ_i). By definition, for any fixed allocation μ, there exist variables (V, W, ζ) to satisfy this primal system if and only if μ is implementable. Such primal problem is a finite-dimensional system of linear inequalities and equations with finitely many variables.

– *Step 2.* By the Theorem of the Alternative, there exist variables (V, W, ζ) to satisfy the primal system of inequalities defined above if and only if the following dual condition holds: given any vector $\lambda \geq 0$, if for every (i, t, c^t, r^t),

$$\sum_{\tau=1}^{t} \sum_{(d_i^\tau, s_i^{\tau-1})} \lambda_{i\tau}(c_i^\tau, d_i^\tau, s_i^{\tau-1}, r_i^{\tau-1})[\Pr(c^t, r^t[s_i^{\tau-1}]|c_i^t[d_i^\tau], r_i^t, \mu) - \Pr(c^t, r^t[s_i^{\tau-1}]|c_i^t[d_i^{\tau-1}], r_i^t, \mu)] + \sum_{(d_i^\tau, s_i^\tau)} \lambda_{i\tau}(c_i^\tau, d_i^\tau, s_i^\tau, r_i^\tau)[\Pr(c^t, r^t[s_i^\tau]|c_i^t[d_i^\tau], r_i^t, \mu) - \Pr(c^t, r^t[s_i^{\tau-1}]|c_i^t[d_i^\tau], r_i^t, \mu)] \;=\; 0, \qquad (*)$$

and for every $(i, t, c_i^t, d_i^t, s_i^t, r_i^t)$,

$$\lambda_{it}(c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1}) \;=\; \sum_{(\delta_i, \rho_i)} \lambda_i(\delta_i, \rho_i)\mathbf{1}\{\delta_{i1}(c_{i1}) = d_{i1}, \rho_{i1}(s_{i1}) = r_{i1}, \ldots, \delta_{it}(c_{it}) = d_{it}\},$$

$$\lambda_{it}(c_i^t, d_i^t, s_i^t, r_i^t) \;=\; \sum_{(\delta_i, \rho_i)} \lambda_i(\delta_i, \rho_i)\mathbf{1}\{\delta_{i1}(c_{i1}) = d_{i1}, \rho_{i1}(s_{i1}) = r_{i1}, \ldots, \delta_{it}(c_{it}) = d_{it}, \rho_{it}(s_{it}) = r_{it}\},$$

where $\mathbf{1}\{E\}$ equals 1 if E is true and 0 otherwise, then

$$\sum_{(i,t,c^t,d_i^t,s_i^{t-1},r^{t-1})} \lambda_{it}(c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1}) \sum_{(c^\tau, r^\tau) \geq (c_i^t, r_i^{t-1})} \{\Pr(c^\tau, r^\tau[s_i^{t-1}]|c_i^\tau[d_i^t], r_i^\tau, \mu)v_{i\tau}(c^\tau[d_i^t], r^\tau[s_i^{t-1}]) - \Pr(c^\tau, r^\tau[s_i^{t-1}]|c_i^\tau[d_i^{t-1}], r_i^\tau, \mu)v_{i\tau}(c^\tau[d_i^{t-1}], r^\tau[s_i^{t-1}])\} + \sum_{(s_{it}, r_{it})} \lambda_{it}(c_i^t, d_i^t, s_i^t, r_i^t) \sum_{(c^\tau, r^\tau) \geq (c_i^t, r_i^t)} \Pr(c^\tau, r^\tau[s_i^t]|c_i^\tau[d_i^t], r_i^\tau, \mu)v_{i\tau}(c^\tau[d_i^t], r^\tau[s_i^t]) - \Pr(c^\tau, r^\tau[s_i^t]|c_i^\tau[d_i^t], r_i^\tau[s_i^t[r_i^{t-1}]], \mu)v_{i\tau}(c^\tau[d_i^t], r^\tau[s_i^t]) \;\leq\; 0.$$

The vector λ collects the multipliers on each of the primal constraints. The multipliers for the first family of equations are denoted by $\lambda_{it}(c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1})$, for the second family by $\lambda_{it}(c_i^t, d_i^t, s_i^t, r_i^t)$, and for the third family by $\lambda_i(\delta_i, \rho_i)$.

– *Step 3.* We will now manipulate this dual condition, applying summation by parts with respect to time to simplify it. Let us begin with the antecedent of the dual condition above, which involves three families of equations. The first family consists of the constraints with respect to which the money payments ζ are multipliers, and the second and third families consist of the constraints with respect to which the V's and W's are multipliers.

The second and third families of constraints imply that

$$\sum_{d_{it+1}} \lambda_{it+1}(c_i^{t+1}, d_i^{t+1}, s_i^t, r_i^t) \;=\; \lambda_{it}(c_i^t, d_i^t, s_i^t, r_i^t) \qquad \forall (i, t, c_i^{t+1}, d_i^t, s_i^t, r_i^t), \tag{A.1}$$

$$\text{and} \sum_{r_{it}} \lambda_{it}(c_i^t, d_i^t, s_i^t, r_i^t) \;=\; \lambda_{it}(c_i^t, d_i^t, s_i^{t-1}, r_i^{t-1}) \qquad \forall (i, t, c_i^t, d_i^t, s_i^t, r_i^{t-1}). \tag{A.2}$$

Substituting (A.2) into (∗) and rearranging, we obtain, for every (i, t, c^t, r^t),

$$-\sum_{d_i^1} \lambda_{i1}(c_i^1, d_i^1) \Pr(c^t, r^t|\mu) +$$

$$\sum_{\tau=1}^{t-1} \sum_{(d_i^\tau, s_i^\tau)} \Pr(c^t, r^t[s_i^\tau]|c_i^t[d_i^\tau], r_i^t, \mu)[\lambda_{i\tau}(c_i^\tau, d_i^\tau, s_i^\tau, r_i^\tau) - \sum_{d_{i\tau+1}} \lambda_{i\tau+1}(c_i^{\tau+1}, d_i^{\tau+1}, s_i^\tau, r_i^\tau)] +$$

$$\sum_{(d_i^t, s_i^t)} \lambda_{it}(c_i^t, d_i^t, s_i^t, r_i^t) \Pr(c^t, r^t[s_i^t]|c_i^t[d_i^t], r_i^t, \mu) \;=\; 0$$

By (A.1), the middle term above disappears, therefore (∗) becomes

$$\sum_{d_i^1} \lambda_{i1}(c_i^1, d_i^1) \Pr(c^t, r^t|\mu) \;=\; \sum_{(d_i^t, s_i^t)} \lambda_{it}(c_i^t, d_i^t, s_i^t, r_i^t) \Pr(c^t, r^t[s_i^t]|c_i^t[d_i^t], r_i^t, \mu) \qquad \forall (i, t, c^t, r^t).$$

By iterating (A.1) and (A.2), it follows that

$$\sum_{d_i^1} \lambda_{i1}(c_i^1, d_i^1) \;=\; \sum_{(d_i^t, r_i^t)} \lambda_{it}(c_i^t, d_i^t, s_i^t, r_i^t) \qquad \forall (i, t, c^t, r^t).$$

Now, dividing both sides by $\sum_{d_i^1} \lambda_{i1}(c_i^1, d_i^1)$ (if it equals zero then there is nothing to prove) and relabeling the ratio of λ's by σ, we finally obtain

$$\Pr(c^t, r^t|\mu) \;=\; \sum_{(d_i^t, s_i^t)} \sigma_{it}(d_i^t, r_i^t|c_i^t, s_i^t) \Pr(c^t, r^t[s_i^t]|c_i^t[d_i^t], r_i^t, \mu) \qquad \forall (i, t, c^t, r^t),$$

where σ_{it} is a stochastic matrix. This takes care of the "undetectable" part of the theorem. The "unprofitable" part follows similar reasoning, and is therefore omitted. □

Theorem 2. The dual of the problem defined by F_i is given by minimizing the inner product of $\zeta^\pm$—the positive and negative parts of ζ—and $\mathbf{z}^\pm$, subject to ζ implementing μ. By the Marginal Value Theorem for linear programming, the subdifferential of the primal with respect to the right-hand side constraints equals the set of solutions to the dual. Finally, the result follows because $F_i(\mathbf{0}|\mu) = 0$ for all i if and only if μ is implementable. □

Corollary 1. Fix any partial history $(\hat{c}^{\hat{t}}, \hat{r}^{\hat{t}})$, and consider the problem of finding an incentive scheme that implements μ subject to $\zeta_{i\hat{t}}(\hat{c}^{\hat{t}}, \hat{r}^{\hat{t}}) = 0$ for every agent i. The dual of this problem is given by $G_i(\mathbf{0}|\mu)$. By standard results (see, e.g., Rockafellar, 1970), the function G_i is differentiable at $\mathbf{0}$ if and only if its subdifferential is a singleton there, which, by the Marginal Value Theorem, is clearly equivalent to revenue equivalence. □

Theorem 3. Follows from Theorem 1. $\square$

Theorem 4. Follows from Theorem 1 and Rahman (2008c, Theorem 2). $\square$

Theorem 5. Consider the following relaxation of (5.1): denote by (η, ξ) any pair such that (i) $\eta = \{\eta_t : C^t \times S^{t-1} \to \mathbb{R}_+\}$ satisfies $\sum_{c_{t+1}} \eta_{t+1}(c^t, c_{t+1}, r^t) = \eta_t(c^t, r^{t-1}) \Pr(r_t | c^t, r^{t-1})$ for all (t, c^t, r^t) and $\eta_0 = 1$, and (ii) $\xi = \{\xi_t : I \times C^t \times S^t \to \mathbb{R}\}$ is a sequence of probability-weighted transfers. Maximize the original objective by choosing (η, ξ) subject to incentive compatibility, except that (i) $\eta_t(c^t, r^{t-1}) \Pr(r_t | c^t, r^{t-1}) \mathrm{Lr}(d_i^t, s_i^t | c^t, r^t)$ replaces every instance of $\Pr(c^t, s_i^t, r_{-i}^t | d_i^t, r_i^t, \mu)$, and (iii) $\xi_{it}(c^t, r^t) \mathrm{Lr}(d_i^t, s_i^t | c^t, r^t)$ replaces every instance of $\zeta_{it}(c^t, r^t) \Pr(c^t, s_i^t, r_{-i}^t | d_i^t, r_i^t, \mu)$. The role of η is to replace μ with unconditional probabilities. This relaxation is now a linear program. Taking its dual yields the Bellman equation above as well as the problem $\min_{\lambda \in \mathscr{U}} J_1$. By hypothesis, there is an incentive compatible mechanism (μ^*, ζ^*) such that $F(\mu^*)$ equals the value of the dual. By changing μ^* into its unconditional probabilities η^* and ζ^* into probability-weighted transfers ξ^*, we obtain a feasible solution to the primal (η^*, ξ^*) that attains the value of the dual. By strong duality, it is also an optimal solution. $\square$

Theorem 6. Maximize f by choosing (η, ξ) as in the proof of Theorem 5, subject to the additional constraint that $\left|\xi_{it}(c^t, r^t)\right| \leq \eta_t(c^t, r^{t-1}) \Pr(r_t | c^t, r^{t-1}) z$. This problem is equivalent to a linear program for each z, with a dual that characterizes virtually implementable optimality as $z \to \infty$ given by the Bellman equation above. $\square$

Theorem 7. The proof is similar to that of Theorem 5, except that now the primal relaxation includes participation constraints. Taking the dual, it follows that the constraints with respect to which the payments are multipliers are given by

$$(\gamma_{it}(c_i^t, r_i^t) - 1) \Pr(r^t | c^t) \;=\; \sum_{(d_i^t, s_i^t)} \lambda_{it}(d_i^t, r_i^t | c_i^t, s_i^t)[\Pr(s_i^t, r_{-i}^t | d_i^t, c_{-i}^t) - \Pr(r^t | c^t)] \quad \forall (i, t, c^t, r^t),$$

where $\gamma \geq 0$ denotes the family of multipliers on the participation constraints. Adding these with respect to (c^t, r^t) it follows that $\sum_{(c^t, r^t)} \gamma_{it}(c_i^t, r_i^t) \Pr(r^t | c^t) = 1$ for all (i, t). This conclusion finally yields the Bellman equation and optimality condition above. $\square$

Theorem 8. Follows from Theorem 1 and Rahman and Obara (2008, Corollary 3). $\square$

Theorem 9. Follows the same lines as the proof of Theorem 8. $\square$

Theorem 10. As in the proof of Theorem 5, we replace (μ, ζ) with (η, ξ) after multiplying the budget constraint by $\Pr(c^t, r^t|mu)$. (Implicitly, we are assuming that agents cannot report being types that have zero probability.) This problem is a linear program with dual described in the statement of the result. The result now follows by strong duality. □

References

Athey, S. and I. Segal (2007): "An Efficient Dynamic Mechanism," Working paper. 1, 2

Bergemann, D. and J. Välimäki (2007): "Dynamic Marginal Contribution Mechanism," Mimeo. 1, 2

Cremer, J. and R. McLean (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247–1257. 1, 7, 8, 9, 10

Gershkov, A. and B. Moldovanu (2007): "The Dynamic Assignment of Heterogeneous Objects: A Mechanism Design Approach," Mimeo. 1, 2

Heydenreich, B., R. Müller, M. Uetz, and R. Vohra (2008): "Characterization of Revenue Equivalence," Mimeo. 2, 6

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73. 14

——— (1986): "Multistage Games with Communication," *Econometrica*, 54, 323–358. 1

Pai, M. and R. V. Vohra (2008): "Optimal Dynamic Auctions," Mimeo. 1, 2

Pavan, A., I. Segal, and J. Toikka (2008): "Dynamic Mechanism Design," Mimeo. 1, 2

Rahman, D. (2008a): "An Alternative Proof of Rochet's Theorem," Mimeo. 7, 18

——— (2008b): "The Alternative to Equilibrium Existence," Mimeo. 18

——— (2008c): "But Who Will Monitor the Monitor?" Mimeo. 2, 10, 11, 12, 23

Rahman, D. and I. Obara (2008): "Mediated Partnerships," Mimeo. 15, 23

Rochet, J. C. (1987): "A Necessary and Sufficient Condition for Rationalizability in a Quasi-Linear Context," *Journal of Mathematical Economics*, 16, 191–200. -1, 1, 7, 8

Rockafellar, R. T. (1970): *Convex Analysis*, Princeton, New Jersey: Princeton University Press. 22