Full Implementation and Belief Restrictions*

Mariann Ollár
University of Pennsylvania,
Econ. Dept. and Warren Center

Antonio Penta
University of Wisconsin-Madison,
Department of Economics

September 22, 2014

Abstract

We introduce a framework to study the problem of full implementation under arbitrary restrictions on agents' beliefs, which we call Δ-Implementation. We show how this framework delivers insights from both a theoretical and an applied perspective. In particular, first we provide a characterization of the properties of the mechanisms that achieve it, and relate it to several notions in the existing literature on full implementation. Then we show how, in environments with transferable utility, using minimal information about agents' beliefs, full implementation can be achieved through the design of simple direct mechanisms. The mechanisms that we use are determined by a payment scheme, that maps agents' reports into transfers. Such mechanisms have a clear economic interpretation, thereby bridging a gap between the literature on partial and full implementation. These results also show how the methodology of the recent literature on robust mechanism design can be extended to address more applied problems of mechanism design, overcoming important limitations of the traditional approach to the full implementation problem.

KEYWORDS: Full Implementation, Robust Mechanism Design, Rationalizability, Moment Conditions, Nice Games, Simple Mechanisms, Uniqueness

1 Introduction

The problem of multiplicity is a key concern for the design of real-world mechanisms and institutions. Unless all the solutions of a mechanism are consistent with the outcome the designer wishes to implement, the designer may not confidently assume that the proposed mechanism will perform well. This is a well known criticism to the *partial implementation* approach to mechanism design, which merely requires that there exists *one* strategy profile consistent with the chosen solution concept that guarantees desirable outcomes. The *full implementation* approach (cf., Maskin, 1999) overcomes the problem of multiplicity. But in pursuit of greater generality, the existing literature has typically adopted rather complicated mechanisms.[1] Thus, while it addresses a practical concern, the full implementation literature overall has provided little insight into how real-world institutions could be designed to avoid the problem of multiplicity.

*We are particularly grateful to Ken Hendricks, Laurent Mathevet, Stephen Morris, Andy Postlewaite and William Sandholm for the helpful comments. We also thank seminar audiences at Stanford, Pompeu Fabra, UW-Madison and at the SAET Conferences in Paris and Tokyo.

[1] See Jackson (1992) for an influential criticism to the tail-chasing mechanisms typically used in this literature.

Another well-known limitation of the classical (Bayesian) approach to mechanism design, particularly important from a practical viewpoint, is the excessive reliance on common knowledge assumptions. This criticism, often referred to as the Wilson doctrine, has recently received considerable attention in the literature on 'robust' implementation. It is fair to say, however, that the promise of the Wilson doctrine, '[...] to conduct useful analyses of practical problems [...]' (Wilson, 1987), is still far from being fulfilled. This is due to two main limitations: On the one hand, most of this literature has focused on environments in which the designer has no information about the agents' beliefs.[2] This extreme assumption represents a useful benchmark to address foundational questions, but significantly limits the relevance of the theory for practical problems of mechanism design. On the other hand, as far as full implementation is concerned (e.g. Bergemann and Morris (2009a), or Penta (2011)), the main focus thus far has been on identifying conditions under which a given mechanism achieves implementation in a robust sense, but it offers little guidance as to *how to design* such a mechanism, *given the objectives* of the designer.

In this paper we address these points pursuing a more pragmatic approach to full implementation. That is, we focus on simple mechanisms that deliver clear practical insights and at the same time are based on more realistic assumptions of common knowledge, intermediate between the classical and the 'belief-free' approaches. More specifically, we study the full implementation problem in environments with transferable utility (TU) and interdependent values, restricting ourselves to using *direct mechanisms*, which only elicit players' payoff-relevant information. To address the concern for robustness, we adopt the solution concept of Δ-Rationalizability (Battigalli and Siniscalchi, 2003), which enables us to study full implementation under general *belief restrictions*, thereby allowing for varying degrees of robustness. Besides the obvious theoretical advantages of maintaining such a level generality, we will show that the flexibility of this solution concept is also extremely convenient from an applied viewpoint.[3]

Direct mechanisms and TU-environments are the main realm of the *partial* implementation literature. Studying *full* implementation in these settings is important for a number of reasons: First, it facilitates the comparison with the literature on partial implementation, by making transparent what features of an incentive compatible mechanism may or may not be problematic from the full implementation viewpoint; Second, because with these restrictions the mechanism design problem boils down to designing transfer schemes which have a clear interpretation and deliver insights that are easily portable to the design of real-world institutions and incentive schemes.[4]

For the sake of illustration, consider the problem of efficient implementation with interdependent values. In environments with single-crossing preferences, the generalized VCG mechanism of Cremer and McLean (1985) guarantees partial implementation of the efficient social choice rule in an ex-post equilibrium, with essentially no restrictions on the strength of the preferences' interdependence. Hence, independent of the agents' beliefs, truthful revelation (hence efficiency) is always achievable as part of a Bayes-Nash equilibrium (e.g., Bergemann and Morris, 2005). The

[2] On the 'belief free' approach to mechanism design, see Bergemann and Morris (2005, 2009a, 2009b, 2011) for static mechanisms and Mueller (2012a,b) and Penta (2011) for dynamic ones. A thorough account of this literature is provided by Bergemann and Morris (2012). We discuss the related literature more extensively in Section 1.1.

[3] For instance, as the belief restrictions are varied, our notion of implementation (viz. Δ-Implementation) includes as special cases both Bergemann and Morris' (2009a, 2011) belief-free implementation and Oury and Tercieux's (2012) ICR-implementation.

[4] D'Aspremont, Cremer and Gerard-Varet (2005) also studied full implementation in environments with transferable utility, but they resort to unbounded mechanisms of the kind criticized above. Duggan and Roberts (2002) fully implement the efficient allocation of pollution via transfers, but under complete information and richer reports.

problem with this mechanism is that it typically admits also inefficient equilibria, which can only be ruled out if the interdependence in agents' valuations is not too strong (e.g., Bergemann and Morris, 2009). However, if the designer does possess some information about agents' beliefs, then it may be the case that a clever modification of the baseline mechanism may guarantee full implementation of the efficient allocation for a larger set of preferences. One may thus consider two kinds of questions: (i) Given certain information about agents' beliefs (e.g., a standard type space, or just a set of beliefs), is it possible to design an efficient and incentive compatible mechanism that guarantees full implementation? And, if so, how would it look like? (ii) What kind and how much information on agents' beliefs is needed to be able to design such an incentive compatible and 'fully efficient' mechanism?

The reason why the VCG mechanism fails to fully implement the efficient decision rule when interdependencies are strong is that under this condition agents' best response functions are strongly affected by the opponents' strategies: different conjectures about the opponents' behavior lead to very different optimal responses. The multiplicity that prevents the full implementation result stems precisely from such strong 'strategic externalities'. If *some* information about agents' beliefs is available, full efficient implementation may be possible. The VCG mechanisn, however, ignores such information. There is thus a tension between the robustness of the partial implementation result (achieved by the VCG mechanism in an ex-post equilibrium), and the possibility to achieve full implementation of the efficient decision rule. The key idea that we pursue is to modify the baseline VCG mechanism using the information about agents' beliefs. Of course, this modification should maintain incentive compatibility (if not in the ex-post sense, at least for the beliefs consistent with the designer's information), but also guarantee that the strategic externalities in the resulting mechanism are not too strong. Depending on the nature and 'amount' of information about agents' beliefs, the strength of the preferences' interdependence need not coincide with the strength of the strategic externalities: the latter can be weakened even though the former are strong. If this can be guaranteed, then full implementation is achievable regardless of the strength of the interdependence in agents' valuations.

While the efficient implementation problem is one of our leading examples, our analysis covers general implementation problems with interdependent values. In Sections 2 and 3 we introduce the notion of Δ-Implementation, and develop a framework to analyze full implementation under general restrictions on agents' beliefs. In Section 3.2 we provide a full characterization of the Δ-Implementable mechanisms. We show that a direct mechanism is Δ-Implementable if and only if it is Δ-Incentive compatible and Δ-Contractive. In Section 6 we relate these concepts to existing notions of incentive compatibility and monotonicity, particularly to those provided by Bergemann and Morris (2009a) and Oury and Tercieux (2012) to characterize belief-free and ICR-Implementation, respectively.

The question of how to obtain Δ-Incentive Compatible and Δ-Contractive mechanisms through the design of simple transfer schemes is considered in Section 4. Our design strategy consists of two steps. First, we derive the 'canonical transfers'. These are obtained from a generalization of well-known necessary conditions for ex-post incentive compatible payment schemes, and generalize the idea leading to the optimal auction mechanisms (Myerson, 1981) and the VCG payments for efficient mechanism design. In environments with single-crossing preferences and increasing decision rules, the canonical transfers induce an ex-post incentive compatible mechanism, but if preference interdependencies are strong they generate strong strategic externalities which lead to a failure of full implementation. The second part of our design then exploits the information about agents' beliefs to correct such strategic externalities, so as to induce a contractive mechanism which guarantees uniqueness. We also consider environments without single-crossing preferences. In these environments, the canonical transfers only guarantee that truthful revelation is an extremal point given the truthful strategy of the opponent, but do not guarantee neither optimality nor uniqueness of the truthful strategy. The information about agents' beliefs in that case is used to correct both the possible non-optimality and the multiplicity.

The conditions that guarantee full implementation relate the strength of the preference interdependence to the information embedded in the belief restrictions, formalizing the idea that belief restrictions are useful in designing contractive mechanisms if they contain information that is 'inversely-aligned' with the agents' preferences, so that it can be used to weaken the strategic externalities induced by the canonical transfers.

In Section 5 we show how these general results can be applied to important special cases, such as the common prior environments that are commonly studied by the classical and applied literature on mechanism design. We show, for instance, that full implementation is possible if types are affiliated or stochastically independent, and we design transfers that achieve it. Our construction suggests a *simple design principle*: start out with the ex-post incentive compatible transfers, and then compensate each agent for a proper measure of the strategic externality he is subject to, given the reports. To avoid that agents misreport their type in order to inflate their compensation, each agent i is asked to pay a fee equal to the expected value of the compensation, given his type:

$$t_i(\theta) = \underbrace{t_i^{EPIC}(\theta)}_{\text{canonical transfers}} + \underbrace{CSE_i(\theta_i, \theta_{-i})}_{\substack{\text{compensation for} \\ \text{strategic externality} \\ \text{(depends on everybody's report)}}} - \underbrace{E(CSE_i|\theta_i)}_{\substack{\text{belief-based adjustment:} \\ \text{expected compensation} \\ \text{(only depends on } i\text{'s report)}}} .$$

The first term we add to the canonical transfers reduces the strategic externalities and ensures that the mechanism is contractive; the last term, derived from the designer's information about agents' beliefs, restores incentive compatibility. Full implementation follows.

Importantly, these implementation results do not rely on the full strength of the common prior assumption: common knowledge of certain summary statistics of the types' distribution suffices. This is true more generally for our construction: our general results achieve full implementation whenever some moments of the types' distribution are common knowledge. Such summary statistics can be estimated from previous data on the performance of the mechanism. As recently argued by Deb and Pai (2013), this is a desirable property for a mechanism, and it is guaranteed here thanks to the combination of the *full* implementation requirement and the robustness entailed by the solution concept. From a different perspective, these results indirectly shed some light on what kind of information on a mechanism past performance may be useful for the designer to disclose. (We discuss this and other extensions in Section 7.)

1.1 Related Literature

Our work is related to several strands of the literature in game theory and mechanism design. We briefly discuss the most closely related literature.

Solution Concept. The solution concept that we use, Δ-Rationalizability, was introduced by Battigalli (2003) and Battigalli and Siniscalchi (2003). It generalizes several versions of rationalizability for incomplete information games, including the 'belief-free' rationalizability studied by Bergemann and Morris (2009) and Dekel, Fudenberg and Morris' (2007) 'interim correlated rationalizability' (ICR). ICR has also been studied by Oury and Tercieux (2012), Penta (2013) and Weinstein and Yildiz (2007, 2011, 2013). Battigalli et al. (2011) provide a thorough analysis of Δ-Rationalizability and its connections with belief-free, ICR and other versions of rationalizability.

Nice Games. At a more technical level, our construction exploits the notion of 'nice mechanisms', which extends Moulin's (1984) idea of nice games to encompass environments with incomplete information. Our implementation results are based on general uniqueness results we establish for 'nice games', more extensively discussed in Ollár and Penta (2013). Nice games are convenient analytical tools, particularly if rationalizability is adopted as solution concept. For a recent application of (complete information) nice games, see Weinstein and Yildiz (2011).

Full Implementation. Within the vast literature on (full) implementation, the closest papers are Bergemann and Morris (2009a) and Oury and Tercieux (2012), which study implementation in 'belief free' rationalizability and ICR, respectively. The first is an important benchmark, in that it represents the most demanding notion of robustness with respect to agents' beliefs. The second characterizes 'continuous implementation', an important property of local robustness for partial implementation (cf. Oury (2013)). Both 'belief free' and ICR-Implementation are special cases of Full Δ-Implementation. Whereas most of our analysis aims at achieving full implementation via simple transfer schemes, and hence focuses on TU-environments, the general framework as well as the results in Sections 3.2 and 6 easily extend to the environments considered by Bergemann and Morris (2009a) and Oury and Tercieux (2012). Their characterization results can be obtained as special cases of ours, with the proviso that Oury and Tercieux (2012) do not restrict attention to direct mechanisms, as we do (cf. Section 6). The restriction to direct mechanisms is also shared by Bergemann and Morris (2009a), while Bergemann and Morris (2011) study belief-free implementation in general mechanisms. Within the classical literature, Jackson (1991) and Postlewaite and Schmeidler's (1986) Bayesian Monotonicity are also connected to Δ-Contractivity (cf. Section 6). From a conceptual viewpoint, our departure from that literature is inspired by Jackson's (1992) critique of unbounded mechanisms. We push the concern for 'relevance' a bit further, requiring that full implementation is achieved via simple transfer schemes.

Mechanism Design in TU-Environments. TU-environments are the typical domain of the partial implementation literature. Within this area, the closest works are those that allow for interdependent values (e.g., Cremer and McLean (1985, 1988), Dasgupta and Maskin (2000), McLean and Postlewaite (2004).[5] In recent years, a growing literature has revisited standard results, imposing extra desiderata inspired by more practical considerations. The already mentioned paper by Deb and Pai (2013) is one such example, which pursues symmetry of the mechanism. Mathevet (2010) and Mathevet and Taneva (2013) instead pursue supermodularity. In those papers, the

[5] McLean and Postlewaite (2002) also explore related ideas and allow interdependent values, but also allow for environments without transferable utility.

extra desiderata are achieved by adding a belief-dependent component to some baseline payments, very much as we attain full implementation appending an extra term to the canonical transfers.[6] One difference is that those papers maintain that types are independently distributed, whereas we allow more general correlations, as well as weaker restrictions on beliefs. At a more technical level, our design results in a contractive mechanism. Given our concern with full implementation, contractivity is a more convenient property than supermodularity. Healy and Mathevet (2013) also pursue contractivity of the mechanism, though in a complete information setting.

Robust Mechanism Design. As already mentioned, most of the literature on robust mechanism design has focused on the belief-free case. See, for instance, Bergemann and Morris (2005, 2009 and 2011) for static mechanism design, and Müller (2012a,b) and Penta (2011) for dynamic mechanism design. Kim and Penta (2012) explore *partial* implementation with interdependent values, maintaining some restrictions on beliefs. Lopomo, Rigotti and Shannon (2013) also explore partial implementation with beliefs restrictions analogous to ours, but focus on single agent problems and consider a different notion of robustness. Artemov, Kunimoto and Serrano (2013) also maintain some restrictions on beliefs, but focus on virtual implementation. Different approaches to robust mechanism design have been recently put forward by Yamashita (2013a,b), Börgers and Smith (2012,2013) and Carroll (2013).

Implementation and the 'price of anarchy': The partial implementation approach often argues that the truthtelling efficient equilibrium is plausible as it is intuitive for players to play that equilibrium instead of some other equilibria. The validity of this argument clearly depends on the particular instance. A growing literature quantifies the potential efficiency loss due to multiplicity in auction environments. Constant bounds are known for the 'price of anarchy' (the welfare ratio between efficient and worst Nash Equilibrium outcomes) in complete information settings. For incomplete information, Roughgarden (2012) shows that the price of anarchy can be arbitrarily large. If the design guarantees full implementation of the efficient outcome, then the price of anarchy is minimized, equal to 1.

2 Model

Environments and Mechanisms. We consider standard environments with transferable utility. We denote by $I = \{1, ..., n\}$ the set of agents, by X the set of (common) social outcomes and by $t_i \in \mathbb{R}$ the private transfer to agent $i \in I$. Agents' preferences depend on the realization of the state of the world $\theta \in \Theta = \times_{i \in I} \Theta_i$. When θ is realized, agent i privately observes the i-th component, $\theta_i \in \Theta_i$. We refer to $\theta_i \in \Theta_i$ as agent i's *payoff type* (or just as 'type'), and let $\theta_{-i} \in \Theta_{-i} = \times_{j \neq i} \Theta_j$ denote the type profile of i's opponents. For each $i \in I$, we let $u_i : X \times \mathbb{R} \times \Theta \to \mathbb{R}$ denote player i's utility function, and we assume that there exists a function $v_i : X \times \Theta \to \mathbb{R}$ such that

$$u_i (x, t_i, \theta) = v_i (x, \theta) + t_i$$

for every $(x, t_i, \theta) \in X \times \mathbb{R} \times \Theta$. We refer to $v_i (\cdot)$ as agent i's *valuation function.* We maintain throughout that $(\Theta_i)_{i \in I}$ are convex and compact subsets of the real line and that X is a convex

[6] Early examples of this principle are the mechanisms of D'Aspremont and Gerard-Varet (1975) and of Cremer and McLean (1985), which append the baseline VCG mechanism with a belief-based component in order to achieve budget balance and surplus extraction, respectively.

and compact subset of a Euclidean space. This model accommodates general externalities in consumption, including both pure cases of private and public goods.

The tuple $\mathcal{E} = \langle I, (\Theta_i, u_i) \rangle$ defines the 'payoff environment', and we assume it is common knowledge among the agents. Payoff types thus represent agents' information about preferences. If v_i is constant in θ_{-i} for every i, then the environment is one of *private values*. If not, the environment has *interdependent values*.

A decision rule (or allocation rule) is a mapping $d : \Theta \to X$, which assigns to each payoff state the social outcome that the designer wishes to implement. We say that an allocation rule is *responsive* if for any i, θ_i and θ_i' such that $\theta_i \neq \theta_i'$, there exists $\theta_{-i} \in \Theta_{-i}$ such that $d(\theta_i, \theta_{-i}) \neq d(\theta_i', \theta_{-i})$. We impose the following assumptions on $(\mathcal{E}, d)$:

Assumption 1 (Smoothness): *(i) $d : \Theta \to X$ is responsive and twice continuously differentiable; (ii) for every i, v_i is three times continuously differentiable in all the arguments.*

In general, a mechanism is a tuple $\mathcal{M} = \left\langle (M_i)_{i \in I}, g \right\rangle$, where $g : \times_{i \in I} M_i \to X \times \mathbb{R}^n$. For the reasons discussed in the introduction, we restrict ourselves to using *direct mechanisms*, in which the sets of messages are $M_i = \Theta_i$, and the common component $x \in X$ is chosen according to d. A direct mechanism is thus uniquely determined by a *transfer scheme* $(t_i)_{i \in I}$, $t_i : M \to \mathbb{R}$, which specifies the (possibly negative) transfer to agent i, for each possible profile of reports $m \in M$ (to distinguish the report from the state, we maintain the notation m even though $M = \Theta$.)

Any direct mechanism $\mathcal{M}$ induces a (belief-free) game $G^{\mathcal{M}} = \left\langle I, (\Theta_i, M_i, U_i)_{i \in I} \right\rangle$, where I is the set of players, Θ_i the set of i's payoff types, M_i is the set of i's actions and ex-post payoff functions $U_i : M \times \Theta \to \mathbb{R}$ are such that

$$U_i(m; \theta) = v_i(d(m), \theta) + t_i(m).$$

For every $\theta_i \in \Theta_i$, $\mu_i \in \Delta(\Theta_{-i} \times M_{-i})$ and $m_i \in M_i$, we let $EU_{\theta_i}^{\mu_i}(m_i)$ denote player i's expected payoff from message m_i, if i's type is θ_i and his conjectures are μ_i:

$$EU_{\theta_i}^{\mu_i}(m_i) := \int_{\Theta_{-i} \times M_{-i}} U_i(m_i, m_{-i}; \theta_i, \theta_{-i})\, d\mu_i.$$

We also define $BR_{\theta_i}(\mu_i) := \arg\max_{m_i \in M_i} EU_{\theta_i}^{\mu_i}(m_i)$.

Belief Restrictions. We model the assumptions on agents' beliefs separately from the environment $\mathcal{E}$. This is because, whereas information about beliefs may be useful in designing a mechanism, agents' beliefs are not directly relevant to the designer's objectives, d. We model 'belief restrictions' as sets of possible beliefs for each type of every player. Formally, $\mathcal{B} = ((B_{\theta_i})_{\theta_i \in \Theta_i})_{i \in I}$ where $B_{\theta_i} \subseteq \Delta(\Theta_{-i})$ for each $\theta_i \in \Theta_i$ and $i \in I$, assumed common knowledge. We maintain throughout that each B_{θ_i} is non-empty, closed and convex. If $\mathcal{B}$ and $\mathcal{B}'$ are such that $B_{\theta_i} \subseteq B_{\theta_i}'$ for all θ_i and i, we write $\mathcal{B} \subseteq \mathcal{B}'$.

This formulation is fairly general. For instance, if B_{θ_i} is a singleton for every θ_i and i, then the pair $(\mathcal{E}, \mathcal{B})$ is a standard Bayesian environment, in which agents' hierarchies of beliefs are uniquely pinned down by their payoff types. The further special case in which $B_{\theta_i} = B_{\theta_i'}$ for all i and all $\theta_i, \theta_i' \in \Theta_i$ corresponds to the case of independent types (cf. Section 2.1). At the opposite extreme, if $B_{\theta_i} = \Delta(\Theta_{-i})$ for every θ_i and every i, then there are no commonly known restrictions on beliefs, and the pair $(\mathcal{E}, \mathcal{B})$ coincides with the belief-free environments that are common in the literature on robust mechanism design.[7] Such 'vacuous restrictions' are thus denoted by $\mathcal{B}^{BF}$. Our model also accommodates settings, intermediate between the Bayesian and belief-free cases, in which some common knowledge restrictions are maintained but not to the point that belief hierarchies are uniquely determined by the payoff types. In those cases, the tuple $\mathcal{B}$ represents the designer's *partial information* about agents' beliefs. Our results apply to all of these cases.

For reasons that will be illustrated in the next example, we will distinguish between the belief restrictions in $\mathcal{B}$ and the beliefs with respect to which full implementation may be obtained. From this viewpoint, it is useful to think of $\mathcal{B}$ as *the most* that the designer is willing assume about agents' beliefs. (Clearly, if $\mathcal{B} \subseteq \mathcal{B}'$, then $\mathcal{B}'$ entails weaker restrictions than $\mathcal{B}$.)

2.1 Example: Full Implementation in a Common Prior Model

Consider an environment with two agents, $i = 1, 2$. The planner chooses some quantity $x \in X \equiv \mathbb{R}_+$ of a public good, with cost of production $c(x) = \frac{1}{2}x^2$. Players' valuation functions are $v_i(x, \theta) = (\theta_i + \gamma\theta_j)x$, where $\gamma \geq 0$ is a parameter of preference interdependence: if $\gamma = 0$, this is a private-value setting; if $\gamma > 0$, values are interdependent. Now, suppose that the planner knows that agents' types are i.i.d. draws from a uniform distribution over $\Theta_i \equiv [0,1]$, denoted by $\bar{v}_{\Theta_i}$, and that this is common knowledge among the agents. This is a standard common prior environment, with independently distributed types and interdependent values. In this model, the planner's information about agents' beliefs is represent by belief restrictions $\mathcal{B} = ((B_{\theta_i})_{\theta_i \in \Theta_i})_{i \in I}$ such that $B_{\theta_i} = \left\{\bar{v}_{\Theta_j}\right\}$ for every i, $j \neq i$ and $\theta_i \in \Theta_i$.

The objective of the social planner is to implement the efficient level of public good, that is $d^*(\theta) = (1+\gamma)(\theta_1 + \theta_2)$. It is well known that this decision rule can be partially implemented through the VCG mechanism, with transfers

$$t_i^{VCG}(m) = -(1+\gamma)\left(0.5m_i^2 + \gamma m_i m_j\right). \quad (1)$$

Given the VCG mechanism, for any pair (θ_j, m_j) of player j's type and report, the *ex-post* best-reply function for type θ_i of player i is

$$BR_{\theta_i}^{VCG}(\theta_j, m_j) = proj_{[0,1]}(\theta_i + \gamma(\theta_j - m_j)). \quad (2)$$

Observe that for any $\gamma \geq 0$ and for any realization of θ, truthful revelation $(m_i(\theta_i) = \theta_i)$ is a best response to the opponent's truthful strategy $(m_j(\theta_j) = \theta_j)$. This is the *ex-post incentive compatibility* of the VCG mechanism. Partial implementation of the efficient allocation is thus guaranteed independent of agents' beliefs. Furthermore, if $\gamma < 1$, equation (2) is a contraction, and its iteration delivers truthful revelation as the only rationalizable strategy. In this case, the VCG mechanism also guarantees full robust implementation (Bergemann and Morris (2009a)). If $\gamma \geq 1$, on the other hand, the VCG mechanism fails to robustly implement the efficient allocation rule. (In the general symmetric case with n agents, it can be shown that *no* mechanism achieves belief-free full implementation if $\gamma \geq 1/(n-1)$.)

[7] See Bergemann and Morris (2005, 2008, 2009a,b, 2011), or Mueller (2011, 2012) and Penta (2011) for dynamic mechanisms.

Hence, with weak interdependence in valuations, the designer need not rely on the common prior the VCG mechanism ensures full implementation in the belief-free model $\mathcal{B}^{BF} \supset \mathcal{B}$. If the interdependence is too strong, however, belief-free implementation is impossible, even under the $\mathcal{B}$-restrictions. For instance, if $\gamma = 2$, the strategy profile $(\hat{m}_1(\theta_1) = 1, \hat{m}_2(\theta_2) = 0)$ is also a Bayes Nash equilibrium (BNE) of the VCG mechanism. To see this, consider the *interim* best reply of type θ_i, given the common prior and the opponent's strategy $\hat{m}_j : \Theta_j \to M_j$:

$$BR^{VCG}_{\theta_i}(\hat{m}_j(\cdot)) = proj_{[0,1]}(\theta_i + \gamma[\mathbb{E}(\theta_j|\theta_i) - \mathbb{E}(\hat{m}_j(\theta_j)|\theta_i)]).$$

If $\gamma = 2$ and j always reports 0 (resp., 1), then i's best reply is to report as high (as low) as possible. Furthermore, this equilibrium is inefficient, as it implements $x = 3$ regardless the state. The source of this multiplicity of equilibria, when interdependence in valuations are strong, is that the VCG mechanism determines strong strategic externalities: if γ is large, players' best reponses are largely affected by changes in the opponents' strategies.

Being designed to achieve ex-post incentive compatibility, the VCG mechanism clearly ignores any information about agents' beliefs. We propose next a different set of transfers, which do exploit some information contained in the common prior. Namely, that $\mathbb{E}(\theta_j|\theta_i) = \frac{1}{2}$ for all θ_i and i:

$$t^*_i(m) := -(1+\gamma)\left(\frac{1}{2}m_i^2 + \gamma m_i \mathbb{E}(\theta_j|\theta_i)\right) = -(1+\gamma)\left(\frac{1}{2}m_i^2 + \gamma m_i \frac{1}{2}\right). \quad (3)$$

These transfers induce the following the best response function:

$$BR^*_{\theta_i}(\hat{m}_j(\cdot)) = proj_{[0,1]}\left(\theta_i + \gamma\left[\mathbb{E}(\theta_j|\theta_i) - \frac{1}{2}\right]\right). \quad (4)$$

Since, under the common prior, $\mathbb{E}(\theta_j|\theta_i) = \frac{1}{2}$ for all θ_i, the term in square brackets cancels out for all types. Truthful revelation therefore is *strictly dominant*, independent of the strength of preference interdependence, γ. In fact, this would be the case for any beliefs that satisfy the condition $\mathbb{E}(\theta_j|\theta_i) = \frac{1}{2}$ for all θ_i. Hence, full implementation is guaranteed not just for the common prior $\mathcal{B}$, but for the weaker restrictions $\mathcal{B}^* = ((B^*_{\theta_i})_{\theta_i \in \Theta_i})_{i \in I}$ defined as $B^*_{\theta_i} := \left\{b_i \in \Delta(\Theta_j) : \int \theta_j \cdot db_i = \frac{1}{2}\right\}$. Furthermore, since truthful revelation is dominant, given $\mathcal{B}^*$, such restrictions need not even be common knowledge among the agents: as long as $\mathbb{E}(\theta_j|\theta_i) = \frac{1}{2}$ for all θ_i and i, full implementation obtains independent of higher order beliefs.

Discussion and Generalizations: This example is a standard Bayesian implementation problem, in which the planner's information is represented by a common prior model, $\mathcal{B}$. As illustrated above, however, the designer may achieve implementation without necessarily relying on the full strength of the common prior model. If $\gamma < 1$, the VCG mechanism ensures belief-free implementation, that is for all beliefs consistent with the vacuous restrictions $\mathcal{B}^{BF} \supset \mathcal{B}$. If $\gamma \geq 1$, the transfers in (3) achieve full implementation for all beliefs consistent with the restrictions $\mathcal{B}^* \supset \mathcal{B}$. The precise definition of $\mathcal{B}^*$ clearly depends on the particular *moment condition* we used to design the transfers (that is, $\mathbb{E}(\theta_j|\theta_i) = \frac{1}{2}$ for all θ_i and i). This is only one of infinitely many conditions that are consistent with the designer's information $\mathcal{B}$. Had we used a different condition, say "$\mathbb{E}(G(\theta_{-i})|\theta_i) = f(\theta_i)$ for all θ_i and i", full implementation may have obtained for different beliefs

$\mathcal{B}' \supset \mathcal{B}$: namely, $\mathcal{B}' = ((B'_{\theta_i})_{\theta_i \in \Theta_i})_{i \in I}$ such that $B'_{\theta_i} = \left\{ b_i \in \Delta(\Theta_{-i}) : \int G(\theta_{-i})\, db_i = f(\theta_i) \right\}$.

Thus, not only the set $\mathcal{B}$, which represents the designer's information, need not coincide with the set of beliefs with respect to which implementation is obtained (such as $\mathcal{B}^{BF}$ or $\mathcal{B}^*$ in the example), but the latter is itself determined by the planner's choice of the mechanism. Both the transfers and the degree of robustness are a *choice* of the designer. For this reason, we distinguish between the designer's assumptions about agents' beliefs, represented by the belief restrictions introduced above, and the beliefs with respect to which implementation is achieved, which will be referred to as 'Δ-restrictions'. Our notion of 'Δ-implementation' treats such Δ-restrictions as a *parameter*, which can be chosen in the design of a mechanism. In Section 3 we characterize the properties of general Δ-restrictions that ensure full implementation. We then use these general results to provide guidelines for the designer's choice of which, of the possibly many Δ-restrictions consistent with $\mathcal{B}$, are useful to design transfers for full implementation. In particular, in Section 4 we develop a general design principle which consists of using properly chosen Δ-restrictions to weaken the strategic externalities of a baseline 'canonical' mechanism. We show that a special kind of Δ-restrictions (namely, *moment conditions*) are particularly suited to this task. In the example above, for instance, the moment condition "$\mathbb{E}(\theta_j|\theta_i) = \frac{1}{2}$ for all θ_i and i" enabled us to completely offset the strategic externalities of the VCG mechanism, thereby ensuring full implementation in dominant strategies. In the general case in which strategic externalities cannot be completely eliminated, our design principle pursues contractivity of the best replies, to ensure that truthful revelation is the unique rationalizable outcome (given the Δ-restrictions). In Section 5 we apply the general results to common prior environments with independent or affiliated types, and show that the key insights of the example fully generalize to these important special cases.

3 Δ-Implementation

3.1 Definitions

Solution Concept. The game $G^{\mathcal{M}}$ is a 'belief-free' game, in that it does not contain any information about agents' beliefs. Restrictions on beliefs are introduced via the solution concept, Δ-Rationalizability (Battigalli and Siniscalchi, 2003). Δ-Rationalizability characterizes the behavioral implications of common certainty of players' rationality and of a set of exogenous restrictions on players' conjectures. The latter are referred to as 'Δ-restrictions'.

Formally, Δ-restrictions are a collection $\Delta = ((\Delta_{\theta_i})_{\theta_i \in \Theta_i})_{i \in I}$ such that $\Delta_{\theta_i} \subseteq \Delta(\Theta_{-i} \times M_{-i})$ for every i and θ_i. Unlike $\mathcal{B}$, which restricts agents' beliefs about the environment, the Δ-restrictions concern agents' *conjectures*, which are also about opponents' behavior in the mechanism. Δ-Rationalizability consists of an iterated deletion procedure in which, for each type θ_i, a given report m_i survives the k-th round of deletion if and only if it can be justified by conjectures in Δ_{θ_i} that are consistent with the previous rounds of deletion:

Definition 1 (Δ-Rationalizability) *Fix a set of Δ-restrictions. For every $i \in I$, let* $R_i^{\Delta,0} =$

$\Theta_i \times M_i$ *and for each* $k = 1, 2, ...,$ *let* $R_{-i}^{\Delta,k-1} = \times_{j \neq i} R_j^{\Delta,k-1}$,

$$R_i^{\Delta,k} = \left\{ (\theta_i, m_i) : m_i \in BR_{\theta_i}(\mu_i) \text{ for some } \mu_i \in \Delta_{\theta_i} \cap \Delta\left(R_{-i}^{k-1}\right) \right\},$$
$$\text{and } R_i^{\Delta} = \bigcap_{k \geq 0} R_i^{\Delta,k-1}.$$

The set of Δ*-rationalizable messages for type* θ_i *is defined as:* $R_i^{\Delta}(\theta_i) := \left\{ m_i : (\theta_i, m_i) \in R_i^{\Delta} \right\}$.

The Δ-Restrictions. A special case of interest is when the Δ-restrictions merely capture the idea that, in the game ensuing from the mechanism, the beliefs represented in the model $\mathcal{B}$ are common knowledge. Formally: the Δ-restrictions are '*equivalent to* $\mathcal{B}$', denoted by $\Delta^{\mathcal{B}} = \left(\left(\Delta_{\theta_i}^{\mathcal{B}} \right)_{\theta_i \in \Theta_i} \right)_{i \in I}$, if for every $i \in I$ and every $\theta_i \in \Theta_i$,

$$\Delta_{\theta_i}^{\mathcal{B}} := \left\{ \mu_i \in \Delta(\Theta_{-i} \times M_{-i}) : marg_{\Theta_{-i}} \mu_i \in B_{\theta_i} \right\}. \tag{5}$$

For the reasons discussed in Section 2.1, we consider general Δ-restrictions, not necessarily equivalent to the designer's information. In particular, it will be useful to consider Δ-restrictions that are '*weaker than* $\mathcal{B}$' : that is, $\Delta_{\theta_i} \supseteq \Delta_{\theta_i}^{\mathcal{B}}$ for all i and θ_i (or $\Delta \supseteq \Delta^{\mathcal{B}}$). Another important special case is when the Δ-restrictions are vacuous, denoted by Δ^{BF} (since they are equivalent to the belief-free model). In general, we maintain the following assumption on Δ-restrictions:

Assumption 2 ('Non Behavioral' Δ-Restrictions): *The* Δ*-restrictions are non-behavioral: that is, there exist belief restrictions* $\mathcal{B}'$ *such that* $\Delta = \Delta^{\mathcal{B}'}$.

To understand the meaning of this assumption, suppose that it is violated. Then, for some θ_i, Δ_{θ_i} restricts beliefs not only about the opponents' types, but also about their behavior in the mechanism. If the designer has information on agents' conjectures about others' behavior, then he would specify Δ-restrictions that violate Assumption 2. This is an interesting idea, unexplored in the implementation literature, though often implicit in more applied work.[8] We point out that such behavioral restrictions can easily be accommodated into our framework, but maintain Assumption 2 throughout. Notice that, if the Δ-restrictions are *non behavioral*, then: (i) the sets $\Delta_{\theta_i} \subseteq \Delta(\Theta_{-i} \times M_{-i})$ are closed, non-empty and convex (this follows from the definition of belief restrictions); and (ii) they are *weaker than* $\mathcal{B}$ if and only if $\Delta = \Delta^{\mathcal{B}'}$ for some $\mathcal{B}' \supseteq \mathcal{B}$.

Full Implementation. Our notion of implementation requires *all* the Δ-Rationalizable profiles of a direct mechanism to induce outcomes consistent with the allocation rule. This notion presents several advantages. First, as the Δ-restrictions are varied, Δ-Rationalizability coincides with various versions of rationalizability, some of which play an important role in the literature on robustness and on implementation.[9] The flexibility of Δ-Rationalizability also allows a unified

[8] Exogenous restrictions on behavior are common in the related literature on market design, where assuming behavioral restrictions such as linear bidding strategies is convenient when comparing performance of different mechanisms (e.g., Ausubel et al. (2013)). Other examples of implicit behavioral Δ-restrictions include the introduction of 'noise traders' in financial market models (e.g., Kyle (1989)), the assumption that players consider the opponents' truthtelling strategies in auctions (e.g., Deb and Pai (2013)) and that bidders do not bid above their true valuations in the empirical auctions literature (e.g., Haile and Tamer (2003)).

[9] In particular, if $\mathcal{B}$ is a Bayesian model and $\Delta = \Delta^{\mathcal{B}}$, then Δ-rationalizability coincides with interim correlated rationalizability (ICR, Dekel, Fudenberg and Morris, 2007). ICR-Implementation has been studied by Oury and Tercieux (2012). If instead the Δ-restrictions are vacuous ($\Delta = \Delta^{BF}$), then Δ-rationalizability coincides with 'belief

analysis of a wide range of different settings, which is convenient from both a theoretical and an applied viewpoint. Second, Δ-Rationalizability in general is a very weak solution concept. This is important because, unlike for 'partial' implementation, full implementation results are stronger if obtained with respect to a weaker solution concept. Sufficient conditions for full Δ-Implementation therefore guarantee full implementation with respect to any refinement of Δ-Rationalizability. Finally, it can be shown that Δ-Rationalizability characterizes the set of all Bayes-Nash equilibria (BNE) in type spaces that are consistent with the Δ-restricitons. Full Δ-Implementation therefore can be seen as a shortcut to analyze standard questions of Bayesian Implementation for general restrictions on beliefs (cf. Appendix A).

Definition 2 (Full Δ-Implementation) *Fix an allocation rule d, a direct mechanism $\mathcal{M} = (d,t)$ and a set of Δ-restrictions. We say that:*

1. *$\mathcal{M}$ (fully) Δ-implements d, if $R^{\Delta}(\theta) \neq \emptyset$ for all θ and $m \in R^{\Delta}(\theta)$ implies $d(m) = d(\theta)$.*
2. *$\mathcal{M}$ truthfully Δ-implements d, if $R_i^{\Delta}(\theta_i) = \{\theta_i\}$ for all θ_i and all $i \in I$.*
3. *$\mathcal{M}$ implements d in (strictly) Δ-dominant strategies, if $R_i^{\Delta,1}(\theta_i) = \{\theta_i\}$ for all θ_i and all i.*

We say that d is Δ-Implementable (resp.: truthfully Δ-Implementable; Δ-DS Implementable) if there exists a direct mechanism that Δ-implements d (resp.: truthfully Δ-implements d; implements d in Δ-dominant strategies).

Consider point 3 first: Δ-DS Implementation. As shown in the example of Section 2.1, if truthful implementation is achieved with only one round of Δ-rationalizability, then truthful revelation is strictly dominant for all the beliefs consistent with the Δ-restrictions. In this case, full implementation actually obtains independent of higher order beliefs, so the Δ-restrictions need not even be common knowledge among the agents. This concept therefore entails a very robust notion of implementation, and it is stronger than the conditions in points 1 and 2. For instance, if $\Delta = \Delta^{\mathcal{B}}$ for some Bayesian model $\mathcal{B}$, then Δ-DS Implementation is equivalent to truthtful revelation being strictly dominant in the interim normal form of the Bayesian game.

It is also easy to see that truthful Δ-Implementation in general is a stronger requirement than Δ-Implementation. The next Proposition, however, shows that the two concepts coincide under the maintained assumptions 1 and 2:

Proposition 1 *If the Δ-restrictions are non-behavioral and d is responsive, then d is (fully) Δ-Implementable if and only if it is Truthfully Δ-Implementable.*
Proof. *The 'if' part is trivial. The 'only if' follows from Lemma 3 in Appendix B.* ■

As discussed in Section 2.1, the general notion of Δ-Implementation is a useful theoretical tool. The next definition instead formalizes the ultimate objective of the designer, which is to achieve full implementation *at least* for all the beliefs in the model $\mathcal{B}$:

Definition 3 (Full $\mathcal{B}$-Implementation) *We say that d is $\mathcal{B}$-Implementable (resp., $\mathcal{B}$-DS Implementable) if it is Δ-Implementable (Δ-DS Implementable) for some $\Delta \supseteq \Delta^{\mathcal{B}}$.*

---

free' rationalizability (e.g., Bergemann and Morris, 2009). See Battigalli, Di Tillio, Grillo and Penta (2011) for a thorough analysis of the connections between Δ-rationalizability and other solution concepts.

Hence, achieving $\Delta^{\mathcal{B}}$-Implementation is the minimum objective for the designer, but this definition also accounts for the possibility of achieving full implementation for a larger set of beliefs $\Delta \supseteq \Delta^{\mathcal{B}}$, which would ensure a more robust result. In the example of Section 2.1, for instance, depending on the parameter γ, full implementation could be obtained with respect to Δ^{BF} or $\Delta^{\mathcal{B}^*}$, both of which are weaker than the designer's information (the common prior, in the example).[10]

3.2 Characterization of Full Δ-Implementability

A close inspection of part 2 of Definition 2 should make it clear that, in order to achieve truthful Δ-implementation, the truthful profile must be a mutual best response for all types, and for all conjectures allowed by the Δ-restrictions. Hence, based on Proposition 1, some notion of incentive compatibility will be necessary for full Δ-implementation. For any direct mechanism, and for every $i \in I$, let $C_i^T \subseteq \Delta(\Theta_{-i} \times M_{-i})$ denote the set of *truthful conjectures* of player i: that is, player i's conjectures that assign probability one to his opponents reporting truthfully. Formally:

$$C_i^T = \{\mu \in \Delta(\Theta_{-i} \times M_{-i}) : \mu(\{(\theta_{-i}, m_{-i}) : m_{-i} = \theta_{-i}\}) = 1\}.$$

Definition 4 *Given Δ-restrictions, a direct mechanism $\mathcal{M}$ is **Δ-incentive compatible (Δ-IC)** if for all $\theta_i \in \Theta_i$ and for all $\mu \in \Delta_{\theta_i} \cap C_i^T$ and $\theta_i' \in \Theta_i$, $EU_{\theta_i}^{\mu}(\theta_i) \geq EU_{\theta_i}^{\mu}(\theta_i')$. It is **strictly Δ-IC** if the inequality holds strictly for all $\theta_i' \neq \theta_i$.*

It is easy to verify that, if $\Delta = \Delta^{BF}$, then Δ-IC coincides with ex-post incentive compatibility (EPIC). If instead $\Delta = \Delta^{\mathcal{B}}$, and $\mathcal{B}$ is a standard type space, then Δ-IC coincides with the standard notion of interim (or Bayesian) incentive compatibility (IIC). Clearly, the weaker the Δ-restrictions, the stronger the Δ-IC condition.

We will use the following notation: For any $\sigma_{-i} : \Theta_{-i} \to \Delta(M_{-i})$ and $b_i \in \Delta(\Theta_{-i})$, we let $\mu^i(b_i, \sigma_{-i}) \in \Delta(\Theta_{-i} \times M_{-i})$ denote the conjectures derived from b_i and σ_{-i}. We will let σ_i^* denote the truthful strategy ($\sigma_i^*(\theta_i) = \theta_i$ for every θ_i) and write $\mu^*(b_i)$ for $\mu^i(b_i, \sigma_{-i}^*)$. Hence, for any $\mathcal{B}$, $\Delta^{\mathcal{B}}$-IC can be written as: $\forall i, \forall \theta_i, \forall b_i \in B_{\theta_i}, \forall \theta_i' : EU_{\theta_i}^{\mu^*(b_i)}(\theta_i) \geq EU_{\theta_i}^{\mu^*(b_i)}(\theta_i')$. Also, let $\Sigma_i : \Theta_i \to 2^{\Theta_i} \backslash \emptyset$ denote an arbitrary non-empty valued correspondence from i's types to his reports. We will write $\sigma_i \in \Sigma_i$ to signify that σ_i is a selection from Σ_i, that is $\sigma_i(\theta_i) \in \Sigma_i(\theta_i)$ for each θ_i. Similarly, we will refer to strategies as selections from R_i^{Δ}.

We provide next a general characterization of Δ-Implementation, which is based on the following property of a mechanism:

Definition 5 *Let the Δ-restrictions be non-behavioral, and $\mathcal{B}' = ((B'_{\theta_i \in \Theta_i})_{i \in I}$ such that $\Delta \equiv \Delta^{\mathcal{B}'}$. A direct mechanism is **Δ-contractive** if for any $\Sigma \neq \sigma^*$ there exists $i \in I, \theta_i \in \Theta_i$, $m_i \in \Sigma_i(\theta_i)$ and $\nu_i \in \Delta(M_i)$ such that: $EU_{\theta_i}^{\mu^i(b_i, \sigma_{-i})}(\nu_i) > EU_{\theta_i}^{\mu^i(b_i, \sigma_{-i})}(m_i)$ for all $\sigma_{-i} \in \Sigma_{-i}$ and all $b_i \in B'_{\theta_i}$.*

Theorem 1 *If the Δ-restrictions are non-behavioral, a responsive allocation rule is (fully) Δ-Implementable by a direct mechanism if and only if there exists a strictly Δ-IC and Δ-Contractive mechanism that truthfully Δ-implements it.*

Proof. *(See Appendix B).* ■

[10] The distinction between the maintained assumptions on beliefs over the environment and the beliefs with respect to which implementation is achieved is not completely new to the literature, though it typically remains implicit. For instance, within the partial implementation literature, ex-post incentive compatibility is often sought even in common prior environments. See, for instance, in Myerson (1981) and Cremer and McLean (1985,1988).

As the Δ-restrictions are varied, Δ-contractivity is related to several notions of monotonicity in the literature on implementation. In particular, in the Δ-restrictions are vacuous, then it coincides with Bergemann and Morris' (2009a) contraction property; if $\Delta = \Delta^{\mathcal{B}}$ and $\mathcal{B}$ is a Bayesian model, then Δ-contractivity is closely related to Oury and Tercieux's (2012) ICR-monotonicity, which in turn is related to robust monotonicity (Bergemann and Morris, 2011) and to Bayesian monotonicity (Jackson (1991) and Postlewaite and Schmeidler (1986)). Discussing such connections requires the introduction of extra concepts and notation. Furthermore, while conceptually important, these notions of monotonicity are not particularly suited to provide insights on the design of transfers for full implementation. We thus postpone that discussion to Section 6, and focus instead on more insightful sufficient conditions for Δ-contractivity of a Δ-IC mechanisms. These conditions will have a clear interpretation: namely, that bounding the strength of the *strategic externalities* is key to ensure Δ-contractivity, hence Δ-Implementation. To this end, we introduce the notion of a *nice mechanism*. Nice mechanisms extend the idea of *nice games* to the incomplete information games induced by a direct mechanism (possibly accompanied by exogenous Δ-restrictions).[11] Besides allowing intuitive and easy-to-check conditions for Δ-contractivity, nice mechanisms are particularly useful to instruct the design of contractive mechanisms through transfers, which we turn to in Section 4.

3.3 Conditions for Δ-Contractivity in 'Nice' Mechanisms

Consider a belief-free game with incomplete information, $G = \left\langle I, (\Theta_i, M_i, U_i)_{i \in I} \right\rangle$, where I is the set of players, Θ_i is the set of i's payoff types, $M_i \subseteq \mathbb{R}$ is the set of i's actions and $U_i : M \times \Theta \to \mathbb{R}$ is the payoff function of player i. We say that a game is 'smooth' if payoff functions are twice continously differentiable. We let $D_{jk} U_i (m, \theta)$ denote the second order partial derivative of i's payoff with respect to strategies of players j and k:

$$D_{jk} U_i (m, \theta) := \frac{\partial^2 U_i (m, \theta)}{\partial m_j \partial m_k}.$$

(Given the maintained assumptions on the environment, smoothness of $G^{\mathcal{M}}$ is guaranteed if the transfers in $\mathcal{M}$ are smooth functions.)

Definition 6 *Fix a set of Δ-restrictions $\Delta = (\Delta_{\theta_i})_{\theta_i \in \Theta_i, i \in I}$. Game $G = \left\langle I, (\Theta_i, M_i, U_i)_{i \in I} \right\rangle$ is* ***'nice' with respect to Δ (or Δ-nice)*** *if it is smooth and for every $i \in I$, $\theta_i \in \Theta_i$ and $\mu \in \Delta_{\theta_i}$, the expected payoff function $EU^{\mu}_{\theta_i} : M_i \to \mathbb{R}$ is strictly concave. A mechanism $\mathcal{M}$ is Δ-nice if $G^{\mathcal{M}}$ is Δ-nice. For convenience, we will use the term 'nice' instead of 'Δ^{BF}-nice'.*

The next proposition provides sufficient conditions for a Δ-IC nice mechanism to be Δ-contractive (hence, by Theorem 1, to guarantee truthful Δ-Implementation.)

Theorem 2 *Let $\mathcal{M}$ be a Δ-nice and Δ-IC direct mechanism. Then: $\mathcal{M}$ is Δ-contractive if one of the following holds:*

[11] The idea of 'nice' games was introduced by Moulin (1984) for games with complete information. For a recent applcation, see Weinstein and Yildiz (2011). A general analysis of nice games with incomplete information and Δ-restrictions is provided in Ollár and Penta (2014).

1. *(Δ-**Self Determination**) for each agent i, for all θ_i, for all $\mu \in \Delta_{\theta_i}$, and for all $m_i, m_i' \in M_i$*

$$\left| \int_{\Theta_{-i} \times M_{-i}} D_{ii} U_i \left(m_i', m_{-i}, \theta_i, \theta_{-i}\right) d\mu \right| > \int_{\Theta_{-i} \times M_{-i}} \sum_{j \neq i} \left| D_{ji} U_i \left(m_i, m_{-i}, \theta_i, \theta_{-i}\right) \right| d\mu. \quad (6)$$

2. *(**Ex-Post Self-Determination**) for each agent i, for all $\theta \in \Theta$, for all $m \in M$, and for all for all $m_i' \in M_i$,*

$$\left| D_{ii} U_i \left(m_i', m_{-i}, \theta\right) \right| > \sum_{j \neq i} \left| D_{ji} U_i \left(m, \theta\right) \right|. \quad (7)$$

Proof. *(See Appendix B).* ■

To understand the meaning of inequalities (6) and (7), consider the first-order condition of the optimization problem of type θ_i, given conjectures $\mu \in \Delta_{\theta_i}$: $\int_{\Theta_{-i} \times M_{-i}} \frac{\partial U_i}{\partial m_i} \left(m_i^*, m_{-i}, \theta_i, \theta_{-i}\right) d\mu = 0$. Because of the the strict concavity assumption implicit in the definition of nice mechanism, this condition is both necessary and sufficient for $m_i^* \in int\left(\Theta_i\right)$ to be a best response to the conjectures $\mu \in \Delta_{\theta_i}$. The second derivative $D_{ji} U_i \left(m_i, m_{-i}, \theta_i, \theta_{-i}\right)$ therefore captures how the report of player j affects the best response of player i, hence (for $j \neq i$) j's 'strategic externalities' on i. Both conditions (6) and (7) require the 'own effect' (the LHS of the equations) to be stronger than the opponents' effects, considered jointly (the RHS of the equations). These conditions therefore capture the idea that strategic externalities should not be too large, extending the main insight underlying the analogous condition in Moulin (1984).[12]

The difference between the two conditions is that the first requires the Self-determination property to hold for all beliefs that are consistent with the Δ-restrictions, while the second is an ex-post requirement. Clearly, the two conditions coincide if $\Delta = \Delta^{BF}$. In general, however, condition (7) is stronger than (6), hence the theorem could be stated in terms of the latter alone. We present (7) nonetheless because it is often easier to check in applications.

4 Designing Contractive Mechanisms through Transfers

In accordance with the literature on implementation, Theorem 1 provides a full characterization of the general properties of the mechanisms that ensure full implementation, but it is not very helpful for understanding how the transfers should be designed to ensure Δ-contractivity. Theorem 2, on the other hand, suggests that nice mechanisms can be used to guarantee Δ-contractivity, provided that the strategic externalities are adequately bounded. In the following we exploit this insight to explicitly construct transfers that achieve full implementation. Assumptions 1 and 2 will be maintained throughout this Section.

In Section 4.1 we consider belief-free implementation, that is, when the $\Delta = \Delta^{BF}$. This is the most demanding notion Δ-Implementation, and in many situations it would not be possible. When

[12] From a technical viewpoint, however, the extension to incomplete information is not straightforward, as Moulin's techniques do not apply to multidimensional strategy spaces. Moulin's (1984) condition guarantees uniqueness of rationalizability in nice games with complete information. Special versions of that condition underly uniqueness results for supermodular games (Milgrom and Roberts (1990)) and submodular games (Jensen (2005)). Moulin (1984) also points out that in nice games, contractive best responses result in stability of the unique rationalizable outcome with respect to best response dynamics.

possible, however, it is convenient to adopt a mechanism that achieves it, because it entails full robustness of the result. We introduce the *canonical transfers*, which characterize the mechanisms that achieve belief-free implementation (whenever possible).[13] Hence, if the canonical transfers induce overly strong strategic externalities, belief-free implementation is impossible. Full implementation may still be possible if information about beliefs is used. In Section 4.2 we introduce a natural class of Δ-restrictions (the '*moment conditions*'), which are particularly suited to design transfers for full implementation. The transfers for full implementation are obtained adding a belief-based term to the canonical transfers, to reduce the underlying strategic externalities. The resulting mechanism is 'nice', and full implementation follows from Theorem 2.

4.1 Canonical Transfers and Belief Free Implementation

In this Section we consider the most demanding notion of robustness, that is when the Δ-restrictions are vacuous. Since Δ-IC coincides with ex-post incentive compatibility, Theorem 1 implies that strict ex-post incentive compatibility (EPIC) is necessary for belief-free truthful implementation (see also Bergemann and Morris, 2009). We thus focus on the question of how to design transfers that fully implements d, if possible. Consider the following transfers: for each $i \in I$ and $m \in \Theta$, let

$$t_i^*(m) = -v_i(d(m), m) + \int^{m_i} \left. \frac{\partial v_i(d(s_i, m_{-i}), \theta_i, m_{-i})}{\partial \theta_i} \right|_{\theta_i = s_i} ds_i. \tag{8}$$

We will refer to $t^* = (t_i^*(\cdot))_{i \in I}$ as the '*canonical transfers*', and to the direct mechanism $\mathcal{M} = (d, t^*)$ as the '*canonical mechanism*'. In the canonical mechanism, agents pay their valuation as entailed by the reports profile (treated as truthful) minus the '*total* own preference effect'. This way, agents' payments fully coincide with the 'total allocation effect' of their report, given the opponents'.[14] Canonical transfers generalize several mechanisms: if d is the efficient allocation rule, then t^* coincides with the VCG transfers; in auction environments, it specializes to the incentive compatible auction of Myerson (1981) for private values and Li (2013) and Roughgarden and Talgam-Cohen (2013) for interdependent values. Proposition 2 below shows that the canonical transfers characterize the direct mechanisms that achieve belief-free full implementation. The proof of this result is based on the following Lemma, which generalizes analogous results for the above mentioned special cases. We report it here because it has intrinsic interest from the viewpoint of partial implementation (proofs are in Appendix C):

[13] The exercise in Section 4.1 is similar to that of Bergemann and Morris (2009a). The important difference is that here we focus on *how to design* transfers for full implementation of a given allocation rule. In Bergemann and Morris (2009a) the transfer scheme was given, implicit in the the social choice function (cf. Section 6).

[14] Let $\varpi_i(\theta) \equiv v_i(d(\theta), \theta)$ and consider its derivative with respect to θ_i,

$$\frac{\partial \varpi_i(\bar{\theta})}{\partial \theta_i} = \frac{\partial v_i(d(\bar{\theta}), \bar{\theta})}{\partial \theta_i} + \frac{\partial v_i(d(\bar{\theta}), \bar{\theta})}{\partial x} \cdot \frac{\partial d(\bar{\theta})}{\partial \theta_i}.$$

The first term represents the 'own preference effect': the variation of i's valuation due to θ_i, holding $d(\bar{\theta})$ constant. The second term is the 'allocation effect': the variation of i's valuation at $\bar{\theta}$, when the allocation changes due to a change in the reported θ_i. Integrating both terms with respect to θ_i, we obtain that ϖ_i can be decomposed as

$$\varpi_i(\bar{\theta}) = \int^{\bar{\theta}_i} \left. \frac{\partial v_i(d(s_i, \bar{\theta}_{-i}), \theta_i, \bar{\theta}_{-i})}{\partial \theta_i} \right|_{\theta_i = s_i} ds_i + \int^{\bar{\theta}_i} \frac{\partial v_i(d(s_i, \bar{\theta}_{-i}), s_i, \bar{\theta}_{-i})}{\partial x} \cdot \frac{\partial d(s_i, \bar{\theta}_{-i})}{\partial \theta_i} ds_i,$$

where the first term is the '*total* preference effect' and the second is the '*total* allocation effect'. Rearranging terms, the canonical transfers can be seen as the negative of the total allocation effect, given the opponents'.

Lemma 1 *Suppose that $\mathcal{M} = (d,t)$ is EPIC and t is differentiable. Then, for every i and for every m, there exists a function $\tau_i : \Theta_{-i} \to \mathbb{R}$ such that $t_i(m) = t_i^*(m) + \tau_i(m_{-i})$.*

Proposition 2 *Allocation rule d is belief-free implementable by a differentiable direct mechanism if and only if the canonical mechanism is belief-free truthfully implementable.*

In many environments of economic interest (e.g., environments with 'single crossing' preferences, as in Section 5) the canonical mechanism induces a nice game. Hence, combining our results, it follows that if in such environments ex-post incentive compatibility is possible, full implementation can only fail due to the canonical mechanism inducing overly strong strategic externalities. We provide next a measure of such strategic externalities. For any $i \in I$, let $V_i : \Theta \times \Theta \to \mathbb{R}$ be such that for any $(m,\theta) \in \Theta \times \Theta$:

$$V_i(m,\theta) := \left(\frac{\partial v_i(d(m),\theta)}{\partial d} - \frac{\partial v_i(d(m),m)}{\partial d} \right) \frac{\partial d(m)}{\partial \theta_i}.$$

For every $i \in I$, define the '*contractivity gap*' as:

$$CG_i := \max_{\theta,m} \left(\sum_{j \neq i} \left| \frac{\partial V_i(m,\theta)}{\partial m_j} \right| - \left| \frac{\partial V_i(m,\theta)}{\partial m_i} \right| \right). \tag{9}$$

Corollary 1 *Suppose that the canonical mechanism is nice. Then: If the allocation rule is EPIC but not belief-free fully implementable, then $CG_i > 0$ for some i.*

To understand this result, notice that $V_i(m,\theta)$ is nothing but the derivative of the ex-post payoff function of the canonical mechanism with respect to i's type, evaluated at state θ, when the reported profile is m. The 'contractivity gap' therefore measures the maximal difference between the opponents' ability to jointly affect this derivative and player i's own effect, evaluated across all possible combinations of states and reports. Hence, $CG_i < 0$ means that i's own effect on the first-order condition of the canonical mechanism always dominates the combined strategic externalities at all states and reports. The result then follows from Theorem 2.

4.2 Full Implementation via Moment Conditions

In environments in which ex-post incentive compatibility is possible and the canonical mechanism is nice, such as environments with single-crossing preferences, failure to achieve belief-free implementation is due to the existence of some positive contractivity gaps. In these cases, information about beliefs may be useful to restore the contraction property that for full implementation. In general, however, ex-post incentive compatibility may also be problematic.[15] In that case, information about beliefs should be used to ensure both contractiveness and incentive compatibility. In the following we consider such general case.

4.2.1 Moment Conditions

The following concept will play an important role in our analysis:

[15] That the very notion of EPIC is too restrictive is a well-known criticism to the belief-free approach. See, for instance, Jehiel et al. (2010).

Definition 7 *A* $\mathcal{B}$-consistent moment condition *is defined by a collection of functions* $(G_i, f_i)_{i\in I}$ *s.t. for every* $i \in I$*,* $G_i : \Theta_{-i} \to \mathbb{R}$ *and* $f_i : \Theta_i \to \mathbb{R}$ *are twice continuously differentiable and for every* i*,* θ_i *and* $b_i \in B_{\theta_i}$*:*

$$\int_{\Theta_{-i}} G_i(\theta_{-i}) \cdot db_i = f_i(\theta_i). \tag{10}$$

We let $\varrho(\mathcal{B})$ *denote the set of* $\mathcal{B}$*-consistent moment conditions.*

In words, a $\mathcal{B}$-consistent moment condition represents the idea that, given $\mathcal{B}$, it is common belief that agent i's conditional expectation (conditional on θ_i) of some moment $G_i(\theta_{-i})$ of the distribution of the opponents' types is described by some function f_i of i's type. Moment conditions are natural objects in a variety of settings. The following examples illustrate how moment conditions are implicit in standard models.

Example 1 (Fundamental Value Models) *Consider a model in which types can be decomposed into a fundamental component and a noise component, i.e.* $\theta_i = \theta_0 + \varepsilon_i$ *where* θ_0 *comes from a normal distribution and* ε_i*'s are i.i.d. across agents and independent of* θ_0*. This model entails many moment conditions. For example, irrespective of further details about the involved distributions, it is common knowledge that* $E(\theta_l - \theta_k|\theta_i) = 0$ *for any* θ_i *and* $l, k \neq i$*. This is represented by setting* $G_i(\theta_{-i}) = \theta_l - \theta_k$ *for some* $l, k \neq i$ *and* $f_i(\theta_i) = 0$ *for any* θ_i*. Examples for such information models include financial models with intrinsic value (e.g., the information structure in Grossman and Stiglitz (1980) and in Hellwig (1980)) and common value auctions.*

Example 2 (Unobserved Heterogeneity) *Consider an auction model in which types* θ_i *are i.i.d. draws from a distribution* F_η*, where* η *is a parameter drawn from some distribution* H*. The realization of* η *is observed by the agents but not by the designer. Then, the moment condition with* $G_i(\theta_{-i}) = \theta_k - \theta_l$ *and* $f_i(\theta_i) = 0$ *holds in this environment, regardless of the specification of* F *and* H*.*

Example 3 (Spatial Values) *Consider an environment with two distinct groups of agents (e.g., by geographic location, technology, etc.). Types for the groups are drawn from distribution* F*, with mean* E_F*. Agents are independently assigned to group 1 with probability* p*. Agents inherit the type of the group. An agent's group and type are his private information. This information structure admits the moment equation* $E(\theta_j|\theta_i) = p(i)\theta_i + (1 - p(i)) E_F(v_j)$*, where* $p(i) = p$ *if* i *belongs to group* 1*, and* $(1 - p)$ *otherwise. The moment condition thus obtains setting* $G_i(\theta_{-i}) = \theta_j$ *for some* $j \neq i$ *and* $f_i(\theta_i) = p(i)\theta_i + (1 - p(i)) E_F$*. This is the information structure considered in Ausubel and Baranov (2010).*

Moment conditions arise even more naturally in many real-world problems of mechanism design. In these problems, knowledge of the expected value of some conditional moments of the distribution is a much more basic and realistic kind of information than the one typically assumed by the standard approach, which requires common knowledge of the prior distribution. Consider the following examples:

Example 4 (Estimation-based Conditions) *Consider a practical situation, in which past data facilitate conditional predictions of the opponents' types in the form of linear regressions. The linear regression gives moment equations, with statistics* $G_{-i}(\theta_{-i}) = \theta_j$ *for* $j \neq i$ *and* $f_i(\theta_i) = a_i\theta_i + b_i$*.*

Alternatively, it could be that past data only report aggregate statistics of the distributions, so that only conditional expectations of the average of the opponents' types may be allowed. In this case, a moment equation is given by $G_i(\theta_{-i}) = \frac{1}{n-1}\sum_{j\neq i}\theta_j$, and so on. In fact, any description of the environment based on some econometric method takes the form of conditional expectation of some moment of the distribution, rather than a 'common prior'. In these cases, the very belief-restrictions $\mathcal{B}$ can be specified as the set of all beliefs consistent with such moment conditions, taken as a primitive.

Observe that in general *any* belief restriction entails common knowledge of some moment conditions (that is, $\varrho(\mathcal{B}) \neq \emptyset$ for any $\mathcal{B}$). At a minimum, condition (10) is trivially satisfied for any constant functions $G_i(\cdot) = f_i(\cdot) = \bar{y}$. If $\mathcal{B}$ is vacuous (i.e., in a belief-free environment) only such trivial moment restrictions are commonly known. The stronger the belief-restrictions, hence the smaller the sets $\mathcal{B}$, the richer the set of moment conditions. In fact, the following remark is immediate:

Remark 1 *The $\varrho(\cdot)$ correspondence is decreasing: $\varrho(\mathcal{B}') \subseteq \varrho(\mathcal{B})$ if $\mathcal{B} \subseteq \mathcal{B}'$.*

For any $\rho = (G_i, f_i)_{i\in I} \in \varrho(\mathcal{B})$, we define the Δ-restrictions derived from ρ as $\Delta^\rho := \left(\left(\Delta^\rho_{\theta_i}\right)_{\theta_i\in\Theta_i}\right)_{i\in I}$ such that for every i and θ_i,

$$\Delta^\rho_{\theta_i} = \left\{\mu \in \Delta(\Theta_{-i}\times M_{-i}) : E^\mu\left(G_i(\theta_{-i})\right) = f_i(\theta_i)\right\}.$$

It is easy to verify that Δ^ρ satisfies Assumption 2 on the Δ-restrictions, hence $\Delta^\rho_{\theta_i} \cap C_i^T \neq \emptyset$ for all θ_i and i. Also, by construction, $\Delta^{\mathcal{B}} \subseteq \Delta^\rho$ if $\rho \in \varrho(\mathcal{B})$.

4.2.2 Designing Transfers for Uniqueness

In the following we will design transfers and provide conditions for full Δ^ρ-Implementation. Since $\Delta^{\mathcal{B}} \subseteq \Delta^\rho$ whenever $\rho \in \varrho(\mathcal{B})$, it follows that achieving Δ^ρ-Implementation for some $\rho \in \varrho(\mathcal{B})$ also guarantees $\mathcal{B}$-Implementation. Given that $\mathcal{B}$-Implementation is the ultimate goal, and that $\varrho(\mathcal{B})$ in general contains several moment conditions, the key question is to understand which $\rho \in \varrho(\mathcal{B})$ is convenient to choose in the design of the mechanism. Our conditions are formulated precisely to infrorm this choice, and provide an explicit formula for the transfers. (The proof is in Appendix C.)

Theorem 3 *Allocation rule $d : \Theta \to X$ is (Fully) $\mathcal{B}$-Implementable by a direct mechanism if there exists a $\mathcal{B}$-consistent moment condition $\rho = (G_i, f_i)_{i\in I} \in \varrho(\mathcal{B})$ that satisfies the following conditions. For all i, for all $\theta_i \in \Theta_i$, for all $m_i, m_i' \in M_i$ and for all $\mu \in \Delta^\rho_{\theta_i}$:*

1. $\int_{\Theta_{-i}\times M_{-i}} \left|\frac{\partial V_i\left(m_i', m_{-i}, \theta_i, \theta_{-i}\right)}{\partial m_i} - f_i'(m_i')\right| d\mu > \sum_{j\neq i}\int_{\Theta_{-i}\times M_{-i}}\left|\frac{\partial V_i(m_i, m_{-i}, \theta_i, \theta_{-i})}{\partial m_j} + \frac{\partial G_i(m_{-i})}{\partial m_j}\right| d\mu$
2. $\int_{\Theta_{-i}\times M_{-i}} \frac{\partial V_i(m_i, m_{-i}, \theta_i, \theta_{-i})}{\partial m_i} d\mu < f_i'(m_i)$.

Moreover, for $\rho = (G_i, f_i)_{i\in I} \in \varrho(\mathcal{B})$ that satisfies the two conditions, the following transfers guarantee Full Δ^ρ-Implementation:

$$t_i^\rho(m) = \underbrace{t_i^*(m)}_{\text{canonical transfers}} + \underbrace{G_i(m_{-i})\, m_i - \int^{m_i} f_i(s_i)\, ds_i}_{\text{moment condition-based term}}. \tag{11}$$

We also provide a stronger, 'ex-post' version of these conditions, which is often easier to check in applications:

Remark 2 *The conditions of Theorem 3 are satisfied if for all i, for all $\theta \in \Theta$, for all $m_{-i} \in M$ and for all $m_i, m_i' \in M_i$:*

1. $\left| \frac{\partial V_i\left(m_i', m_{-i}, \theta\right)}{\partial m_i} - f_i'\left(m_i\right) \right| > \sum_{j \neq i} \left| \frac{\partial V_i\left(m_i, m_{-i}, \theta\right)}{\partial m_j} + \frac{\partial G_i\left(m_{-i}\right)}{\partial m_j} \right|.$

2. $\frac{\partial V_i(m, \theta)}{\partial m_i} < f_i'\left(m_i\right)$

Theorem 3 states two properties of moment conditions that are useful to guarantee full implementation. To understand what these are, let us consider the stronger versions stated in Remark 2. First, notice that if the contractivity gap (defined in the previous Section) is negative for all i, then Condition 1 is satisfied by any trivial moment condition, in which functions f_i and G_i are constant. Since such trivial moment conditions are consistent with any belief restrictions, the full implementation result is guaranteed by the canonical mechanism in the belief-free sense of Bergemann and Morris (2009, 2011). Now, suppose that the contractivity gap is positive for some agent. Condition 1 clarifies which properties of beliefs can be used to create a wedge between the strength of the strategic externalities and the interdependence in agents' valuations: a moment condition in which the derivative of f_i has the opposite sign as $\partial V_i / \partial m_i$ can be used to increase the 'own effect', whereas moment functions on the opponents' types with derivatives that contrast the externality in the canonical mechanism can be used to weaken the 'external effects'. Condition 2 instead requires that the 'own effect' in the canonical mechanism is bounded above by the derivative of the f_i function.

To gain further insights on how these conditions contribute to the full implementation result, it is useful to consider the exact transfers that guarantee full implementation (eq. 11). Under these transfers, the first order derivatives of the expected utility of type θ_i is:

$$\frac{\partial EU_{\theta_i}^{\mu}\left(m_i\right)}{\partial m_i} = \int_{\Theta_{-i} \times M_{-i}} \left(\frac{\partial v_i\left(d\left(m\right), \theta\right)}{\partial d} - \frac{\partial v_i d\left(m\right), m}{\partial d} \right) \frac{\partial d\left(m\right)}{\partial m_i} + G_i\left(m_{-i}\right) - f_i\left(m_i\right) \; d\mu.$$

This shows that for any truthtelling conjecture $\mu \in \Delta_{\theta_i}^{\rho} \cap C_i^T$, the report $m_i = \theta_i$ is an extremal point. Notice that this does not necessarily result in Δ^{ρ}-IC, as that depends on the second order conditions as well. Condition 1 in Theorem 3, however, guarantees that the ensuing mechanism is nice, which ensures that the second order conditions are satisfied. This mechanism therefore is Δ^{ρ}-IC (hence $\mathcal{B}$-IC). The full implementation follows from the fact that Condition 1 in Theorem 3 also guarantees the Self-Determination condition of Theorem 2.

Theorem 3 is constructive in the sense that it pins down a precise design principle: the designer shall start out with the canonical transfers, and then add a new term which is based on suitable moment-conditions. 'Suitable' here means that the term added to the canonical transfers ought to guarantee niceness of the mechanism and reduce the strategic externalities. We illustrate the versatility of this general principle in the next section.

5 Applications: Single-Crossing Environments

We consider next different special cases of economic interest, with different assumptions on preferences and beliefs, to illustrate how the general results of Theorem 3 can be applied to these special settings.

An important class of environments are those that satisfy the following *single-crossing condition* (SCC): for every $i \in I$, and for every (x, θ), $\partial^2 v_i (x, \theta) / \partial x \partial \theta_i > 0$. Single-crossing conditions are known to facilitate ex-post incentive compatibility.

Lemma 2 *Suppose that the environment satisfies the single-crossing condition. Then the canonical mechanism is EPIC if and only if the allocation rule is strictly increasing: $\partial d (\theta) / \partial \theta_i > 0$ for every θ and every i.*
Proof. *(See Appendix D)* ■

This lemma generalizes many standard results on ex-post (partial) implementation. Joint with Corollary 1, this Lemma implies that in SCC-environments with increasing allocation rules, failure to achieve belief-free full implementation is possible only if the canonical mechanism is not nice or some players have positive contractivity gaps. In those cases, information about beliefs may be used to achieve full implementation. In these environments therefore there is an interesting tension between the robustness of the partial implementation result, obtained by the canonical mechanism in a belief-free sense, and the possibility of achieving full implementation: the mechanism that achieves the latter will necessarily exploit information about beliefs, and therefore reduce the robustness of the partial implementation result.

To simplify the analysis, we consider quadratic SCC-environments:

Definition 8 *An 'SCC-environment' satisfies the following conditions: (SCC.1) $0 < \partial d (\theta) / \partial \theta_i < \infty$ for every θ and every i and (SCC.2) $0 < \partial^2 v_i (x, \theta) / \partial x \partial \theta_i < \infty$ for all i, x and θ. An environment is 'quadratic' if (Q.1) for all $i, j, k \in I$ and all $(x, \theta) \in X \times \Theta$, $\partial^3 v_i (x, \theta) / \partial^2 x \partial \theta_j = \partial^3 v_i (x, \theta) / \partial x \partial \theta_j \partial \theta_k = 0$ and (Q.2) $\partial^2 d (\theta) / \partial \theta_i \partial \theta_j = 0$ for all $i,j \in I$ and $\theta \in \Theta$.*

Under these assumptions, for any $i,j \in I$ and $\theta, m \in \Theta$:

$$\frac{\partial V_i (m, \theta)}{\partial m_j} = - \left(\frac{\partial^2 v_i}{\partial d \partial \theta_j} (d (m), m) \right) \frac{\partial d (m)}{\partial \theta_i} < 0.$$

(For $j = i$, this condition ensures that the canonical mechanism is nice.)

Conditions (SCC.1) and (SCC.2) are standard assumptions for general SCC-environments. Conditions (Q.1) and (Q.2) are satisfied, for instance, by environments in which agents' valuations are (at most) quadratic functions of players' types and the allocation rule maximizes a linear functional of agents' valuations (the efficient rule would be a special case). While rather special in the context of our paper, these assumptions are extremely common in the applied literature on unit demand auctions, divisible good auctions, finance, etc. Conditions (Q.1) and (Q.2), however, are not essential to our analysis, and can be relaxed (see Section 5.3).

5.1 Common Prior Models

5.1.1 Independent Types

Independent common prior models are such that $\mathcal{B}$ is such that for every j, there exists $p^j \in \Delta(\Theta_j)$ such that for all i and for all θ_i, $B_{\theta_i} = \{\otimes_{j \neq i} p_j\} \subseteq \Delta(\Theta_{-i})$. In an independent common prior model, for any $G_i : \Theta_{-i} \to \mathbb{R}$, the condition $E(G_i(\theta_{-i})|\theta_i) = f_i(\theta_i)$ holds true with $f_i : \Theta_i \to \mathbb{R}$ s.t. $f_i' = 0$ (just by the definition of independence). Since (under the common prior assumption) G_i can be chosen freely, independence leaves us huge leeway in manipulating the external effects on the RHS of Condition 1 of Theorem 3. The LHS, on the other hand, will not be affected, as the f_i function is constant. The ex-post condition of Remark 2 therefore can be rewritten as:

$$min_{m_i \in M_i} \left| \left(\frac{\partial^2 v_i}{\partial d \partial \theta_i} (d(m), m) \right) \right| > \sum_{j \neq i} \left| \left(-\frac{\partial^2 v_i}{\partial d \partial \theta_j} (d(m), m) \right) + \frac{\frac{\partial G_i(m_{-i})}{\partial m_j}}{\frac{\partial d(m)}{\partial \theta_i}} \right| \tag{12}$$

For any i and $j \in I$, let $g_{ij}(m) \equiv -\frac{\partial^2 v_i}{\partial d \partial \theta_j}(d(m), m)$ and for any $l \in I$, consider

$$\frac{\partial g_{ij}(m)}{\partial m_l} = -\frac{\partial^3 v_i}{\partial^2 d \partial \theta_j}(d(m), m) \frac{\partial d}{\partial \theta_l}(m) - \frac{\partial^3 v_i}{\partial d \partial \theta_j \partial \theta_l}(d(m), m).$$

Condition (Q.1) in Definition 8 guarantees that these derivatives are zero, hence there exists $K_{ij} \in \mathbb{R}$ such that $g_{ij}(m) = K_{ij}$ for all m. Inequality (12) therefore simplifies to:

$$|K_{ii}| > \sum_{j \neq i} \left| K_{ij} + \frac{\frac{\partial G_i(m_{-i})}{\partial m_j}}{\frac{\partial d(m)}{\partial \theta_i}} \right|. \tag{13}$$

Since the assumption of independent common priors allows full freedom in the specification of the G_i functions, in this case we can choose the moment condition to completely neutralize the strategic externalities, setting the RHS of (13) equal to zero. That is, if $\bar{G}_i$ is such that

$$\frac{\partial \bar{G}_i}{\partial m_j}(m_{-i}) = -K_{ij} \frac{\partial d}{\partial \theta_i}(m) \text{ for all } m \text{ and } j \neq i. \tag{14}$$

Under (Q.2), this condition is satisfied by the following function:

$$\bar{G}_i(m_{-i}) = -\sum_{j \neq i} K_{ij} \frac{\partial d}{\partial \theta_i}(m) \cdot m_j \tag{15}$$

(Equations (14) and (15) are well-defined because Condition (Q.2) guarantees that $\frac{\partial d}{\partial \theta_i}(m_i, m_{-i})$ is constant in m_i). Hence, the following Proposition holds:

Proposition 3 *Full implementation is always possible in quadratic SCC-environments with independent common prior. In particular, let $\mathcal{B}$ be an independent common prior model. For any $i \in I$, let $\bar{G}_i : \Theta_{-i} \to \mathbb{R}$ be defined as in (15) and define $\bar{f}_i(\theta_i) = E(\bar{G}_i(\theta_{-i})|\theta_i) \equiv C_i$. Then the* *transfers in (11) with $\rho = \left(\bar{f}_i, \bar{G}_i\right)_{i \in I}$, that is*

$$t_i^{\rho}(m) = \underbrace{t_i^*(m)}_{\text{canonical transfers}} + \underbrace{\bar{G}_i(m_{-i})\, m_i}_{\substack{\text{compensation for} \\ \text{strategic externality} \\ \text{(depends on } m_i \text{ and } m_{-i})}} - \underbrace{\int^{m_i} \bar{f}_i(s_i)\, ds_i}_{\substack{\text{belief-based adjustment:} \\ \text{expected compensation} \\ \text{(only depends on } m_i)}}\,, \tag{16}$$

guarantee full Δ^{ρ}-implementation. Since $\rho = \left(\bar{f}_i, \bar{G}_i\right)_{i \in I} \in \varrho(\mathcal{B})$, full $\mathcal{B}$-Implementation follows. Furthermore, in this mechanism truthful revelation is strictly Δ^{ρ}-dominant for every type.

To understand the logic of the mechanism, first notice that the function $\bar{G}_i(m_{-i})$ constructed above is nothing but a measure of the strategic externality the other players impose on i.[16] The transfers in (16) therefore are such that, starting from the canonical mechanism, player i is compensated for the total strategic externality he is subject to. The last term in (16) is nothing but the expected value of such compensation, when i reports m_i. This term is added to prevent the agent from misreporting his type, in order to inflate the implied compensation for the strategic externality. Thus, the compensation for the strategic externality ensures that the mechanism is 'contractive', but possibly upsets incentive compatibility; incentive compatibility is restored by the last term, which only depends on i's report, and therefore it entails no further strategic externality. Full implementation follows.

5.1.2 Affiliated Types

Suppose that $\Theta_i = \Theta_j$ for all i and j, and that types are *affiliated* (Milgrom and Weber, 1982). For any i, let $\bar{G}_i : \Theta_{-i} \to \mathbb{R}$ be defined as in equation (15). Under the maintained assumptions for quadratic SCC-environments, $\frac{\partial \bar{G}_i(m)}{\partial m_j} > 0$ for all m and $j \neq i$. If types are affiliated, Theorem 5 in Milgrom and Weber (1982) implies that $\mathbb{E}\left(\bar{G}_i(\theta_{-i})\,|\theta_i\right)$ is an increasing function of θ_i. Hence, letting $\bar{f}_i(\cdot) \equiv \mathbb{E}\left(\bar{G}_i(\theta_{-i})\,|\cdot\right)$, we have that the moment condition $\rho = \left(\bar{G}_i, \bar{f}_i\right) \in \varrho(\mathcal{B})$ satisfies $\bar{f}_i' > 0$ for all i. By construction, $\bar{G}_i$ is such that the RHS of Condition 1 in Theorem 3 is equal to zero. Since $\bar{f}_i' > 0$, SCC implies that the LHS of the same condition is (strictly) positive. The condition for full Δ^{ρ}-Implementation therefore is satisfied, which implies the following:

Proposition 4 *Full implementation is always possible in quadratic SCC-environments with affiliated types. In particular, for any $i \in I$, let $\bar{G}_i : \Theta_{-i} \to \mathbb{R}$ be defined as in (15) and and define $\bar{f}_i : \Theta_i \to \mathbb{R}$ as $f_i(\theta_i) \equiv \mathbb{E}(G_i(\theta_{-i})\,|\theta_i)$ for each θ_i. Then the transfers in (11) with $\rho = \left(\bar{f}_i, \bar{G}_i\right)_{i \in I}$, guarantee full Δ^{ρ}-implementation. Since $\rho = \left(\bar{f}_i, \bar{G}_i\right)_{i \in I} \in \varrho(\mathcal{B})$, full $\mathcal{B}$-Implementation follows. Furthermore, in this mechanism truthful revelation is strictly Δ^{ρ}-dominant for every type.*

[16] Substituting for the constant $K_{ij} = -\frac{\partial^2 v_i}{\partial d \partial \theta_j}(d(m), m)$ in (15), we obtain:

$$\bar{G}_i(m_{-i}) = \sum_{j \neq i} \left(\frac{\partial^2 v_i}{\partial d \partial \theta_j}(d(m), m) \cdot m_j \right) \cdot \frac{\partial d}{\partial \theta_i}(m).$$

The term in parenthesis represents the effect of j's report on i's marginal utility for the public good, and is multiplied by the impact of i's report on the allocation. Overall, this is the total strategic externality that player i is subject to, for each increment of his own report.

5.1.3 Equivalence of EPIC and DS-Implementation

The construction above can also be used to derive an interesting equivalence between ex-post and (interim) dominant strategy incentive compatibility:

Proposition 5 *Under assumptions Q.1-2 and SCC.2, in a common prior environment with independently distributed or affiliated types, an allocation function is EPIC-Implementable if and only if it is interim DS-implementable.*
Proof. *(See Appendix D.)* ■

The proof of this result is very simple. First, we show that an allocation rule is iDSIC only if it is increasing. The 'only if' part of the proposition then follows immediately from Lemma 2. The 'if' direction follows from the discussion above: if the allocation rule is EPIC, Lemma 2 implies that it is increasing, hence condition SCC.1 is satisfied. Propositions 3 and 4 in turn imply that the allocation rule is iDSIC.

This result is somewhat related to an important equivalence result provided by Manelli and Vincent (2010), recently generalized by Gershkov et al. (2013). Those results show that, in Bayesian environments with *private values*, for any Bayesian Incentive Compatible mechanism there is an 'equivalent' mechanism that is Dominant Strategy Incentive Compatible. Given the restriction to private values, one way of interpreting this result is as an equivalence between 'partial' and 'full' implementation in direct mechanisms. From this viewpoint, Proposition 5 can be seen as a generalization of that insight to Bayesian environments with *interdependent values*.[17] We should point out, however, that Manelli and Vincent's notion of 'equivalence' is different from ours. In particular, Manelli and Vincent (2010) define two mechanisms as 'equivalent' if they deliver the same interim expected utilities for all agents and the same ex-ante expected social surplus. Here instead we maintain the traditional notion of equivalence, which requires that the mechanisms induce the same ex-post allocation. (As shown by Gershkov et al. (2013), equivalence results à la Manelli and Vincent do not extend beyond environments with linear utilities and independent types.)

5.2 Common Knowledge of 'Conditional Moments'

As discussed in Section 4.2, the designer's information about the distribution of types in the environment often takes the form of conditional moment condition. In that case, the model of beliefs $\mathcal{B}$ should specify such primitive information of the designer, rather than imposing a single common prior. For instance, suppose that the only information available to the designer consists of a set of linear regressions estimated on the environment. Then, the model of beliefs $\mathcal{B}$ is just a set of moment conditions (cf. Example 4 in Section 4.2). Theorem 3 is useful for these settings too. As an illustration, suppose that the only information available to the designer concerns the conditional averages $\mathbb{E}(\theta_j|\theta_i)$. For each i, j, let $\varphi_{ij} : \mathbb{R} \to \mathbb{R}$ be such that, for each $\theta_i \in \Theta_i$, $\varphi_{ij}(\theta_i) := \mathbb{E}(\theta_j|\theta_i)$. For simplicity, assume that these functions φ_{ij} are differentiable. Then, the designer's information is represented by belief restrictions $\mathcal{B}^{ave} = \left(\left(B^{ave}_{\theta_i}\right)_{\theta_i \in \Theta_i}\right)_{i \in I}$ such that $B^{ave}_{\theta_i} = \{\beta \in \Delta(\Theta_{-i}) : \mathbb{E}_\beta(\theta_j) = \varphi_{ij}(\theta_i)\}$, for each $i \in I$ and $\theta_i \in \Theta_i$.

[17] We are grateful to Stephen Morris for this insight.

For any i, and for any $G_i : \Theta_{-i} \to \mathbb{R}$, define $f_i^G : \Theta_i \to \mathbb{R}$ such that $f_i^G(\theta_i) = G_i((\varphi_{ij}(\theta_j))_{j\in I\setminus\{i\}})$. It is easy to see that, for any collection of *affine* functions $G_i : \Theta_{-i} \to \mathbb{R}$, the collection $\left(G_i, f_i^G\right)_{i\in I}$ defines a moment condition consistent with $\mathcal{B}^{ave}$. Next notice that, in quadratic SCC environments, the function $\bar{G}_i : \Theta_{-i} \to \mathbb{R}$ (eq. 15) is an increasing linear function of the types of i's opponents. Hence, $\rho = \left(\bar{G}_i, f_i^{\bar{G}}\right)_{i\in I} \in \varrho(\mathcal{B}^{ave})$. Furthermore, if the functions φ_{ij} are non-decreasing, then $f_i^{\bar{G}}$ is also non-decreasing. Then, the next result follows from Theorem 3 for the same reasons as Proposition 4 did:

Proposition 6 *Consider a quadratic SCC-environment, and $\mathcal{B}^{ave}$ is such that, for each i, j, the functions φ_{ij} are non-decreasing. Then, the mechanism defined in (11) with $\rho = \left(\bar{G}_i, f^{\bar{G}}\right)_{i\in I}$ and $\bar{G}_i$ defined as in (15), ensures Δ^ρ-DS Implementation. Since $\rho = \left(\bar{G}_i, f_i^{\bar{G}}\right)_{i\in I} \in \varrho(\mathcal{B}^{ave})$, $\mathcal{B}^{ave}$-DS Implementation follows.*

5.3 Extensions

The logic of Propositions 3, 4 and 5 extends well beyond the cases of common prior models with independent or affiliated types. To see this, notice that for $\bar{G}_i : \Theta_{-i} \to \mathbb{R}$ defined as in equation (15), the maintained assumptions for quadratic SCC-environments guarantee that the RHS of Condition 1 in Theorem 3 is equal to zero. Affiliation or independence further guarantee that the conditional moment $E\left(\bar{G}_i(\theta_{-i})|\theta_i\right)$ is (weakly) increasing in θ_i, hence the moment condition $\rho = \left(\bar{f}_i, \bar{G}_i\right)_{i\in I}$ can be used with no risk of upsetting the LHS of that condition. This argument, however, remains valid whenever $E\left(\bar{G}_i(\theta_{-i})|m_i\right) < \frac{\partial V_i(m,\theta)}{\partial m_i}$ for all m, which ensures that both conditions for contractions are still statisfied by $\rho = \left(\bar{f}_i, \bar{G}_i\right)_{i\in I}$. In Proposition 6, the assumption that functions φ_{ij} are non decreasing plays the same role as the assumptions of independence and affiliation in the common prior models, and can be weakened similarly.

Conditions (Q.1) and (Q.2) in may also be weakened in Propositions 3, 4 and 6. In the argument above, we used these conditions to ensure that $\frac{\partial V_i(m,\theta)}{\partial m_j} < 0$ and that $\bar{G}_i$ could be designed to completely neutralize the strategic externalities of the canonical mechanism. Clearly, $\frac{\partial V_i(m,\theta)}{\partial m_j} < 0$ can be guaranteed by weaker conditions. If (Q.2) is violated, however, then we may not be able to choose $\bar{G}_i$ to completely offset the strategic externalities: if $\frac{\partial d}{\partial \theta_i}(m)$ varies with m_i, the same $\bar{G}_i(\cdot)$ cannot set the RHS identically equal to zero. But if $\left|\frac{\partial^2 d}{\partial\theta_i\partial\theta_j}\right|$ and the variations of the valuations' concavity are small relative to $\left|\frac{\partial V_i(s_i,m_{-i},\theta)}{\partial m_i}\right|$, then $\bar{G}_i$ can be chosen so that the RHS is bounded above by $\left|\frac{\partial V_i(s_{i,},m_{-i},\theta)}{\partial m_i}\right|$, and the argument goes through essentially unchanged. The only difference is that now the full Δ^ρ-Implementation results would not be obtained with one round of Δ-rationalizability only. That is, we would have Full Δ-Implementation, but not in 'Δ-dominant' strategies.

6 Δ-Contractivity and Monotonicity

To compare Δ-contractivity to the related notions introduced in the implementation literature, it is useful to reformulate some of our concepts in the context of that literature, which typically considers general (non quasilinear) environments. In such frameworks, the space of outcomes is given by an abstract set A, agents' preferences are defined as $u_i : A \times \Theta \to \mathbb{R}$, and social choice functions are $f : \Theta \to A$. A mechanism is a tuple $\left((M_i)_{i \in I}, g\right)$ s.t. $g : M \to A$, and it is direct if $g = f$ and $M_i = \Theta_i$ for all i. To see the connection with our framework, let $A = X \times \mathbb{R}^n$. Given $d : \Theta \to X$, any direct mechanism $\mathcal{M} = (d, t)$ in our setting induces a function $f^{\mathcal{M}} : \Theta \to A$ in the general setting: $f^{\mathcal{M}}$ is such that $f^{\mathcal{M}}(\theta) = (d(\theta), t(\theta))$ for each $\theta \in \Theta$. That $\mathcal{M}$ truthfully Δ-implements d in our model therefore is the same as saying that $f^{\mathcal{M}}$ is truthfully Δ-Implemented by a direct mechanism in the general setting. For the sake of the argument, we will slightly abuse notation as follows: for any $\alpha \in \Delta(A)$, we write $u_i(\alpha, \theta)$ instead of $\int u_i(a, \theta)\, d\alpha(a)$ and for any $f : \Theta \to A$ and $\nu_i \in \Delta(\Theta_i)$ we write $f(\nu_i, \theta_{-i})$ for the measure over A induced by v_i, given θ_{-i}.[18]

Most of the literature on full implementation does not impose the restriction that it should be achieved through direct mechanism. An important exception is Bergemann and Morris (2009a), who consider rationalizable implementation through direct mechanisms in belief-free environments, under some restrictions on agents' preferences $u_i : A \times \Theta \to \mathbb{R}$ (in particular, preferences are required to admit monotone aggregators and satisfy a strict single crossing property). As already mentioned, belief-free environments in our model correspond to the special case in which $\Delta = \Delta^{BF}$. It is easy to shown that, under Bergemann and Morris' (2009a) extra restrictions on preferences, Δ^{BF}-contractivity is equivalent to Bergemann and Morris' (2009) contraction property.

Another special case of interest is when the $\Delta = \Delta^{\mathcal{B}}$ and $\mathcal{B}$ is such that $B_{\theta_i} = \{b_{\theta_i}\}$ for every θ_i, so that the belief-restrictions induce a standard Bayesian environment. As shown by Battigalli et al. (2011), in that case Δ-Rationalizability coincides with Dekel et al. (2007) interim correlated rationalizability (ICR). Full Δ-Implementation therefore coincides with ICR-implementation, which has been studied by Oury and Tercieux (2012). Oury and Tercieux (2012) introduce the notion of ICR-monotonicity, which is a necessary and (with minimal extra assumptions) sufficient condition for ICR-Implementation. ICR-monotonicity is a Bayesian version of a notion introduced by Bergemann and Morris (2011) in the ex-post environment, and strengthens Postlewaite and Schmeidler (1986) and Jackson's (1991) Bayesian monotonicity to account for the weaker solution concept. For the sake of the argument, we recall both Jackson's and Oury and Tercieux's notions of monotonicity.

Definition 9 *A* deception *is a collection of mappings $\Sigma = (\Sigma_i)_{i \in I}$, where each $\Sigma_i : \Theta_i \to 2^{\Theta_i}$. A deception Σ is acceptable for $f : \Theta \to A$ if, for all $\theta \in \Theta$, and for all $\theta' \in \Sigma(\theta)$, $f(\theta) = f(\theta')$. A deception Σ is unacceptable if it is not acceptable.*

Definition 10 (Bayesian Monotonicity) *The SCF $f : \Theta \to A$ is Bayesian monotonic if, for every unacceptable single-valued deception $\sigma : \Theta \to \Theta$, there exist i, $\theta_i \in \Theta_i$ and $y : \Theta \to A$ such that*

$$\begin{aligned} &\int u_i\left((y \circ \sigma)(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i}(\theta_{-i}) \\ &> \int u_i\left((f \circ \sigma)(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i}(\theta_{-i}) \end{aligned}$$

while

$$\int u_i\left(y(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i}(\theta_{-i}) \leq \int u_i\left(f(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i}(\theta_{-i}).$$

[18]Formally, $f(\nu_i, \theta_{-i}) \in \Delta(A)$ such that for any measurable $E \subseteq A$, $f(\nu_i, \theta_{-i})[E] = \nu_i[\{\theta_i : f(\theta_i, \theta_{-i}) \in E\}]$.

Similar to Maskin's monotonicity, Bayesian monotonicity ensures elimination of undesirable equilibria. For instance, suppose that some unacceptable deception β is played, so that $f \circ \beta \neq f$. Under Bayesian monotonicity, there would be at least one player who, given β, has the incentives to signal the deception and induce an outcome according to function y rather than f. The second condition ensures that the same player would not have an incentive to induce y if the opponents were reporting truthfully.

Definition 11 (ICR-Monotonicity) *The SCF $f : \Theta \to A$ is ICR-monotonic if, for every unacceptable deception Σ, there exist i, $\theta_i \in \Theta_i$ and $\theta'_i \in \Sigma_i(\theta_i)$ such that, for every $\mu_i \in \Delta(\Theta_{-i} \times \Theta_{-i})$ that satisfies the properties: (i) $(\theta_{-i}, \theta'_{-i}) \in supp(\mu_i) \Rightarrow \theta'_{-i} \in \Sigma_{-i}(\theta_{-i})$; (ii) $marg_{\Theta_{-i}}\mu_i = b_{\theta_i}$; there exists $y^* : \Theta_{-i} \to \Delta(A)$ such that*[19]

$$\int u_i\left(y_i^*\left(\theta'_{-i}\right), \theta_i, \theta_{-i}\right) d\mu_i > \int u_i\left(f\left(\theta'_i, \theta'_{-i}\right), \theta_i, \theta_{-i}\right) d\mu_i \tag{17}$$

while for all $\theta_i \in \Theta_i$,

$$\int u_i\left(y_i^*\left(\theta'_{-i}\right), \theta_i, \theta_{-i}\right) db_{\theta_i}(\theta_{-i}) \leq \int u_i\left(f\left(\theta_i, \theta_{-i}\right), \theta_i, \theta_{-i}\right) db_{\theta_i}(\theta_{-i}). \tag{18}$$

ICR-monotonicity extends the idea of Bayesian monotonicity, except that rather than giving players the incentives to signal a deviation from truthful reporting, some player has the incentive to signal *sets* of non-truthful reports whenever it is the case (17), and only if it is the case (18).[20] The existence of the function $y^* : \Theta_{-i} \to \Delta(A)$ in the definition ensures that at least one type of some player would find it profitable to signal the deception and choose according to y^* rather than f.

In the definitions above, functions y^* and y need not have any relation with f beyond those specified by the above inequalities. The reason is that both notions of implementation do not require that direct mechanisms are used. Using richer message spaces, and letting the outcome function change from f to the designated y^*, then (under typically weak extra assumptions) such monotonicity conditions are also sufficient for the implementation results. If the restriction to direct mechanism is imposed, however, the 'reward function' y^* should itself be achievable as part of the direct mechanism, through a unilateral misreport of some of the players from the candidate deception.

Definition 12 *Function $y^* : \Theta_{-i} \to \Delta(A)$ is a* 'directly feasible' reward *for type θ_i if there exists $\nu_i \in \Delta(\Theta_i) : f(\nu_i, \theta_{-i}) = y^*(\theta_{-i})$ for all $\theta_{-i} \in supp(b_{\theta_i})$. Let $Y^*(\theta_i) \subseteq [\Delta(A)]^{\Theta_{-i}}$ denote the set of type θ_i's directly feasible rewards.*

This notion provides precisely the missing link between ICR-Monotonicity and Δ-Contractivity:

Definition 13 (direct ICR-Monotonicity) *The SCF $f : \Theta \to A$ is* direct ICR-monotonic *if, for every unacceptable deception Σ, there exist i, $\theta_i \in \Theta_i$ and $\theta'_i \in \Sigma_i(\theta_i)$ such that, for every*

[19] In Oury and Tercieux (2012), the reward y^* is required to induce a degenerate lottery over A (that is, $y^* : \Theta_{-i} \to A$). We allow for the possibility of non-degenerate lotteries, as done by Bergemann and Morris (2011). Clearly, this change does not affect the necessity result, as the version with non-degenerate lotteries is weaker. Degenerate lotteries in general make sufficiency arguments easier. Given the restriction to direct mechanisms, however, the weaker version with non-degenerate lotteries is also sufficient in our setting.

[20] Notice that in Definition 11 the deception Σ need not be single valued, as in Definition 10.

$\mu_i \in \Delta(\Theta_{-i} \times \Theta_{-i})$ *that satisfies the properties: (i)* $(\theta_{-i}, \theta'_{-i}) \in supp(\mu_i) \Rightarrow \theta'_{-i} \in \Sigma_{-i}(\theta_{-i})$*; (ii)* $marg_{\Theta_{-i}} \mu_i = b_{\theta_i}$*; there exists* $y^* \in Y^*(\theta_i)$ *that satisfies the inequalities in (17) and (18).*

Proposition 7 *Let* $d : \Theta \to X$ *be responsive,* $\mathcal{M} = (d, t)$ *and the* Δ*-restrictions be derived from* $\mathcal{B}$ *that satisfies* $|B_{\theta_i}| = 1$ *for every* θ_i *and* i*. Then the following are equivalent:*

1. $\mathcal{M}$ *is* Δ*-IC and* Δ*-contractive*
2. $f^{\mathcal{M}}$ *satisfies direct ICR-Monotonicity.*

Proof. *(See Appendix D.3).* ■

7 Discussion and Extensions

A central goal of the theory of mechanism design is to identify conditions under which there exist institutions (or mechanisms) which guarantee socially desirable outcomes through the decentralized interaction of self-interested economic agents. Assumptions on the form of such interactions are embodied in some solution concept, which is used to solve the game that results from the mechanism. The relevance of the theory therefore depends on the nature of the solution concepts and the mechanisms that are used to achieve these results, as well as the auxiliary assumptions that may be explicit or implicit in the model used to formalize the question.

The theory of partial implementation has been very successful in providing useful insights for the design of real world institutions (e.g. auctions, contracts, compensation schemes, etc.). The problem of multiplicity, however, is a serious limitation of this approach, both from a practical and a theoretical viewpoint. The full implementation literature overcomes this problem, but often adopting rather complicated mechanisms which have no clear economic interpretation. Such mechanisms guarantee great generality of the results, but provide little insight on how real-world institutions could be designed to avoid the problem of multiplicity. Thus, despite the stronger results, the practical relevance of the full implementation approach is more limited.

In this paper we introduced a new approach to full implementation under arbitrary restrictions on beliefs. Besides generalizing several existing results, such as those on belief-free implementation by Bergemann and Morris (2009) and on ICR-implementation by Oury and Tercioux (2012), we showed that this approach is also useful from a more applied perspective. Namely, to achieve uniqueness through the design of simple transfer schemes. Our approach therefore solves the problem of multiplicity using mechanisms that have a clear economic interpretation and are as simple as those developed by the partial implementation literature. These results show how the methodology put forward by the recent literature on robust mechanism design can be extended to overcome important limitations of the traditional approach to full implementation.

While largely inspired by the literature on belief-free mechanism design, we departed from it in important ways. The requirement that a mechanism should perform well independent of the agents' beliefs is often unnecessarily demanding. By treating belief restrictions as parameters, our approach improves on that literature, enabling us to show that often minimal assumptions on agents' beliefs may suffice to overcome the limitations of the belief-free approach. Another important difference with respect to that literature is that the main focus thus far has been on identifying conditions under which a given transfer scheme achieves implementation in a robust sense (e.g., Bergemann and Morris (2009) and Penta (2011)), offering very little guidance on how to design such a mechanism, given the objectives of the designer. The results we provided in Sections 3.2 and 6 are analogous to those provided by that literature. In Section 4, however, we went beyond such mere possibility results and pursued a more constructive approach. The transfers we propose provide explicit design principles that can easily be applied to a variety of important problems, some of which we explored in Section 5.

The notion of a nice mechanism, as well as the conditions that ensure contractivity of best replies (further developed in Ollár and Penta (2013)), are useful game theoretic concepts that extend beyond the present context. (For instance, they may be used to extend the analysis of Weinstein and Yildiz (2011) beyond complete information games). Besides ensuring uniqueness, contractivity of the best replies which we pursued has other important properties, such as (1) small sensitivity to small changes in moment conditions, (2) small sensitivity to lower levels of common belief in rationality, (3) small sensitivity to small misspecifications of the domain. All these properties point to a broader concept of robustness outlined in the Wilson doctrine. The first point, in particular, is important because by assuming common knowledge of the moment conditions, considerable robustness was given up compared to the belief-free case. Yet, the contractiveness of our design ensures that small changes or small mistakes in the moment conditions result in an outcome that is proportionately close to the desired outcome. It is important to note that 'closeness' here is in terms of the natural allocation space, as opposed to the probabilistic notion of the virtual implementation literature. Further developing these results is part of ongoing research.

The weakness of the solution concept and the generality of our framework, both in terms of preferences and belief restrictions, ensure that our results accommodate several important special cases. For instance, for the case in which the belief restrictions are derived from a standard type space, our notion of implementation coincides with ICR-Implementation. Since most of the classical literature on implementation focuses on Bayesian environments, this is an important special case of our analysis. Clearly, since ICR is a superset of Bayes Nash equilibrium (BNE), our sufficient conditions also guarantee full implementation in BNE (cf. Jackson, 1991), as well as an explicit design of the transfer scheme that achieves that. Also, Oury and Tercieux (2012) have shown an important connection between partial and full implementation. Namely, say that a social choice function is 'continuously partially implemented' at a given type t_i of the universal type space, if it is incentive compatible for all types in the neighborhood of t_i. Building on important work by Weinstein and Yildiz (2007), Oury and Tercieux (2012) show that this is possible if and only if the SCF is fully implemented in ICR. Hence, our sufficient conditions for Δ-Implementation are also sufficient for 'continuous (partial) implementation'.

Finally, we note that our construction sheds some light on at least two novel important research questions. In particular, Deb and Pai (2013) recently argued that an important topic for future research is the design of mechanisms that only use properties of the distribution which can be estimated from previous performances of the mechanism. The moment conditions we use in our design can be thought of precisely as summary statistics of the designer data. Our mechanisms therefore satisfy the property advocated by Deb and Pai, and the notions of Δ-Incentive Compativility and Δ-Implementation provide the analytical framework to address that problem, from the viewpoint of partial and full implementation, respectively. The question raised by Deb and Pai is related to a second one, that is understanding what kind of information on a mechanism's past performance may be useful for the designer to disclose. This question is extremely relevant for practical mechanism design. For instance, the recent development of online trading platforms has provided the designers of those platforms with a huge quantity of data on the distribution of users' preferences, strategies, etc. Some of this information is clearly used by the platforms to better target the advertising campaigns, but very little is understood on the potential to use such information to shape the very interaction in such mechanisms. If, as in these examples, the designer of the mechanism has information on the distributions that may not be commoly known by the agents, he or she may decide to publicly disclose part of that information if this could improve the outcome of the mechanism. The conditions we provide to achieve contractivity of the mechanism point at the properties of the moment conditions that can be conveniently made public: once common knowledge of particular moment conditions is established, transfers may be suitably designed to guarantee full implementation.

Appendix

A Equivalent Approaches to Full Δ-Implementation

We define the set of type spaces consistent with $\mathcal{B}$ as follows:

Definition 14 *A (Θ-based)* type space *is a tuple* $\mathcal{T} = \left(T_i, \hat{\theta}_i, \tau_i\right)_{i\in I}$ *such that* T_i *is a compact set of player i's types;* $\hat{\theta}_i : T_i \to \Theta_i$ *is a measurable function assigning to each type a payoff type, and* $\tau_i : T_i \to \Delta(T_{-i})$ *is a belief function. A type space is* consistent with belief restrictions $\mathcal{B}$ *(or, $\mathcal{B}$-consistent) if the belief functions* $(\tau_i)_{i\in I}$ *satisfy the following: for every* $i \in N$ *and for every* $t_i \in T_i$*, there exists* $\beta \in \mathcal{B}_{\hat{\theta}(t_i)}$ *such that for any measurable* $E \subseteq \Theta_{-i}$*,*

$$\tau_i(t_i)\left[\left\{t_{-i} \in T_{-i} : \hat{\theta}_{-i}(t_{-i}) \in E\right\}\right] = \beta(E). \tag{19}$$

Equation (19) is a consistency condition, which requires that type t_i's beliefs about the opponents' types are consistent with the belief restrictions $\mathcal{B}$. It can be shown that any hierarchy of beliefs consistent with the $\mathcal{B}$ restrictions can find an implicit representation as a type in an $\mathcal{B}$-consistent type space.

Proposition 8 *Fix an environment* $\mathcal{E}$*, belief restrictions* $\mathcal{B}$ *and a direct mechanism* $\mathcal{M}$*. Let the Δ-restrictions be derived from $\mathcal{B}$. For every* θ_i*,* $R_i^{\Delta}(\theta_i)$ *characterizes the set of messages that payoff type* θ_i *would play in some BNE for some $\mathcal{B}$-consistent type space.*

(This Proposition generalizes an analogous result by Battigalli and Siniscalchi's (2003), who proved it for games with finite actions. The argument, however, is essentially the same.)

B Omitted Proofs From Sections 2 and 3.2

B.1 Proof of Theorem 1

The proposition follows immediately from the following lemmata.

Lemma 3 *Let the Δ-restrictions be non-behavioral. Then: If d responsive and Δ-implementable by a direct mechanism, then it is truthfully Δ-implementable by a strictly Δ-IC mechanism.*

Proof. *Fix the mechanism $\mathcal{M} = (d, t)$, and suppose that it Δ-Implements d, and let R^{Δ} denote the set of Δ-rationalizable reports in mechanism $\mathcal{M}$. Then: (a) for all i and θ_i, $R_i^{\Delta}(\theta_i) \neq \emptyset$, and (b) for all $\theta, \theta' \in \Theta$, $\theta' \in R^{\Delta}(\theta)$ implies $d(\theta) = d(\theta')$.*

Step 1:*. We show that for any $\theta_i \neq \theta_i'$, $R_i^{\Delta}(\theta_i) \cap R_i^{\Delta}(\theta_i') = \emptyset$. Suppose not, and let $m_i' \in R_i^{\Delta}(\theta_i) \cap R_i^{\Delta}(\theta_i')$. By definition of implementation, for any $\sigma_{-i} \in R_{-i}^{\Delta}$, we must have that $d(m_i', \sigma_{-i}(\theta_{-i})) = d(\theta_i, \theta_{-i})$ and $d(m_i', \sigma_{-i}(\theta_{-i})) = d(\theta_i', \theta_{-i})$ for any θ_{-i}. But $d(\theta_i, \theta_{-i}) = d(\theta_i', \theta_{-i})$ for all θ_{-i} contradicts the fact that d is responsive.*

Step 2: *We show that for any i, for any $\theta_i \in \Theta_i$: $|R_i^{\Delta}(\theta_i)| = 1$. Suppose not, then there exist $i \in I$ and $\theta_i^*, \theta_i, \theta_i' \in \Theta_i$ such that $\theta_i, \theta_i' \in R_i^{\Delta}(\theta_i^*)$ and $\theta_i \neq \theta_i'$. For each $i \in I$, let $Y_i := \bigcup_{\theta_i \in \Theta_i} R_i^{\Delta}(\theta_i)$. Then, by Step 1, for every $i \in I$ there exists an onto function $f_i : Y_i \to \Theta_i$ such that: (1) $\theta_i \in R_i^{\Delta}(f_i(\theta_i))$ for any $\theta_i \in Y_i$; (2) for any $\theta \in Y$, $d(\theta) = d(f(\theta))$. Under the absurd hypothesis, $\theta_i, \theta_i' \in Y_i$ are such that $\theta_i \neq \theta_i'$ and $f_i(\theta_i) = f_i(\theta_i')$. But then $d(f_i(\theta_i), f_{-i}(\theta_{-i})) = d(f_i(\theta_i'), f_{-i}(\theta_{-i}))$ for any $\theta_{-i} \in Y_{-i}$. Using (2), it follows that $d(\theta_i, \theta_{-i}) = d(\theta_i', \theta_{-i})$ for all $\theta_{-i} \in Y_{-i}$. This would contradict Responsiveness unless there exists $\theta_{-i}'' \in \Theta_{-i} \backslash Y_{-i} : d(\theta_i, \theta_{-i}'') \neq d(\theta_i', \theta_{-i}'')$. But because f_{-i} is onto, there exists $\theta_{-i}''' \in Y_{-i} : f_{-i}(\theta_{-i}''') = \theta_{-i}''$, but if $f_i(\theta_i) = f_i(\theta_i')$, then $d(f_i(\theta_i), f_{-i}(\theta_{-i}''')) = d(f_i(\theta_i'), f_{-i}(\theta_{-i}'''))$, which by (2) implies that $d(\theta_i, \theta_{-i}'') = d(\theta_i', \theta_{-i}'')$, a contradiciton.*

Step 3: *We show that there exists a mechanism that truthfully Δ-implements d. From step 2, for every $i \in I$ there exists a one-to-one function $\iota_i : \Theta_i \to \Theta_i$ such that $R_i^{\Delta}(\theta_i) = \{\iota_i(\theta_i)\}$ and $d(\theta) = d(\iota(\theta))$ for each $\theta \in \Theta$, where $\iota(\theta) \equiv (\iota_i(\theta_i))_{i \in I}$. We let $\iota_i(\Theta_i) = \bigcup_{\theta_i \in \Theta_i} \iota_i(\theta_i)$. Let $\bar{R}^{\Delta}$ denote the set of rationalizable strategies in the mechanism $\bar{\mathcal{M}}$ that is identical to $\mathcal{M}$ except that each i's action space is set equal to $\iota_i(\Theta_i)$. Clearly, $\emptyset \neq \bar{R}_i^{\Delta} \subseteq R_i^{\Delta}$ hence $\bar{R}_i^{\Delta}(\theta_i) = \{\iota_i(\theta_i)\}$ for every i and θ_i. Now consider the mechanism $\hat{\mathcal{M}} = (d, t')$ s.t. $t' = t \circ \iota$, and let $\hat{R}^{\Delta}$ denote the Δ-rationalizable strategies in $\hat{\mathcal{M}}$. Clearly, this mechanism is strictly Δ-IC, hence $\theta_i \in \hat{R}_i^{\Delta}(\theta_i)$ for each i and θ_i. We will show that, in fact, for any $k = 0, 1, ...,$ $\iota\left(\hat{R}^{\Delta,k}\right) \subseteq \bar{R}^{\Delta,k}$, that is $\theta_i' \in \hat{R}_i^{\Delta,k}(\theta_i)$ implies $\iota_i(\theta_i') \in \bar{R}_i^{\Delta,k}$. Once this is proven, it follows that $\bar{R}_i^{\Delta}(\theta_i) = \{\theta_i\}$, for if there exists $\theta_i' \neq \theta_i$ s.t. $\theta_i' \in \hat{R}_i^{\Delta}(\theta_i)$, then $\iota_i(\theta_i') \in \bar{R}_i^{\Delta}(\theta_i)$ and $\iota_i(\theta_i') \neq \iota_i(\theta_i)$, contradicting $\bar{R}_i^{\Delta}(\theta_i) = \{\iota_i(\theta_i)\}$. The proof is by induction. For $k = 0$ the condition $\iota\left(\hat{R}^{\Delta,k}\right) \subseteq \bar{R}^{\Delta,k}$ is trivially satisfied. For the inductive step, suppose that the statement is true up to $k-1$. We show that $\iota_i\left(\hat{R}_i^{\Delta,k}(\theta_i)\right) \subseteq \bar{R}_i^{\Delta,k}(\theta_i)$. Let $\theta_i' \in \hat{R}_i^{\Delta,k}(\theta_i)$, then there exists $\hat{\mu} \in \Delta_{\theta_i} \cap \hat{R}_{-i}^{\Delta,k-1}$:*

$$\theta_i' \in \arg\max_{\theta_i''} \int \left[u_i\left(d\left(\theta_{-i}', \theta_i''\right), \theta_{-i}, \theta_i\right) + t_i'\left(\theta_{-i}', \theta_i''\right)\right] d\hat{\mu}\left(\theta_{-i}', \theta_{-i}\right) \tag{20}$$

$$= \arg\max_{\theta_i''} \int \left[u_i\left(d\left(\iota\left(\theta_{-i}', \theta_i''\right)\right), \theta_{-i}, \theta_i\right) + t_i\left(\iota\left(\theta_{-i}', \theta_i''\right)\right)\right] d\hat{\mu}\left(\theta_{-i}', \theta_{-i}\right)$$

Now, let $\mu \in \Delta(\Theta_{-i} \times \Theta_{-i})$ s.t. $\mu(\iota(\theta'), \theta) = \hat{\mu}(\theta', \theta)$. Under the inductive assumption, and if Δ_{θ_i} entails no behavioral restrictions, $\mu \in \Delta_{\theta_i} \cap \bar{R}_{-i}^{\Delta,k-1}$. We want to show that

$$\iota_i(\theta_i') \in \arg\max_{\theta_i'' \in \iota_i(\Theta_i)} \int \left[u_i\left(d\left(\theta_{-i}', \theta_i''\right), \theta_{-i}, \theta_i\right) + t_i\left(\theta_{-i}', \theta_i''\right)\right] d\mu\left(\theta_{-i}', \theta_{-i}\right)$$

Suppose not. Then $\exists \theta_i'' \in \iota_i(\Theta_i)$:

$$\int \left[u_i\left(d\left(\theta_{-i}', \theta_i''\right), \theta_{-i}, \theta_i\right) + t_i\left(\theta_{-i}', \theta_i''\right)\right] d\mu\left(\theta_{-i}', \theta_{-i}\right)$$
$$> \int \left[u_i\left(d\left(\theta_{-i}', \iota_i\left(\theta_i'\right)\right), \theta_{-i}, \theta_i\right) + t_i\left(\theta_{-i}', \iota_i\left(\theta_i'\right)\right)\right] d\mu\left(\theta_{-i}', \theta_{-i}\right)$$

But, letting $\theta_i''' = \iota_i^{-1}(\theta'')$, *we can write the two sides of this inequality as follows:*

$$\text{LHS:} \int \left[u_i\left(d\left(\theta_{-i}', \theta_i''\right), \theta_{-i}, \theta_i\right) + t_i\left(\theta_{-i}', \theta_i''\right)\right] d\mu\left(\theta_{-i}', \theta_{-i}\right)$$
$$= \int \left[u_i\left(d\left(\iota_{-i}\left(\theta_{-i}'\right), \iota_i\left(\theta_i'\right)\right), \theta_{-i}, \theta_i\right) + t_i\left(\iota_{-i}\left(\theta_{-i}'\right), \iota_i\left(\theta_i'''\right)\right)\right] d\hat{\mu}\left(\theta_{-i}', \theta_{-i}\right)$$
$$= \int \left[u_i\left(d\left(\theta_{-i}', \theta_i'\right), \theta_{-i}, \theta_i\right) + t_i'\left(\theta_{-i}', \theta_i'''\right)\right] d\hat{\mu}\left(\theta_{-i}', \theta_{-i}\right)$$

$$\text{RHS:} \int \left[u_i\left(d\left(\theta_{-i}', \iota_i\left(\theta_i'\right)\right), \theta_{-i}, \theta_i\right) + t_i\left(\theta_{-i}', \iota_i\left(\theta_i'\right)\right)\right] d\mu\left(\theta_{-i}', \theta_{-i}\right)$$
$$= \int \left[u_i\left(d\left(\iota_{-i}\left(\theta_{-i}'\right), \iota_i\left(\theta_i'\right)\right), \theta_{-i}, \theta_i\right) + t_i\left(\iota_{-i}\left(\theta_{-i}'\right), \iota_i\left(\theta_i'\right)\right)\right] d\hat{\mu}\left(\theta_{-i}', \theta_{-i}\right)$$
$$= \int \left[u_i\left(d\left(\theta_{-i}', \theta_i'\right), \theta_{-i}, \theta_i\right) + t_i'\left(\theta_{-i}', \theta_i'\right)\right] d\hat{\mu}\left(\theta_{-i}', \theta_{-i}\right)$$

Hence,

$$\int \left[u_i\left(d\left(\theta_{-i}', \theta_i'\right), \theta_{-i}, \theta_i\right) + t_i'\left(\theta_{-i}', \theta_i'''\right)\right] d\hat{\mu}\left(\theta_{-i}', \theta_{-i}\right)$$
$$> \int \left[u_i\left(d\left(\theta_{-i}', \theta_i'\right), \theta_{-i}, \theta_i\right) + t_i'\left(\theta_{-i}', \theta_i'\right)\right] d\hat{\mu}\left(\theta_{-i}', \theta_{-i}\right),$$

which contradits 20. Thus, $\hat{R}_i^{\Delta}(\theta_i) = \{\theta_i\}$ *for each* θ_i, *which is truthful Full* Δ*-Implementation. That* $\hat{\mathcal{M}}$ *is strictly* Δ*-IC follows trivially.* ■

Lemma 4 *Under Assumption 2, if* d *is responsive and truthfully* Δ*-implementable, then the* Δ*-contraction property holds.*

Proof. *Truthful implementation implies that for any* $\theta_i' \in \Theta_i$, $m_i \neq \theta_i'$,*and* $b_i \in \vartheta_{\theta_i'}$.

$$EU_{\theta_i'}^{\mu^*(b_i)}\left(\theta_i'\right) > EU_{\theta_i'}^{\mu^*(b_i)}\left(m_i\right) \tag{21}$$

Fix $\Sigma \neq \sigma^*$, *and let* $\hat{k}$ *be the largest* k *such that for every* i, θ_i, $\theta_i' \in \Sigma_i(\theta_i)$, $R_i^{\Delta}(\theta_i') = \{\theta_i'\} \subseteq R_i^{\Delta,k}(\theta_i)$. *Such* $\hat{k}$ *exists, because* $R_i^{\Delta,0}(\theta_i) = M_i$ *(by definition) and* $R_i^{\Delta}(\theta_i) = \{\theta_i\}$ *(by truthful implementation) for all* $\theta_i \in \Theta_i$. *Hence,* $\exists i, \theta_i, \theta_i' \in \Sigma_i(\theta_i)$:

$$\theta_i' \in R_i^{\Delta,\hat{k}}(\theta_i) \text{ and } \theta_i' \notin R_i^{\Delta,\hat{k}+1}(\theta_i),$$

but this means that for all $b_i \in \vartheta_{\theta_i}$ *and for all* $\sigma_{-i} \in R_{-i}^{\Delta,\hat{k}}$, *there exists* $m_i : EU_{\theta_i}^{\mu^i(b_i,\sigma_{-i})}(m_i) > EU_{\theta_i}^{\mu^i(b_i,\sigma_{-i})}(\theta_i')$. *In other words: define the operator* L *s.t.* $\mu \mapsto L(\mu)$ *where* $L(\mu) : M_i \to \mathbb{R}$ *is such that* $L(\mu)(m_i) = EU_{\theta_i}^{\mu}(\theta_i') - EU_{\theta_i}^{\mu}(m_i)$; *then there exists no* $\mu \in \Delta_{\theta_i}$ *such that* $L(\mu)(m_i) \geq 0$ *for all* $m_i \in M_i$. *Notice that* L *is a linear operator from* Δ_{θ_i} *to* $\mathbb{R}^{M_i}$. *Since* Δ_{θ_i} *is closed and* *convex, the image* $L(\Delta_{\theta_i}) := \bigcup_{\mu\in\Delta_{\theta_i}} L(\mu)$ *is also closed and convex. Moreover, it is disjoint from the positive orthant of* $\mathbb{R}^{M_i}$*. Hence there is a nonzero functional* ϕ *separating these two convex sets, such that*

$$\phi(L(\mu)) < 0 \text{ for all } \mu \in \Delta_{\theta_i} \text{ and}$$
$$\phi(T) > 0 \text{ for all } T \in \mathbb{R}^{M_i}_{++}.$$

By these two properties, the normalization of ϕ *gives a nonegative probability measure* $\nu_i \in \Delta(M_i)$ *such that* $L(\mu) < 0$ *for all* $\mu \in \Delta_{\theta_i}$*. That is,* $EU^{\mu^i(b_i,\sigma_{-i})}_{\theta_i}(\nu_i) > EU^{\mu^i(b_i,\sigma_{-i})}_{\theta_i}(\theta_i')$ *for all* $b_i \in B_{\theta_i}$ *and for all* $\sigma_{-i} \in R^{\Delta,\hat{k}}_{-i}$*. This claim remains true also for* $\sigma^*_{-i} \in R^{\Delta,\hat{k}}_{-i}$*, that is*

$$EU^{\mu^*(b_i)}_{\theta_i}(\nu_i) > EU^{\mu^*(b_i)}_{\theta_i}(\theta_i') \text{ for all } b_i \in \vartheta_{\theta_i}.$$

Furthermore, by definition of $\hat{k}$*, we have that* $\sigma_{-i}(\theta'_{-i}) \in R^{\Delta,\hat{k}}_{-i}(\theta_{-i})$ *for any* $\sigma_{-i} \in \Sigma_{-i}$ *and for any* θ'_{-i}*. Thus, we have established that:*

$$\forall \Sigma \neq \sigma^*, \exists i, \theta_i, \theta_i' \in \Sigma_i(\theta_i): $$
$$EU^{\mu^i(b_i,\sigma_{-i})}_{\theta_i}(\nu_i) > EU^{\mu^i(b_i,\sigma_{-i})}_{\theta_i}(\theta_i') \text{ for all } \sigma_{-i} \in \Sigma_{-i} \text{ and } b_i \in \vartheta_{\theta_i}. \tag{22}$$

which is precisely the contraction property. ■

Lemma 5 *If d is strictly* Δ*-IC and satisfies* Δ*-contractivity, then it is truthfully* Δ*-Implementable.*
Proof. *Clearly, strict* Δ*-IC implies that* $\sigma^* \in R^\Delta$*. We neet to show that in fact* $R^\Delta = \{\sigma^*\}$*. Suppose not, then* Δ*-contractivity implies that there exists* $i, \theta_i, \theta_i' \in R^\Delta_i(\theta_i)$ *and* $\nu_i \in \Delta(\Theta_i)$ *such that* $EU^{\mu^i(b_i,\sigma_{-i})}_{\theta_i}(\nu_i) > EU^{\mu^i(b_i,\sigma_{-i})}_{\theta_i}(\theta_i')$ *for all* $b_i \in \vartheta_{\theta_i}, \sigma_{-i} \in R^\Delta_{-i}$ *and* $b_i' \in \vartheta_{\theta_i'}$*. But this implies that* θ_i' *is dominated by* v_i *for all beliefs concentrated on* $\Delta_{\theta_i} \cap R^\Delta_{-i}$*, hence* $\theta_i' \notin R^\Delta_i(\theta_i)$*, a contradiction.* ■

B.2 Proof of Theorem 2

Let us assume that the longest distance between the fixed strategy and another rationalizable one, $l := \max_{i,\theta_i}\left\{\max_{m_i\in R^\Delta_i(\theta_i)} |m_i - \theta_i|\right\} \neq 0$. We prove the result by contradiction.

Let i, θ_i^* and $m_i^* \in R^\Delta_i(\theta_i^*)$ be s.t. $|m_i^* - \theta_i^*| = l$. Since $m_i^* \in R^\Delta_i(\theta_i^*)$, $\exists \mu \in \Delta\left(R^\Delta_{-i}\right)$: $m_i^* \in \arg\max_{m_i} EU^\mu_{\theta_i}(m_i)$. By Δ-IC we also know that $\theta_i \in R^\Delta_i(\theta_i)$ for all θ_i and i, hence $C_i^T \subseteq R^\Delta_{-i}$. Let $\mu^* \in C_i^T$:$marg_{\Theta_{-i}}\mu^* = marg_{\Theta_{-i}}\mu$. By the assumed Δ-niceness of the mechanism, best responses are unique and described as the minimizer of the absolute value of the derivative of the expected utility function. We examine the difference in the first derivative of the expected utility that justifies m_i^* and the first order condition that justifies θ_i^*:

$$\left.\frac{\partial EU^{\mu}_{\theta_i^*}(m_i)}{\partial m_i}\right|_{m_i=m_i^*} - \left.\frac{\partial EU^{\mu^*}_{\theta_i^*}(m_i)}{\partial m_i}\right|_{m_i=\theta_i^*} \tag{23}$$

$$= \int_{M_{-i}\times\Theta_{-i}} \left.\frac{\partial U_i(m_{-i}, m_i, \theta_i^*, \theta_{-i})}{\partial m_i}\right|_{m_i=m_i^*} d\mu$$

$$- \int_{M_{-i}\times\Theta_{-i}} \left.\frac{\partial U_i(m_{-i}, m_i, \theta_i^*, \theta_{-i})}{\partial m_i}\right|_{m_i=\theta_i^*} d\mu^* \tag{24}$$

We consider two cases:

Case 1: $m_i^* > \theta_i^*$. Since $EU^{\mu}_{\theta_i^*}(m_i)$ is strictly concave and maximized at m_i^*, whereas $EU^{\mu^*}_{\theta_i^*}(m_i)$ is strictly concave and maximized at θ_i^*, if $m_i^* > \theta_i^*$ it follows that $\left.\frac{\partial EU^{\mu}_{\theta_i^*}(m_i)}{\partial m_i}\right|_{m_i=m_i^*} \geq 0$ and $\left.\frac{\partial EU^{\mu}_{\theta_i^*}(m_i)}{\partial m_i}\right|_{m_i=\theta_i^*} > 0$, whereas: $\left.\frac{\partial EU^{\mu^*}_{\theta_i^*}(m_i)}{\partial m_i}\right|_{m_i=\theta_i^*} \leq 0$ and $\left.\frac{\partial EU^{\mu^*}_{\theta_i^*}(m_i)}{\partial m_i}\right|_{m_i=m_i^*}$. Hence:

$$\left.\frac{\partial EU^{\mu}_{\theta_i^*}(m_i)}{\partial m_i}\right|_{m_i=m_i^*} - \left.\frac{\partial EU^{\mu^*}_{\theta_i^*}(m_i)}{\partial m_i}\right|_{m_i=\theta_i^*} \geq 0.$$

Letting $b_{\theta_i}(\theta_{-i}, m_{-i}, m_i) = \left(\left.\frac{\partial U_i(m_{-i}, s_i, \theta_i^*, \theta_{-i})}{\partial m_i}\right|_{s_i=m_i}\right)$, this can be rewritten as:

$$\int_{M_{-i}\times\Theta_{-i}} b_{\theta_i^*}(\theta_{-i}, m_{-i}, m_i^*)\, d\mu - \int_{M_{-i}\times\Theta_{-i}} b_{\theta_i^*}(\theta_{-i}, m_{-i}, \theta_i^*)\, d\mu^* \geq 0$$

Next, we separate the own effect from external effects by adding and subtracting $\int_{M_{-i}\times\Theta_{-i}} b_{\theta_i^*}(\theta_{-i}, m_{-i}, \theta_i^*)\, d\mu$. Rearranging terms, we obtain

$$\begin{aligned} B_i &:= \int_{M_{-i}\times\Theta_{-i}} b_{\theta_i^*}(\theta_{-i}, m_{-i}, \theta_i^*)\, d\mu - \int_{M_{-i}\times\Theta_{-i}} b_{\theta_i^*}(\theta_{-i}, m_{-i}, \theta_i^*)\, d\mu^* \\ &\geq \int_{M_{-i}\times\Theta_{-i}} b_{\theta_i^*}(\theta_{-i}, m_{-i}, \theta_i^*)\, d\mu - \int_{M_{-i}\times\Theta_{-i}} b_{\theta_i^*}(\theta_{-i}, m_{-i}, m_i^*)\, d\mu =: A_i \end{aligned}$$

By the mean value theorem, there exists $m_i' \in [\theta_i^*, m_i^*]$ such that

$$A_i = \left(\int_{M_{-i}\times\Theta_{-i}} \left.\frac{\partial b_{\theta_i^*}(\theta_{-i}, m_{-i}, m_i)}{\partial m_i}\right|_{m_i=m_i'} d\mu\right) \cdot (\theta_i^* - m_i^*)$$

Since $(\theta_i^* - m_i^*) < 0$ and the strict concavity of the expected payoffs, both terms are negative, thus A can be written as

$$A_i = \left|\int_{M_{-i}\times\Theta_{-i}} D_{ii}U_i(m_{-i}, m_i', \theta_i^*, \theta_{-i}) \cdot d\mu\right| \cdot l$$

Since $marg_{\Theta_{-i}}\mu^* = marg_{\Theta_{-i}}\mu$ and $\mu^* \in C_i^T$, the term B can be written as:

$$B_i = \int_{M_{-i}\times\Theta_{-i}} b_{\theta_i^*}(\theta_{-i}, m_{-i}, \theta_i^*)\, d\mu - \int_{M_{-i}\times\Theta_{-i}} b_{\theta_i^*}(\theta_{-i}, \theta_{-i}, \theta_i^*)\, d\mu$$

which by a mean-value Cauchy-Schwarz inequality, is bounded by

$$\begin{aligned} B_i &\leq \left(\left| \int_{M_{-i}\times\Theta_{-i}} \sum_{j\neq i} \frac{\partial b_{\theta_i^*}\left(\theta_{-i}, m_{-i}, m_i\right)}{\partial m_j}\Bigg|_{m_i=\theta_i^*} \right| \cdot \left|\theta_j - m_j\right| \right) d\mu \\ &\leq \left| \int_{M_{-i}\times\Theta_{-i}} \sum_{j\neq i} D_{ij} U_i\left(m_{-i}, m_i', \theta_i^*, \theta_{-i}\right) \cdot d\mu \right| \cdot l \end{aligned}$$

Since $A_i \leq B_i$, we have that

$$\begin{aligned} &\left| \int_{M_{-i}\times\Theta_{-i}} D_{ii} U_i\left(m_{-i}, m_i', \theta_i^*, \theta_{-i}\right) \cdot d\mu \right| \\ &\leq \left| \int_{M_{-i}\times\Theta_{-i}} \sum_{j\neq i} D_{ij} U_i\left(m_{-i}, m_i', \theta_i^*, \theta_{-i}\right) \cdot d\mu \right|, \end{aligned}$$

which contradicts the Δ-Self Determination condition for i.

Case 2: If $m_i^* < \theta_i^*$, proceed similarly until contradicting the Self Determination condition.

C Proofs from Section 4

Lemma 1: Suppose that $\mathcal{M} = (d, t)$ is EPIC and t is differentiable. Then, for every i and for every m, there exists a function $\tau_i : \Theta_{-i} \to \mathbb{R}$ such that $t_i(m) = t_i^*(m) + \tau_i(m_{-i})$.

Proof: A necessary condition for truthful revelation to be a best response to the opponent trutful revelation at every state (that is, EPIC) is that the following first-order condition is satisfied for every i and every θ:

$$\begin{aligned} \frac{\partial v_i\left(d(\theta), \theta\right)}{\partial x} \cdot \frac{\partial d(\theta)}{\partial \theta_i} + \frac{\partial t_i^*(\theta)}{\partial \theta_i} &= 0 \\ \text{hence, } \frac{\partial t_i^*(\theta)}{\partial \theta_i} &= -\frac{\partial v_i\left(d(\theta), \theta\right)}{\partial x} \cdot \frac{\partial d(\theta)}{\partial \theta_i}. \end{aligned}$$

Integrating over m_i, it follows that, for any $\theta = (\theta_i, \theta_{-i})$

$$t_i^*\left(\theta_i, \theta_{-i}\right) = -\int_0^{\theta_i} \frac{\partial v_i\left(d\left(s, \theta_{-i}\right), s, \theta_{-i}\right)}{\partial x} \cdot \frac{\partial d\left(s, \theta_{-i}\right)}{\partial \theta_i} ds + K \tag{25}$$

Now, for every i, define the function $\varpi_i : \Theta \to \mathbb{R}$ s.t. $\forall (\theta_i, \theta_{-i}) \in \Theta_i \times \Theta_{-i}$, $\varpi_i(\theta_i, \theta_{-i}) = v_i\left(d(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}\right)$, and notice that

$$\frac{\partial \varpi_i\left(\theta_i, \theta_{-i}\right)}{\partial \theta_i} = \frac{\partial v_i\left(d\left(\theta_i, \theta_{-i}\right), \theta_i, \theta_{-i}\right)}{\partial x} \cdot \frac{\partial d\left(\theta_i, \theta_{-i}\right)}{\partial \theta_i} + \frac{\partial v_i\left(d\left(\theta_i, \theta_{-i}\right), \theta_i, \theta_{-i}\right)}{\partial \theta_i},$$

hence (25) can be rewritten as

$$t_i^*(\theta_i,\theta_{-i}) = -\int_0^{\theta_i} \frac{\partial \varpi_i(s,\theta_{-i})}{\partial \theta_i} ds + \int_0^{\theta_i} \frac{\partial v_i(d(s,\theta_{-i}),s,\theta_{-i})}{\partial \theta_i} ds + K \tag{26}$$

$$= -v_i(d(\theta_i,\theta_{-i}),\theta_i,\theta_{-i}) + \int_0^{\theta_i} \frac{\partial v_i(d(s,\theta_{-i}),s,\theta_{-i})}{\partial \theta_i} + K + v_i(d(0,\theta_{-i}),0,\theta_{-i}). \tag{27}$$

The result follows letting $\tau_{-i}(\theta_{-i}) = K + v_i(d(0,\theta_{-i}),0,\theta_{-i})$ for every θ_{-i}.

Proposition 2: Allocation rule d is belief-free implementable by a differentiable direct mechanism if and only if the canonical mechanism is belief-free truthfully implementable.

Proof: The 'if' part is trivial. For the 'only if', suppose that d is truthfully belief-free implemented by $\mathcal{M} = (d,t)$. Results in Bergemann and Morris (2009) imply that $\mathcal{M}$ is EPIC, hence by Lemma 1 transfers t can be written as $t_i(m) = t_i^*(m) + \tau_i(m_{-i})$ for some $\tau_i : \Theta_{-i} \to \mathbb{R}$. It follows that the ex-post best-responses generated by $\mathcal{M}$ and by the canonical mechanism are identical, but this implies that also the sets of (belief-free) Rationalizable strategies are identical for the two mechanisms. Hence, if $\mathcal{M}$ truthfully implements d, so does the canonical mechanism.

C.1 Proof of Theorem 3

Consider the mechanism with transfers as in eq. (11). Observe that Condition 2 in the Theorem guarantees niceness of the mechanism. By strict concavity, truthelling is best response to any allowed conjecture concentrated on the truthtelling profile, thus the mechanism is Δ-IC. Condition 1 in the Theorem implies the Δ-Self Determination Condition of Theorem 2. The result thus follows from Theorem 2.

D Proofs from Sections 5 and 6

D.1 Proof of Lemma 2

Lemma 2: If the environment satisfies the single-crossing condition. Then: (1) The canonical mechanism is EPIC if and only if the allocation rule is strictly increasing: $\partial d(\theta)/\partial\theta_i > 0$ for every θ and every i.

Proof: To prove (1), notice that truthful revelation satisfies the (necessary) first-order conditions in the canonical mechanism, in that $V_i(\theta,\theta) = 0$ for all $\theta \in \Theta$. Taking the second order derivative of the ex-post payoff function, and simplifying, we obtain:

$$\frac{\partial^2 U_i^*(d(\theta_i,\theta_{-i}),\theta_i,\theta_{-i})}{\partial^2 m_i} = -\frac{\partial^2 v_i(d(\theta),\theta)}{\partial x \partial \theta_i} \frac{\partial d(\theta)}{\partial \theta_i}.$$

Since the SCC implies that $\frac{\partial^2 v_i(d(\theta),\theta)}{\partial x \partial \theta_i} > 0$, truthful revelation is uniquely optimal only if $\frac{\partial d(\theta)}{\partial \theta_i} > 0$.

D.2 Proof of Proposition 5

Proposition 5 Under assumptions Q.1-2 and SCC.2, in a common prior environment with independently distributed or affiliated types, an allocation function is EPIC(BIC?)-Implementable if and only if it is iDSIC-implementable.

Proof: As explained in the text, the proof of the result follows from Lemma 2 and Propositions 3 and 4, provided that we prove the following Lemma.

Lemma 6 *Under assumptions Q.1-2 and SCC.2, if d is iDSIC, then it is strictly increasing.*
Proof. $\forall\theta_i \in \Theta_i$, $\forall m \in \Theta$, *define*

$$U_i(m,\theta_i) := \int_{\theta_{-i}} v_i(d(m),\theta_i,\theta_{-i}) \cdot B_{\theta_i}(\theta_{-i})$$

A necessary condition for truthful revelation to be an (interim) best response independent of the opponents' strategies is: $\forall\theta_i \in \Theta_i$, $\forall\theta_{-i} \in \Theta_{-i}$,

$$t_i^*(\theta_i,\theta_{-i}) = -\int_0^{\theta_i} \frac{\partial U_i(s_i,\theta_{-i},\theta_i)}{\partial m_i} ds + K.$$

Substituting these transfers and taking the first order conditions of the θ_i*'s optimization problem in the resulting mechanism, it is easy to see that truthful revelation satisfies the (necessary) first-order conditions. Under the maintained assumptions Q.1 and Q.2, the second order derivative of the interim payoff, for each* $\theta_i \in \Theta_i$ *and* $m_{-i} \in \Theta_{-i}$*, simplifies to:*

$$\text{(S.O.C.)} \int_{\theta_{-i}} \frac{\partial^2 U_i^*(d(m_i,m_{-i}),\theta_i,\theta_{-i})}{\partial^2 m_i} \cdot B_{\theta_i}(\theta_{-i}) = -\frac{\partial d(m)}{\partial\theta_i} \cdot \int_{\theta_{-i}} \frac{\partial^2 v_i(d(m),\theta_i,\theta_{-i})}{\partial x \partial\theta_i} B_{\theta_i}(\theta_{-i}).$$

Since SCC.2 implies that $\int_{\theta_{-i}} \frac{\partial^2 v_i(d(m),\theta)}{\partial x\partial\theta_i} B_{\theta_i}(\theta_{-i}) > 0$*, truthful revelation is uniquely optimal only if* $\frac{\partial d(m)}{\partial\theta_i} > 0$. ■

D.3 Proof of Proposition 7

Lemma 7 *If* Σ *and* Σ' *are acceptable for* f*, then* $\Sigma^* := \Sigma \cup \Sigma'$ *defined as* $\Sigma^*(\theta) = \Sigma(\theta) \cup \Sigma'(\theta)$ *for each* θ *is also acceptable for* f*.*
Proof. *This is trivial. Let* $\theta,\theta' : \theta' \in \Sigma^*(\theta)$*. By definition of* Σ^**, it must be* $\theta' \in \Sigma(\theta)$ *or* $\theta' \in \Sigma'(\theta)$*. One way or the other, if* Σ *and* Σ' *are both acceptable,* $f(\theta') = f(\theta)$*, hence* Σ^* *is acceptable* ■

Lemma 8 *If* f *is responsive, any deception* $\Sigma \neq \sigma^*$ *is unacceptable for* $f^{\mathcal{M}}$*.*
Proof. *The proof is based on three steps. The first two are essentially the same as Steps1 and 2 in Lemma 3.*

Step 1:. *For any* $\theta_i \neq \theta_i'$, $\Sigma_i(\theta_i) \cap \Sigma_i(\theta_i') = \emptyset$*. Suppose not, and let* $\theta_i^* \in \Sigma_i(\theta_i) \cap \Sigma_i(\theta_i')$*. Then, for any* θ_{-i} *and* $\theta'_{-i} \in \Sigma_{-i}(\theta_{-i})$*,* $f(\theta_i^*,\theta_{-i}) = f(\theta_i,\theta_{-i})$ *and* $f(\theta_i^*,\theta_{-i}) = f(\theta_i',\theta_{-i})$*.But* $f(\theta_i,\theta_{-i}) = f(\theta_i',\theta_{-i})$ *for all* θ_{-i} *contradicts responsiveness.*

Step 2: *For any* i*, for any* $\theta_i \in \Theta_i$*:* $|\Sigma_i(\theta_i)| = 1$*. Suppose not, then there exist* $i \in I$ *and* $\theta_i^*,\theta_i,\theta_i' \in \Theta_i$ *such that* $\theta_i,\theta_i' \in \Sigma_i(\theta_i^*)$ *and* $\theta_i \neq \theta_i'$*. For each* $i \in I$*, let* $Y_i := \bigcup_{\theta_i\in\Theta_i} \Sigma_i(\theta_i)$*. Then, by Step 1, for every* $i \in I$ *there exists an onto function* $\gamma_i : Y_i \to \Theta_i$ *such that: (1)* $\theta_i \in \Sigma_i(\gamma_i(\theta_i))$ *for any* $\theta_i \in Y_i$*; (2) for any* $\theta \in Y$*,* $f(\theta) = f(\gamma(\theta))$*. Under the absurd hypothesis,* $\theta_i,\theta_i' \in Y_i$ *are such that* $\theta_i \neq \theta_i'$ *and* $\gamma_i(\theta_i) = \gamma_i(\theta_i')$*. But then* $f(\gamma_i(\theta_i),\gamma_{-i}(\theta_{-i})) = f(\gamma_i(\theta_i'),\gamma_{-i}(\theta_{-i}))$ *for any* $\theta_{-i} \in Y_{-i}$*. Using (2), it follows that* $f(\theta_i,\theta_{-i}) = f(\theta_i',\theta_{-i})$ *for all* $\theta_{-i} \in Y_{-i}$*. This* *would contradict Responsiveness unless there exists* $\theta''_{-i} \in \Theta_{-i} \backslash Y_{-i} : f\left(\theta_i, \theta''_{-i}\right) \neq f\left(\theta'_i, \theta''_{-i}\right)$*. But because* γ_{-i} *is onto, there exists* $\theta'''_{-i} \in Y_{-i} : \gamma_{-i}\left(\theta'''_{-i}\right) = \theta''_{-i}$*, but if* $\gamma_i(\theta_i) = \gamma_i(\theta'_i)$*, then* $f\left(\gamma_i(\theta_i), \gamma_{-i}\left(\theta'''_{-i}\right)\right) = f\left(\gamma_i(\theta'_i), \gamma_{-i}\left(\theta'''_{-i}\right)\right)$*, which by (2) implies that* $f\left(\theta_i, \theta''_{-i}\right) = f\left(\theta'_i, \theta''_{-i}\right)$*, a contradiciton.*

Step 3: *Suppose that* Σ *is acceptable. Since* σ^* *is trivially acceptable, Lemma 7 implies that* $\Sigma' = \Sigma \cup \{\sigma^*\}$ *is also acceptable. But if* $\Sigma \neq \sigma^*$*,* Σ' *is such that* $\exists i, \theta_i : |\Sigma'_i(\theta_i)| > 1$*, contradicting Step 2.* ■

Proposition 7: Let $d : \Theta \rightarrow X$ be responsive, $\mathcal{M} = (d, t)$ and the Δ-restrictions be derived from $\mathcal{B}$ that satisfies $|B_{\theta_i}| = 1$ for every θ_i and i. Then the following are equivalent:

1. $\mathcal{M}$ is Δ-IC and Δ-contractive

2. $f^{\mathcal{M}}$ satisfies direct ICR-Monotonicity.

Proof: First of all, notice that if d is responsive, so is $f^{\mathcal{M}}$, hence any deception $\Sigma \neq \sigma^*$ is unacceptable for $f^{\mathcal{M}}$ (Lemma 8).

(1) $\Rightarrow$ (2) is trivial: if $\mathcal{M}$ is Δ-contractive, then for any $\Sigma \neq \sigma^*$ there exists $i, \theta_i, \theta'_i \in \Sigma_i(\theta_i)$ and $\nu^*_i \in \Delta(\Theta_i)$ such that: $EU_{\theta_i}^{\mu^i\left(b_{\theta_i}, \sigma_{-i}\right)}(\nu_i) > EU_{\theta_i}^{\mu^i\left(b_{\theta_i}, \sigma_{-i}\right)}(\theta'_i)$ for all $\sigma_{-i} \in \Sigma_{-i}$, but this is just inequality (17) for $y^* \in Y^*(\theta_i)$.obtained as $f(\nu^*_i, \theta_{-i}) = y^*(\theta_{-i})$. Inequality (18) follows trivially from Δ-IC.

(2) $\Rightarrow$ (1) direct ICR-Monotonicity implies that there exists $y^* \in Y^*(\theta_i)$ that satisfies (17), but this implies that $\exists \nu^*_i \in \Delta(\Theta_i)$ such that: $EU_{\theta_i}^{\mu^i\left(b_{\theta_i}, \sigma_{-i}\right)}(\nu_i) > EU_{\theta_i}^{\mu^i\left(b_{\theta_i}, \sigma_{-i}\right)}(\theta'_i)$ for all $\sigma_{-i} \in \Sigma_{-i}$, which is Δ-contractivity. Furthermore, for any i, θ_i, $\theta'_i \neq \theta_i$, consider the deception

$$\Sigma_j(\theta_j) = \begin{cases} \{\theta_i, \theta'_i\} & \text{if } (j, \theta_j) = (i, \theta_i) \\ \theta_j & \text{otherwise} \end{cases} .$$

This deception is unacceptable, hence by direct ICR-monotonicity, there exists $y^* \in Y^*(\theta_i)$ such that for any θ_{-i},

$$\int u_i\left(y^*_i(\theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i} > \int u_i\left(f(\theta'_i, \theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i}$$

$$\text{and} \tag{28}$$

$$\int u_i\left(f(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i} \geq \int u_i\left(y^*_i(\theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i}, \tag{29}$$

Together, this implies:

$$\int u_i\left(f(\theta_i, \theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i} > \int u_i\left(f(\theta'_i, \theta_{-i}), \theta_i, \theta_{-i}\right) db_{\theta_i}$$

Since this holds for any i,θ_i, θ'_i, Δ-IC follows.■

REFERENCES

1. Ausubel, L. M., O. V. Baranov (2010) "Core-Selecting Auctions with Incomplete Information," *mimeo.*

2. Ausubel, L. M., P. Cramton, M. Pycia, M. Rostek and M. Weretka (2013) "Demand Reduction and Inefficiency in Multi-Unit Auctions," *mimeo.*

3. Artemov, G., T. Kunimoto and R. Serrano (2013) "Robust Virtual Implementation with Incomplete Information: Towards a Reinterpretation of the Wilson Doctrine," *Journal of Economic Theory* 148(2), 424-447.

4. d'Aspremont, C., J. Cremer and L-A. Gerard-Varet (2005), "Unique Implementation in Auctions and in Public Goods Problems," *economie publique* 17, 125–139.

5. Battigalli, P. (2003), "Rationalizability in infinite, dynamic games with incomplete information," *Research in Economics,* 57, 1–38.

6. Battigalli, P. and M. Siniscalchi (2003), "Rationalization and Incomplete Information ," *The B.E. Journal of Theoretical Economics* 3(1).

7. Battigalli P., A. Di Tillio, E. Grillo and A. Penta (2011), "Interactive Epistemology and Solution Concepts for Games with Asymmetric Information," *The B.E. Journal of Theoretical Economics: Vol. 11 (Advances),* Article 6.

8. Bergemann D. and S. Morris (2005), "Robust Mechanism Design," *Econometrica* 73, 1521–1534.

9. Bergemann, D. and S. Morris (2009a) "Robust Implementation in Direct Mechanisms," *Review of Economic Studies*, **76**, 1175–1204.

10. Bergemann, D. and S. Morris (2009b) "Robust virtual implementation," *Theoretical Economics*, 4(1).

11. Bergemann, D. and S. Morris (2011) "Robust Implementation in General Mechanisms," *Games and Economic Behavior,* 71(2), 261–281.

12. Bergemann, D. and S. Morris (2012), "Robust Mechanism Design," *World Scientific Publishing*, Singapore.

13. Börgers, T. and D. Smith (2012) "Robustly Ranking Mechanisms", *American Economic Review Papers and Proceedings*, (2012) 325-329

14. Börgers, T. and D. Smith (2013), "Robust Mechanism Design and Dominant Strategy Voting Rules," *Theoretical Economics,* forthcoming.

15. Carroll, G. (2013) "Robustness and Linear Contracts," *mimeo*, Stanford.

16. Cremer, J. and R.P. McLean (1985), "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica,* 53(2), 345–361.

17. Cremer, J. and R.P. McLean (1988), "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica,* 56(6), 1247–1257.

18. Dasgupta, P. and E. Maskin (2000), "Efficient Auctions," *Quarterly Journal of Economics,* 115, 341–388.

19. Deb, R. and M. M. Pai (2013), "Symmetric Auctions," *mimeo,* UPenn.

20. Dekel, E., D. Fudenberg and S. Morris (2007), "Interim Correlated Rationalizability," *Theoretical Economics,* 2, 15–40.

21. Duggan, J. and J. Roberts (2002), "Implementing the Efficient Allocation of Pollution," *American Economic Review*, 92(4), 1070-1078

22. Gershkov, A., J.K. Goeree, A. Kushnir, B. Moldovnu and X. Shi (2013), "On the Equivalence of Bayesian and Dominant Strategy Implementation," *Econometrica*, 81(1), 197-220.

23. Grossman, S., J. Stiglitz (1980), "On the Impossibility of Informationally Efficient Markets," *American Economic Review,* 70, 393–408.

24. Haile, P. A. and E. Tamer (2003), "Inference with an Incomplete Model of English Auctions,"*Journal of Political Economy,* 111(1), 1–51.

25. Healy, P. J. and L. Mathevet (2012), "Designing stable mechanisms for economic environments," *Theoretical Economics,* 7(3).

26. Hellwig, M. F. (1980), "On the Aggregation of Information in Competitive Markets," *Journal of Economic Theory,* 22, 477–498.

27. Jackson, M. O. (1991), "Bayesian Implementation," *Econometrica,* 59(2), 461–477.

28. Jackson, M. O. (1992), "Implementation in Undominated Strategies: A Look at Bounded Mechanisms," *Review of Economic Studies,* 59(4), 757–75.

29. Jensen, M. K. (2005), "Existence, Comparative Statics, and Stability in Games with Strategic Substitutes," *mimeo,* University of Birmingham.

30. Jehiel, P., M. Meyer-ter-Vehn, B. Moldovanu and W.R. Zame, "The Limits of Ex Post Implementation," *Econometrica,* 74(3), 585–610.

31. Kim, K. and A. Penta (2012) "Efficient Auctions and Robust Mechanism Design: A New Approach, "*mimeo*, UW-Madison.

32. Kyle, A. S. (1989) "Informed Speculation With Imperfect Competition " *Review of Economic Studies,* 56, 317-356.

33. Li, Y. (2013), "Approximation in Mechanism Design with Interdependent Values, " *mimeo,* U-Penn.

34. Lopomo, G., L. Rigotti and C. Shannon (2013) "Uncertainy in Mechanism Design," *mimeo.*

35. Manelli, A. and D.R. Vincent (2010) "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model", *Econometrica*, 78(6), 1905-1938.

36. Maskin, E. (1999) "Nash equilibrium and welfare optimality.", *Review of Economic Studies*, 66, 23-38. (Working Paper version circulated in 1977)

37. Mathevet, L. A. (2010), "Supermodular Mechanism Design, " *Theoretical Economics,* 5(3).

38. Mathevet, L. A. and I. Taneva (2013), "Finite Supermodular Design with Interdependent Valuations," *Games and Economic Behavior*, forthcoming.

39. McLean, R. and A. Postlewaite (2002), "Informational Size and Incentive Compatibility," *Econometrica,* 70, 2421-2453.

40. McLean, R. and A. Postlewaite (2004), "Informational Size and Efficient Auctions," *Review of Economic Studies,* 71, 809-827.

41. Milgrom, P. and J. Roberts (1990), "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica,* 58.

42. Milgrom, P. and R. J. Weber (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, Vol. 50, No. 5. pp. 1089-1122.

43. Moulin, H. (1984), "Dominance Solvability and Cournot Stability," *Mathematical Social Sciences,* 7.

44. Mueller, C. (2012a), "Robust Virtual Implementation under Common Strong Belief in Rationality" *mimeo*, Carnegie-Mellon.

45. Mueller, C. (2012b), "Robust Implementation in Weakly Rationalizable Strategies", *mimeo*, Carnegie-Mellon.

46. Myerson, R. B. (1981), "Optimal Auction Design," *Mathematics of Operations Research,* 6.

47. Ollár, M. and A. Penta (2014) "Nice Games with Incomplete Information and Belief Restrictions," *mimeo*, UW-Madison.

48. Oury, M. and O. Tercieux (2012), "Continuous Implementation," *Econometrica*, 80, 1605-1637.

49. Oury, M. (2013), "Continuous Implementation: a Full Characterization Result," *mimeo*, Paris-Dauphine.

50. Penta, A. (2011), "Robust Dynamic Mechanism Design," *mimeo*, UW-Madison.

51. Penta, A. (2013), "On the Structure of Rationalizability on Arbitrary Spaces of Uncertianty", *Theoretical Economics*, 8, 405–430.

52. Postlewaite, A. and D. Schmeidler (1986), "Implementation in Differential Informati Economieson," *Journal of Economic Theory*, 39, 14–33.

53. Roughgarden, T. (2012), "The price of anarchy in games of incomplete information." In *Proceedings of the 13th ACM Conference on Electronic Commerce, pages 862—879. ACM.*

54. Roughgarden, T. and I. Talgam-Cohen (2013), "Optimal and Near-Optimal Mechanism Design with Interdependent Values," *EC'13 ACM Conference on Electronic Commerce*.

55. Schmeidler, D. (1989), "Subjective Probability and Expected Utility Without Additivity," *Econometrica*, 57, 571–587.

56. Weinstein, J. L. and M. Yildiz, (2007) "A Structure Theorem for Rationalizability With Application to Robust Predictions of Refinements," *Econometrica*, 75, 365–400.

57. Weinstein, J. L. and M. Yildiz, (2011) "Sensitivity of Equilibrium Behavior to Higher-order Beliefs in Nice Games.," *Games and Economic Behavior,* 72(1), 288–300.

58. Weinstein, J. L. and M. Yildiz, (2013) "Robust Predictions in Infinite-horizon Games–An Unrefinable Folk Theorem," *Review of Economic Studies,* forthcoming.

59. Wilson, R., (1987) "Game-Theoretic Analysis of Trading Processes," *Advances in Economic Theory, ed. by Bewley*, Cambridge University Press.

60. Yamashita, T. (2013a) ""Robust trading mechanisms to strategic uncertainty", *mimeo*, Toulouse School of Economics

61. Yamashita, T. (2013b) "Robust Auctions with Interdependent Values", *mimeo*, Toulouse School of Economics.