Subsidising the Private Sector for Development: Lessons from Mechanism Design

Paddy Carter, Francesco Decarolis and Nathaniel Young*

August 2, 2017

Abstract

The paper discusses insights from the theory of mechanism design in the context of development finance institutions (DFIs) who wish to allocate subsidies for private investments, to achieve development objectives. A dominant theme in donors post-2015 development finance strategies is the idea of using aid to 'catalyse' other financial flows, motivated by the observation that estimated investments exceed donor budgets by an order of magnitude. However, DFIs face a problem of dealing with private investors who have an incentive to pretend they need financial support. DFIs need a mechanism to select projects that: a) will deliver development outcomes b) will be delivered efficiently c) genuinely require a subsidy to be commercially viable. The paper discusses circumstances under which the mechanism design approach indicates that this asymmetric information might hinder the activities of DFI.

*Carter, Center for Global Development, *pcarter@cgdev.org*; Decarolis, Boston University and EIEF, *fdc@bu.edu*; Young, European Bank for Reconstruction and Development, *youngna@ebrd.com*. This paper was commissioned by the Overseas Development Institute, London, with funding provided by the Bill and Melinda Gates Foundation. We are grateful to Paddy Carter for his feedback. All opinions expressed are those of the authors alone and do not necessarily reflect the opinions of the Bill and Melinda Gates Foundation, the Center for Global Development, the European Bank for Reconstruction and Development or the Overseas Development Institute.

1 Introduction

In the last thirty years, mechanism design has passed from being an abstract analytical framework used by mathematical economists to being a tool actively applied in designing better institutions across a broad range of economic activities.

The basic idea of mechanism design is to find 'rules' that induce participants to reveal their private information through their actions, or to minimize the impact of their private information on target outcomes. Its focus is thus on problems where the presence of private information can lead to inefficient outcomes - such as public support being given to a private investment project that does not warrant it.

Successful applications of mechanism design have entered everyday life, shaping the procedures several governments use to allocate scarce resources (like oil drilling rights or spectrum space for telecommunications), how firms like Facebook or Google choose the advertisements that consumers see, or how students are allocated to schools, or organ donors to their matches.[1]

Promoting private sector development involves dealing with some of the same problems that characterize mechanism design. Namely, attempts to boost private investment to promote sustainable development will often be hindered by the presence of private information on the side of the entrepreneur seeking financing. Development finance institutions (DFIs), and other public sector entities typically seek to avoid crowding out private finance. This is referred to as achieving 'additionality' - making something happen that would not have happened otherwise. Sometimes that is achieved by selecting investment projects that would not be viable without some support from the public sector. But this objective will often be hindered by the presence of private information on the side of the entrepreneur. If the true financing cost is only observed by the entrepreneur, she or he may have the incentive to misrepresent it. For instance, he or she might try to request a subsidized loan, while the project would have been viable at market rates. Relatedly, when the final quality of the project on completion is only observed by the entrepreneur, she or he may have the incentive to deliver

[1] See Roth (2015) for a survey of the many areas where mechanism design has been applied.

a low quality outcome. The problems posed by asymmetric information are traditionally categorized in two types: (i) hidden information (or adverse selection) is a situation in which an agent may not reveal the true state of affairs; (ii) hidden action (or moral hazard) is a situation in which an agent may not deliver on promises, due to imperfect monitoring.

Beyond the asymmetric information between the entrepreneur and the financing institutions, many other aspects of the activities involving financing the private sector for development can be understood using the same theoretical tools. For instance, DFIs might seek "crowding in" other investors via syndicated (or shared) loans. However, here again, the asymmetric information between the DFI that has performed the selection of the project and the other potential investors might adversely influence whether, and to what extent, they will be willing to co-finance the project. As before, the potential presence of both hidden information and hidden actions will obstruct the ease with which the additionality goal can be reached. Another example can involve what happens within the DFI itself. For instance, the DFI officer who is following the project in the developing country might have incentives that are not perfectly aligned with the goals of the DFI in terms of the monitoring effort he exercises: since his effort in monitoring is not perfectly observable by the DFI, he might shirk and reduce his effort at the cost of potentially worse outcomes for the project.

Different rules will obviously shape how asymmetric information effects outcomes. In the examples above, the project subsidy, the achievement of additionality and the quality of the final project all depend on the set of rules under which the *agents* (i.e., DFI, the entrepreneurs or both) operate. The purpose of mechanism design is to establish - for given goals of the designer - the optimal set of rules leading to the fulfillment of such goals. Clearly, there are limits to this approach and they can be of three kinds: (i) the mechanism identified might not be implementable through any feasible real world contract or arrangement, (ii) the problem is too complex and an optimal mechanism cannot be identified and (iii) in certain situations it might be hard to perfectly nail down what is *the* goal and what are the relevant constraints. The latter problem is particularly troubling as small changes to the objective or the constraint might lead to the identification of very different systems. In response to this weakness, a recent strand in the literature addresses some of these problems by taking a less rigid view - instead of attempting to locate the 'optimal mechanism', the objective is to find 'good mechanisms' in the sense of exhibiting certain desirable properties, despite not necessarily being optimal. This latter approach is sometimes referred to as 'market design'.

Given both the broad spectrum of problems to which the mechanism design approach can be applied as well as its various methodological declinations, we focus this essay on two specific issues: the choice of projects and of financing tools. Moreover, we mostly analyze environments where there exists potential for competition between the agents possessing private information. This is an important restriction because, clearly, competition will not always be present in all DFIs' activities. For instance, DFIs will often be involved in confidential bilateral negotiations.

In the absence of competition, the problems created by asymmetric information can be partially dealt with through the careful design of incentive schemes and contractual arrangements.[2] We will discuss how, in absence of competition, other options exist in terms of screening via the choice of terms and conditions relative to the market. However, competition would typically substantially expand the scope of what can be achieved. Its presence in some of the DFI problems, involving the choice of projects and of financing tools, will allow us to describe what prescriptions can be drawn from the mechanism design literature.

This essay is intended as a review of how some key principles of the mechanism design literature can be of use for the activities of DFIs. In particular, with regard to additionality, we will discuss conditions under which a certain design of the rules of the game should bolster the chances of selecting projects that the private market would not finance. To be clear with the terminology, our use of the term additionality implies having a project that would not exist otherwise *in the form wanted by the DFI.* This means, for instance, that the construction of a certain infrastructure satisfies additionality if the DFI wants this project to fulfill some green goals, but the project that could receive financing from the private sector is one where these green goals are not satisfied. In this sense, whenever the DFI provides a subsidy to a project that changes its nature relative to what would be achievable without the DFI, it achieves additionality. Of course, mechanism design cannot say what is the right

[2]See the contract theory literature surveyed in Bolton and Dewatripont (2004).

goal but rather serves as a tool to accomplish implementing a predetermined targeted goal.

In addition to positive results about what the design of contracts can achieve and how, mechanism design also has a few important negative results on what trades will not happen, despite being potentially beneficial for all parties. This type of market failure is the negative consequence of asymmetric information: since each party is unsure about the other's true value, trade will sometimes fail and no mechanism with certain desirable properties can be designed to always overcome this problem. The presence of such impossibility results also means that a mechanism designer often cannot implement their targeted goal and must instead construct rules that implement the next best thing. Further, how close the designer can get to their goal directly depends on the extent of asymmetric information. As such, institutions that acquire, and provide the designer with, otherwise private information are often integral in achieving efficiency. For instance, institutions providing effective project monitoring can be a key element to foster the outcomes that proper contracting can achieve.

Needless to say, mechanism design gives solution tools conditional on the 'information structure' of the situation, the preferences of the participants, market composition, and other contextual details. As such, statements about what can and cannot be achieved with mechanism design in the context of DFIs depends directly on what the analyst is willing to assume and, of course, how well those assumptions correspond to reality.

The rest of this document is organized as follows. Section 2 summarizes the key ideas from the mechanism design approach. Section 3 describes more formally the theoretical foundations of the mechanism design approach and can be skipped by less technically-oriented readers. Section 4 presents a series of examples and cases from the activities of DFIs discussing how they could be handled from a mechanism design perspective. These examples involve such diverse situations as how to choose projects to finance, how to select contractors and how to deal with different financing instruments. Section 5 concludes.

2 The Mechanism Design Approach

A key pillar of how economists look at economic transactions among small groups of agents (as opposed to the whole economy) is the 'Coase theorem.' This states that an efficient trade between two or more parties is always possible in the absence of transaction costs. Stated otherwise, we shall expect economic agents to be able to trade effectively (for instance, DFI selecting deserving entrepreneurs) whenever there are no transaction costs. These latter costs can take many forms, but are essentially linked to two types of phenomena: bounded rationality and incentives. Incentives are a problem for efficient trade whenever the parties involved in the transaction have *asymmetric information*.[3] As mentioned in the introduction, asymmetric information can arise in a number of activities involving DFIs and can involve the presence of adverse selection (hidden information, for instance on the likelihood of success of a project), moral hazard (hidden action, for instance on the amount of effort that the entrepreneur receiving investment will put into the realization of development goals) or any combination of the two.

Contract theory, which encompasses mechanism design, is a branch of economics that seeks to analyze which types of trades are possible in the presence of asymmetric information and which institutions are desirable (or optimal) to help parties achieve trades despite the presence of transaction costs. A key result is that in general no mechanism exists to guarantee that every desirable trade takes place, which implies that DFIs will not be able to identify and finance absolutely every deserving project whilst excluding the undeserving, but how close they can get to that goal depends on the context. It may be worth funding some undeserving projects so that deserving projects can be implemented.

The main tool of contract theory is game theory, "the study of mathematical models of conflict and cooperation between intelligent rational decision-makers" (Myerson, 1991).[4]

[3]The problems driven by bounded rationality are likely no less important than those due to asymmetric information, but are less well understood in the economics literature. Thus, this essay will not address issues related to bounded rationality.

[4]Games are mathematical models that specify who are the agents (players), what actions they can take, what is their information and what are their payoffs. Using these elements together with some notion of how agents should behave, the model can be solved to find a equilibrium, which represents a prediction of what strategies agents will follow and what payoffs will ensue. For instance, the well known Nash equilibrium is

Mechanism design is a particular incarnation of contract theory in which, thanks to the knowledge of how agents will behave under certain 'rules of the game' (or 'institutions') these rules are optimally designed in such a way to ensure that the optimizing behavior of the agents will produce exactly what the designer of the rules seeks to achieve. So for example, a mechanism might be designed to minimize the subsidy needed to finance a deserving entrepreneur. Hence, while game theory seeks to predict how agents behave given the institutions (positive economics), mechanism design seeks to state which institutions are optimal given how agents behave (normative economics).

An example of the mechanism design approach that can help to clarify its application to DFIs is as follows. Suppose a DFI seeks to procure a project. There are N firms that can execute it, but the DFI does not know what is the right price to pay since the production cost is idiosyncratic to each firm and privately known only to it. A fundamental result established by Myerson (1981) is that, under certain conditions, an auction with a reserve price (i.e., the maximum price the procurer is willing to pay for the project) is the optimal mechanism that ensures both an efficient allocation of the contract to the socially most efficient firm (i.e., the one with the lowest production cost) and the minimization of the DFI's expenditures. Furthermore, all conventional auction formats (first price, second price, ascending or descending) are all equally good in implementing such a mechanism (the so called 'revenue equivalence' result).

In cases where competition is present, this type of approach is potentially of great importance to DFIs. Indeed, a recent example where the influence of this type of idea can be identified is the World Bank pilot approach to incentivize green projects in developing countries though the *Pilot Auction Facility for Methane and Climate Change Mitigation (PAF).* In 2016, the pilot held its second auction, allocating $20 million in funding directly to the private sector for projects reducing methane emissions. We return to this example in section 4.

A noteworthy feature, however, is that both the optimal auction and the revenue equiv-

the solution concept wherein each player, knowing what the other players are doing, does not wish to change their choice.

alence results hold only under certain conditions. In any particular DFI setting these conditions may or may not hold. Therefore, in addition to knowing what the market design literature has found in terms of optimal mechanisms and impossibility results, it is also important to know what contract theory more generally has been able to say in terms of what can be achieved through simpler, but more realistic, types of contracts. This is important since the assumptions necessary to implement an optimal mechanism can be difficult to test, or in some context are empirically obviously false.[5] The next section describes some of the general results of the mechanism design and contract theory approaches. Section 4, then discusses how these results can be applied in real wold environments.

3 An illustrative model

In this section we introduce a stylized model, that tries to be as general as possible to encompass multiple instances that might be relevant for DFIs' activities. The discussion opens with a simple case where a principal can contract with many agents without the terms offered to one agent depending on those offered (or accepted) by others.[6] The first-best scenario with full information is discussed, before introducing asymmetries to explore the costs imposed in the contracting frame. The concepts of incentive compatibility and the revelation principle are explored. We then approach the analysis of asymmetric information when there exists competition between multiple principals, then consider the possibility of strategic interactions between multiple rival agents to introduce basic tenants from the mechanism design literature. We discuss direct revelation mechanisms, impossibility results, Bayesian incentive compatibility and revenue equivalence. While the section aims for a non-technical approach to these concepts, portions of the discussion can become technical to motivate these ideas.[7] Therefore, the rest of this section can be skipped by the less technically oriented readers. Section illustrates several applications of the concepts from

[5]For example Li (2015) shows that two allocation mechanisms that should preform identically in theory do not empirically. He attributes the difference in in performance to one being is easier to understand by participants then the other, despite both having identical best strategies and expected payoffs/allocations.

[6]This case is refereed to as that of 'non-rival' agents.

[7]For an even more in-depth discussion, presented along the same lines followed here, see Mas-Colell et al. (1995).

this section using examples from imagined and real life scenarios facing DFIs.

3.1 One principal and one (or more, but non-rival) agents

Suppose we are in a situation where there is a unique investor - referred to as the 'principal' - that could finance a one-time project proposed by a project sponsor - referred to as the 'agent'. Suppose that the project's gross profits depend on the agent's execution of some tasks, in many contexts we can think of the task as 'effort' such as attention to project detail, monitoring of subcontractors or anything that is relevant to the project goal, but which may not be perfectly observed by the mechanism designer. In particular, denote t as the task level and $\pi(t)$ an (increasing) function describing how profitability depends on the task level.[8] Suppose that $t \geq 0$ is fully observable after the contract is signed and while the project is in execution. For the agent, however, performing each given task level can be more or less costly depending on his or her *type*, which is only privately observed by the agent. In particular, the monetary cost of performing task level t is given by the function $g(t, \theta)$ that maps both the task and the *type*, θ, onto costs.

One may also look at θ, as is frequently followed in the literature, as an ex-post, privately observed realization of an unpredictable 'state of the world' that accounts for how in different situations performing each task level has different implied costs.[9] For example a project may have failed because the agent worked hard despite but the project was too being difficult, high t but low θ, or it may have failed because the agent shirked on an easy project, low t high θ. Obviously we would like to reward the hard working agent and penalize the agent who shirks, but the principal can only observe project outcomes not the underlying cause of failure. Further, as much as possible, we would like to choose an agent with a high θ.

While not explicit in this model, we could think of high θ projects as those likely to occur even without the support of a DFI, low θ projects as those who cost more then their social value and hence should not be undertaken at all, and we can imagine an intermediate range

[8]The following results pertain to a situation with $\pi'(t) > 0$ and $\pi''(t) < 0$ for all t.

[9]The following discussion is based on an environment where the following restrictions on $g(t, \theta)$ apply: $g(0, \theta) = 0$, $g_t(t, \theta) > 0$, $g_{tt}(t, \theta) > 0$, $g_\theta(t, \theta) < 0$ and $g_{t\theta}(t, \theta) < 0$.

of θ where the social value of the project exceeds its cost but the private value to the agent does not, and thus for these values of θ there is room for DFI additionality. In what follows, it will be convenient to consider a case where θ can take one of two values only: high (θ_H), with probability λ, or low (θ_L), with probability $1 - \lambda$.[10]

The agent seeks to obtain the highest private return r from the project's execution. This quantity is conceptually distinct from the project profitability π, for example even if the project completely fails, $\pi = 0$, the agent will still need to be paid something for their work and when projects are successful the investor will keep a share of the profits, $\pi > r$. If we assume that the agent aims to maximize expected utility, meaning the difference between his private return r and his cost $g(t, \theta)$, then the problem of the investor is to propose a private remuneration to the agent, r, such that it induces him to achieve a task level that maximizes the project's profitability.[11] All the difficulties in choosing r derive from the fact that the asymmetric information on θ might induce an agent to pretend to require a high remuneration even though his realized cost is low. That's analogous to the agent pretending the project requires support from a DFI.

One additional constraint arises from the outside options of the agent (opportunities for his or her time and resources that do not involve contracting with the principal), such that a minimum remuneration level exists denoted by value $\bar{r}$, below which the agent is unwilling to participate. This is sometimes referred to as the reservation value, and the need for the private return to not be lower than it is defined as the *participation constraint* as it accounts for the need to make the parties willingly enter the contract. A project that offers a return of less then $\bar{r}$ will not be undertaken regardless of the value of the project. A project that is socially valuable but where the private market will only pay less then $\bar{r}$ is the ideal project for DFI funding as, by definition, its completion is pure additionality.

The proposed contract must satisfy two separate and somewhat conflicting objectives.

[10]So, $Pr(\theta_H) = \lambda \in (0, 1)$.

[11]Formally, we can write this utility as a function of (r, t, θ) as follows: $u(r, t, \theta) = v(r - g(t, \theta))$. The shape of the function $v(.)$ is particularly relevant. For risk neutral agents, the incentive problem can be easily solved with a contract that makes the agent internalize the marginal returns from his chosen task level. The problem is more interesting when the agent is strictly risk averse, $v''(.) < 0$. This is the case we consider here. Note that under the stated assumptions the high and low type agents' indifference curves show the so called *single crossing* property.

On the one hand, it shall ensure that the agent will earn a sufficient return to justify participation, despite the ex post realization of θ. The second is the need for the contract to create the incentives for the agent to act in a way that maximizes the project's profitability, despite only the agent observing θ. Stated otherwise, while a contract can certainly specify a value of t, since this is assumed to be observable, the 'optimal contract' must accomplish something more sophisticated: make the task level selected by the agent responsive to the realization of θ, essentially inducing the agent to reveal the realization of θ that only they observe.[12] It turns out that satisfying this latter objective necessarily implies wasting some resources, in the sense that the agent is able to extract some private 'rents' (excess returns) from the fact that they are the only one observing θ.

Before entering into the details of how the optimal contract looks and how it solves this complex problem, it is worth presenting as a benchmark case the situation in which the state θ is observable.

3.2 What if the state is observable?

The (unrealistic) situation where θ is observable is useful as a benchmark to understand what limits asymmetric information imposes on contracting. With full information the private return r can now directly be linked to the task level t, making the contract outcome conditional on the realization of θ. That is, the contract can now be: (r_H, t_H) for state θ_H and (r_L, t_L) for state θ_L. An *optimal contract* in this context is defined as the quadruplet of values (r_H, t_H, r_L, t_L) that maximizes the project's profitability, net of any private return and subject to this private return being no less than the reservation value $\bar{r}$. In the solution to the problem of finding the optimal contract, the agents reservation value plays a key role: the optimal contract fully insures the agent in the sense that in both states, θ_H and θ_L, the agent's value is equalized and set equal to $\bar{r}$.[13]

[12] In the context of DFIs, the problem of inducing optimal effort from the agent may at first seem less salient, as the project sponsor is typically an equity investor with 'skin in the game', but that should be seen as a contractual solution to the problem of inducing effort, not as a sign that inducing effort is irrelevant to DFIs.

[13] Formally, the optimal contract $(r_H^*, t_H^*, r_L^*, t_L^*)$ is such that: $v(r_i^* - g_i(t_i^*, \theta_i)) = \bar{r}$ for $i \in L, H$.

Thus, for instance, a DFI would leave the agent/investor with 'market returns' (i.e., no rents, but just the minimum amount required to induce participation). Furthermore, the optimal contract requires a higher task level from the agent when the state is θ_H than when it is θ_L. This requirement is necessary to induce the incentives that ensure that the marginal benefit from task execution equals its marginal cost.[14] This type of solution is known as the *first best*, to differentiate it from what one can achieve from the optimal contract in the case of incomplete information. In cases of incomplete information, the optimal contract achieves at most a *second best*, characterized by inefficient productivity and distortions in the task levels.

A graphical representation of this contract is presented in Figure 1 the agent's private return is on the vertical axis and his task level is on the horizontal axis. The agent is better off as for pairs (t, r) that are more toward the top-left corner (higher private return and lower task level), but profitability increases by moving toward the bottom-right corner. The optimal contract entails a tangency of the agent's value and project's profitability curves. Moreover, the risk elimination for the agent is represented from the two values achieved at (r_H^*, t_H^*) and (r_L^*, t_L^*) being the same and equal to $\bar{r}$.

3.3 What if the state is observed only by the agent?

We now consider the case in which the agent perfectly observes the state θ after the contract is signed, but this state is not observable to others. To understand the incentive problems created by this form of asymmetric information consider again Figure 1. Now, suppose that the same quadruplet of values $(r_H^*, t_H^*, r_L^*, t_L^*)$ were offered in a situation with asymmetric information. It is clear from our assumptions that the agent would have an incentive to always declare that the state is θ_L in order to produce a low required task level of t_L, even in case when the realized state is θ_H, as this increases his own payoff. Indeed, $g(t_L^*, \theta_H) < g(t_H^*, \theta_H)$ so that falsely declaring the state allows the agent to increase his own utility: $v(r_L^* - g(t_L^*, \theta_H)) > \bar{r} = v(r_H^* - g(t_H^*, \theta_H))$. This, clearly, comes at the cost of a reduced project's profitability as $\pi(t_L) < \pi(t_H)$.

[14] Formally, the optimal contract is such that $\pi'(t_i^*) = g_t(t_i^*, \theta_i)$.

Figure 1: The Optimal Contract with Observable States

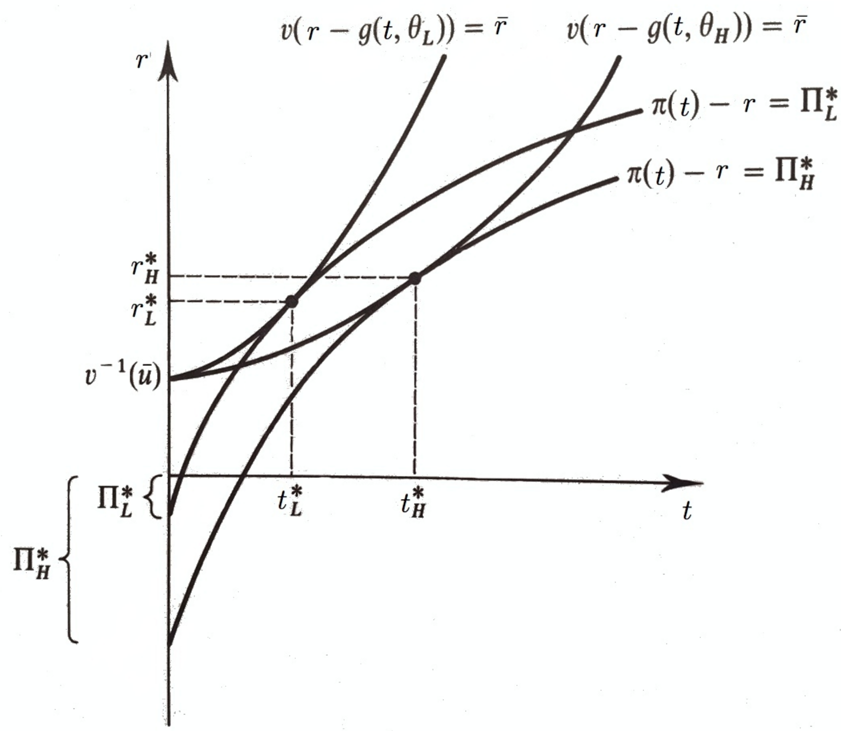

In the context of private investment, for some projects it may be that no contract (r_L^*, t_L^*) exists that can satisfy the agent's incentive compatibility constraint, i.e. provide the agent at least $\bar{r}$, but it may still be socially optimal to implement the project. In such case there would be no project in the absence of DFI, which - although not explicit in this illustrative model - is able to offer a contract that delivers the agent at least $\bar{r}$ and thus bring about a project that would otherwise not occur. However, as the true state is not observable, and the informed agent will always claim θ_L is the true state of the world, in the absence of a mechanism to prevent it, the DFI would also fund projects when the true state is θ_H which provides no additionality and crowd out private funding.

Given this incentive to lie created by asymmetric information, the optimal contract in this situation has to take a different form. Identifying this contract could in principle resemble a nearly impossible task given the plethora of contract types possible: r could be made dependent on any combination of the observable quantities, for example π and or t. Nevertheless, a major result in the mechanism design literature is that for this problem (as well as for most problems of asymmetric information) to identify the optimal contract one can restrict attention to *revelation mechanisms*. It is important to note that for any mechanism there is an infinite number of equivalent mechanisms, many of which at first glance look nothing alike.[15] Revelation mechanisms are contracts (or, more generally, binding rules/procedures) such that: *i)* after the state θ is realized, the agent announces a realization for it, $\hat{\theta}$; *ii)* the contract requires a pair of task level and private remuneration for each announced state (so, if the announced state is $\hat{\theta}$, then the contract entails a pair $(t(\hat{\theta}), r(\hat{\theta})$ for each $\hat{\theta}$); *iii)* in every state the agent finds it optimal to reveal truthfully: $\hat{\theta} = \theta$.

Revelation mechanisms are *incentive compatible* only if it is indeed the case that truthful revelation is optimal for the agent possessing private information. There are two very general and non obvious aspects related to the fact that to identify an optimal contract one can restrict attention to truthful mechanisms. The first aspect relates to the logic behind this result. While the formal mathematical argument is complex, the intuition is straightforward: for any chosen contract that can be written (based on the realization of either t or π, or any combination of them) there is always a revelation mechanism that produces exactly the same outcome. Thus, instead of thinking of all possible contract types, it is sufficient to think about a revelation mechanism, and then find at least one such contract that can be implemented in practice.

To see this, consider devising a revelation mechanism that replicates a contract linking the private benefit to profitability $r(\pi)$. The latter contract leaves t unspecified and the agent will thus pick t_L and t_H optimally in each of the two states (as this choice is unconstrained). But, if the revelation mechanism is designed so that those two chosen task levels are the same as would be specified in a contract if θ was observable, then evidently the agent has no incentive to lie about θ under that mechanism. The same type of logic can be applied to nearly all types of contracts. Methodologically, the result that one can limit attention to truthful mechanisms is a result known as the *revelation principle* and it is a founding pillar of modern microeconomics.

[15] An interested reader can look to Larsen and Zhang (2017) for examples of an implementable mechanism that is observationally equivalent to the revelation mechanism we describe here.

The second general result is that the optimal contract must entail some 'distortions' relative to the case of the first best achievable outcome without asymmetric information. Put differently there are always trade-offs when information is asymmetric. It is not possible to eliminate the incentive to lie (as a truthful mechanism does) whilst also achieving the same outcome as that described in Figure 1 - something has got to give. The distortion will typically entail the agent earning a 'rent' - a return above what he or she would have earned given full information, in at least some states of the world. Moreover, a typical result is that the incentive to lie will not affect all states so that distortions relative to the first best do not characterize all realized states.

What the best contract looks like in the presence of asymmetric information depends on contextual details. Suppose, for example, that the nature of the project a principal wishes to undertake would depend on the agent's type, and denote these two potential project types by S (small project) and B (big project). Suppose that project type S creates value of π_S and that project B creates value π_B. Now suppose that the agent is either type θ_L or type θ_H and that the type-specific costs for the two projects are (c_{SL}, c_{BL}) and (c_{SH}, c_{BH}) respectively. Suppose that the probability of the agent being type θ_L is γ and that the agent is aware his or her type but that the principal is not.[16] In this example, the principal would like to have the agent undertake the small project if the agent is type θ_L and undertake the big project if the agent is type θ_H. The principal must then select a price for each project with the goal of maximizing social value, net the cost of project implementation and subject to incentives. As such we can represent the principal's problem as follows:[17]

$$\max_{p_S, p_S} \gamma(\pi_S - p_S) + (1-\gamma)(\pi_B - p_B)$$

subject to the agents being willing to participate: $p_B > c_{BH}$, $p_S \geq c_{SL}$ and that the agents incentives induce them into selecting the contract targeted at their type

$$p_B - c_{BH} \geq p_S - c_{SH}$$

[16]Without loss of generality suppose $\gamma \in (01)$, $\pi_B > \pi_S$, $c_{Bj} > c_{Sj}$ for $j \in \{L, H\}$ and $c_{iL} > c_{iH}$ for $i \in \{S, B\}$.

[17]The principal would like to contract with both type so long as $\gamma(\pi_S - c_{SL}) + (1-\gamma)(c_{SH} - c_{SL}) \geq 0$.

$$p_S - c_{SL} \geq p_B - c_{BL}$$

Note that only two of the 4 conditions above actually matter. Specifically any price for the big project that will only be selected by the type θ_H will be larger then the type θ_H cost of production and thus the individual rationality constraint $p_B \geq c_{BH}$ is trivially satisfied. Similarly the big project will be too costly for the type θ_L agent to want to pretend to be type H and thus the incentive compatibility constraint $p_S - c_{SL} \geq p_B - c_{BL}$ will be trivially satisfied. As such the solution to this maximization problem is to set $p_S = c_{SL}$ and $p_B = c_{BH} - (c_{SH} - c_{SL})$. Note that compared to the perfect information case the θ_H type now earns an addition value of $(c_{SH} - c_{SL})$ which mechanism design terms 'information rent.' Also note that this information rent is subtracted from the principal's value and thus the total profitability of the project is lower then in the first best case.□

But, can we do better, if we relax the assumption that there is only a single principal and a single agent (or non-rival agents)?

3.4 Asymmetric Information and Competition between Principals

Suppose now that there are a multiplicity of potential principals who compete and all want to select agents to perform a project. The agents have different unobservable productivity levels: $u(r, t|\theta) = r - g(t, \theta)$. This situation is isomorphic to the one illustrated above with the only difference that we can think of λ not as the probability of state θ_H realized as before, but as the share of agents with type θ_H. In this environment, principals offer a menu of contracts that try to distinguish, or *screen*, between agents' types by inducing them to self select into a contract based on their type. Screening settings of this kind have been extensively studied in the literature since the 1970s (see Rothschild and Stiglitz (1976) and Wilson (1977)). A general tenant of this literature is that competition between principals by itself does not guarantee efficient outcomes in environments with asymmetric information. In principle, two scenarios can emerge in this market: a *pooling* scenario, in which all agents receive the same contract, or a *separating* scenario, where the two types accept different contracts. However, it can be shown that only a separating scenario is compatible with an equilibrium outcome.

Indeed, since a pooling contract is necessarily a compromise between the ideal outcomes of the two different types of agents, if a principal tries to offer a single contract to all, a different principal will profit by offering a separating contract that appeals only to the most productive agents. In this way, the principal offering the pooling contract would sign up only unproductive agents, which would offer them an excessive private return. Since this is clearly not sustainable, only a separating equilibrium is possible. However, a major problem highlighted by the literature is that separating equilibria will often also fail to exist. Here too, it can be that for any pair of contracts offered by one principal, another principal will have an incentive to offer something different. The lack of any sort of equilibrium in theory means we may expect to observe instability in markets where private information matters in practice. [18] This suggests that, in many cases, it will not be possible to completely separating a DFI market and a private market and that when the markets cannot be separated behavior is out of equilibrium and thus behavior may not be well described by a static mechanism.[19]

3.5 Asymmetric Information and Multiple Agents

In this paragraph, we return to a situation with only one principal, but with multiple agents. The major difference from the discussion above is that we now allow for the possibility of *strategic interactions* between agents (i.e., interactions in which agents take into account each others' responses to their actions). In the analysis above, we implicitly assumed that the same contract could have been offered to all agents potentially present in the market, without any need for agents to compete against each other. However, it is often the case that, for instance due to the scarcity of financing resources, two or more agents will be in competition to obtain the one contract from the principal. There is an extremely broad set of situations where strategic interactions between agents might matter. The field of *mechanism*

[18] A similar market failure can also occur in environments with *signaling* (Spence, 1973, 1974). In this case, we do not have principals taking actions (i.e., designing contracts) to screen agents, but agents themselves taking actions to signal to the principals their (privately observed) type. The signal is typically assumed as an action imposing a cost on the agent, but that lacks any benefits by itself for either the principal or the agent. Thus, although in signaling markets might -under certain conditions - have both pooling and separating equilibria, they also entail a welfare loss relative to first best due to the waste associated with the signaling cost.

[19] Dynamic mechanisms are complex and well beyond the scope of this paper.

design is mostly devoted to analyze what kind of contracts (or, more generally, rules and procedures) are optimal in this environment. The essential concept underlying the analysis is that of equilibrium. In the paragraph above, we stated that the lack of equilibrium in the competitive screening market is due to the presence of at least one principal that, for any possible menu of contracts offered by the various principals, will find it optimal to change its offered contract after observing the contracts offered by the other principals. In this section where we consider only one principal, the notion of equilibrium refers to the agents. Thus, we have an equilibrium if we are in a situation where each agent does not want to change actions, given what the other agents are doing.

In the presence of asymmetric information, this general description of an equilibrium needs to be refined to account for whether the lack of any incentive to deviate from the prescribed action is ex ante (before agents observe the realization of θ), interim (after each agent has privately observed his own θ, but before the θ of all agents are commonly revealed) or ex post (after the θ of all are commonly revealed). Depending on which of these three notions of equilibria one would like to achieve, different mechanisms might be available or, alternatively, might fail to exist. Clearly, achieving an ex post equilibrium would imply a very strong notion of stability and, moreover, mechanically implies that the ex ante and interim equilibria are also obtained. Accordingly, however, there is a limited set of instances where an ex post equilibrium can be achieved. We start by discussing these situations and then move to analyze what is achievable if one seeks only to obtain an interim equilibrium.

Regardless of the equilibrium notion adopted, the key finding in the literature that allows the characterization of optimal mechanisms in a complex environment with multiple competing agents is the revelation principle discussed above. In this context, a mechanism is defined as a contract (or rule or procedure) such that, for every possible collection of agents' actions (one action per agent), the contract specifies which outcome will correspond to these actions. For instance, consider a scenario where multiple agents compete by bidding their task level t to a principal that will select only one of them. Then one possible rule is to state that for whatever combination of task levels offered, the agent who offers the highest task level is the one who gets signed up and that his required contractual task level is what he bid. Alternatively, the contract can say that the agent offering the highest task level is signed up, but that his required effort level equals the task level of the second highest bidder. These are just two possible examples of mechanisms, but how to identify the optimal one? The result from the revelation principle is, once again, that we can restrict attention to study simple mechanisms of the following form: all of the N agents are required to each report a type and, based on these announcements $(\hat{\theta}_1, \hat{\theta}_2, ..., \hat{\theta}_N)$, the outcome $f(\hat{\theta}_1, \hat{\theta}_2, ..., \hat{\theta}_N)$ is implemented. This class of mechanisms are known as *direct revelation mechanisms.* By the revelation principle, any mechanism that is truthfully implementable (or incentive compatible) in dominant strategies (i.e., by taking actions that are optimal regardless of the actions taken by the other players) must produce an outcome that is an ex post equilibrium under truth telling of a direct revelation mechanism. As above, despite the complexity of a formal proof, the logic behind this result is simple: if some mechanism is able to induce an ex post equilibrium, then it is always possible to construct a direct revelation mechanism in which, to the truthful revelation of the state by each agent, the contract associates exactly the outcome of the ex post equilibrium.

Despite the appeal of the ex post equilibrium notion, as it entails substantial stability of the outcome to agents' incentives to deviate, a main finding in the literature is that it is typically impossible to find a mechanism with desirable properties in this context. The Gibbard-Satterthwaite theorem states that whenever there are more than two possible contractual outcomes, then truthful implementation in *dominant strategies* is possible if and only if the mechanism is *dictatorial.* A mechanism is dictatorial whenever there is one agent i such that the contractual outcome chosen is always one of i's top-ranked alternatives. Clearly, contractual outcomes with this *dictatorial* property are not satisfactory at all as they might leave all other agents worse off. To overcome this 'impossibility' result, the literature has identified two solutions. The first is to focus on situations where the agents' utility can be represented in a specific way (quasilinear-utility models). The second is to switch the focus to the less demanding notion of interim equilibrium.

The case of quasi-linear utility is certainly special, but potentially of major relevance for DFI's activities. Indeed, its only restriction is that the agents' utility can be written as:

$u_i(x, \theta_i) = v_i(k, \theta_i) + m_i$, where k is a contractual outcome, m is a numeraire commodity ('money') and $x = (k, m_1, m_2, ..., m_N)$. In this situation, it has been shown that truthful implementation in dominant strategies is possible if and only if the numeraire m that is assigned to each agent in response to any profile of announcements $(\hat{\theta}_1, \hat{\theta}_2, ..., \hat{\theta}_N)$ takes a very special form: $m_i(\theta) = [\sum_{j \neq i}(k^*(\theta), \theta_i)] + h_i(\theta_{-i})$, where $h_i(\theta_{-i})$ is an arbitrary function depending on the reports of all agents other than i. Despite appearing complex, this formula has a straightforward intuition: dominant strategy implementation is feasible only if for each agent i the monetary transfer he gets is independent of his own reported state and only depends on the final contractual outcome and on the reports of the other agents. An example is useful: in the so called *second price auction* mechanism, the winner of the auction is the agent offering the highest bid, but payment is equal to the second highest bid. This is an instance of a mechanism satisfying the above requirement because the monetary transfer is independent of the winners' bid (which only determines the winning agent) and, instead, it is equal to the sum of the utilities of the other players had the winner abstained from participating. Indeed, the second bid is the bid of the agent who would have won absent the current winner and, moreover, it can be shown that his bid truthfully reveals its true state. Mechanism of the form described above are known as Groves mechanisms and their payment rule reflects exactly the externality that each agents imposes on the others because of his presence.[20]

Despite their desirable properties, Groves mechanisms have two limitations. The first is that their payment rule is both complex, thus limiting its practical usefulness, and prone to collusion between agents, as payments are interlinked. Nevertheless, a form of this mechanism has recently been adopted by one of the largest social media platforms (Facebook) to sell its banner space to advertisers. The second, and more technical problem, is that - as shown by Green and Laffont (1979) - it is typically impossible to achieve truthful implemen-

[20]For a simple example of the externality that arises when competing for a scarce good, suppose there are 2 agents who want a good and that the first values the good at \$1 and the second at \$5. And further suppose that if not for the second agent, the first would have received this good for free, thereby obtaining a \$1 net benefit. If the good is allocated to the second agent, who places a higher value upon it, then the first agent foregoes that \$1 benefit and the presence of the second agent is said to impose an 'externality' of \$1. Accordingly, a Groves mechanism that prices the externality allocates the good to the second agent but requires this agent to pay \$1 (note that with two bidders the Groves mechanism is equivalent to a second price auction).

tation in dominant strategies without the mechanism creating some waste of the numeraire. This is often referenced as the inefficiency, or lack of budget balancedness, of this type of mechanisms. More positive results are instead achievable when one looks at interim equilibria as we do next.

The second solution to the Gibbard-Satterrthwaite's impossibility theorem is to look at interim equilibria. These are also known as Bayesian-Nash equilibria. A Bayesian-Nash equilibrium is a mapping from privately observed states θ to actions such that each agent, knowing his own type θ and the actions of all agents (but not their privately observed states) has no incentive to individually deviate from the prescribed action. For this different notion of equilibrium, the revelation principle discussed above applies once again so that any contractual outcome is truthfully implementable in Bayesian-Nash equilibrium (or Bayesian incentive compatible) if this outcome can be achieved as the Bayesian-Nash equilibrium of a direct revelation mechanism. The notion of Bayesian implementation is weaker relative to that of dominant strategy implementation in the sense that if something can be implemented in dominant strategies, it can certainly be implanted in a Bayesian-Nash equilibrium, but not the other way around. A major result of looking at this weaker concept is that, contrary to the earlier impossibility results, in this environment we have that, if we are also willing to assume quasilinear preferences, then whenever agents' types are statistically independent of one another at least one efficient (i.e., budget balanced) outcome is always Bayesian-Nash implementable.

One of the most remarkable results in this context is the so called *revenue equivalence* theorem. It states that in an auction setting with N risk-neutral buyers, under some mild restrictions on the distribution of types θ, then any pair of auction mechanisms satisfying the following two properties will produce the same expected revenue for the seller: *i)* each buyer i has an identical probability of winning in the two auctions, for each possible realization of $(\hat{\theta}_1, \hat{\theta}_2, ..., \hat{\theta}_N)$ and *ii)* when an agent happens to have the lowest possible valuation, then his expected utility level is the same across the two mechanisms. This important result likely explains, at least partially, why there is not one prevailing auction system in the real world, but instead a plethora of auction procedures: under rather general conditions they all produced the same (expected) revenues. This is a practically useful result as it can guide the design of real world allocation mechanisms. However, its validity stops to hold when, for instance, the environment is characterized by certain forms of asymmetries between buyers or by the lack of independence in their types.

One limitation of the discussion above is that it has (implicitly) assumed that the agents have no option but to participate in the mechanism. This is not the case when participation is voluntary and agents can step out if unsatisfied with what the mechanism imposes on them. This apparently small change in the characterization of the environment has rather profound effects on the set of mechanisms that are implementable. To be more specific, the well-known Myerson-Satterthwaite *impossibility to trade* theorem shows that in this situation there might not be any implementable outcome. This theorem states that, in a bilateral trade setting where gains from trade are possible but not certain because each party is unsure about the true value that the other party assigns to the good, there is no ex post efficient mechanism that is both Bayesian incentive compatible and that satisfies the constraint of (interim) voluntary participation.

It is worth stressing that when multiple mechanisms producing different Bayesian-Nash equilibrium are all implementable, then it might be desirable to devise criteria to select among those. The literature on *optimal* Bayesian mechanism is broad (starting from Holmstrom and Myerson (1983), Myerson (1991) and Fudenberg and Tirole (1991). The optimal contract that we described earlier for the case of one principal and one agent is indeed an instance of an optimal mechanism developed in this literature. While optimal mechanisms can be constructed for several environments, the results presented above also stress the fundamental limits imposed by asymmetric information on how contracts can induce desirable behaviors. In dominant strategy implementation, the incentives to deviate are often hard to curb, outside special cases (quasi-linear utility) and very specific and often complex mechanism (Clark mechanisms). The situation is more encouraging for Bayesian-implementation, but when participation constraints are accounted for then impossibility results are likely to emerge (the Myerson-Satterthwaite theorem). In the next section, we discuss how these abstract results can be more concretely applied to inform the actions of DFIs.

Finally, although the above discussion has focused on problems of *hidden information*, as discussed in the introduction, asymmetric information can also take the form of *hidden action* (or moral hazard). Similar, albeit not identical considerations to the ones developed above can be obtained for the case of moral hazard or in mixed environments where both hidden action and hidden information are present. While mixed models are appealing on the ground of their greater realism, their great complexity often makes it hard to analyze them and to learn general lessons from their results. Similarly, while dynamic considerations are often relevant as in many situations principals and agents are likely to repeatedly interact over time, we have preferred to discuss above the simpler, static models of one-shot interactions leaving the discussion of dynamic mechanism design to the technical literature.

A natural question arises: why are auctions the most common form of mechanism design implementation? What is gained by being able to design a mechanism under which agents compete against one another? To some extent, competition can be thought of as placing a bound on the maximum attainable information rents. For example, observing auction outcomes allows for inference about the agents' valuations of the good, as implied through their bids, and this in turn reduces the principal's uncertainty over the agents types. More generally strategic interaction between competing agents implicitly provides information to the principal on the agents types and thus reduces the asymmetry in information.

4 Examples and Cases

In this section, we present a series of examples and cases. They serve to illustrate how the theoretical considerations developed above can impact various aspects of the DFIs' activities.

First we present a numerical example to help the reader visualize the key insights from mechanism design in the context of DFIs. To keep the example simple we can think of projects having a benefit that is potentially independent from profitability. We will also define the 'principal' as the agency that determines DFI subsidy rules and is interested in additionality. Suppose that construction firms that implement contracts pay for the construction using credit from the private market and DFI subsidies.[21] For example a new park in a residential area may cost 2 million dollars to create, may provide value to the neighborhood of 10 million dollars but may only sell on the private market for 1 million dollars. Regardless of the information structure it is clear in this context that the private sector would not fund this project. It is also clear that were there to be a DFI subsidy of at least 1 million dollars then the private sector would find this project profitable and, by definition, DFI subsidization would result in additionality. Whether crowding out of private lending or investing occurs depends on the size of the subsidy; specifically, a subsidy of exactly 1 million dollars will result in no crowding out and deliver the project. However, if the subsidy is larger then 1 million dollars, then there will be crowd out. Now consider the previous example, but suppose that the cost of the project is either $500,000, 1 or 2 million dollars, and without loss of generality that the private sector is aware of the true cost.[22] Further suppose that the principal does not know the cost of the project when it decides on a subsidy amount, which can also take three values $0, $500,000 and 1 million dollars.[23]

This example is extremely simple. It serves to show how the size of public subsidy should depend on the believed distribution of underlying costs and the importance that the DFI places on additionality versus crowding out. In this example, additionality and crowding out are not binary alternatives - the DFI can make a project happen that would not otherwise (achieve additionality) whilst also partially crowding out the private sector.

First suppose the cost realized cost is 2 million dollars:

- With a $0 or $500,000 subsidy the project is not implemented. Trivially there is no crowd out.
- With 1 million dollar subsidy the project is implemented, DFI creates additionally and there is no crowd out.

Now suppose the realized cost is 1.5 million dollars.

[21] Realistically firms will only borrow to cover costs in excess of subsidies and for this example we shall assume the same.

[22] It is sufficient for the private sector to be "more informed" than the principal

[23] For simplicity consider 'choosing not to subsidize' as choosing a subsidy of $0.

- With a \$0 subsidy the project is not implemented. Trivially there is no crowd out.
- With a \$500,000 subsidy the project is implemented, DFI creates additionally and there is no crowd out.
- With a 1 million dollar subsidy the project is implemented and DFI provides additionally. However, *the private sector is partially crowded out.*

Now suppose the realized cost is \$ 500,000.

- With a \$0 subsidy the project is implemented, there is no DFI and there is no crowd out.
- With any positive subsidy the project is implemented. However, the DFI does not create additionality.

The next logical question is which of the three subsidy rules would the principal be best off choosing? The answer depends both on how likely the costs are and on the principal's value of additionality versus aversion to crowd-out.

Let us first ignore crowd-out and consider how the distribution of costs impacts the likelihood of additionality for the three subsidy rules. Specifically let us define γ_1 as the probability of the cost being \$ 500,000, γ_2 be the probability of the cost being 1 million dollars and $\gamma_3 = 1 - \gamma_1 - \gamma_2$ as the probability of the cost being 2 million dollars. As such the expected cost is $C = \$1,000,000(\gamma_1.5 + \gamma_2 + \gamma_3 2)$ and, depending on the values of γ_1 and γ_2, can take on any value between \$500,000 and 2 million dollars. With a risk neutral principal, the DFI will, in expectations, generate additionality whenever $C > 1$ million dollars.

Now let us consider crowd-out. Crowd-out from DFI financing occurs whenever the cost of the project is below the private value of the project or when the value of the project plus the subsidy is greater then the cost. As such, the only way there can be both additionality, no crowd out and no 'waste' of funds (i.e., extra rents from excessive funding), is when the cost of the project is above the private value and the subsidy is exactly equal to the difference.

For tractability we shall introduce a particular 'principal objective function'. This can be thought of as a mathematical representation of the trade-offs between funding projects that would not have otherwise been funded and crowding out private sector financing. The function we shall introduce should not be thought of as the 'correct' objective function, as it is up to the principal to decide how they value additionality versus crowd-out, rather this reasoning can be applied to any objective function of the principal and we shall choose one that is mathematically simple for the reader's convenience. For notational convenience let s be a subsidy amount, c be the realized cost, and p be the private sector value.

Let the principal's objective function be the added value of projects that would not otherwise exist less a multiple of the square of the crowd out. Recall that the neighborhood valued the park in this example at 10 million dollars, then the value of additionality will be 10 million less the cost of production, 10,000,000-c. Also if the subsidy more than covers the difference between the cost and the private value then every dollar in excess of the difference crowds out private sector funding and as such induces an offset burden of the form $\alpha(\min\{c, p\})^2$. In words, α is the relative weight the principal puts on the disutility of crowd out versus the benefit of additionally. If $\alpha = 0$, then the principal does not care at all about crowd out and only values additionality. By constant if α is large then the principal does not care at all about additionality and wants to ensure no crowd out.

Notice that increasing s makes the first component of the objective function larger, since larger subsidies make it more likely that the project occurs. By contrast increasing s has a negative effect on the principal's value though the second component, as a larger s will induce more crowd out. Suppose that each of the three costs are equally likely, then if $\alpha > 0.132$ the principal should choose not to subsidize the project, for such parameters the cost of crowd out is larger than the benefit of additionality. For $\alpha < 0.132$ the benefits of additionality outweigh the cost of crow out. Similarly for $\alpha < \frac{2}{25}$ the principal should offer the subsidy of 1 million dollars rather than the subsidy for \$500,000.

4.1 Non-rival agents, one principal

The above example abstracts from the size or number of agents, but this is a trivial abstraction so that all discussed above applies unchanged to the case of multiple, non-rival agents. Indeed, we can think of the project cost discussed above as the ex-post cost of the agent selected. Thus, the example above directly applies. Other aspects of interest in this case can be explored, although it will depend on the specific features of the DFI to what extent it will deal with these additional aspects.

Suppose, for instance, that the environment is as follows: a DFI has selected a target area to finance (such as wind or solar panel farms to support renewable energy, or small food processors in a certain region). Financing would be available for all projects in the target area. Bankers would approach potential borrowers, and selected which projects would be vetted for sound-banking. Assume the set of projects in question meet the sound banking requirements, and further, satisfy the impact assessment (does it deliver suitable impact on social, transition, poverty alleviation criteria, etc.) to be eligible for funding by the DFI, for a selected range of tasks. The problem is now to devise a contract that provides the right incentives to the borrowers in order to induce them to choose the task level maximizing the project's profitability (note: profitability can instead be thought of as the social value of the project) when only the borrowers know their private cost associated with given task levels (or *effort*). Finally, and most importantly, the DFI knows that borrowers are potentially heterogenous, but cannot know for certain the type (or ability or effort cost) of each potential borrower.

If the same contract is offered to all potential borrowers, extensive forgone opportunities to maximize the project outcome may occur as some borrowers might elect an inefficiently high/low task level. This situation very closely resembles what is described in section 3 above. Accordingly, we know that the optimal contract must be a menu of options designed in such a way that borrowers have an incentive to self select into the right option based on their privately known cost. This optimal contract must also entail a distortion relative to the first best outcome that would be achievable absent any asymmetric information. Stated otherwise, in the presence of asymmetric information, it is impossible to obtain the same profitability as in the case of full information, even when an ideal *optimal contract* is in place to steer the incentives of the borrowers in the right direction. When importing the lessons from mechanism design to consideration in the real world of DFIs, additional concerns arise. We know that the DFI, in order to design an optimal contract, must be able to properly assess a series of quantities that, especially in a development context, might be hard to estimate. These include the outside option of the borrower (determining his reservation utility $\bar{r}$), the shape of the cost and profit functions, the source and extent of asymmetric information and so on.

Finally, even when all the necessary information is available, and the optimal contract can be designed, its enforcement might be impossible. That may occur when either the contract requirements fall outside what is allowed under the local legal system, or because there are other obstacles arising from the local institutions that prevent the parties from trusting the binding nature of the contract. Far from operating within a vacuum, optimal contracts and the outcomes from mechanism design depend acutely on the environments in which they operate.

4.2 Screening and signaling examples

In the example at the beginning of section 4, assuming that firms are perfectly informed of the cost, but that the principal cannot observer this cost may seem like a strong assumption. Surely communication exists and it would seem reasonable that the principal could attempt to acquire information. However this assumption simply generates mathematically tractable and easily readable solutions.

Consider an instance where the private sector instead gets a signal about the cost of the project, and that firms decide to participate according to that signal, while the principal receives less information on the true cost. For reasonable signal spaces and under appropriate beliefs with respect to the relationship between the state of the world, signals and the actions of the other market participants responding to the signal, then there can still exist cases in which private firms choose to offer financing and the principal would choose to fill the expected financing gap for the project.

To make this more concrete, note that agents in this context are businesses and not all of their business activity is unobservable. Thus, as an example on how to set up a screening mechanism, consider a situation in which firms managed by female managers are more likely to execute projects in ways that are more likely to generate high social value than are those managed by male owners, but female managed firms tend to be underfunded by the private sector. If this were the case then the prevalence of female managers in a firm can serve as a signal to the principal of what projects are most likely to provide additionality. This kind of screening can help direct DFI subsidies towards projects that are more likely to provide additionality. Practical mechanism to implement this type of policies in terms of project set asides, quotas or bid subsidies are discussed for instance in Athey, Coey and Levin (2013).

The competitive screening setting illustrated above has multiple, competing principals and non-rival agents. We argued that the principals' incentive to undercut each other to maximize their own profits can lead to perverse situations of market failure where no separating equilibria are sustainable. If private banks are considered to be the competing principals, then a positive view of the role of DFIs is that they might be able to solve the cause of this market failure by absorbing some types of borrowers that the market cannot cover. The logic is the same for which the presence of a public insurance plan might enhance the effectiveness of the private insurance market: if the public plan can absorb some of the high risk types, then the private market can more effectively offer its products to a more homogenous pool of enrollees. In this kind of situation, it appears that additionality in the action of the DFI would go well beyond not crowding out the private market, but actually allow for such a market to function.

On the other hand, the model for signaling considered above can illustrate situations DFIs may face that present potentially severe problems. Consider a circumstance where borrowers undertake costly, but unproductive signaling actions to demonstrate a need for funds. These actions can amount to paying for certifications to verify their need for funds or actions to alter their structural (or legal) condition to qualify for subsidization. As argued above, in addition to the wasteful signal, a second problem can emerge in this type of environment if the equilibrium result then entails the pooling of all borrowers at the same contractual conditions and, hence, an inefficient allocation of agents to tasks. Furthermore, a different problem, related to the pooling of agents in a single contract and the competition between potential financers, occurs if inexperienced private donors compete with DFIs but provide generous subsidies. This might lead to a pooling equilibrium emerging around the generous terms while these donors operate, followed by a complete market collapse once those funds are exhausted (since they were unsustainable to begin with). An extensive literature on aid dependency exists, related to these concerns. Firms, in seeking out subsidies, may inefficiently chase signaling behavior undermining the long term growth of an industry or the economy as a whole.

4.3 Auction examples

The problem for the DFI in the example at the beginning of this section was a lack of information on the cost to the firm of implementing the park construction. The example abstracts from any allocation mechanisms, thus asserting the existence of a single "price" for construction. Now consider extending the assumptions to a case of multiple firms competing for the park contract where the DFI is free to select an allocation mechanism. The DFI may now be able to elicit the unknown cost. Given a set of potential firms competing for a contract and an auction format, there exist private valuations such that we can think of the contract's 'cost' as the expected winning bid. When the winning bid is below the reservation value, the DFI can efficiently assign the contract to a firm *at the true cost*, under appropriate conditions.

A large and growing interest exists in encouraging firms to adopt higher safety standards, energy saving and otherwise 'green' technologies by governments, public institutions and DFIs. Introducing and enforcing such policies via regulations in emerging markets can be politically and pragmatically challenging. DFIs might wish to incentivise firms to reduce their pollution levels, but the question is then which are the best firms to approach, and what level of incentives (i.e., at what cost, known to firms but unknown by the DFI, to reduce pollution) would be sufficient? To answer these questions, The World Bank Group's

Pilot Auction Facility for Methane and Climate Change Mitigation (PAF) has designed a successful approach drawing on the lessons from mechanism design in creating markets for reducing harmful pollutants. The group holds online 'forward auctions' for put-options by private firms to sell emission reductions in the future at a strike price.[24]

The design of the auction and creation of the put-options overcomes several challenges DFIs often face when attempting to contract for production or behavior changes. First, DFIs might face restrictions on the set of firms that they may contract with,[25] and worry that holding an auction is writing a 'contract' first, before knowing who will participant and at what terms. However, prescreening potential auction participants can address these issues, as the PAF does through a certification and refundable deposit scheme. Further, should winners of the auction later decide they cannot deliver on the pollution reduction commitments (or find it unprofitable to do so), then a secondary market for the put-options exists so that the winner can trade with other private sector firms.[26]

As discussed above, auctions typify the type of mechanisms that the mechanism design literature suggests to use to efficiently allocate a scarce resource between agents (in this case firms) that are privately informed of their benefit (in this case their need to produce toxic emissions as a consequence of their production activities). An important aspect mentioned above is that while under ideal conditions a simple first price auction with reserve price might be optimal to ensure efficiency and maximize the auctioneer's revenues and, as implied by the revenue equivalence theorem, most conventional auctions tend to deliver the same outcome, in complex environments these results are likely to be violated. Thus, while the basic principle of using an auction to elicit privately held information is still valid, the design of the specific auction rules requires a thorough understanding of the market and surrounding institutions. The design of the PAF is a clear instance where the toolbox of mechanism design was explicitly used to try to steer the players' incentives toward a behavior more in line with

[24]See http://blogs.worldbank.org/climatechange/using-auctions-support-climate-and-development-outcomes. Details on auction parameters for the forward auction: http://www.pilotauctionfacility.org/content/auction-parameters-second-auction.

[25]Institutions might not be allowed or want to work with firms tied to organized crime, engaging in questionable business practices, etc.

[26]This feature is an example how the DFI can relax the participation constraint of firms by carefully constructing the allocation mechanism.

the social good. Still, it is essential to acknowledge the presence of frictions in the market that prevent the straight application of the abstract findings from the literature.

4.4 Choosing Contractors

The leading example from this section abstracts from selection over contractors, indeed if there is heterogeneity in contractor performance, then the principal will want to select the best contractor. Suppose there are 2 contractors who have different likelihood of costs. Specifically suppose that the first contractor has a $\frac{1}{2}$ probability of costing \$500,000, a $\frac{1}{4}$ probability of costing 1.5 million dollars and a $\frac{1}{4}$ probability of costing 2 million dollars. Suppose the second contractor has a a $\frac{1}{3}$ probability of costing \$500,000, a $\frac{1}{3}$ probability of costing 1.5 million dollars and a $\frac{1}{3}$ probability of costing 2 million dollars. Assume the likelihoods, but not the realized costs, are known by the principal. Then the principal can simply select to undergo the same process as before, separately for each agent. Specifically, given their cost likelihoods, the principal can compute the additionality values and crowd out costs for each subsidy level using the same objective function from before and then select the optimal subsidy for each agent. In an additional final step, the principal selects the contractor that provides the highest value out of the two, giving the winner the subsidy that maximizes the principal's objective function for that specific contractor. For example if $\alpha = 0.1$ then the optimal subsidy level for the first contractor is \$500,000 and gives the principal a value of $\frac{7}{8}$, while the optimal subsidy level for the second contractor is \$500,000 giving the principal a value of 2. As such the principal should set the subsidy level at \$500,000 and contract with the second contractor.

However, conditional on any particular contractor being selected, the same incentive problems are present so long as the contractors cost (or the private sector value) is not observable to the the principal. Other considerations follow.

Although for some projects that a DFI might be interested in financing there will not be any form of competition, it is easy to imagine at least two environments where competition might matter. The first one is the case in which, given a certain project, like the construction of an electrical plant, multiple firms are competing to be selected for financing. The second one is the case in which entrepreneurs or projects from potentially different sectors and regions are put in competition due to the limited resources (time and/or money) of the DFI. The two cases are closely related, but not identical and for both a series of lessons can be drawn from mechanism design. Indeed, the presence of multiple, competing agents is at the heart of the mechanism design approach.

Identifying the contractor for a project. Consider the problem of a DFI offering a loan to perform a certain project, and firms competing by asking for a different interest on this loan. Should the loan be awarded to the firm asking the lowest interest and should it be indeed charged that interest? We have argued above that mechanism design can give an answer to this question by exploiting the powerful revelation principle, which allows us to focus on simple revelation mechanisms. In the context of choosing a contractor for a project, the 'optimal auction' result of Myerson (1981) tells us that the optimal mechanism tends to have the form of a first price auction with reserve price (such that bids failing to exceed the reserve price cannot win the auction). Relatedly, the revenue equivalence theorem tells us that we can concentrate, in a broad set of situation, on the analysis of simpler mechanisms as all mechanism within certain classes share the same expected payment rule. However, this result shall be qualified in a few ways. Two examples will help. The first involves a situation where, to boost the development goals, certain entrepreneurs are preferable to others (say to enhance equality, such as firms owned, managed or employing young, minorities or women). In this case, auctions with set-asides, quotas and bid-subsidies might all be preferable to a simple first price auction, depending on the nuances of the environment. The second case involves the possibility of 'defaults' (broadly defined) by the financed entrepreneur. When bids are not binding commitments and ex post the entrepreneur can default on his bid, simple first price auctions have the perverse effect of exacerbating the default risk by inducing strong price competition in the awarding phase followed by low profit margins and, hence, a high probability of default ex post. Mechanisms that can fix this situation involve more complex auction designs that might exploit: reputation, ex ante screening and ex post monitoring (Decarolis, 2017). Some of these latter institutions, might also be relevant to curb a special form of market failure often observed in public contracts: that involving the relationship between firms' collusion and their corruption.

4.5 Challenge funds and choosing projects

Similar to the example in section 4.4, instead of multiple contractors consider a single agent of the DFI evaluating two distinct projects, potentially from different sectors, with different distributions of costs as before. The same numerical example as in 4.4 can apply here, substituting multiple contractors with multiple projects. A plausible difference may, however, be that each project carries a different social benefit, such that the first project is for the park from the original example valued at 10 million dollars, where the second project is for a water treatment plant that a neighborhood in different city values at 15 million dollars.

DFIs often face the challenge of operating in markets where there may be few participants or potential entrants. Leveraging competition within such sectors or localities can be problematic if not impossible. Still, the scarcity of resources within the DFI can be utilized to help achieve efficient outcomes. Challenge funds can be used to foster competition between projects from different sectors or even countries by creating a competition for funding involving particular themes, such as energy efficiency. Such funds typically hold a call for business or project proposals from private sector companies in need of funding.[27] Finalists selected from the pool of applicants can then be reviewed more closely by agents from the fund, with the winners agreeing to specified repayment schedules and monitoring by the fund. Since candidate projects are put in competition, the social impact per funding amount can be compared across projects and provides incentive for the applicants not to greatly exaggerate their funding needs. While individual agents of DFIs often face similar resource restrictions over their energy and time, which are then passed down to firms designing their projects, challenge funds may offer an attractive alternative by drawing from wider communities or sectors, and making the competition between projects more salient.

Challenge funds are just one example of how DFIs can deal with the common problem of choosing between projects. A mechanism with good properties identified by the literature

[27]See The African Enterprise Challenge Fund (AECF) as one example of such funds.

for this purpose is the scoring rule auction/mechanism. These mechanisms can provide the principal with greater flexibility in targeting projects with certain features while retaining the transparency and selection discipline of auctions.

4.6 Search for Financial Partners

Most of the action in the model comes from the lack of credible communication between the principal, the agents and the financial sector. If the principal and the financial sector can communicate effectively, then the principal could make the subsidy value conditional on the private sector's willingness to finance. As such relationships between the principal and financial partners can help mitigate the impact of asymmetric information.

To map this discussion to the earlier example, we can think of a partnership with the financial sector as a cost incurred to reduce the asymmetry in information. Clearly, the example illustrates that the difference between the values in the first and second best contracts can be substantial. Thus, for any amount less then this difference, the principal would be willing to incur costs to remove asymmetric information.

Since financial institutions rely on high quality predictions of market outcomes they tend to have a better idea of what projects are more or less likely to be undertaken by the private sector. By providing the principal with better information on the likelihood of successful private funding the principal can choose only to offer DFI subsidies on projects that the financial sector does not think will be likely to be privately funded.

More broadly speaking, one last set of instances where mechanism design might speak to the activities of the DFI is in the choice between financing tools. DFIs may select to directly finance projects, provide funds for on-lending through other financial institutions, partner with other financial institutions via syndicated loans, or engage in cofinancing with private sector firms through equity investment, to highlight popular arrangements observed in practice. The asymmetric information problem influences both project selection, as addressed in the cases above, as well as the decision between modes of finance. The choice of financing tools might help minimize the subsidy rent extracted due to asymmetric information. As pointed out in Hainz and Hakens (2012), no de facto optimal choice as to which institutions should administer lending exists, but depends importantly on the market and institutional characteristics.

Market characteristics, such as the opacity of a market or industry can affect the screening and monitoring capacity of lenders. Equity financing might constitute a greater commitment of resources (or risk) to a project, but allows for closer monitoring. Alternatively, debt financing requires less exposure of the investing institution to risk, but might leave greater opportunities for mismanagement or taking actions that lead to higher potential rewards for the borrower but lower expected returns for the lender.

DFIs can seek to crowd-in private sector partners regardless of choosing equity or debt financing. DFIs of any size ultimately face resource limitations that may limit the projects in which they participate. By utilizing their relative expertise in screening or monitoring in certain markets, DFIs might attract commercial banks to participate in syndicated loans. Such projects can often achieve additionality due to their size and the reluctance of commercial banks to underwrite such large projects without partnering with the DFI. In other cases, the DFI may choose to instead co-invest in a firm with a third party firm via equity. The source of risk in the environment and confidence in institutions to support and enforce contracts define the trade-offs presented by each mode of finance.[28]

[28]Syndicating a loan does not guarantee additionality. However, this can happen when, for instance, the private sector lacks the reviewing and/or monitoring technology to properly evaluate the project that the DFI is interested in. Traditionally, the DFI might finance the whole project. Since the private sector cannot make the loan, there is no crowding out. However, the DFI can take on the review or monitoring bank role and syndicate out shares of the loan to the private sector. Now we have additionality and crowding-in. Nevertheless, in general nothing guarantees that the private sector could not do everything without a DFI. But many of the same mechanisms already discussed earlier apply to this case as well. The DFI has limited resources, so choosing between potential projects to syndicate can generate internal competition. Furthermore, syndication also allows the DFI to go after projects that may be too large for it to finance independently, which in itself is additionality if the private sector also lacks the size and motivation to create a syndicate. Finally, as allured earlier on, syndication is also a way for the DFI to 'rent out' its sector expertise (or other expertise) to banks in order to bring in the private sector, thus potentially reducing the asymmetric information that the lack of this expertise can generate and, through that, allowing the creation of trades that would be impossible otherwise.

5 Conclusions

This paper provides a discussion of mechanism design as it applies to situations commonly encountered by development finance institutions. More broadly, the issues discussed here apply to any investor operating in emerging or other markets where asymmetric information may pose significant hurdles impeding the efficient setting of terms. The greater lessons from mechanism design, rather than prescribing a set of rules to be applied broadly across circumstances, instead highlights the importance of understanding the market environment and institutions involved in selecting potential projects, setting the terms of finance and determining the enforcement of agreements.

Certain settings of hidden information, such as the ability of a firm to adopt new practices, might be resolved through carefully constructing a choice of debt finance contracts and allowing the entrepreneurs to self sort into the option best suited to their ability level. When environments include costly signals, agents may be able to separate themselves from others as being high achieving, but the equilibria may be costly in terms of efficiency compared with the achievable outcomes with the full disclosure of information. When competition between projects or agents may be levered, then a larger set of tools from mechanism design may be employed. Still, as discussed in above, many results are fragile and depend on strong assumptions for most environments. Rather than viewing these outcomes as failures, the theory instead directs one to consider the concessions that must instead be made with their associated costs from information rents reluctantly awarded to the participants.

The outcomes from applications of the mechanism design are pervasive in everyday life, from the advertisements one sees on search engines to the costs taken on by the government to procure construction firms to build new roads. Examples from DFIs in emerging markets, such as the forward auctions implemented by the World Bank for pollution reductions and the challenge funds awarding project finance in Africa, constitute just two cases benefiting from the mechanism design literature. By their very nature of operating in emerging markets that may lack the financial markets or other institutions intended to facilitate investment and trade, DFIs are likely to face opportunities to increase the efficiency of their funds by recognizing the environment and tools that may be available to them in eliciting truthful revelation of funding needs and deliverable products. The reduction of information rents and support for improved institutions can help contribute to the greater sustainability of financing with donor and public funds.

References

Athey, S., D. Coey, J. levin, (2013). "Set-Asides and Subsidies in Auctions," *American Economic Journal: Microeconomics*, 5 (1), 1-27.

Bolton, P., Dewatripont, M., (2005). *Contract theory.* MIT press.

Decarolis, F., (2017). "Comparing Procurement Auctions," forthcoming at the *International Economic Review.*

Fudenberg, D., Tirole, J., (1991). "Perfect Bayesian equilibrium and sequential equilibrium," *Journal of Economic Theory*, 53(2), 236-260.

Green, J., Laffont, J. J., (1979). "On coalition incentive compatibility," *The Review of Economic Studies*, 46(2), 243-254.

Hainz, C., Hakenes, H., (2012). "The politician and his banker - How to efficiently grant state aid," *Journal of Public Economics*, 96(1), 218-225.

Holmstrom, B., Myerson, R. B., (1983). "Efficient and durable decision rules with incomplete information," *Econometrica*, 1799-1819.

Larsen, B., A.L. Zhang, (2017). "A Mechanism Design Approach to Identification and Estimation," mimeo.

Li, S., (2015). "Obviously strategy-proof mechanisms," mimeo.

Mas-Colell, A., Whinston, M. D., Green, J. R., (1995). *Microeconomic theory.* New York: Oxford university press.

Myerson, R. B., (1981). "Optimal auction design," *Mathematics of operations research*, 6(1), 58-73.

Myerson, R. B., (1991). *Game theory: analysis of conflict.* Harvard University.

Roth, A. E., (2015). *Who Gets What - and Why: The New Economics of Matchmaking and Market Design.* Houghton Mifflin Harcourt, Chicago.

Rothschild, M., Stiglitz, J., (1976). "Equilibrium in competitive insurance markets: An essay on the economics of imperfect information," in *Foundations of Insurance Economics* (pp. 355-375). Springer, Netherlands.

Wilson, C. (1977). "A model of insurance markets with incomplete information," *Journal of Economic Theory*, 16(2), 167-207.