Participation Rights and Mechanism Design

Peter S. Faynzilberg[*]

September 1997
Revisions: March 1998, July 1999

[*] I would like to thank for discussions and helpful comments Sugato Bhattacharyya, Tracy Lewis, Roger Myerson, Richard Romano, Ilya Segal, Asher Wolinsky, the participants of the Econometric Society Meeting, the Economic Theory Seminar at the University of Florida, and the Economic Seminar Series at Northwestern University.

Abstract

This paper is concerned with the procedural aspects of collective choice and the impact of the parties' participation rights on the optimal mechanism. We ¯nd that the mechanism designer generally bene¯ts from the selective engagement of the agents | the exclusion of some agent- types from the choice process. We show that optimization of mechanisms with voluntary participation involves two mutually dependent instruments: the scope of the agents' engagement, and the functional form of the social choice function. The bene¯ts of selective engagement, as well as two optimization methodologies, are illustrated on principal-agent models.

We ¯nd that the participation constraint is redundant and generally leads to suboptimal mechanisms. Contrary to its general interpretation, this restriction does not re°ect the voluntary aspect of the agents' participation. Rather, it gives them an additional entitlement: to force their involvement in the collective choice.

We formulate a free-exit constraint that is devoid of incentives and fully accounts for the voluntary aspect of participation. It also leads to an equivalent representation of incentive-compatibility that explicates incentives and speci¯es the feasibility of a mechanism.

Key words: Participation rights, voluntary participation, economics of information, incentives, incentive compatibility, principal-agent model.

1 Introduction

In this paper, we are concerned with the problem of collective choice, i.e., selection from a specified set of alternatives undertaken jointly by a given group of agents. The resolution of this problem is effectively facilitated by the well-known device of a mechanism that is constructed by the designer, a party introduced expressly for this purpose. Given the ubiquity of collective-choice problems in economics, the framework of mechanism design unifies a vast multitude of applications, such as provision of public goods, auctions, principal-agent relationships, bilateral trading, etc.

The conflict between the social and the private values of the choice alternatives is aggravated whenever the distribution of information among the parties lacks uniformity. To mitigate the impact of the informational disparities, the designer supplies the agents with incentives that induce them to reveal their private information. Thus, in a principal-agent relationship with hidden actions (Holmström [14]), the wage function is constructed to induce the agent's compliance with the action allocated to him by the principal. In the model with hidden information (e.g., Baron and Myerson [1], Maskin and Riley [18], Guesnerie and Laffont [13]), the agent is similarly motivated to reveal truthfully his productivity or preference parameter. In the hybrid model with both hidden actions and hidden information (e.g., Laffont and Tirole [15], Faynzilberg and Kumar [9]), additional incentives dissuade the agent from a simultaneous misreporting of his type and deviation from the socially optimal action.

The literature has widely recognized that the choice made by the voluntarily participating agents requires additional considerations. The very ability of the agents to abstain may have profound consequences for the optimal mechanism. The Myerson-Satterthwaite Theorem demonstrates, for instance, that the voluntary aspect of bilateral trading renders the ex post efficiency of the allocation unfeasible (Myerson and Satterthwaite [25]).

The task of securing the agents' participation may be accomplished by incentives. Formally, these incentives are usually described by a restriction — the participation, or individual rationality constraint — on the form of the social choice function. The imperfectly informed designer who facilitates the collective choice of voluntarily participating agents is confined, then, to mechanisms that satisfy both the participation and the incentive-compatibility constraints.

In the present paper, we show that, contrary to the common assumption, the participation constraint neither follows from nor reflects the voluntary aspect of the agent's participation. The incentive argument is, of course, valid: the agents who have the right to abstain from choice need inducements to forego the exercise of this right. We find, however, that the incentive-compatibility already accounts for the participation incentives. Consequently, the participation constraint is, at best, redundant, and generally leads to an oversupply of incentives with deleterious results for the social welfare.

The agent's involvement in collective choice is governed by procedural rules, to which we refer as participation rights. Whereas property rights are associated with the outcomes of the collective choice, the participation rights of the parties regulate the process rather than outcomes. Within the category of participation rights, we distinguish the privilege to abstain and the entitlement to participate. Accordingly, the agent's right of free exit allows him to decline unilaterally any part in the choice process. When imputed to the agent, this right makes the agent's involvement voluntary for him, and involuntary otherwise. The agent's entry right creates an entitlement to participate. When imputed to the agent, it allows him to force upon other parties his involvement in the process.

When the rights of exit and entry are allocated to the agent, they impose their dual duties on the designer. He must not interfere with the agents' pursuit of exogenous opportunities if they choose to abstain from choice, and, respectively, o®er them at least one choice outcome if they prefer to participate. Naturally, these duties limit the designer in his selection of the mechanism. Formally, they manifest themselves as the free-exit and the free-entry constraints. The free-exit constraint is devoid of incentives and re°ects fully the voluntary aspect of the agents' participation.

The present framework allows us to reinterpret the participation constraint: it is an entry right granted to the agent rather than a device dissuading him from abstention. The traditional use of the participation constraint has, therefore, a possibly unintended consequence: it gives the agent entitlement beyond his ability to abstain from choice. The participation constraint models a guaranteed, rather than voluntary, participation. As any of the agent's entitlements, this privilege is costly and reduces the welfare value of the mechanism.

The legal environment may, of course, expressly grant all participation privileges to the agents, in which case both the participation and the nonparticipation become voluntary for them. In other words, the agents may exogenously enforce their abstinence from the choice process as well as their participation in it. Dually, the designer has the respective duties to comply with the involvement-related choices of the agents. Such institutional settings have been extensively studied by the extant literature, so the present paper focuses mainly on choice problems with voluntary participation only.

In mechanisms with voluntary participation, the designer is free from the duty to engage an agent and may prefer to exclude some of his types from participation. Formally, this selective engagement of the agents becomes possible because the designer is no longer limited by the participation constraint. This is in contrast to the complete engagement, prescribed as a rule by the extant models.[1] In some multi-agent situations, such an auction with a random number of bidders in McAfee and McMillan [19], provisions are made for the variable involvement of the agents. These settings construe the act of an agent's (non)participation as, respectively, the (non)participation of all types of that agent. In contrast, we allow the agent's involvement to be type-contingent. Even in single-agent mechanisms, the designer must consider and, when optimal, adopt the selective-engagement regime. Below, we illustrate the bene¯ts of selective engagement on two broad classes of principal-agent models, which have attracted much attention in the mechanism-design literature (e.g., Mirrlees [20], Mussa and Rosen [21], Baron and Myerson [1], Maskin and Riley [18], Guesnerie and La®ont [13]).

The fact that selective engagement generally improves the social welfare raises the question of its feasibility. This question is answered in the positive by the explicit form of the incentive-compatibility of a mechanism with voluntary participation (Equivalence Theorem 2.1 below).

We find that the incentive-compatibility of the social choice function is equivalent to a combination of: (i) the truth-telling constraint, satisfied on the participating types only; (ii) the participation constraint, also satisfied only on the participating types; and (iii) the nonparticipation constraint satisfied only on the nonparticipating agent-types. It follows that the need for incentives, including those for participation, leaves the scope of participation undetermined, hence under the discretion of the designer.

Given that selective engagement is both beneficial and feasible, mechanism optimization involves an endogenous choice of the scope of participation. The designer has therefore two instruments at his disposal: the functional form of the social choice function, and the scope of participation. The nature of the second instrument complicates optimization of the mechanism: the scope of participation is a set, which must be chosen simultaneously with a social choice function.

The general mechanism-optimization methodology is addressed elsewhere (Faynzilberg [5, 6, 7]). The companion paper Faynzilberg [5] contains a complete characterization of the principal-agent model with hidden information (adverse selection), which builds on the presently developed framework. As in Holmström [14], where the same task for the hidden-actions model (moral hazard) is addressed, the characterization of the optimal contract is an effective means of mechanism selection. The hidden-information contract requires a more elaborate characterization than its hidden-actions counterpart. This is because the additional variable is presently a set (the scope of participation) rather than a number (the agent's effort) — a markedly greater freedom of choice. The companion paper also shows that, generically, the designer has multiple pooling options at his disposal. Since each of the pooling regimes extends to its own allocation path, there exists a quantized spectrum of locally optimal social choice functions. This spectrum may be rather large even in relatively simple settings. Consequently, the selection from it of the globally optimal, second-best mechanism may be far from trivial. We leave these effects outside the scope of the present paper, however, and merely illustrate the characterization methodology in Section 4.2.

Thus, the scope of agents' participation is an optimization instrument in its own right. As such, it should be built into the mechanism design from the outset, before the optimization is undertaken. The fact that ex post, in the optimal mechanism, some agents abstain from participation has been recognized, of course, by many authors, most notably by Baron and Myerson [1]. The two most common alternatives to the present framework are discussed below (Section 5). It appears that their validity is premised on certain idiosyncrasies, such as the equivalence of the pre-choice status quo and the outside options of the parties. In addition, the initial inclusion of all agent-types into the scope of participation presents methodological difficulties as well, and may lead to suboptimal results. In comparison, the presently proposed framework appears to be more general. The explicit treatment of the parties' involvement decisions allows it to serve as a foundation for a consistent optimization methodology.

The objective of the present paper is therefore threefold: (i) to analyze the impact of participation rights on the choice of the mechanism, (ii) to demonstrate that selective engagement is both beneficial and feasible in mechanisms with voluntary participation, and (iii) to illustrate two optimization methodologies that allow for the endogenous determination of the scope of participation jointly with the form of the social choice function.

The rest of the paper is organized as follows. Section 2 is devoted to participation rights and the impact they have on the selection of the mechanism. We begin with a discussion and de¯nitions of the participation rights in Section 2.1 and give their formal description in Section 2.2. Focusing on mechanisms with voluntary participation, we present a characterization of incentive-compatibility in Section 2.3 that leads to a convenient speci¯cation of feasibility. Optimization issues are addressed in Section 3. Here we formalize the designer's optimization problem (Section 3.1) and identify the economic e®ects of selective engagement (Section 3.2). Optimization methodologies and applications are the subject of Section 4, where we solve in closed form two principal-agent models. Two methodologies are presented: a step-wise optimization procedure (Section 4.1), and a characterization of the optimal contract (Section 4.2). Next, in Section 5, we revisit the alternate approaches to selective engagement. The concluding remarks are gathered in Section 6. Technical details and proofs may be found in the Appendix.

2 Participation Rights

In the present section, we revisit the issue of agents' involvement in the collective choice. We extend its analysis to include both of its facets | participation and abstinence. The corresponding rights of free exit and free entry are formulated in Section 2.1 and formalized in Section 2.2 as constraints (2.6){(2.7) on the choice of the mechanism. The Equivalence Theorem 2.1 shows that voluntary participation leaves the selective engagement feasible, and that incentive- compatibility accounts for all inducements needed to secure the agents' participation.

2.1 Collective Choice: The Rights of Exit and Entry

Consider a group of agents who collectively perform a selection from a given set of alternatives. The parties to bilateral trading, for instance, select a pair of transfers | of the object of trade and of the numeraire. In the agency setting, the principal and the agent jointly select a contract comprised of wages to be paid to the agent, and an action he is supposed to take in behalf of the principal. In a single-unit auction, the participants jointly determine the party to whom the good is allocated, and the compensatory transfer to the auctioneer.

The selection made by the parties amounts to an assignment of property rights. These rights may be merely redistributed, such as when exiting goods are being traded. Alternatively, the process may end with an allocation of claims to the objects that are yet to come into existence. In agency settings, for instance, the agent contributes to the output of the technology that belongs to the principal. Here the property rights are both reallocated and newly assigned: the principal is a residual claimant to the output, net of the wages paid to the agent.

In contrast to property rights, the redistribution of which is both the objective and the result of the collective choice, we are concerned with the procedural rights and duties that govern the very process of outcome selection. Oftentimes, these rules take the form of participation- related restrictions that limit the consequences of the agents' involvement in the process. To illustrate, consider a corporate capital- allocation committee voting on a project. The ¯rm's bylaws may contain a rule requiring unanimous participation for the committee to reach a decision. The members' involvement decisions may then severely limit the outcome: the absence of one of them renders the decision infeasible. The bylaws may alternatively prescribe that the decision be reached by a dichotomous vote, in which case the presence of an odd number of members guarantees the existence of a majority and thereby resolves the collective- choice problem at hand.

In terms of the mechanism design, the participation-related restrictions are duties imposed on the designer who facilitates the choice process. For the agents, it a®ects neither the freedom to participate in nor the freedom to abstain from the process. This is because the right, which corresponds to the designer's duty, lies outside of the mechanism. In the above-given example of the capital-allocation committee, it belongs to the shareholders of the ¯rm, so the right-duty pair is not fully endogenous to the mechanism. In contrast, participation rights | the focus of the present paper | are entirely endogenous to the choice problem.

We refer as participation rights and duties to those privileges and obligations that govern the ability of one party to exclude another from participation in the collective choice. Such rights and duties are in the usual legal duality: a right of one party implies a duty for another, namely, for the party being excluded to abstain from choice. Here and elsewhere in the paper participation refers to the involvement of a party in the collective choice rather than the mechanism. Every party takes part in the mechanism | either by taking part in the choice process (participation), or by abstaining from it (nonparticipation).

Participation rights and their dual duties specify to the corresponding parties their courses of action. Depending on how they are distributed, these rights may prescribe or proscribe either or both actions | the participation in or the abstinence from choice. The mechanism designer, for instance, prefers to engage all agents when their involvement leads to greater social welfare. When the legal environment grants him the privilege to impose this decision on the agents, the agents' participation becomes involuntary: the designer may proscribe their abstinence from choice. In other situations, where an agent's contribution is negative, the designer may prefer to prescribe the agent to abstain from choice. When this course of action is made feasible by the legal environment, it is the agent's nonparticipation that becomes involuntary. In line with the usual practice, we say that a party's participation is voluntary if it can unilaterally decline its part in the choice process. By extension, the nonparticipation is voluntary if the party can unilaterally decline not to take part. Such a party can force its participation in the collective choice.

We concentrate below on the subclass of individual participation rights | the rights of free entry, or entry rights, and rights of free exit, or exit rights. We shall focus, furthermore, on the rights that govern the interactions between an agent and the designer rather than those between two agents. A participation right granted to one agent creates a dual duty for the designer but does not directly a®ect other agents. Clearly, this type of duality is not the only one economically relevant: one agent may grant, for instance, a right to another (or impose a duty on him) to join him in the resolution of the choice problem. In Anglo- American law, the class status of a legal action has such flavor. Further, participation rights need not be individual: a coalition of agents may be entitled to abstain unanimously and unilaterally from taking part in the process. Although such inter- and multi-agent privileges are well- constructed participation rights, they lie outside the scope of the present paper.

The rights of free exit and free entry arise in certain symmetry to each other. Specifically, we refer as the exit right to the privilege of a party to decline unilaterally its participation in the collective choice. Let us focus first on an agent and suppose that the exit right is imputed to him. The agent's participation is voluntary: even when contrary to the wishes of the designer, his abstinence from choice can be exogenously enforced. Alternatively, the designer in possession of the exit right can prescribe the agent's involvement, thereby making it involuntary for the agent.

To illustrate, consider the monopoly of Maskin and Riley [18] that sells its output to a homogeneous population of consumers with an uncertain taste. The parties' abilities to abstain from market interactions may be governed in one of four ways. The consumers are typically endowed with their exit right and may therefore refuse to buy the good offered to them by the monopolist. If this right is given to the monopolist instead, he can extend a binding offer to each type of consumer. The monopolist, too, may or may not be free to exit the market. He may be compelled to withdraw if his costs make production prohibitively expensive. Whether or not he may take this course of action is predicated on the imputation of the right of free exit.

In symmetry with the exit right, the entry right of a party is a privilege to decline unilaterally its nonparticipation in the collective choice. Stated differently, it is the privilege of the party to force upon the designer its involvement in the choice process. Naturally, this right expands the agent's freedom only if it lies with that agent. The designer endowed with this right is absolved from any duty to engage the agent, and can proscribe the agent's participation altogether.

In the aforementioned case of the monopoly pricing its good, a consumer's entry right is the entitlement to at least one allocation from the monopolist. Whereas the consumer is still free not to buy the good, the firm is not: its withdrawal from the market is proscribed. When the consumer's entry right is imputed, alternatively, to the firm, the monopoly has no duty to the consumer. It is free not to extend any offer at all to some consumer types. Our choice of words is purposeful: the absence of the firm's offer is economically distinct from an offer to buy zero quantity of the good. A detailed discussion of this distinction may be found in Section 5.

What the rights do not forbid or prescribe, they allow. A party to the collective choice has three courses of action: it may decline, accept, or force its participation in the process. Which of these avenues is open to the party is determined by the distribution of participation rights, as summarized in Table 2.1. An interpretation of the table elements that involve the parties' utilities requires the formal developments of the next section. We shall find there that two pairs of scenarios — namely, Cells 2 and 4, and 2b and 4b — are behaviorally equivalent.

It suffices therefore to consider only the three remaining situations depicted as 1, 2a; and 3 in Table 2.1. The earlier literature on mechanism design focused on the first of these settings by imputing all participation rights to the designer. It has been subsequently recognized by many

Exit Right \ Entry Right	Agent	Designer
Agent	2 Participation: voluntary. Nonparticipation: voluntary. Exclusion of the agent may be: Proscribed: No. Prescribed: Yes. Guaranteed attainability of: Choice outcome: Yes. Exogenous sources: Yes. Utility attainable: Exogenously, U_0: Endogenously, U_1: Rational behavior: (a) If U_1 ¸ U_0; the agent participates. (b) If $U_1 < U_0$; the entry right is inconsequential.	3 Participation: voluntary. Nonparticipation: involuntary. Exclusion of the agent may be: Proscribed: No. Prescribed: No. Guaranteed attainability of: Choice outcome: No. Exogenous sources: Yes. Utility attainable: Exogenously, U_0:
Designer	4 Participation: involuntary. Nonparticipation: voluntary. Exclusion of the agent may be: Proscribed: Yes. Prescribed: No. Guaranteed attainability of: Choice outcome: Yes. Exogenous sources: No. Utility attainable: Endogenously, U_1: Rational behavior: (a) If exit actions are not o®ered or o®ered but U_1 ¸ U_0; the agent participates. (b) Otherwise, the entry right is inconsequential.	1 Participation: involuntary. Nonparticipation: involuntary. Exclusion of the agent may be: Proscribed: Yes. Prescribed: Yes. Guaranteed attainability of: Choice outcome: No. Exogenous sources: No.

Table 2.1: Imputations of the Participation Rights

authors (e.g., Myerson and Satterthwaite [25]) that mechanisms so developed lack credibility in situations where the agents are free to abstain from the choice process. To account for the voluntary aspect of participation, extant models impose on the designer the individual-rationality, or participation constraint. We shall see shortly that this entitles the agent to a choice outcome and entails his complete engagement. In Table 2.1, this case of guaranteed participation is depicted by cell 2a:

Given that the imputations 1 and 2a have been extensively studied in the literature, in the remainder of the paper we concentrate mostly on the remaining distribution of the participation rights | mechanisms with voluntary but not guaranteed participation.

2.2 A Formal Description of the Participation Rights

As we discussed in the preceding section, participation rights in possession of an agent create duties on the part of the designer and thereby restrict his choice of the mechanism. Formally, these restrictions appear below as the free-exit and the free-entry constraints (2.6) and (2.7). We begin their derivation with the notational conventions and a brief statement of the mechanism-design problem (see, e.g., Myerson [24], Fudenberg and Tirole [10], Mas-Colell, Whinston and Green [17] for a complete exposition). We do so in su±cient detail to accommodate explicitly the presently examined participation-related issues.

The choice problem being facilitated by a mechanism deals with selection of a choice outcome from a given set of alternatives, X^0; made jointly by a set of agents, I: It may be resolved with no selection from X^0 having been made. With such no-choice ending ν; the set of allocations is

$$X = X^0 \cup \{\nu\}. \tag{2.1}$$

The information sets of each agent $i \in I$ are indexed by the type of that agent — the real numbers in a given set T_i. The common and commonly known ex ante beliefs regarding the likelihood of types are described by a probability P on the set of type combinations, $T = \prod_{i \in I} T_i$. The agents' type-contingent preferences over the outcomes in X^0 are represented by a vector of their utility functions, $u : X^0 \times T \to R^I$. If instead of a choice outcome an agent i pursues exogenous opportunities, he attains his interim reservation utility U_{i0}.

The mechanism designer attempts to induce the revelation of information held privately by the agents. Toward this end, he offers them to participate in a game, Γ, to be played with him and Nature. Nature chooses random realizations, and the strategy of the designer is an arbitrary allocation-valued function $\mu : S \to X$. We refer to μ as mechanism [2] whenever its domain has a product form $S = \prod_{i \in I} S_i$, in which case the elements of a factor S_i are viewed as actions available in Γ to player i. A Bayesian strategy of an agent i in Γ is a function $\xi_i : T_i \to S_i$ that specifies an action $\xi_i(\tau) \in S_i$ at each information set $\tau \in T_i$. For a type realization $t \in T$, the allocation resulting from a strategy combination $\xi = \prod_{i \in I} \xi_i$ is described by the social choice function $\Phi : T \to X$ defined as a composition: $\Phi = \mu \circ \xi$. This function is said to be implemented by a mechanism μ in a Bayesian equilibrium ξ if, given μ, the strategy ξ_i of each agent i is a best reply to the strategy ξ_{-i} of all others.

To illustrate, consider a firm contracting with two independent agents. As usual, we think of a contract as a wage-action pair, in which the specified action is to be taken in consideration of the wage offered by the firm. Let C_i be the sets of the institutionally feasible contracts that may be awarded to an agent $i \in I = \{1, 2\}$. Then the choice set, $X^0 = C_1 \cup C_2 \cup (C_1 \times C_2)$, has a three-component structure: the outcomes in $C_1 \times C_2$ obtain with the unanimous participation of the agents, whereas C_1 and C_2 contain all contractual obligations that may be accepted by the respective agent when the other abstains. An agent's abstinence from choice manifests itself formally as his exit action, which for $i \in I$ we denote as a_i^e. To fix ideas further, suppose that in the game Γ each agent i makes an offer $c \in C_i$ to the firm. The strategy space is, then, a disjoint union: $S = (C_1 \times C_2) \cup (C_1 \times \{a_2^e\}) \cup (\{a_1^e\} \times C_2) \cup (\{a_1^e\} \times \{a_2^e\})$. Since it has a product form, any allocation function $\mu : S \to X$ is a mechanism. The no-trade ν ending is allocated by μ whenever both agents refuse to contract with the firm, that is, $\mu(\{a_1^e\} \times \{a_2^e\}) = \nu$. All other components of S are mapped to the choice set X^0. Given a strategy combination ξ in Γ, the sets $\xi_1^{-1}(\{a_1^e\} \times (C_2 \cup \{a_1^e\}))$ and $\xi_1^{-1}(C_1 \times (C_2 \cup \{a_1^e\}))$ comprise respectively the nonparticipating an the participating types of the first agent.

More generally, agents' participation in the collective choice enters the formal description of the mechanism as follows. Whenever the legal environment of the mechanism allows an agent to abstain from choice, his action set is comprised of two components: for each voluntarily participating agent $i \in I$,

$$S_i = S_i^0 \cup S_i^e, \tag{2.2}$$

where the disjoint subsets S_i' and S_i^e contain the actions that amount to the participation in and, respectively, the abstinence from the collective choice.

The strategy space S of Γ is partitioned by the "participating coalitions." If the agents in a set $A \subseteq I$ participate and the rest abstain, let S_A be the set of their action combinations: $S_A = \prod_{i \in A} S_i'$. For any two different groups A' and A'', the sets $S_{A'}$ and $S_{A''}$ are disjoint, and the sets of the form $S_A \times \prod_{i \notin A} \{a_i^e\}$ collectively cover S. Hence, the collection of all sets S_A is, up to an isomorphism, a partition of the state space S.

The participation-related restrictions discussed in the preceding section limit the mechanism in terms of the outcomes of X' it can allocate. The designer may be required, for instance, to effect the no-choice ending ν unless one "participation coalition" from a given collection G takes part in the process. The unanimous involvement, for instance, obtains with $G = I$. Conversely, if the designer allocates a choice outcome, he observes some action combination in the set $\mu^{-1}(X') = \bigcup_{A \in G} S_A$ that is taken by a type realization in the set

$$T_G^p = \Phi^{-1}(X') = \xi^{-1}\left(\bigcup_{A \in G} \prod_{i \in A} S_i'\right). \tag{2.3}$$

The pullback of (2.2) to T along a strategy combination ξ allows us to classify the agent-types according to their involvement: for each agent i, the inverse-image sets

$$T_i^p = \xi_i^{-1}(S_i'), \qquad T_i^n = \xi_i^{-1}(S_i^e) \tag{2.4}$$

are respectively the participation set and the nonparticipation set of types of that agent. Naturally, either of these sets may be empty. We shall say that the mechanism is characterized by the complete engagement of the agent i whenever $T_i^p = T_i$. Alternatively, we shall say that the selective engagement of i takes place if the set T_i^n is not empty, and T_i^p is a proper subset of T_i. The choice-effecting type combinations (2.3) are uniquely specified by the scope of participation as follows:

$$T_G^p = \bigcup_{A \in G} \prod_{i \in A} T_i^p. \tag{2.5}$$

Before concluding the formal treatment of participation, we note that the designer, in complete similarity with the agents, may also be free to decline his part in the collective choice. This feature is standard in principal-agent models, where the principal acts as mechanism-designer and may forego contracting with the agent if extensive exogenous opportunities are available to him. In the setting of Baron and Myerson [1], for example, the regulatory agency may choose not to control a monopolist whose technology it deems to be more useful elsewhere. Formally, the designer's strategy set, too, has a two-component structure: besides the class of X-valued mechanisms, it contains the designer's exit function, a_D^e. Thus, the no-choice allocation ν may result not only from the agents' abstinence exhibited in the interim but also from the designer's ex ante denial of a game Γ.

The formal description of participation rights, to which we devote the remainder of the section, reflects the earlier discussed symmetry between them. To see the specifics, consider first the exit right and assume that it is imputed to the designer, so that the agent's participation is involuntary. Since the designer is able to enforce the agent's involvement exogenously, he is free to eliminate from Γ all exit actions by letting $S_i^e = \emptyset$. The action set (2.2) with two components in general is thereby reduced to having at most one. Whether or not the designer will utilize his freedom to make S_i^e empty is the subject of mechanism optimization. The role of the exit right is to make this option feasible for him.

Suppose next that the exit right of an agent i lies with the agent himself who is thereby empowered to take the exit action unilaterally. The component S_i^e of the action set (2.2) is granted to the agent exogenously, as part of the institutional setting. Thus, the imputation of the exit right to an agent i creates a duty for the designer, namely, to make the exit component S_i^e in the game Γ nonempty.

The duty not to interfere with the agent's pursuit of exogenous opportunities has a simple but essential implication for the social choice function φ: if some allocation is reached with a given agent abstaining, it must yield to that agent his reservation level of utility. To formalize, let $t \in T$ be a type realization, σ a strategy combination, and τ_i a nonparticipating under σ type of an agent $i \in I$, so that the allocation $x = \mu(\sigma_i(\tau_i), \sigma_{-i}(t_{-i}))$ obtains with i abstaining. In terms of the interim expectation of his utility, $U_{i0}(t_i)$, the "exogeneity-of- exit" property of φ may be expressed as follows: for each agent $i \in I$ whose participation is voluntary,

$$\forall t_i \in T_i,\ \forall \tau_i \in T_i^n, \qquad E\left[u_i\left(\varphi(\tau_i, t_{-i}), t\right) \mid t_i\right] = U_{i0}(t_i). \tag{2.6}$$

This condition, to which we shall henceforth refer as the free-exit constraint, is imposed on the social choice function, in contrast to (2.2) which restricts implementations.

In similarity with the right of free exit, the entry right creates a duty for the designer only when it is imputed to the agent. Recall from the preceding section that such empowerment of the agent makes his nonparticipation voluntary; he can force his participation onto the designer. The agent is assured therefore of at least one choice outcome. Whereas the allocation specified by φ belongs to the extended set (2.1), the designer is presently bound to make available at least one outcome in the smaller set X^0. Since allocations in X^0 are reachable, by definition, only by a participating agent-type, the component S_i^0 of the action set (2.2) is nonempty. An implication for the social choice function is also immediate: for each agent $i \in I$ endowed with the entry right,

$$\forall t_i \in T_i^p, \qquad E\left[u_i\left(\varphi(t), t\right) \mid t_i\right] \geq U_{i1}(t_i), \tag{2.7}$$

where $U_{i1}(t_i)$ is the utility expected in the interim of the guaranteed choice outcome. To the restriction (2.7) we shall refer as the free-entry constraint. An alternative imputation of the entry right to the designer imparts no such duty on him.

The foregoing makes explicit the aforementioned symmetry between the participation rights of exit and entry. Indeed, the exit right deals with the participation of the agent, whereas the entry right addresses his nonparticipation. When the agent has the right of free exit (entry), his participation (nonparticipation) in the collective choice is voluntary for him. Further, the exit right determines whether in the implementation Γ the action-set component S_i^e is empty, while the entry right similarly regulates the size of the component S_i^0. And, whereas the exit right imputed to the agent ensures via the free-exit condition (2.6) that the exogenously derived reservation utility is attainable, the entry right guarantees him via the free-entry constraint (2.7) a certain level of utility from the endogenous sources.

The entry-exit symmetry is reflected in Table 2.1, in which the scenarios characterized with both utilities, U_{0i} and U_{1i}, require further attention. The fact that the utility level U_{1i} is endogenously available reflects the entry right in the possession of the agent, as in Cells 2 and 4 of the table. The two cases differ in the reasons for the availability of the reservation utility U_{0i}: whereas in Cell 2 the reservation utility is made available exogenously by the exit right, in Cell 4 it arises endogenously as an act of the designer. In the latter case, the exit actions are neither prescribed nor proscribed, hence under the discretion of the designer. Observe that the rationality of both parties makes both situations behaviorally equivalent. Specifically, scenarios 2a and 4a always entail the agent's participation: either he has no other options, such as when the exit actions are not provided in 4a; or he prefers to do so because it improves his welfare. In the remaining two cases, 2b and 4b, the entry right is inconsequential since a better (exogenous) source of utility is available to the agent; the agent's behavior remains ambiguous but unaffected by his entry right.

The foregoing analysis enables us to interpret the individual-rationality constraint. It is easy to see that it obtains with $U_{i1} = U_{i0}$ (2.7). This is the source of the earlier-mentioned conclusion: rather than reflecting the voluntary aspect of an agent's participation, the individual-rationality constraint guarantees it to him. It creates an entitlement to receive the reservation level of utility endogenously, from a choice outcome. Scenarios 2a and 4a of Table 2.1 make the implications clear as well: when the participation constraint is required to hold for all agent-types, the designer's engagement of the agent is complete.

In mechanisms with voluntary but not guaranteed participation, the designer is free to engage the agents selectively. As the first step toward optimization of such mechanisms, we provide in the next section a convenient characterization of mechanism feasibility.

2.3 Incentive-Compatibility under Voluntary Participation

The agents' participation rights (2.6)–(2.7) apply to an arbitrary strategy combination ξ adopted by the agents in the game Γ offered to them by the designer as part of the mechanism μ. From the Revelation Principle (Gibbard [11], Green and Laffont [12], Myerson [22], Dasgupta, Hammond and Maskin [3]) we know that, when ξ is an equilibrium in Γ, the implemented by μ social choice function $\varphi = \mu \circ \xi$ is Bayesian incentive-compatible:

$$\forall i \in I, \qquad \forall t_i, \tau_i \in T_i, \qquad E\left[u_i\left(\varphi(t), t\right) \mid t_i\right] \geq E\left[u_i\left(\varphi(\tau_i, t_{-i}), t\right) \mid t_i\right]. \tag{2.8}$$

This circumstance allows the well-known replacement of the original choice problem addressed by the mechanism μ by an equivalent one that is facilitated by a so-called direct mechanism, μ^d. The choice set of the extended problem is the set of allocations (2.1). The mechanism μ^d is based on the direct-revelation game Δ, in which the agents' actions are their reports, perhaps false, of their respective types to the designer: $S_i^d = T_i$; and the strategies are $\xi_i^d : T_i \to S_i^d$. With the action set of the mechanism being T, it is by definition direct. From (2.8) and the Revelation Principle it follows that whenever an incentive-compatible function Φ is implementable by μ in an equilibrium ξ of the original game Γ, it is also implementable by a direct mechanism μ^d in truthful revelation of type $\xi^d = \iota_T$, where ι_T denotes the identity on T.

The two choice problems and their respective mechanisms μ and μ^d have a notable distinction in terms of the agent's participation: although it may be voluntary in the original game Γ, the agent's involvement in the direct-revelation game Δ is always involuntary. Clearly, the action-sets T_i of Δ have no exit components; every agent- type reports to the designer. Since involvement in Δ is involuntary, no incentives beyond those for truthful revelation of type are needed. It is for this reason that the free-exit constraint (2.6) is void of incentives. By satisfying (2.8) alone the designer provides to the agents all the inducements he finds necessary.

To see explicitly the sufficiency of the incentive-compatibility, we specialize it below to the participating and the nonparticipating types. Both the type realization and the report of an agent may be of either kind, so we consider four possibilities.

Participating Agent-Types: $t_i \in T_i^p$. If the social choice function Φ is incentive-compatible, the agent-type t_i cannot gain by reporting itself falsely. This is so regardless of whether he misreports himself as another participating type $\tau_i \in T_i^p$, or a nonparticipating $\tau_i \in T_i^n$. In the former case, we have the following (optimality-of-)truth-telling constraint, which is merely a restatement of (2.8):

$$\forall t_i, \tau_i \in T_i^p, \qquad E\left[u_i(\Phi(t), t) \mid t_i\right] \geq E\left[u_i(\Phi(\tau_i, t_{-i}), t) \mid t_i\right]. \tag{2.9}$$

We reserve the term truth-telling for (2.9) to differentiate it from its antecedent (2.8). They differ only in scope: whereas (2.8) is imposed on all of the agent-type pairs, $t_i, \tau_i \in T_i$, the truth-telling condition (2.9) holds only on the participating ones: $t_i, \tau_i \in T_i^p$. The intuition behind the reduction in scope is clear: once the participation of t_i is secured (via an appropriate restriction; see below), he is not tempted to mimic the abstaining types. The only remaining ones are of the participating kind, and the role of the truth-telling constraint is to make these reports unprofitable.

When the type $t_i \in T_i^p$ mimics a nonparticipating τ_i, we may utilize the free-exit constraint (2.6) in the right hand side of (2.8). The inequality no longer involves τ_i, and becomes the participation constraint:

$$\forall t_i \in T_i^p, \qquad E\left[u_i(\Phi(t), t) \mid t_i\right] \geq U_{i0}(t_i). \tag{2.10}$$

Our use of the name "participation constraint" should not cause any confusion: it is merely a special case of the incentive- compatibility constraint. Rather than being logically independent, it is a joint consequence of the incentive-compatibility constraint (2.8) and the free-exit constraint (2.6). Note also that it acts on the participating type values only.

The reduction of scope in (2.10) is intuitive. Suppose that the truth-telling has been secured by the restriction (2.9). Then t_i has no incentives to mimic the participating types. The only remaining ones are of the nonparticipating kind, and (2.10) precludes gains from such false reports.

Thus, the truth-telling and participation incentives are decoupled from each other, and their formal descriptions are given by (2.9) and (2.10). Note that it is the scope reduction from T to T^p in both restrictions that makes selective engagement feasible.

Nonparticipating Agent-Types: $t_i \in T_i^n$: A nonparticipating agent-type t_i; too, may misreport himself as a participating type or another nonparticipating one. The counterpart of (2.9) is trivially satisfied: the free-exit constraint (2.6) is presently applicable in both sides of the inequality, which turns it into an identity.

The remaining condition, to which we refer as the nonparticipation constraint, removes incentives for a nonparticipating type t_i to misrepresent himself as a participating type τ_i : for all agents $i \in I$ and $t_i \in T_i^n$;

$$\forall t_i \in T_i^n;\ \forall \tau_i \in T_i^p; \qquad U_{i0}(t_i) \geq E\,[u_i\,(\varphi(\tau_i; t_{-i}); t) \mid t_i]: \tag{2.11}$$

The preceding analysis leads to an equivalent representation of the incentive-compatibility:

Equivalence Theorem 2.1 Let a social choice function $\varphi : T \to X$ describe the allocations to a set of agents I with type combinations in a set $T = \prod_{i \in I} T_i$: Suppose that the agent's participation in the collective choice is voluntary, and an abstaining agent i derives his reservation utility $U_{i0} : T_i \to R$ from exogenous sources.

Then the function φ is incentive-compatible if and only if for every agent $i \in I$ it satisfies:

(i) the truth- telling constraint (2.9),

(ii) the participation constraint (2.10), and

(iii) the nonparticipation constraint (2.11).

The theorem explicates the incentives provided by the designer: to induce participation (from those whom he wants to engage), nonparticipation (from those whom he wants to abstain), and truth-telling (by the participating agent-types). The designer's reliance on incentives emerges in (2.9)–(2.11) for rather different reasons. The truth- telling, as is well known, is necessitated by the informational asymmetry between the designer and the agent. Formally, this may be seen from the fact that a perfectly informed designer can easily satisfy this condition by the so-called forcing mechanism that inflicts severe disutilities on the agents who falsify their types. In fact, not only (2.9) but also (2.11) is satisfied by a forcing mechanism. However, infliction of disutilities for abstinence from choice is expressly forbidden by the exit right. Consequently, even the mechanism of a perfectly informed designer is subject to the participation constraint (2.10). The incentives provided to the agents induce them to forego the perfectly effective exit right in their possession. As we can see, the incentive-compatibility is required even when the designer is perfectly informed.

As an alternative, suppose that all participation rights lie with the designer. When perfectly informed, he is then entirely unrestricted in his choice of mechanism: a forcing mechanism applies to all three constraints (2.9)–(2.11). With only imperfect information, however, the designer finds himself at the other extreme — his choice is subject to all of these constraints. Of course, having the exit rights in his possession, he is still able to force the agents to participate. If given also the entry rights, he can exclude them as well. However, the lack of information reduces the efficacy of the participation rights: although their enforcement is still possible on the agent-specific basis, the type-specific enforcement is not, and the designer resorts to incentives.

The operationalization of incentive-compatibility offered by the Equivalence Theorem affords us a convenient criterion of mechanism feasibility. We shall exploit it in Section 3.1 in the formulation of the designer's second-best optimization problem. Our present objective is to isolate the restrictions stemming from two sources: participation rights and information.

Towards this end, we shall say that a mechanism is perfectly feasible if it can be implemented by a perfectly informed designer in possession of all participation rights. Perfect feasibility is a convenient benchmark in that it leaves the designer most free in his choices (he remains a subject to the idiosyncratic restrictions of the institutional setting). The earlier discussed participation-related restrictions may be recalled as an example. In particular, the unanimity of participation is required in all one-agent choice problems — most notably, the principal-agent relationships. Some multi-agent settings, such as bilateral trading, place this restriction on the designer as well.

Further, we shall say that a mechanism is participation-feasible if the possession of the participation rights alone makes it implementable. Similarly, a mechanism is informationally feasible if a perfectly informed designer can implemented it even when all participation rights lie elsewhere. Mechanisms in each of these two classes are, of course, perfectly feasible.

Under voluntary participation, we distinguish between the first- and the second-best mechanism depending on whether the designer is perfectly or partially informed. We shall refer to the corresponding feasibility as being of the first and, respectively, second degree. The impact of voluntary participation on mechanism feasibility may be summarized as follows.

Proposition 2.2 For collective-choice problems with voluntary participation:

(i) A mechanism is first-degree feasible if it is perfectly feasible and, in addition, satisfies the free-exit condition (2.6).

(ii) A mechanism is second-degree feasible if it is first-degree feasible and incentive-compatible or, equivalently, if it satisfies the conditions of free-exit (2.6), truth-telling (2.9), participation (2.10), and nonparticipation (2.11).

To summarize, the impact of voluntary participation is entirely localized by the free-exit condition. The reduced in scope participation constraint merely reflects incentive-compatibility and allows for the selective engagement of the agents. For mechanisms with voluntary participation, the Equivalence Theorem presents the designer with an alternative: either (a) the incentive-compatibility alone should be satisfied on all type values, or (b) a combination of truth-telling and participation be required to hold for the participating types only, in which case they are complemented by the nonparticipation condition for the remaining type values.

We presently adopt the second approach: as the next section indicates, it improves methodological tractability and provides for a parsimonious formalization of the designer's optimization problem.

3 Optimal Mechanisms

In this section, we formulate the mechanism-optimization problem (3.2)–(3.5) and identify the generic sources of welfare gains from selective engagement of the agents (Section 3.2).

3.1 The Mechanism-Optimization Problem

The optimization task of the designer is to maximize, by a judicial selection of a feasible mechanism, the social value of the allocation made by the mechanism. This value is represented by the designer's von Neumann-Morgenstern utility function, $u^D : X \times T \to \mathbb{R}$. In particular, $u^D(\nu, t)$ is the social value of the exogenous alternatives.

When comparing mechanisms and their social choice functions, the designer is guided by the induced preferences. For a perfectly informed designer, they are type-contingent and defined on the set $X^T \times T$ by $V(\Phi, t) = u^D(\Phi(t), t))$. The preferences of a less informed designer are represented on X^T by his ex ante expectation of V; namely, $U^D(\Phi) = \int_T u^D(\Phi(t), t) P(dt)$.

To illustrate, consider the principal-agent model with hidden information, where the choice set X^0 is comprised of action-wage pairs (contracts) $(\hat{x}, \hat{y})$, so that $X = \mathbb{R}^2 \cup \{\nu\}$. A social choice function Φ allocates either a contract from a menu $(x, y) : T \to \mathbb{R}^2$, or the “no-trade” ending ν. Both the contracting and the abstaining agent-types contribute to the social welfare, which has the following ex ante expectation in terms of the sets (2.4):

$$\int_{T^p} u^D((x(t), y(t)), t) P(dt) + \int_{T^n} u^D(\nu, t) P(dt).$$

After the subtraction of the decision-independent quantity $\int_T u^D(\nu, t) P(dt)$, we find that the principal's welfare may be equivalently viewed as the expectation of the contract value net of the opportunity costs:

$$U^D(\Phi) = \int_{T^p} \left[u^D((x(t), y(t)), t) - u^D(\nu, t) \right] P(dt). \tag{3.1}$$

In more general, multi-agent choice problems, the domain of integration is the defined in (2.3) set T^p_G rather than T^p.

We continue to focus on the settings with voluntary but not guaranteed participation, the feasibility of which is specified in Proposition 2.2. The designer's optimization problem is thus to find

$$\max_{\Phi, T^p} \int_{T^p_G} \left[u^D(\Phi(t), t) - u^D(\nu, t) \right] P(dt) \tag{3.2}$$

subject to the truth-telling, the participation, and the nonparticipation constraints: for all $i \in I$,

$$\forall t_i \in T_i^p, \quad E[u_i(\varphi(t); t) \mid t_i] \ge \sup_{\tau_i \in T_i^p} E[u_i(\varphi(\tau_i; t_{-i}); t) \mid t_i]. \tag{3.3}$$

$$\forall t_i \in T_i^p, \quad E[u_i(\varphi(t); t) \mid t_i] \ge U_{i0}(t_i) \tag{3.4}$$

$$\forall t_i \notin T^p, \quad U_{i0}(t_i) \ge \sup_{\tau_i \in T_i^p} E[u_i(\varphi(\tau_i; t_{-i}); t) \mid t_i]. \tag{3.5}$$

The form of the designer's problem (3.2)–(3.5) shows that the economically relevant domain of the social choice function is the set T_G^p. Since $\varphi(t) = \emptyset$ for all $t \notin T_G^p$, there are no choices to be made at such type values. The set T_G^p is uniquely specified by the scope of participation T^p, as may be seen from (2.5). The designer has therefore two independent optimization instruments:[3] the scope of participation T^p, and the form of the social choice function φ on the set T_G^p.

The traditionally assumed complete engagement obtains from (3.2)–(3.5) with a fixed $T^p = T$ and optimization with respect to φ only. The nonparticipation constraint (3.5) is then vacuously satisfied, and the remaining constraints (3.3)–(3.4) become respectively the incentive-compatibility condition and the participation constraint in their traditional forms. Since the optimization with respect to φ is partial, the complete-engagement mechanism may be suboptimal.

It is instructive to compare the present optimization problem with that of the designer to whom the agents' actions rather than their types are unknown. Concentrate for the moment on the single-agent case, that is, two principal-agent models with, respectively, hidden information (adverse selection) and hidden actions (moral hazard). It is easy to see what makes the former considerably richer in terms of economic content. In both cases the designer has two optimization instruments at his disposal, of which one is a function: the social choice function φ in hidden information, and the wage schedule for hidden actions. The two differ significantly in the nature of the second instrument: whereas the hidden action is a number, the scope of participation T^p is a set. In the case of hidden information, therefore, the principal has infinitely greater flexibility in mechanism optimization.

The additional degrees of freedom in the hidden-information setting are reflected in the characterization of the optimal contract. Recall from Holmström [14] that the optimal hidden-actions contract is characterized by two first-order conditions: the Euler equation, which governs the optimality of the wage function; and the Fermat condition of the usual calculus, which determines optimality of the action allocation. A complete characterization of the hidden-information contract, which is presented in the companion paper Faynzilberg [5], requires many more conditions of optimality. A relatively simple example of Section 4.2 shows that besides the Euler equations, multiple occurrences of transversality, corner, and free-end conditions are also generally needed. This is in addition to the Kuhn-Tucker and the rent-extraction conditions, which are non-differential in nature. For the model of Section 4.2, the entire ensemble takes the form of the equations (4.28)–(4.33). In comparison to the hidden-actions case, optimization becomes far more complex.

Observe that the present framework is rather general in that it does not rely on any structural features of the choice model. The continuity properties of the probability P have been left unrestricted. More importantly, the proposed approach applies equally to the discrete and continuous cases: both the choice set X^0 and the set T may contain as components continua as well as discrete subsets.

When the set T is discrete, the optimization problem (3.2)–(3.5) acquires a combinatorial flavor. Interestingly, even when T is continuous, the possible pooling ambiguities make a discrete-optimization step necessary as well. More specifically, recall that the designer often pools the types by assigning to them the same allocation. The analysis in Faynzilberg [5] shows that the pooling regime is not generally unique. The corresponding social choice functions are quantized and form a discrete spectrum of the locally optimal mechanisms. The methods of continuous optimization, which produce the spectrum, must therefore be followed by a discrete optimization within the mechanism. The second stage of this process may be nontrivial because the spectrum may comprise dozens of members, even when the informational asymmetry has relatively simple characteristics.

To conclude, optimization of mechanisms with voluntary participation, which is formally represented by (3.2)–(3.5), calls for the selective engagement of the agents. The scope of participation T^p is determined endogenously and jointly with the functional form of the social choice function Ψ. While optimization may result in complete engagement, the result $T^p = T$ may not be presumed without loss of optimality.

3.2 The Economic Impact of Selective Engagement

A compete assessment of the impact of selective engagement on the mechanism requires, naturally, the knowledge of the optima in both settings — with and without the participation constraint. We postpone this line of analysis until Section 4 and presently provide some intuition behind the welfare gains from selective engagement. An inspection of the designer's optimization problem (3.2)–(3.5) reveals that they stem from three generic sources.

To begin, the designer's choice of the participation scope is an informational signal to all the parties to the mechanism. Taken together, the sets (2.4) form a dichotomous classification scheme. As is well known, classification reduces uncertainty: for a given scope of participation T^p, the likelihood of a participating type is given by the conditional on T^p probability P, whereas the prior uncertainty is expressed by P itself, hence generally greater.

The reduction in uncertainty may be quantified in terms of informational measures. Recall (from, e.g., Martin and England [16], Faynzilberg [4], Faynzilberg and Kumar [9]) that the information function $I(p) = -\log p$ and its expectation, entropy, are natural measures of information in a probabilistic experiment. Here p is the density of the probability P. Conditioning reduces entropy (see, e.g., Cover and Thomas [2]). We conclude that selective engagement is always informative. Further, any reduction in uncertainty makes the observer informationally closer to the perfectly informed agent, so the informational asymmetry between them is reduced, hence the distortions of the mechanism subside as well. The reduced incentive-compatibility pressures mitigate the information-related welfare loss.

Secondly, the mere exclusion of the agent-types that contribute negatively to the social welfare improves the mechanism. Consider, for example, a monopoly pricing its good for consumers with uncertain tastes. The firm may find it unprofitable to serve some parts of the market: even the optimally chosen price need not offset the production costs. Under these circumstances, a mere withdrawal from the unprofitable segments improves the expected profits of the monopolist. In general, the removal from T of a P-positive set of types on which the integrand of U^D is negative, leads to a strictly positive change in the value of welfare function U^D: Observe that such deletion makes it easier to satisfy the truth-telling and the participation constraints. The nonparticipation constraint tightens, however, so the deletion need not be feasible.

Finally, we observe that selective engagement may qualitatively change the informational structure of the collective-choice problem. To illustrate, return to the type-pooling phenomenon. Recall that the designer resorts to this costly device only when his ideal, that is, first-order incentive-compatible allocation does not have appropriate monotonicity that oftentimes results from the second-order considerations. Most of the settings in the extant literature, as well as the principal- agent models of Section 4, admit effective type $E(t)$ (Faynzilberg [5]), the monotonicity of which determines the need for pooling. In the companion paper we find that the type premium, $E(t) - t$; is proportional to the shadow costs of incentive-compatibility, which depend strongly on the participation scope T^p: Hence, the latter may significantly alter, via the shadow costs, the functional forms of the effective type $E(t)$ and of the ideal allocation. Although the complete-engagement mechanism may have involved pooling because the ideal allocation was infeasible, the exclusion from participation of some types may obviate the need for pooling of the rest.

The three aforementioned benefits of selective engagement do not come without costs. The explicit form of (3.2)–(3.5) shows that these costs are of two kinds: the loss of the contribution of the excluded types, when this contribution is positive; and the incentives reflected in the nonparticipation constraint (2.11).

The models of Section 4 suggest that, when the preferences satisfy the usually assumed sorting condition, the nonparticipation constraint is rather mild: its effect is localized to the marginal types on the participation-nonparticipation boundary. As for the positive contribution of types foregone by their exclusion, it may be more than offset by the easing of incentive-compatibility pressures on the remaining types.

The balance between the benefits and costs of selective engagement is resolved by the simultaneous optimization of the welfare value (3.2) with respect to the social choice function, and the scope of agents' participation in the collective choice.

4 Optimization Methodology and Applications

In the present section, we extend the analysis of the principal-agent model with hidden information that has been extensively studied in the economic literature (e.g., Mirrlees [20], Mussa and Rosen [21], Baron and Myerson [1], Maskin and Riley [18]). Following the rather general approach of Guesnerie and La®ont [13] to such single-agent mechanisms, we distinguish and illustrate two classes of models | depending on whether or not the principal is indi®erent with respect to the monetary transfer.

In both cases, we exhibit explicitly the optimal contract. In doing so we pursue three objectives. Firstly, we quantify the bene¯ts of selective engagement outlined in Section 3.2. By comparing the value of the optimal contract to the traditionally used complete-engagement regime, we determine the social cost of allocating the entry right to the agent. Alternatively, the drop in the value of the mechanism may be viewed as the cost of imposing the participation constraint on the designer. Secondly, we show that the informational asymmetry causes the mechanism to undergo two kinds of distortions: in the form of the optimal contract ©; and in the scope of participation T^p: Thirdly, we present two optimization methodologies, the ¯rst of which involves a two-stage optimization and has the advantage of showing explicitly the interdependence of © and T^p: The second methodology is more e±cient. It adopts and illustrates the complete characterization of the optimal contract developed in Faynzilberg [5].

The agency settings considered below in Sections 4.1 and 4.2 have the following common features. The principal acts as mechanism designer and to those agent-types that he wishes to engage he o®ers a contract, $(\hat{x}; \hat{y})$; in which $\hat{x}$ is an action the agent must take in consideration of a pecuniary transfer, $\hat{y}$: The set of allocations is $X = R^2$ [fºg; where the ¯rst component contains contracts and º is the no-trade ending that leads to the reservation utilities U_0^P and U_0^A of the respective parties. The parties' preferences for contracts are separable in the action and money. By P we shall henceforth denote the cumulative distribution function of the respective probability rather than the probability itself. Its density with respect to the Lebesgue measure is denoted p:

4.1 The Principal-Agent Model: A Risk-Neutral Principal

We postpone characterization issues to the next section and presently rely on a two-stage optimization procedure. Besides being e®ective in its own right, it allows us to demonstrate explicitly the mutual dependence between the form of the optimal contract and the set of types to which it is o®ered. In Section 5 we shall compare this method with the step-wise methods developed in the auctions (Myerson [23]) and agency literatures (Baron and Myerson [1] and Guesnerie and La®ont [13]).

As we mentioned earlier, two of the objectives are to assess the information-related distortions of the mechanism, and to quantify the cost of the entry right. Toward this end, we construct two benchmarks: the ¯rst-best optimum (Proposition 4.3), and the best complete-engagement regime (Proposition 4.2). They are subsequently compared to the second-best contract exhibited in Proposition 4.3. In the notation for these contracts we employ superscripts: zero for the ¯rst-best, star for the complete-engagement by the imperfectly informed principal, and double star for the second- best.

Following the classi¯cation of Guesnerie and La®ont [13], we ¯rst consider the following setting the with risk-neutral parties:

$$u^P(x, y, t) = x - y \tag{4.1}$$

$$u^A(x, y, t) = y - \frac{x^2}{t}. \tag{4.2}$$

The agent's type t, which is a real number in the unit interval $T = [0, 1]$, describes his aversion to action. We first leave the ex ante uncertainty regarding the type unspecified: this allows us to recognize the economically important effective type of the agent. To find the contract in closed form, we subsequently specialize P to the uniform distribution.

The perfectly-informed principal As we discussed in Section 3.1, the perfectly informed principal can rely on forcing contracts to elicit truth-telling and nonparticipation, but not participation, from the agent. Of the three kinds of inducements explicated in Theorem 2.1, only the participating incentives are utilized by the principal with the following results:

Proposition 4.1 Let the parties' respective reservation utilities be U^P and U^A and their preferences for contracts given by (4.1)–(4.2). Then the first-best contract,

$$\forall t \in T^{0p}, \qquad \psi^0(t) = \left(x^0, y^0\right) = \left(\frac{t}{2}, U_0^A + \frac{t}{4}\right), \tag{4.3}$$

is formed with the agent-types in the participation set

$$T^{0p} = T \cap \left[4(U^P + U^A), +1\right). \tag{4.4}$$

It yields to the principal the indirect expected utility $U^{P0} = t/4 - U_0^A$.

Not surprisingly, even the perfectly informed principal benefits from selective engagement. The size of the participation set in (4.4) may range from complete withdrawal of the principal ($T^{0p} = \emptyset$) to complete engagement of the agents ($T^{0p} = T$). Whenever at least one of the parties has lucrative exogenous opportunities, so that $U_0^P + U_0^A > \inf T = 0$, contracting with the agent-types $t < 4(U_0^P + U_0^A)$ becomes prohibitively costly for the principal. Being perfectly informed of the type value, the principal excludes these values from T^{0p}, i.e., engages the agent selectively. The extreme case of selective engagement is withdrawal, which the principal prefers whenever $4(U_0^P + U_0^A) > \sup T = 1$. Here the principal extends to the agent no offer whatsoever, which is formally indicated in (4.3) by the empty domain of the function ψ. The principal's exit action a^{eP} leads to an ex ante breakdown in contracting, i.e., the no-choice allocation ξ.

The principal's optimization problem The imperfectly informed designer need not choose a simply connected scope of participation (Faynzilberg [7]): the set T^p may be comprised of several disjoint intervals of the set T. Given the objectives of the present paper, we omit the details which rule out such possibility and assume that the participation set is an interval: $T^p = [a, b] \subseteq [0, 1]$ for some $0 \le a \le b \le 1$. Optimization with respect to T^p involves, then, the selection of two real variables, a and b.

Proposition A.1 and Remark A.2 of the Appendix show that, instead of the transfer y, the principal may optimize with respect to a constant, $U(a)$, which is the utility of the least participating type. With this modification, the principal's problem (3.2)–(3.5) is to maximize the welfare function,

$$U^P(x;a;b) = \int_a^b L(x(t);t;a;b)p(t)dt; \tag{4.5}$$

the Lagrangian density of which,

$$L(x;t;a;b) = \left(x - \frac{x^2}{E(t;b)} - U(a) - U_0^P \right) p(t); \tag{4.6}$$

is expressed in terms of the effective type $E(t;b)$:

$$E(t;b) = \frac{t^2 p(t)}{tp(t) + P(b) - P(t)}. \tag{4.7}$$

As the first step in solving the principal's problem, we fix an arbitrary scope of participation and optimize U^P with respect to the rest of degrees of freedom: the contract $(x;y)$ and the constant $U(a)$. This stage results in a partially optimized contract $(x^*(t;a;b); y^*(t;a;b))$ that yields the scope-contingent utility $U^{P*}(a;b) = U^P(x^*;a;b)$ to the principal. The explicit form of this contract is given in (4.8)– (4.9) below. At the second stage, we determine the optimal scope $[a^{**};b^{**}]$ by maximizing $U^{P*}(a;b)$ with respect to the boundaries a and b. Proposition 4.3 presents the fully optimal, second-best contract in terms of its functional form, $(x^{**};y^{**}) = (x^*(t;a^{**};b^{**}); y^*(t;a^{**};b^{**}))$, and domain, $T^{p**} = [a^{**};b^{**}]$.

Stage 1: partial optimization Optimality of the contract with respect to $U(a)$ is attained by the complete rent extraction from the type $t = a$, i.e., $U(a) = U_0^A$. This outcome of optimization is far from general: although there always exists a type with zero rent from private information, it may lie anywhere within the participation set, and must be determined endogenously (Faynzilberg [5, 6]).

The ideal action allocation, x^*, may be found by inspection. Observe that the Lagrangian (4.6) is the principal's welfare function with the actual type t replaced by the effective type (4.7). Hence, the ideal allocation may be expressed in terms of its first-best counterpart (4.3) as follows:

$$x^*(t;a;b) = x^0 \circ E(t) = \frac{E(t)}{2}. \tag{4.8}$$

It is well known that the ideal allocation need not be feasible: while (4.8) satisfies the first-order optimality of truth-telling, it must also be monotonically increasing to satisfy the second-order condition of optimality. On the intervals where the effective type $E(t)$ is a decreasing function, the designer deviates from the ideal allocation and resorts to pooling.

The general theory of the effective type is presented in the companion paper Faynzilberg [5] which contains, in particular, an existence criterion for E. Not all choice problems admit effective type. When they do, however, the optimization task for the imperfectly informed designer is easily fulfilled: his ideal allocation $x^*(t)$ to type t is the first- best allocation to the effective type $E(t)$.

Most of the parameterizations encountered in the extant literature admit effective type. The results are usually expressed in terms of a related quantity, the so-called virtual type, which in the current setting appears in the generalized form

$$J(t; b) = \frac{t^2}{E(t; b)} = t + \frac{P(b) - P(t)}{p(t)}.$$

The special case, $J(t; 1)$, which corresponds to the complete engagement, has been encountered in a variety of mechanisms, such as auctions (Myerson [23]), bilateral trading (Myerson and Satterthwaite [25]), monopoly pricing (Maskin and Riley [18]), and regulation (Baron and Myerson [1]). Although the effective and the virtual types are related, it is the effective type rather than the virtual that affords the unified treatment of the general mechanism-design framework.

Returning to the parameterization at hand, we substitute (4.8) into (A.1) and (4.2) to find the ideal monetary transfer:

$$y^*(t; a; b) = U_0^A + \frac{E^2(t; b)}{4t} + \int_a^t \left(\frac{E(\tau; b)}{2\tau} \right)^2 d\tau. \tag{4.9}$$

The pair (4.8)–(4.9) is the optimal contract for an arbitrary but heretofore fixed scope $[a; b]$. An inspection of these expressions reveals two of the three effects of selective engagement, which we discussed in Section 3.2. The first of those effects manifests itself in the range of integration in (4.9) being reduced to $[a; t]$ by the exclusion of types below a. The third of the discussed effects may be seen in the very dependence on b of the effective type E.

Below, we use the partially optimized contract (4.8)– (4.9) in two ways. First, we specialize it to $[a; b] = [0; 1]$ and thereby find the best complete-engagement contract (Proposition 4.2). And, second, we optimize $(x^*; y^*)$ with respect to the scope of participation $[a; b]$ and obtain the second-best contract (Proposition 4.3). We henceforth specialize P to the uniform distribution: $P(t) = t$ and $E = t^2/b$. Since the ideal contract is feasible, it is also optimal.

The best complete-engagement contract is a special case of (4.8)–(4.9) with $a = 0$ and $b = 1$. It has the following form:

Proposition 4.2 Let the parties' respective reservation utilities be U^P and U^A; their preferences for contracts given by (4.1)–(4.2), and the distribution of types be uniform on the unit interval $T = [0; 1]$. Then the complete-engagement contracting is feasible if and only if $U_0^P + U_0^A \le 1/12$, in which case the best such contract has the form

$$\forall t \in [0; 1]; \qquad x^*(t; 0; 1) = \frac{t^2}{2}; \quad y^*(t; 0; 1) = U_0^A + \frac{t^3}{3}; \tag{4.10}$$

and yields to the principal the indirect expected utility $U^{P*}(0; 1) = 1/12 - U_0^A$.

It is easy to see that the complete-engagement contract (4.10) is suboptimal. To begin, observe that it calls for the employment of the agent even when he not contractible in the perfect- information situation. To fix ideas, let $U_0^A = 1/12$ and $U_0^P = 0$. According to Proposition 4.1, the principal excludes the agent-types in the interval $[0, 1/3)$; offers the first-best contract $(x^0, y^0) = (t/2, t/4 + 1/12)$ to the rest, and attains the utility $t/4 - 1/12$ from the agent of type t. Under the complete-engagement (4.10) the imperfectly informed principal is forced to employ the previously excluded agent-types $t < 1/3$. With less information available to the principal these types can become only more costly, and their engagement is clearly deleterious to social welfare.

A direct examination of the contract (4.10) leads to the same conclusion. The ex post contribution of type t to the principal's welfare is

$$u^P(x^*, y^*, t)\Big|_{\substack{a=0\\ b=1}} = x^*(t; 0, 1) - y^*(t; 0, 1) = \frac{t^2}{2} - \frac{t^3}{3} - U_0^A, \tag{4.11}$$

which is negative when the reservation utility is positive and the values of t and small. In the preceding example of $U_0^A = 1/12$; half of the agent-types — those in the interval $[0, 1/2]$ — contribute negatively. The exclusion of these types will have a positive effect on the principal's welfare. Below we show that such exclusion is, in fact, feasible, and the second-best contract presented in Proposition 4.3 is formed only with the types $t \geq 2/\sqrt{12}$. It results in the principal's utility $U^{P**} = 0.032$ — strictly better than $U^{P*}(0, 1) = 0$ from the complete-engagement contract (4.10).

Stage 2: the second-best contract The gradient of the partially optimized welfare function $U^{P*}(a, b) = U^P(x^*, a, b)$ determines the optimal scope of participation. Whenever both of the following inequalities are satisfied,

$$\frac{\partial U^{P*}(a, b)}{\partial a} \leq 0, \quad \frac{\partial U^{P*}(a, b)}{\partial b} \geq 0, \tag{4.12}$$

the engagement is complete and selective otherwise. In computing the gradient of $U^{P*}(a, b)$, we invoke the Envelope Theorem and ignore the dependence on a and b in the partially optimized contract (x^*, y^*) :

$$\begin{aligned}
\frac{\partial U^{P*}}{\partial a} &= -L(x^*(a; b); a; b) = p(a)\left(U_0^A - \frac{x^*(a; b)}{2}\right) = U_0^A - \frac{a^2}{4b}; \\
\frac{\partial U^{P*}}{\partial b} &= L(x^*(b; b); b; b) + \int_a^b \frac{\partial L(x; t; b)}{\partial b}\Bigg|_{x = x^*(t; b)} dt \\
&= p(b)\left[\frac{x^*(b; b)}{2} - U_0^A - \int_a^b \frac{x^{*2}(t; b)}{t^2} dt\right] = \frac{b}{6} + \frac{a^3}{12b^2} - U_0^A.
\end{aligned}$$

It is easy to see that the optimal pair (a^{**}, b^{**}) is not in the interior of the unit square T^2: the sum of the above derivatives is equal to $(2b^3 - 3a^2b + a^3)/(12b^2)$ and positive for all $b > a$. In determining the component of the boundary on which the maximum is attained, we concentrate on the economically more interesting case of a positive reservation utility.[4] The Fermat condition (4.12) has a unique solution, $a^{**} = 2\sqrt{U_0^A b}$; hence the selective engagement is optimal. The proof of the next proposition indicates (see Appendix) that the derivative $\partial U^{P*}/\partial b$ is positive, so the optimal is attained on the upper boundary of the set T.

Proposition 4.3 Let the respective reservation utilities of the contracting parties be U_0^A and U_0^P; their contracting preferences given by (4.1)–(4.2), and the agent's types distributed uniformly on the unit interval $T = [0, 1]$. Then contracting is feasible if and only if $U_0^P + U_0^A - (4/3)\left(U_0^A \vee 0\right)^{\frac{3}{2}} \leq 1/12$, in which case the second-best contract

$$x^{**}(t) = x^{*}(t; a^{**}, b^{**}) = \frac{t^2}{2}, \qquad y^{**}(t) = y^{*}(t; a^{**}, b^{**}) = U_0^A + \frac{4t^3 - a^{**3}}{12}, \tag{4.13}$$

is formed with the agent-types in the set

$$T^{p**} = [a^{**}, b^{**}] = \left[2\sqrt{U_0^A \vee 0}, 1\right] \tag{4.14}$$

and yields to the principal the indirect expected utility $U^{P**} = \frac{1}{12} - U_0^A + \frac{4\left(U_0^A \vee 0\right)^{\frac{3}{2}}}{3}$.

The conclusions of Propositions 4.2 and 4.3 allow us to assess quantitatively the social cost of placing the entry right with the agent and thereby imposing on the designer the participation constraint. To see the details, we shall assume for the remainder of the section that the agent's reservation utility is positive. The second-best contract (4.13) is clearly superior: it exists even when the complete engagement fails, as may be deduced by a comparison of the respective efficiency frontiers $U_0^P + U_0^A = 1/12$ and $U_0^P + U_0^A - (4/3)\left(U_0^A \vee 0\right)^{3/2} = 1/12$. Moreover, whenever both regimes are feasible, the utility attained under selective engagement is greater by the amount $(4/3)\left(U_0^A\right)^{2/3}$, which is, then, the social cost of the entry right.

The earlier discussed effects of the selective engagement are quantifiable as well. Observe, specifically, that the negatively contributing low type values are excluded altogether: the contract is offered only to $t \geq 2\sqrt{U_0^A}$. The functional distortions for the remaining types are reduced due to the diminished disparity in information between the parties: the transfer y^{**} is reduced by $a^{**3}/3$ in comparison to the complete-engagement contract (4.10), while the action allocation remains the same.[5]

In order to see the full extent of distortions, we compare the second-best contract of Proposition 4.3 with its first-best counterpart of Proposition 4.1. The functional form of the action allocation changes from linear, $x^0 = t/2$, to quadratic, $x^{**} = t^2/2$ on the common domain $t \geq a^{**} = 2\sqrt{U_0^A}$, while the transfer is distorted by $y^{**} - y^0 = (4t^3 - 3t - 2{U_0^A}^{\frac{3}{2}})/12$.

- The scope of participation is also distorted. It is reduced in size: the agents in the interval $[4(U^P + U^A), a^{**}]$ contracted by the perfectly informed principal are not engaged by the lesser informed one. Observe that the complete-engagement contract aggravates the principal's difficulties: it expands rather than reduces the scope of participation, and forces the principal to form a contract even with the prohibitively costly agent-types.

4.2 The Principal-Agent Model: A Transfer-Indifferent Principal

The presently analyzed agency setting is somewhat more representative than that of the preceding section: it leads to an optimal contract with both separating and pooling features. We utilize it to illustrate the characterization of the optimal contract developed in Faynzilberg [5].

The adopted parameterization of preferences places this model into the second class of agency settings in the classification of Guesnerie and Laffont [13]. The principal has an "ideal point" — a type-contingent action allocation,

$$x_0(t) = (t - 1)^2,$$

in terms of which

$$u^P(x, y, t) = 2\left[1 - 40(x - x_0(t))^2\right] \tag{4.15}$$

$$u^A(x, y, t) = 5\left[y - x^2(3 - t)\right]. \tag{4.16}$$

The type values are distributed uniformly on the interval $T = [0, 2]$ of real numbers.

For a perfectly informed principal, the complete engagement of the agent is best; he allocates to all types the action $x_0(t)$. A less informed principal is better off under the selective-engagement regime, as reported in Proposition 4.5 below. As for the complete engagement, we have the following:

Proposition 4.4 Let the parties' respective reservation utilities be U^P and U^A, their preferences for contracts given by (4.15)–(4.16), and the distribution of types be uniform on the unit interval $T = [0, 2]$.

Then the complete-engagement contracting is feasible if and only if $U_0^P \le -5/2$, in which case the optimal contract has the form $\left(x_0(t \wedge \frac{3}{2}), y^*(t \wedge \frac{3}{2})\right)$, where x_0 is the ideal allocation,

$$y^*(t) = y_0 - 12t + 20t^2 - 16t^3 + 6t^4 - \frac{4t^5}{5}, \tag{4.17}$$

and y_0 is a constant that is bounded below by $U_0^A/5 + 3$ but arbitrary otherwise. The contract yields to the principal the indirect expected utility $U^{P*}(0, 1) = -5/2$.

The type value $t = 3/2$ delineates the pooled and the separated regions. Observe that the contribution to U^P of the separated types $t \ge 3/2$ is positive (and equal 2), whereas the contribution of a pooled type t,

$$u^P\left(x_0(\frac{3}{2}), y^*(\frac{3}{2}), t\right) = 2 - 5\left(t - \frac{1}{2}\right)^2\left(t - \frac{3}{2}\right)^2, \tag{4.18}$$

is negative for most of the values $t \leq 3/2$; which suggests that selective engagement is potentially beneficial for the principal.

That the complete-engagement mechanism may be improved upon may be quickly deduced by considering $T^p = [1; 2]$ as an alternative. Such selective engagement, if feasible, yields to the principal the utility 1 — strictly better than $U^{P*} = -2.5$ under the complete-engagement conditions of Proposition 4.4. Indeed, the action $x_0(t)$ may be allocated to all participating types because x_0 is increasing for $t \geq 1$. Every agent-type t contributes, then, 2 units of utility, the probability of the event $t \geq 1$ is 1/2, so the expected value $U^P = 1$. Full optimization reported as Proposition 4.5 below improves this utility value even further, to $U^P = 1.37$.

The principal's optimization problem under selective engagement is to choose a participation set $[a; b] \subseteq [0; 2]$ and a defined on $[a; b]$ contract $(x; y)$ that jointly maximize his expected utility,

$$U^P = \int_a^b u^P(x(t); y(t); t)\frac{dt}{2}; \tag{4.19}$$

subject to the truth-telling, the participation, and the nonparticipation constraints. As most of the extant literature, we utilize the second-order approach to truth telling, that is, replace it with the first- and second-order condition of optimality. Feasibility is defined, therefore, by the following constraints:

$$\forall t \in [a; b]; \quad \frac{\partial u^A}{\partial x}\dot{x} + \frac{\partial u^A}{\partial y}\dot{y} = 0 \tag{4.20}$$

$$\forall t \in [a; b]; \quad \dot{x} \geq 0 \tag{4.21}$$

$$\forall t \in [a; b]; \quad u^A(x(t); y(t); t) \geq U_0^A \tag{4.22}$$

$$\forall t \notin [a; b]; \quad U_0^A \geq \sup_{\tau \in [a;b]} u^A(x(\tau); y(\tau); t). \tag{4.23}$$

The monetary transfer is decoupled from the rest of the problem: because the principal is indifferent with respect to the transfer, the latter is present only in the equation (4.20), which may be chosen to be continuous:

$$\forall t \in [a; b]; \qquad y(t) = y(a) + \int_a^t (3 - \tau)\frac{d}{d\tau}x^2(\tau)\, d\tau. \tag{4.24}$$

We shall relax the principal's problem further by ignoring at first both the participation and the nonparticipation constraints: they impose a restriction only on the integration constants (see the proof of Proposition 4.5 on the Appendix).

To summarize, the remaining principal's problem is to maximize (4.19) with respect to x; a; and b; subject to the condition (4.21). The classical method of undetermined multipliers allows us to replace it with an unconstrained maximization of $\int_a^b L(x(t); t)dt$; where

$$L(x; \dot{x}; t) = 1 - 40[x - x_0(t)]^2 + \lambda(t)\dot{x}(t) \tag{4.25}$$

is the Lagrangian obtained by adding to the integrand of (4.19) the left hand side of the constraint (4.21) multiplied by the yet to be determined factor $\eta(t)$. The latter is chosen to satisfy the Kuhn-Tucker condition

$$\forall t \in [a, b], \qquad \eta(t)\underline{x}(t) = 0, \tag{4.26}$$

the form of which allows for both the separation ($\underline{x} > 0$) and the pooling ($\underline{x} = 0$) of types. The boundary value between the pooling and separating regions is denoted t_0.

A complete characterization of solutions The first-order necessary conditions of optimality express the continuity of the Lagrangian along the optimal path (Faynzilberg [8]). In the problem at hand, the potential discontinuities of the Lagrangian (4.25) may arise[6] in the interior of the participation set $[a, 2]$, at the point $t = t_0$, and at the boundaries $t = a$ and $t = 2$. Accordingly, we have the following first-order conditions.

In the interior of the participation set $[a, 2]$, the Euler equation,

$$\frac{\partial L}{\partial x} - \frac{d}{dt}\frac{\partial L}{\partial \underline{x}} = 0,$$

which in the present case has the form

$$\underline{\eta}(t) = -80[x(t) - x_0(t)], \tag{4.27}$$

must be satisfied. For the separated types $\underline{x} > 0$, the allocation x_0 is feasible: by the Kuhn-Tucker condition (4.26), $\eta = 0$ identically, and (4.27) becomes

$$\forall t \in [t_0, 2], \qquad x(t) = x_0(t). \tag{4.28}$$

The pooled types receive a constant allocation, say C. For these values, the Euler equation yields the shadow cost of incentive compatibility. Once the continuity of η is ascertained (see below and Faynzilberg [5] for details),

$$\forall t \in [a, t_0], \qquad \eta(t) = \eta(t_0) - 80\int_{t_0}^{t} [C - x_0(\tau)]d\tau. \tag{4.29}$$

At the boundary $t = 2$, the momentum of the allocation must vanish by the transversality condition:

$$\left.\frac{\partial L}{\partial \underline{x}}\right|_{t=2} = 0.$$

In the present case it is trivially satisfied since the momentum is equal to the shadow cost η, which vanishes for all separated types.

At $t = t_0$, the contract x has a kink, so both the momentum and the Hamiltonian of the allocation must be continuous (Erdmann corner conditions):

$$\left[\frac{\partial L}{\partial \underline{x}}\right]_{t_0} = 0, \qquad \left[L - \underline{x}\frac{\partial L}{\partial \underline{x}}\right]_{t_0} = 0.$$

Here the brackets denote a jump at the specified point: for any function $f(t)$, the notation $[f]_\tau$ stands for $\lim_{t \to \tau+0} f(t) - \lim_{t \to \tau-0} f(t)$. For the present parameterization these conditions take the form

$$[\eta]_{t_0} = 0, \qquad [1 - 40(x(t) - x_0(t))^2]_{t_0} = 0.$$

Thus, both x and η must be continuous at t_0; hence

$$x_0(t_0) = C, \qquad \eta(t_0) = 0. \tag{4.30}$$

At the boundary $t = a$, we impose the free-end conditions,

$$\left.\frac{\partial L}{\partial \dot{x}}\right|_{t=a} = 0, \qquad L(x(t), t)|_{t=a} = 0, \tag{4.31}$$

that require the momentum of the allocation and the contribution of the marginally participating type to vanish. They have the following form:

$$\eta(a) = \eta(t_0) - 80 \int_{t_0}^{a} [C - x_0(\tau)d\tau \;=\; 0 \tag{4.32}$$

$$1 - 40[x(a) - x_0(a)] \;=\; 0. \tag{4.33}$$

Equations (4.28)–(4.33) jointly characterize the second-best contract. Conditions such as transversality are often viewed as technical in nature and useful only in fixing the constants of integration of the Euler equation. We have just seen, however, that they govern the optimality of the principal's choices with respect to their respective degrees of freedom — to the same extent as the Euler equations do that in the interior of T^p.

As a consequence of (4.30) and (4.32), we obtain an instance of the general quantization condition that determines the optimal pooling regimes (Faynzilberg [5]):

$$\int_{a}^{t_0} [x_0(t_0) - x_0(\tau)]\, d\tau = 0. \tag{4.34}$$

In general, this condition has multiple solutions. When this is the case, there are multiple contracts that satisfy the first-order conditions (4.28)–(4.33). In the space of social choice functions, these contracts do not lie close to each other. Rather there is a quantized spectrum of optimal contracts, each of which is a candidate for the second best. An additional step is required therefore to select from the spectrum a global optimum. This global optimum need not be unique.

In conclusion of this section, we solve the system (4.28)–(4.33) for the optimal contract. Under the assumptions made, it has a unique solution, which puts us in a position to assess the benefits of selective engagement and the informational distortions.

Proposition 4.5 Let the respective reservation utilities of the contracting parties be U_0^A and U_0^P, their contracting preferences given by (4.15)–(4.16), and the agent's types distributed uniformly on the unit interval $T = [0, 2]$.

Then contracting with the agent of an arbitrary type is feasible if and only if $U_0^P \le U^{P**} = 1 + (8/5)360^{-\frac{1}{4}}$. In this case U^{P**} is the indirect expected utility of the principal derived from the second-best contract $\left(x_0(t \vee \frac{3}{2}), y^{**}(t \vee \frac{3}{2})\right)$; where x_0 is the ideal allocation, $t_0 = 1 + 360^{-\frac{1}{4}}$; and

$$y^{**}(t) = \frac{1}{5}\left[U_0^A + 14\left(1 + 360^{-\frac{5}{4}}\right)\right] - 12t + 20t^2 - 16t^3 + 6t^4 - \frac{4t^5}{5}. \qquad (4.35)$$

The corresponding numerical values are as follows:

$$a^{**} = 0.54; \quad t_0 = 1.23; \quad x_0(t_0) = \sqrt{10}/60 = 0.05; \quad U^{P**} = 1.37.$$

A comparison of the second-best contract of Proposition 4.5 with the best complete-engagement contract of Proposition 4.5 demonstrates once more the benefits of selective engagement. The exclusion of the types $t < a^* = 0.54$ improves the principal's welfare from $U^{P*}(0, 1) = -5/2$ to $U^{P**} = 1.37$. This improvement is made possible by the two earlier discussed effects: the exclusion of the unproductive types, and the reduction in the informational asymmetry. Observe, specifically, that selective engagement permits the principal to allocate the ideal action x_0 to a greater proportion of types: all separated types are in the interval $[1.23, 2]$ vs. $[1.5, 2]$ under complete engagement. The action allocated to the pooled types lies more close to the ideal point $(t - 1)^2$ as well: it is $x_0(t_0) = 0.05$ whereas the complete-engagement allocation is $x_0(3/2) = 0.25$. As in the preceding model, we observe the second kind of distortion: the set of participating types subsides from the first-best, complete-engagement value $T = [0, 2]$ to $T^p = [0.54, 2]$ for the second best.

5 Other Approaches to Selective Engagement

The extant literatures on auctions, contracting, etc., contain numerous examples of mechanisms where the involvement of agents is incomplete. For the purpose of comparison with the present framework, we briefly revisit the alternative methods.

Two methodologies of addressing agents' nonparticipation appear to be prominent. The first, to which for brevity we refer as the equivalence argument, attempts to make the agents' nonparticipation endogenous to choice. According to this argument, the mechanism allocates quantities of goods to the agents, so whenever the designer wants to exclude some agent- type, he merely allocates to it the goods in quantity zero. Therefore, the agents' abstinence need not be considered separately from choice outcomes, as we have done in the present framework. The second approach relies on what we shall call exclusion functions. It is most prominently exemplified by the monopoly-regulation model of Baron and Myerson [1] where a "shutdown" variable at the disposal of the regulator allows him to exclude explicitly the firm of a given cost type. This method appears to be more closely related to our framework, which relies on the participation sets rather than functions.

Although intuitively appealing, the two aforementioned approaches suffer from certain methodological difficulties. In addition, their validity appears to depend on some idiosyncrasies of the mechanism, so the properties of the resulting mechanisms may not be generalizable.

To see why the equivalence argument may not be generally valid, consider it first in further detail. To fix ideas, we present it in the context of bilateral trade as in Myerson and Satterthwaite [25], where the buyer and the seller exchange the object of trade and the numeraire. The outcome of their collective choice is a pair of the respective transfers, (q, y), where $q \in \{0, 1\}$ is quantity of the good sold, and $y \in R$ is the payment to the seller. The choice set is comprised of two copies of the real line, $X^0 = (\{0\} \times R) \cup (\{1\} \times R)$ and includes the zero transfer, $\star = (0, 0)$. The preferences in Myerson and Satterthwaite [25] are trilinear in q, y, and type: the respective utilities of the buyer and the seller are $u^B(q, y, t_B) = t_B q - y$ and $u^S(q, y, t_S) = -t_S q + y$, in which the agents' types t_B and t_S parameterize their tastes for the good.

When applied to bilateral trade, the equivalence argument would suggest that the zero transfer $\star = (0, 0)$ effectively excludes the agent-type to which it is allocated. Hence, if the mechanism designer wishes to prevent the trade between the buyer of type t_B and the seller of type t_S, he selects a social choice function Φ so that $\Phi(t_B, t_S) = \star$. Although $\star$ is a transfer and therefore a choice outcome, it is economically equivalent to the no-trade ending ν. The abstinence from choice — an exogenously specified feature of the model — is handled endogenously, as a choice outcome, by the equivalent transfer $\star$. If no loss of generality is incurred, this approach should lead to the same results as the present framework, perhaps with even greater efficiency.

The argument is not generally valid, however. Contrary to its postulate, the allocation $\star$ is not economically equivalent to the no-trade ending ν. Some of the differences between the two are as follows.

To begin, observe that the allocations ν and $\star$ are informationally incommensurate. Consider the buyer and the seller who attempt to predict their indirect utilities in two cases: ex ante, before receiving an incentive-compatible message μ from the designer; and in the interim, upon receiving this message.[7] For the nonparticipating types the message μ carries no new information: they take exit actions, the consequences of which they knew ex ante. In contrast, for the parties of a participating type combination t, the message μ is informative. Their truthful reports $\iota_T(t) = t$ lead them to expect the specified by μ allocation $\star = \mu \circ \iota_T(t)$. This allocation, although random in the interim, yields certain expectations — $E\left[u^B(x, t_B) \mid t_B\right]$ for the buyer and similarly for the seller. These expectations are dependent on μ, so this message is indeed informative.

Further analysis of the equivalence argument reveals additional difficulties. To see them, suppose that this argument is valid: there exists a choice outcome $\Psi \in X^0$, not necessarily zero, that is economically equivalent to the breakdown of trade ν. Then such $\Psi = (\Psi_q, \Psi_y)$ must be type-contingent because two allocations should yield the same utilities:

$$E\left[u^B(\Psi_q, \Psi_y, t) \mid t_B\right] = U_0^B(t_B); \qquad E\left[u^S(\Psi_q, \Psi_y, t) \mid t_S\right] = U_0^S(t_S). \tag{5.1}$$

A specialization to a uniform distribution shows at once that a constant Ψ does not satisfy this system with arbitrary reservation utilities. Thus, the outcome Ψ; which is posited by the equivalence argument, is not even an outcome — it is a social choice function.

Of course, the idiosyncrasies of the setting may result in this function being constant. The system (5.1) is solved by $\Psi(t) = \bar{x} = (0; 0)$ whenever the reservation utilities vanish identically, as assumed in Myerson and Satterthwaite [25]. Observe that this choice of the reservation values is rather special: it assumes that, in economic terms, the exogenous opportunities are equivalent to the status quo. The pre-trade conditions in which the parties find themselves are entirely different and independent from the outside options they may have, so the assumed parity between the two is coincidental.

Further, even when the equivalent social choice function Ψ exists, it need not be feasible. Recall that the second-order considerations usually require the allocations to be appropriately monotone. The derivative of the reservation utility is arbitrary, however, and the expectations $E[\Psi_q(t) \mid t_B]$ and $E[\Psi_q(t) \mid t_S]$ in (5.1) need not have the appropriate monotonicity. Of course, additional idiosyncrasies may still make this possible. In Myerson and Satterthwaite [25], for instance, the assumed risk- neutrality of both parties allows Ψ to be incentive- compatible.

Furthermore, not only is the incentive-compatibility of Ψ in question, but it may not be even first-degree feasible. With arbitrary reservation utilities, the path $\Psi(t)$ may lie outside the region in R^2 defined by the initial endowments and the indivisibility of the goods being traded. In system (5.1), for instance, Ψ_q may take only the values 0 and 1. For some values of the reservation utilities, this may be too restrictive and render Ψ infeasible.

Moreover, even the very existence of the Ψ is problematic due to the curse of dimensionality. The equivalence of Ψ and ν requires, in addition to (5.1), the designer's indifference as well: $u^D(\Psi_q; \Psi_y) = U_0^D$: The resulting system of three equations cannot be generally solved with only two degrees of freedom in $\Psi = (\Psi_q; \Psi_y)$:

Finally, the economic interpretation may preempt the equivalence argument altogether. To illustrate, consider the agency setting, for which zero allocation $\bar{x}$ may be suggested as being equivalent to the nonparticipation of the agent. Recall that in the moral-hazard model, the agent is called upon to take action, a; in exchange for wages, w: Here $\Phi = (a; w)$ and under the allocation $\bar{x} = (0; 0)$; the agent is not paid for the action $a = 0$ taken by him. We speak of a as "action" and often interpret it as the agent's effort. This is no more than an abbreviation device, however. The value of a is actually a factor in the production technology, and the agent's action amounts to the selection of a level of that factor. The distinction is presently relevant: as factors of production, all values of a; including $a = 0$; are part of the technological specification of the agency setting. When a contract is not formed, however, the agent does not use the principal's technology. Thus, working for the principal with zero effort under the allocation $\bar{x}$ is distinct from not working for him at all. The agent's abstinence has no equivalents among the factors of the production technology.

We conclude that the equivalence argument lacks general validity. It may be applicable in some settings due to their additional, idiosyncratic features. Such exceptional situations are characterized in the following proposition, which summarizes the preceding analysis.

Proposition 5.1 The endogenous treatment of nonparticipation is valid if and only if there exists a choice-valued function $\Psi : T \to X^0$ that satisfies all feasibility requirements of the choice problem at hand and, in addition, provides the reservation utilities to all parties:

$$\forall i \in I;\ \forall t_i \in T_i; \qquad E\left[u_i(\Psi(t); t) \mid t_i\right] = U_i^0(t_i) \qquad (5.2)$$

$$E\left[u^D(\Psi(t); t)\right] = E\left[u^D(\nu; t)\right] : \qquad (5.3)$$

If a function Ψ has the properties annunciated in the proposition, the designer is free to choose a social choice function Φ so that it coincides with Ψ for some types. The allocations to these types are economically indistinguishable form their exogenous opportunities. The effect on the mechanism is the same as if these types were excluded, i.e., selectively engaged by the designer.

The regulatory setting of Baron and Myerson [1], to which we turn for the remainder of the section, exemplifies the second approach to selective engagement, by means of the exclusion functions. The regulated monopolist receives an allocation $(r; p; q; s)$ that consists of a price p to be charged by the firm, a prescribed output level q; a subsidy s to be provided to the firm by the regulator, and r — "the probability that the regulator will permit the firm to do business at all." When $r = 0$; the authors refer to the firm as being shut down, hence nonparticipating in the market; the other extreme, $r = 1$; corresponds to the firm operating with certainty.

In comparison to the mechanism-optimization problem of Section 3.1, this formulation has two additional layers. Although the regulator shuts down the firm of some cost types, these type values are initially included into the scope of participation. This is done by subjecting the regulator to the participation constraint that must be satisfied by all types, including those that will be ultimately shut down. The extraneous types t are subsequently removed by the "shutdown" variable: $r(t) = 0$: In comparison, the formulation (3.2)–(3.5) of the regulator's problem appears to be more parsimonious and efficient: it avoids the nonparticipating types altogether and makes the exclusion variable r unnecessary.

The presence of the shutdown variable creates, in fact, certain difficulties in both interpretation and selection of the regulatory policy. Optimization of the policy is carried out in a two-stage procedure, which is an adaptation of the auction- design method developed by Myerson [23]. This procedure is very similar to the one we have offered in Section 4.1, and we compare the two in more detail below. Baron and Myerson begin with finding, for the firm of each cost type t; the optimal operating regime $(p^*(t); q^*(t); s^*(t))$ and compute its welfare value. Next, the optimal scope of participation is determined by setting $r(t) = 0$ if type t contributes negatively, and $r = 1$ otherwise. This yields the participation and nonparticipation sets $T^p = r^{-1}(1)$ and $T^n = r^{-1}(0)$; respectively.

It is not always clear how the monopolist should interpret the prescriptions of the regulator. In what follows we shall omit the price and view the remaining triple $\Phi = (r; q; s)$ as the policy of the regulator. Indeed, for a given level of output, the price of the good is determined by market forces (the demand function is assumed to be known), hence beyond the regulatory control. Observe that the first stage of the Baron-Myerson method results in a path $(q^*(t); s^*(t))$ with positive q and s: Consider now a type value t being shut down by $r(t) = 0$ at the second stage.

How should the monopolist of this type interpret the regulator's orders? His allocation is of the form $(0; q^*(t); s^*(t))$; so he is required to seize operations ($r = 0$) and reach the production level $q^*(t) > 0$ at the same time. The allocations to the nonparticipating types $t \in T^n$ do not appear to be consistent.

The monopolist may attempt to regain consistency by an alternate interpretation. He may view q^* and s^* in $(0; q^*; s^*)$ as the quantity and the subsidy that would have been realized had the firm been allowed to operate, i.e., as if they where a part of the triple $(1; q^*; s^*)$ instead. In actuality, every non-operating firm receives the allocation $\bar{x} = (0; 0; 0)$:

Unfortunately, this interpretation does not eliminate the aforementioned difficulties: it replaces them with the methodological ones. The latter stem from the discontinuities in the optimal policy, which experiences a jump every time it changes from $(1; q^*; s^*)$ to $(0; 0; 0)$ and conversely. Far from being merely technical, this property is of considerable economic importance. As we have seen in Section 4.2, the regulator acquires additional optimization freedoms whenever the allocation-subsidy path lacks continuity or smoothness. This additional flexibility is gainfully exploited by the regulator both at the aforementioned jumps of the policy (discontinuities), and at the boundaries of the pooling regions (kinks). Although the additional degrees of freedom improve the optimization results, they require special handling during the process of policy optimization.

Discontinuities along the regulatory path may be handled within the framework of the calculus of variations or optimal control, as we have done in Section 4.2. The policy must be assumed, then, to be only piece-wise continuous. Any imperfections along the path call into existence the corresponding first-order conditions — transversality, quantization, free-end, etc. A construction that does not utilize such conditions makes a tacit assumption regarding the optimality of the regulator's decisions and may be, therefore, suboptimal.

Alternatively, the discontinuities of the policy may be handled by the calculus of generalized functions, or distributions in the sense of L. Schwartz. The differentiation of a generalized function leads, however, to an additional term for each jump of the function.[8] Thus, differentiation without regard to these terms amounts to a supposition about their values, which may be suboptimal.

Baron and Myerson rely on heuristic considerations based on convex approximations of functions. Since the aforementioned effects are not addressed explicitly, it is difficult to ascertain the optimality of the conjectured second-best policy. The generalizability of the Baron-Myerson method is also called in to question by the rather unlikely feature of the optimal policy: the optimal output-subsidy allocations, $(q^*(t); s^*(t))$; depend only on the informational asymmetry between the firm and the regulator. We have seen, however, both in the general case of Section 3 and in the models of Section 4, that the policy $(q; s)$ and the scope of participation T^p are mutually dependent. In fact, their interdependence is one of the generic sources of gains from selective engagement. The regulator's task (3.2)–(3.5) is the problem of joint optimization with respect to T^p and Φ precisely because the two cannot be chosen independently.

To see the source of the aforementioned feature of the Baron- Myerson solution, we explicate their procedure in further detail and compare it to the step-wise procedure of Section 4.1. To ease the comparison, we specialize to $T = [0,1]$ and use the same notation.

The two approaches diverge at the first step of the optimization process. The Baron-Myerson procedure begins with the complete engagement of the firm, $[a, b] = T = [0,1]$ and, by an elaborate convexity-related construction, computes the optimal quantity- subsidy regulation, $(q^*(t;0,1), s^*(t;0,1))$. At the second step, the assumption $b = 1$ is relaxed. The authors proceed to maximize with respect to b the expected contribution of the firm,

$$\int_0^b u^D(q^*(t;0,1), s^*(t;0,1))dt. \qquad (5.4)$$

In comparison, the procedure of Section 4.1 begins with an entirely unrestricted participation set $[a, b]$. The first step results in the partially optimized policy $((q^*(t;a,b), s^*(t;a,b))$ that, unlike the preceding solution, depends on the scope of participation. The expected welfare value,

$$\int_a^b u^D(q^*(t;a,b), s^*(t;a,b))dt, \qquad (5.5)$$

is now optimized with respect to both a and b.

Baron and Myerson assume that it is always desirable to employ the most efficient firm and, accordingly, set $a = 0$. The alternative, $a > 0$, cannot be ruled out automatically. Whether the firm should be shut down is determined by the balance between the consumer surplus of the good it produces and its informational rents. The balance may be positive for the less efficient cost types, which will be allowed to operate, and negative for the most efficient type, which is then shut down. An example of such situation is presented in Faynzilberg [6].

The principal difference between (5.4) and (5.5) is not, however, in the value of a. For even with $a = 0$ the utility (5.5) is affected by both the quantity and productivity of “labor.” Formally, they stem from the presence of b in the range of integration and, respectively, in the integrand. Since the Baron-Myerson procedure accounts only for the first of the two effects, it may be suboptimal. The aforementioned independence of the optimal policy from the scope of participation appears in that solution as an assumption made at the first step of optimization.

The foregoing does not imply, of course, that the policy conjectured by the authors is not optimal in the specific illustration they have chosen. Recall that besides the informational disparity between the firm and the regulator — the main focus of analysis — Baron and Myerson make the following assumptions: (i) the cost function of the firm is bilinear in type and output, $C(q,t) = k_0 + k_1 + q(c_0 + c_1)$; (ii) the coefficients of proportionality are nonnegative, $\min\{k_0, k_1, c_0, c_1\} \geq 0$; and (iii) the reservation utilities are constant and equal to zero. In economic terms, the later assumption posits equivalence between the exogenous opportunities and the status quo — the pre-production state with zero consumer surplus and zero monopoly profits. By employing variational methods, we show in the companion paper (Faynzilberg [5]) that the solution conjectured in Baron and Myerson [1] is valid for the stated parameterization.

Thus, the commonly used “normalization” to zero of the reservation utilities appears to be essential for both the equivalence argument and the exclusion functions. As we have seen, this amounts to a supposition that the status quo, pre-choice state and the exogenous opportunities are economically equivalent.

The methodological di±culties arise, not surprisingly, in the treatment of nonparticipating agent-types. They may be traced to the imposition on the designer of the individual-rationality constraint, which leads him to include all agent-types into the scope of participation. Once this is done, the removal of the undesirable types at the subsequent stages becomes di±cult, and may not be resolved fully. In comparison, the redundant types do not even enter the designer's problem (3.2){(3.5). The presently proposed formulation appears to be more conducive to a consistent optimization methodology.

6 Conclusions

We have seen that the results of collective redistribution of property rights are substantially a®ected by the procedural rules that govern this process. Speci¯cally, the manner in which the participation rights and duties are distributed among the parties has a signi¯cant impact on the feasibility of the mechanism.

The involvement of an agent in the collective choice is often described as dichotomous: he may either participate or abstain from the process. We have seen, however, that it is helpful to think of involvement as having three facets rather than two, and di®erentiate the agent's ability to decline, accept, and force his participation. In order to govern three alternatives, at least two rights are needed, and we identi¯ed them as the rights of free exit from and free entry into the collective- choice process. Formally, these rights of the free-exit and free-entry enter the picture as constraints (2.6){(2.7) on the designer's choice of the mechanism.

The proposed approach to mechanism design accommodates all possible imputations of the participation rights of exit and entry, which we summarized in Table 2.1. For the agent, involvement in the choice process may be: (i) altogether involuntary, if all participation rights lie with the designer; (ii) altogether voluntary, if he himself is in possession of all participation rights; and, (iii) partly voluntary wherein he may decline but not force his involvement. Most of our analysis involved the last scenario, i.e., mechanisms with voluntary but not guaranteed participation.

The rights of free exit and entry are void of incentives. The Equivalence Theorem 2.1 shows that the incentive-compatibility alone is su±cient to induce the cooperation of the agent. The cooperation includes, of course, the participation of those whom the designer wishes to engage, and the abstention of those whom he wishes to exclude from the choice process. The corresponding incentives arise because the e±cacy of the participation rights, even when all of them are imputed to the designer, is reduced by the lack of perfect information regarding the type realizations. We have seen, more speci¯cally, that it is in the interest of the designer to exercise his rights in a type-speci¯c fashion. With less than perfect information, however, he is unable to do so, and resorts to incentives.

The type-speci¯c exercise of the designer's rights, i.e., selective engagement of the agents, requires less participation and truth-telling incentives, which increases the value of the mechanism. In Section 3.2, we have traced these welfare gains to a decrease in the designer's uncertainty. Extra information is acquired, however, at a cost in the form of the nonparticipation constraint. In addition, the designer foregoes the contribution of the excluded types, which itself may be beneficial as well as costly. Whether the net effect of the selective engagement is positive is determined, of course, by the solution of the optimization problem (3.2)–(3.5).

The solution of the designer's problem (3.2)–(3.5) appears to be more complex, in comparison to both the traditionally assumed complete-engagement regime, and hidden-action equivalent. The complication stems from a substantial gain in flexibility from an additional optimization instrument — the scope of participation. We encountered the added degrees of freedom and saw their impact on the optimization process in the two methodologies we employed in Section 4: the step-wise procedure, and the characterization of optimal contract.

The step-wise procedure of Section 4.1 allowed us to measure the welfare impact of selective engagement. We have seen that the scope of participation changes both the productivity and quantity of the agent's involvement. The former manifests itself through the presence of the participation set in the designer's utility, the integrand of the expected welfare value. The latter comes from the range of integration, which itself is the agents' participation set.

The characterization of the optimal mechanism in Section 4.1 has been also complicated by the selective engagement. When the set of types is disconnected, the problem (3.2)–(3.5) must be solved by discrete-optimization methods, which may require complete enumeration. In the continuous case, the endogenous treatment of the participation set entails free- end conditions of the calculus of variations. The generally multiple instances of these restrictions augment numerous other first-order conditions, and increases further the already substantial complexity of the characterization task. Nonetheless, the characterization methodology appears to be more efficient than the step-wise procedure. In addition, it quantifies the shadow costs of incentive-compatibility, which furthers the economic intuition and allows the comparative- statics analysis of the optimal mechanism.

Mechanisms with selective engagement merely reflect the overall economic richness of the collective-choice problem. Most applications have exploited the freedom of mechanism design stemming from the great variety of social choice functions. Our analysis shows that a second avenue — the scope of agents' participation is generically available to the designer. The substantive and methodological implications of this additional freedom remain to be fully explored.

A Proofs

Proof of Proposition 4.1. Let $v(t) = u^A(x(t); y(t); t)$: With the change of variables from y to v; the principal may be viewed as maximizing his utility $x - (v + x^2/t)$ with respect to v and x: The participation constraint $v \ge U_0^A$ must be satisfied, which makes $v = U_0^A$ the optimal choice. The first-order condition for x is $1 - 2x/t = 0$; and results in $x^0(t) = t/2$: The associated transfer is given by $y^0(t) = U_0^A + x^{0\,2}(t)/t = U_0^A + t/4$; as stated. The indirect utility of the principal is equal to $x^0 - y^0 = t/4 - U_0^A$ and, in order to ensure participation of the principal, must exceed the reservation level U_0^P: Thus, the agent is contractible if and only if $U_0^P + U_0^A \le t/4$; which defines the first-best scope of participation (4.4). Q.E.D.

Proposition A.1 If the preferences of the agent are given by (4.2) and the allocation $x(t)$ is continuous, then the agent's indirect utility $U(t)$ is continuously differentiable, hence absolutely continuous, and given by the following expression:

$$U(t) = U(a) + \int_a^t \frac{x^2(\tau)}{\tau^2} d\tau. \tag{A.1}$$

Proof of Proposition A.1. The truth-telling condition (2.9) implies the following for $t, \tau \in [a, b]$:

$$U(t) = y(t) - \frac{x^2(t)}{t} \geq y(\tau) - \frac{x^2(\tau)}{t} = U(\tau) + x^2(\tau)\frac{t-\tau}{t\tau}$$
$$x^2(t)\frac{t-\tau}{t\tau} \geq U(t) - U(\tau) \geq x^2(\tau)\frac{t-\tau}{t\tau};$$

where the second inequality is obtained from the first by an interchange of t and τ.

The assumption that the set T^p is an interval and hence connected allows one to take limits $\tau \to t - 0$ and $t \to \tau + 0$ in the preceding inequalities, which gives the following bounds on the left and right derivatives $D^{\pm}U$ of U:

$$\frac{x^2(t)}{t^2} \geq D^-(t) \geq \frac{x^2(t-0)}{t^2}; \qquad \frac{x^2(t)}{t^2} \leq D^+(t) \leq \frac{x^2(t+0)}{t^2}.$$

The assumed continuity of x implies now that these one-sided derivatives exist. Moreover, $D^{\pm} = x^2(t)/t^2$; and the function U is both continuous and continuously differentiable (should an idiosyncratic discontinuity at $t = 0$ exist, the convergence $x \to 0$ must be rapid enough to ensure the continuity.) Under these circumstances, the Fundamental Theorem of the calculus applies and leads to A.1. Q.E.D.

Remark A.2 on the form of the Lagrangian (4.6) in principal's problem. By eliminating y from (4.1)–(4.2) with the help of (A.1) we obtain

$$U^P = \int_a^b \left[x(t) + \frac{x^2(t)}{t} - U(a) - \int_a^t \frac{x^2(\tau)}{\tau^2} d\tau \right] p(t)dt.$$

The last term may be integrated by parts:

$$p(t)\int_a^t \frac{x^2(\tau)}{\tau^2} d\tau = \frac{d}{dt}\left[P(t)\int_a^t \frac{x^2(\tau)}{\tau^2} d\tau \right] - P(t)\frac{x^2(t)}{t};$$

which leads to (4.5) and (4.6).

Proof of Proposition 4.2 amounts to verifying the feasibility of the stated mechanism. The expression for x^* follows from (4.8) with $a = 0$ and $b = 1$. With the transfer y^* given by (4.10), the agent of type t who reports it as τ attains the utility $U_0^A + \tau^3(4t - 3\tau)/(12t)$. This function has a maximum $\tau = t$, where the value is $U_0^A + \tau^3/12 \geq U_0^A$. Thus, the mechanism (4.10) is feasible and satisfies the Euler equations. A substitution of the mechanism (4.10) into the expression (4.5) yields the stated in the proposition value $1/12 - U_0^A$ of the principal's utility. Q.E.D.

Proof of Proposition 4.3. With $a = 2\sqrt{U_0^A b}$ obtained from the Fermat condition $\partial U^P/\partial a = 0$, the derivative $\partial U^P/\partial b$ in (4.13) takes the following form,

$$\frac{\partial U^P}{\partial b} = \frac{b}{6} + \frac{2U_0^{A\,3/2}}{3b^{1/2}} - U_0^A = U_0^A\left(\frac{z}{6} + \frac{1}{3z^{1/2}} - 1\right);$$

after a substitution $b = zU_0^A$. The polynomial in z is always positive in the relevant region $z \geq 4$ defined by the nonemptiness condition $b > a$. Hence, $b^{**} = 1$ and $a^{**} = 2\sqrt{U_0^A}$. With this choice of the participation set, the partially optimized solution (4.8)–(4.9) leads to the contract (4.13), a substitution of which into the objective function U^P yields the stated in the proposition expression for the principal's utility.

It remains to verify that the menu (x^{**}, y^{**}) in (4.13) is feasible, that is, the nonparticipation constraint (2.11) is satisfied: for all $t \in [a^{**}, b^{**}]$, we must have $\sup_{\tau \in [a^{**}, b^{**}]} u^A(x^{**}(\tau), y^{**}(\tau), t) \leq U_0^A$. The utility attained by the agent of a type t who reports it as τ,

$$u^A(x^{**}(\tau), y^{**}(\tau), t) = U_0^A + \frac{t(4\tau^3 - a^3) - 3\tau^4}{12t},$$

is increasing in t and decreasing in τ for all pairs (τ, t) in the interior of the set $[a^{**}, b^{**}] \subset [0, a^{**}]$. Therefore its supremum is attained at $t = \tau = a^{**}$, and the maximal utility gained by misreporting the type is

$$\sup_{t \in [a^{**}, b^{**}]} \sup_{\tau \in [a^{**}, b^{**}]} u^A(x^{**}(\tau), y^{**}(\tau), t) = u^A(x^{**}(a^{**}), y^{**}(a^{**}), a^{**}) = U_0^A.$$

We conclude that the nonparticipation constraint (2.11) is satisfied, the exclusion of the types $t < a^{**}$ is incentive-compatible, and the contract (x^{**}, y^{**}) is second-best. Q.E.D.

Proof of Proposition 4.4 is based on the computations utilized in the proof of Proposition 4.5 below. Specifically, with $a = 0$ we ascertain from (A.2) that $t_0 = 3/2$, and $C = 1/4$. As before, the optimal transfer is obtained by (4.24) and has the form (A.3). The participation condition necessitates the constant y_0 to be sufficiently large (and arbitrary otherwise) so that $\min_t u^A(x(t), y(t), t) \geq U_0^A$, which leads to the stated in the proposition lower bound on y_0. Q.E.D.

Proof of Proposition 4.5. The quantization condition (4.34), after integration and omission of the extraneous solution $a = t_0$, yields the following expressions:

$$t_0 = \frac{3 - a}{2}, \qquad C = \left(\frac{1 - a}{2}\right)^2. \tag{A.2}$$

The free-end a condition (4.33) reads: $1 = 40(t_0 - a)^2(t_0 + a - 2)^2$. After a substitution of expression (4.24) for t_0, it reduces to $(a - 1)^4 = 2/45$. Discarding once again the irrelevant (negative) roots, we arrive at the stated in the proposition value of a^{**}. A substitution of $a = a^{**}$ into (A.2) results in the stated value of t_0. The fact that $x^{**} = x_0$ follows from the statement of the Euler equation (4.28), by using which in (4.24) we obtain

$$y^*(t) = \frac{y_0}{5} + -12t + 20t^2 - 16t^3 + 6t^4 - \frac{4t^5}{5}, \tag{A.3}$$

for some constant y_0. By subjecting this solution to both the participation and the nonparticipation constraints (4.32) and (4.23) we find that y_0 must be, respectively, not smaller and not greater than, hence equal to, $U_0^A + \frac{1}{4}\left(1 + 360^{-\frac{5}{4}}\right)$. The indirect utility U^P is now obtained by integration. Q.E.D.

References

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

[2] T.M. Cover and J.A. Thomas. Elements of Information Theory. Wiley Series in Telecommunications. John Wiley & Sons, New York, 1991.

[3] P. Dasgupta, P. Hammond, and E. Maskin. The implementation of social choice rules. Review of Economic Studies, 46:185–216, 1979.

[4] P.S. Faynzilberg. Statistical mechanics of choice: MaxEnt estimation of population heterogeneity. Annals of Operations Research, 68:161–80, 1996.

[5] P.S. Faynzilberg. Optimal selection of a regulated monopoly. Working paper, J.L. Kellogg Graduate School of Management, Evanston, Illinois, (forthcoming).

[6] P.S. Faynzilberg. Rents of a regulated monopolist. Working paper, J.L. Kellogg Graduate School of Management, Evanston, Illinois, (forthcoming).

[7] P.S. Faynzilberg. The scope of regulatory engagement. Working paper, J.L. Kellogg Graduate School of Management, Evanston, Illinois, (forthcoming).

[8] P.S. Faynzilberg. Necessary conditions of optimality in the calculus of variations: A uni¯ed approach. Working paper 1994{18, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, Pennsylvania, July 1994.

[9] P.S. Faynzilberg and P. Kumar. Optimal contracting of separable production technologies. Games and Economic Behavior, 21:15{39, 1997.

[10] D. Fudenberg and J.-J. Tirole. Game Theory. The MIT Press, Cambridge, Mass., 1993.

[11] A. Gibbard. Manipulation of voting schemes. Econometrica, 41:587{602, 1973.

[12] J. Green and J.-J. La®ont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. Econometrica, 45:427{38, 1977.

[13] R. Guesnerie and J.-J. La®ont. A complete solution to a class of principal-agent problems with an application to the control of a self-managed ¯rm. Journal of Public Economics, 25:329{69, 1984.

[14] B. HolmstrÄom. Moral hazard and observability. The Bell Journal of Economics, 10:74{91, 1979.

[15] J.-J. La®ont and J. Tirole. Using cost observation to regulate ¯rms. Journal of Political Economy, 94:614{41, 1986.

[16] N.F.G. Martin and J.W. England. Mathematical Theory of Entropy, volume 12 of Encyclopedia of Mathematics and Its Applications. Addison-Wesley, Reading, Mass, 1981.

[17] A. Mas-Colell, M.D. Whinston, and J.R. Green. Microeconomic Theory. Oxford University Press, New York, 1995.

[18] E. Maskin and R. Riley. Monopoly with incomplete information. Rand Journal of Economics, pages 171{96, 1984.

[19] R.P. McAfee and J. McMillan. Auctions with a stochastic number of bidders. Journal of Economic Theory, 43:1{19, 1987.

[20] J.A. Mirrlees. An exploration in the theory of optimum income taxation. Review of Economic Studies, 38(1):175{208, 1971.

[21] M. Mussa and S. Rosen. Monopoly and product quality. Journal of Economic Theory, 18:301{17, 1978.

[22] R.B. Myerson. Incentive compatibility and the bargaining problem. Econometrica, 46:61{73, 1979.

[23] R.B. Myerson. Optimal auction design. Mathematics of Operations Research, 6:58{73, 1981.

[24] R.B. Myerson. Game theory: Analysis of con°ict. 1991.

[25] R.B. Myerson and M.A. Satterthwaite. E±cient mechanisms for bilateral trading. Journal of Economic Theory, 28:265{81, 1983.

Remarks

1. See, however, Faynzilberg and Kumar [9] who in the context of the generalized principal-agent model allow for the endogenous determination of the participation set. Some models, such as Baron and Myerson [1], provide for the exclusion of agent-types after their participation has been initially secured (see below).

2. To ease exposition, we concentrate on deterministic mechanisms. The mechanism-design framework incorporated the designer's randomizations on some set - by the extensions ¹ : S £ - ! X; © : T £ - ! X; P on T - -; etc.

3. One could formally say that the designer optimizes with respect to one instrument, ©; if © is viewed as a partial function, that is, a single-valued binary relation from T to X: Its domain T_G^p may be smaller than T: In the traditionally employed complete engagement, © is a true function with full domain.

4. The case of a negative reservation utilities is, not surprisingly, simple: faced with poor exogenous opportunities, the agents have strong incentives to accept the principal's o®er, and their employment does not involve large transfers. As expected, it is optimal for the principal to employ all available agent-types. A direct calculation con¯rms this intuition: the derivative $@U^P=@a$ is negative, so that $a^{¤¤} = 0$; which in turn implies that in (4.13) the derivative $@U^P=@b$ is positive, and $b^{¤¤} = 1$:

5. Independence of the action allocation from the choice of the participation set appears to be idiosyncratic to the chosen parameterizations (c.f., Section 4.2).

6. The discontinuities of the function `(t) = L(x(t); x̰(t); t) may arise: at the boundaries | because L vanishes identically outside the interval [a; b]; in the interior of [a; b] | if L lacks continuity; and, at $t = t_0$ where `(t) may have a jump inherited from x̰ even when the function L is continuous.

7. The order of moves by Nature and the designer is immaterial. To ease exposition, we assume that they move simultaneously, so the type-related timing references are the same as those related to the message ¹:

8. Speci¯cally, if f(x) is discontinuous at $x = x_0$; and df(x)=dx is its (classical) derivative, where it exists, and then the derivative of f as a distribution is $Df=Dx = df=dx + [f]_{x_0}±(x ¡ x_0)$; where ± is the Dirac distribution.