

The Subsidy Sorting Hat

Paddy Carter and Mark Plant*

In the Harry Potter novels, a magic hat decides which of four school houses new pupils should join. Development finance institutions (DFIs) need something like that when trying to decide which private firms to subsidise, although applicants only need sorting into two groups: firms that are doing something socially valuable and which genuinely require a subsidy, and firms that are merely trying their luck to get a subsidy for a project they would undertake in any case.

If we want to enable socially worthwhile investments that financiers motivated by private returns will not back, then we need to offer firms some sort of a subsidy. But if we offer subsidies to private firms, they will all want one. So, unless we can sort genuine need from opportunism, public funds will be wasted.

Government R&D funding agencies have existed for decades. Offering explicit subsidies (sometimes called blended concessional finance) is relatively new for DFIs, who have traditionally invested on commercial (i.e., profit-based) terms. The current approach to subsidy allocation in the context of development finance is for specialist funding agencies to vet applicants as best they can, by looking at project details and figuring out whether it looks like a subsidy is needed.

Economists think there is another way. Rather than rely on expertise, judgement, and good faith, contracts could be designed so that private firms sort themselves into the right groups—like magic!

ODA RAISES THE STAKES

Everyone understands that investing means taking bets, some of which will not pay off. Most people probably understand that the optimal amount of waste in government spending is not zero—the only way to achieve that would be to do too little. When the US Department of Energy lost \$528m backing the renewable energy firm Solyndra, opposition politicians portrayed it as a scandal. The more sensible reaction was that if you are not getting it wrong occasionally you are not trying hard enough, and that the more relevant number was total profits in the order of \$5bn generated by the funding program overall.¹

But fishing deeper in the pool of projects that do not offer commercially appealing financial returns, in an attempt to have more development impact, could entail losing money on average, not making it. Many DFIs are funded from the foreign aid budget. Wasting public money is not ideal at the best of times but allocating subsidies to firms that do not really need them, resulting in excess profits for private investors who are probably already wealthy, would be even more painful when the money comes from a budget that is supposed to be saving lives and putting children in school.

¹ <https://krugman.blogs.nytimes.com/2011/09/25/the-solyndra-scandal/>

* Paddy Carter is writing in a personal capacity and this note does not represent the views of CDC Group

The case for taking scarce public money from the foreign aid budget and using it to subsidise the private sector is motivated by a more general logic: some firms generate benefits to society that are far larger than the private returns captured by investors. Economists call these benefits “positive externalities.” When that happens, firms will not invest enough, or will invest in the wrong things, from society’s point of view. Some of these externalities align with the objectives of foreign aid: to reduce poverty and raise the quality of life and—increasingly—tackle climate change.

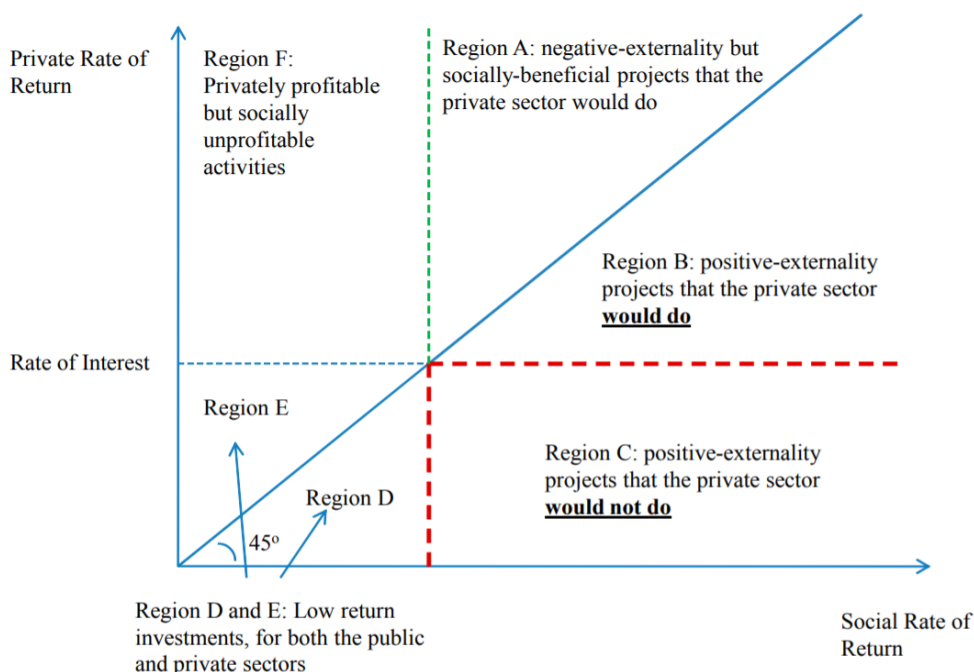
A firm could pioneer an innovation with a huge positive impact, [greenhouses that use sunshine to desalinate water](#) for example, but might choose not to if it fears that it will fail to recoup its investment after the market is flooded with copycat firms that can do it cheaper after learning from its mistakes. It is easy to understand why we might want to subsidise private firms to pioneer green technologies, but social benefits can go beyond the obvious; for example, in economies with widespread under-employment, [good jobs are an externality](#) to firms’ investment decisions that could justify a subsidy.

Weighing the costs and benefits of subsidies is not easy, but if the good-things-for-development achieved by using aid to subsidise private investment are comparable, dollar for dollar, with the good-things-for-development achieved by traditional grant-funded aid, then subsidising private investment is a sensible use of aid.

THE OUTCOME WE WANT

Let’s say we have a fair idea of which firms are going to generate positive social benefits. Our knowledge about that may be imperfect, but it is a good starting point for analysis. However, that only gets us halfway there. Next, we need to sort those who need a subsidy from those who do not. Why? We want to support projects which generate positive externalities, but we don’t want public support to be redundant.

The figure below, taken from an [IMF working paper by Andrew Warner](#), shows what we want to achieve. We need our magic hat—or clever contract—to sort socially beneficial firms that are asking for



a subsidy into regions B and C.

The potential solution comes from a branch of economics that is all about inducing actors to reveal their private information, called mechanism design. Its better-known cousin, game theory, involves taking a strategic situation and figuring out what the players will do; mechanism design starts from what you want the players to do, and then figuring out what game to make them play.

Often the mechanisms that economists devise to solve these problems are auctions of one sort or another that exploit competition between bidders to get them to truthfully disclose their type. [Charles Kenny at the Center for Global Development has urged donors to find competitive mechanisms for allocating subsidies](#). Auctions are likely to work best in more mature sectors with relatively low transaction costs that will attract bidders, such as procuring solar parks. It is also not clear how auctions would work if instead of starting with something you want to build, and are looking for someone to build it, you are starting with a project associated with a unique promoter and are looking to finance it with minimal public subsidy. Auctions are an important mechanism to get potential suppliers to reveal their private information about their costs etc., and can minimise subsidies when they are payments for outcomes like [methane abatement](#). When the subsidy is in the form of concessional capital, persuading project sponsors to participate in a process designed to minimise the concession (maximise their cost of capital) is a trickier proposition. One lesson to draw from this could be that we want to find ways of allocating subsidies that look more like payments for outcomes, and less like concessional capital.

But when concessional capital is called for, there is an alternative to a competitive process called “screening.” Actors can be induced to reveal their private information by contracts cleverly designed in such a way that those who genuinely need a subsidy accept it, and those who do not reject it.

SCREENING CONTRACTS

An early attempt at this problem was the paper “[The politician and his banker](#)” by the economists Christa Hainz and Hendrik Hakenes. The authors assume it is possible to discover what kind of project has applied for finance, but if a politician delegates that task to a development bank, it won’t always use that information as we would want.

The model is simple; there are three types of projects, all of which make the same private financial return if successful, but which have different chances of success: low, medium, and high. The private sector will fund high-type projects. Everyone rejects low-type projects, which fail too often to be worth subsidising. The problem comes with the medium-type projects: politicians would like the development bank to fund them, but because only the bank learns project types, staff at the development bank face an incentive to say yes to high-type applications and pretend they are medium.

The solution is simple too: the politician imposes a rule that the development bank must charge a higher interest rate than private banks. That means we do not have to worry about the development bank sneakily funding high types, because entrepreneurs with high-type projects will prefer to go to commercial banks, where the lower interest rate will leave them with surplus returns. Medium-type projects are rejected by commercial banks and must turn to the development bank where the higher rate of interest will leave them with no surplus when projects succeed (just the bare minimum they need to participate). Despite charging a higher interest rate, because the medium-type projects fail more often, the development bank will lose money and hence require a public subsidy. The private bank charges a lower interest rate but is repaid more often and so covers its costs.

That solution is neat, and works in a simple model, but the model is not general enough to include the possibility of funding socially worthwhile projects that even if successful could not afford to repay that higher interest rate.

MECHANISM DESIGN

A new paper by the economists Saul Lach, Zvika Neeman, and Mark Schankerman takes the screening idea further. Unlike in the Hainz and Hakenes model, in which the government can delegate decisions to a development bank capable of observing a project's true type, the authors tackle the harder problem in which a project's true type is unobservable to the public sector. They also do not assume that the supply of applicants is given. That creates a problem with two dimensions: as before, we do not want to offer redundant subsidies to firms that could have obtained private finance, but now we also want to attract as many applicants as possible so that we do not miss out on opportunities to create benefits for society. These two objectives are in tension, and the paper uses the mechanism design approach to derive a contract that achieves the trick of helping as many unprofitable projects as possible if they generate positive social value, whilst also subsidising as few opportunistic applicants as possible.

The paper is written with government subsidies in the form of loans for R&D in mind, but its lessons could be applied to any project with social returns that exceed private returns. Instead of competing with private banks, in this model the government is up against venture capitalists. Venture capitalists differ from banks, because in addition to finance they offer expertise that raises the probability of a project's success. The government has two tools: the interest rate it charges and the level of co-financing it demands. Government R&D subsidy schemes often require the entrepreneur to put up some of their own money, on the theory that "skin in the game" (co-financing) is needed to induce effort, and government schemes also usually charge zero or very low interest to encourage as much R&D as possible. Surprisingly, Schankerman and co-authors show that the optimal contract looks very different: (almost) no co-payment by the entrepreneur but high interest rates in most economic environments they study in their simulation analysis.

In this model, the loan is only repaid if the project is successful. This contractual arrangement is common when governments finance R&D, where the technical success or failure of a project is easier to identify. Such a contract insures the researcher against failure, and so increases willingness to undertake risky projects. More generally, we could think of this as a limited liability loan without collateral, where the lender will not be repaid when if the project fails. The element of the contract that requires a co-payment works in the other direction: any money the entrepreneur must pay upfront is lost if the project fails.

Venture capitalists, on the other hand, take equity stakes in projects in return for their money, and they would reject projects once they can't break even, even if they are entitled to 100 percent of the returns.

Projects can differ in their probability of success, size of the positive externality, and the magnitude of expected private returns but many of the main results emerge from a simpler version of the model, in which projects only differ in terms of their probability of success. Like Hainz and Hakenes, the government is trying to "target the middle." If projects are a good bet, VCs will back them; if they are too unlikely to succeed, nobody should; the government wants intermediate projects with a decent change of generating positive externalities, but which VCs would reject.

The model includes various parameters to reflect real world considerations, such as the cost of public funds, size of the project externalities, and the effectiveness of the VC market, which we might expect to affect the optimal contract in the following ways. First, the higher the cost of public funds, the less tolerant we should be of wasted subsidies. Secondly, the more effective VCs are at raising the probability of project success, the less we want to give public money to projects that VCs would have funded. Lastly, the larger the positive externality, the more important it is to attract as many applicants as possible.

To their surprise, the authors found that for most configurations of these parameters (cost of capital, effectiveness of VCs, size of externality) the optimal contract looked the same: almost-zero co-payment by the entrepreneur and an interest rate that takes all the returns away from the entrepreneur, save the minimum needed to induce their participation and effort. The absence of substantial co-payment spreads the net for projects as widely as possible. Because it exposes the entrepreneur to almost no risk, almost nobody is dissuaded from applying because they cannot afford it. The interest rate is set to keep private returns to a minimum when the project succeeds, which means that anyone who can get funded by a VC prefers to do so. Only for a very few parameter combinations does the optimal contract flip to another extreme of requiring a co-payment but setting the interest rate to -1 (e.g., a full grant)—the cases where this happens to be optimal require that cost of public funds is very low and the size of the project externalities very high.

We should not get hung up on the specific solutions in a stylized model. The almost-zero co-payment result, for example, is probably an artefact of a theoretical environment in which one only needs a tiny co-payment to screen out projects with a low probability of succeeding, because throwing away money is never rational, no matter how small the sum. It might require rather more ‘skin in the game’ to signal that the entrepreneur has a decent chance of succeeding. What matters are the insights and “direction of travel.” The model suggests that asking for too much money upfront from entrepreneurs is a bad idea because it shrinks the set of applicants, and some of those projects may be high-social-impact projects that the government would want to fund. At the same time, entrepreneurs need a minimal level of return to induce participation and effort, but the ideal contract should take away any returns above that, to ensure nobody will take public money unless they really need it.

FROM THE BLACKBOARD TO REALITY

How does this theoretical ideal contract compare to what DFIs currently do, and what they could plausibly do?

The problem of subsidising innovation is not the core business of DFIs, which is to provide finance on close-to-commercial terms to finance the expansion of more established businesses. That roughly translates to providing finance whilst making commercial levels of profit.² Commercial pricing is important, both because DFIs want to create commercially sustainable enterprises that will survive once they exit, and because they want to encourage private financiers to follow their lead into frontier markets. DFIs must be additional, which means doing things the private market would not, but they also want to set pricing (returns) benchmarks so that commercial investors might contemplate following their example. Furthermore, DFIs are often dealing with businesses whose options include

² The idea of supplying finance in segments of markets where the private market does not, but on terms market participants would offer if they did, is complicated and would take us too far from this note’s topic. One of the authors of the mechanism design paper, Mark Schankerman, was the co-author of a seminal paper on that question, with the then chief economist at the EBRD, Willem Buiter: “Blended Finance and Subsidies,” [EBRD Working Paper \(2002\)](#).

“do something else entirely” and “do nothing at all,” so making their finance as unattractive as possible is not a tactic they will willingly adopt.

To fix ideas in a more realistic context, we could think of a DFI that wishes to embark on a new business line, financing pioneering green technology start-ups that the private VC industry would not, and which are too risky to be funded from the DFI's main balance sheet. But the DFI does not want to go to market offering finance on such attractive terms that it starts to poach start-ups that could have gone to VCs.

Inspired by mechanism design theory, the ideal contract would capture all returns above the minimum needed to secure the entrepreneurs participation and effort. But we are not in a simple model where the financial returns to projects are identical if they succeed, so the interest rate charged would need to vary across projects. And because even for a given project “success” is not a binary but is a continuous variable, ideally payments under the contract should respond to realised returns, and not be fixed in advance like the interest on a loan. In which case the first problem is that monitoring financial performance and writing a contract that charges better-performing businesses a higher interest rate would probably be too difficult. It would introduce an incentive for the entrepreneur to conceal profits.

That suggests something more like an equity investment, in which the project sponsor also holds equity so that their incentives are more closely aligned with those of the DFI. Both the entrepreneur and the DFI will gain from higher reported profits. But rather than share the residual returns equally across equity investors (in proportion to shareholding) as normal, it would need to have an unusual feature that it adjusts the distribution of returns in response to the commercial success of the investment. There are some precedents in the market—we have come across some structures that cap the return of public sector investors to create greater potential upside for private participants. We want to do something similar, but in reverse.

Limiting (though not fully removing) the entrepreneur's upside in that way should encourage them to take private finance, when it is available, and thus avoid crowding-out. The theoretical result that zero co-payment is optimal might suggest that DFIs should not require project sponsors to stump up a large share of the equity finance themselves, but could instead consider structures that promise the sponsor returns without the initial equity investment that would usually be required to get it. But a contract in which the DFI owns most of the equity would have implications for control rights. If DFIs do not want to find themselves having to manage the enterprises that they invest in, an innovative structure that separates ownership, control rights, and the distribution of financial returns might be needed.

The role of co-payment in bringing some assurance that the project is not a duffer has already been discussed, but DFIs might also look for other investors willing to put money on the line. The idea here is that if somebody has a good idea and a credible team, then they should be able to raise some external equity somewhere. DFIs gain confidence when other investors are willing to back a project, so when they encounter a project sponsor with no ability to raise some money themselves, it would be a red flag. A potential concern could be that this mechanism alone will not screen out low quality projects – perhaps because irrationally optimistic entrepreneurs and Angel investors are willing to take risks with their own money that the public sector should not – so some additional due diligence may be required.

A final potential problem could be that because the entrepreneur requires some upside to incentivise

effort, and that probably needs to be negotiated case by case, DFIs would find themselves back in the world of not knowing when they have been too generous and have crowded-out private finance.

All that said, people who are familiar with conventional VC contracts should be able to tell when they are doing something different, and a contract that varies the distribution of returns under different scenarios should be easily distinguishable from the norm. To cast the net as wide as possible, we want a contract that set lowers barriers to participation and insures the entrepreneur against failure in a way that a standard equity co-investment would not, but which also takes away (much of the) returns from the entrepreneur when things go well, in a way that standard equity co-investment would not.

No doubt designing and implementing such an innovative contract would be difficult, and practical considerations will likely require some compromises. But the perfect should not get in the way of the good. Improvements are possible in the way governments, foundations, and others structure contracts that would increase their effectiveness and reduce the redundancy in impact-investments. The holy grail of an innovative contract that will allow DFIs to use public funds to subsidise private enterprise in a way that solves the additionality problem—which will meaningfully increase the probability that any subsidy is genuinely warranted, not wasted—seems sufficiently attractive to be worth considerable effort exploring.

At the Center for Global Development, figuring out ways in which development finance can be used more effectively to make progress towards the Sustainable Development Goals is [one of our core priorities](#). We are hoping to facilitate conversations between creative academics and development finance practitioners to see if these ideas can be turned into action. Please contact us if you are interested in joining that conversation.

The authors are grateful to Professor Schankerman for very helpful discussion and suggestions.

REFERENCES

Warner, Andrew M. (2013), “A Framework for Efficient Government Investment,” IMF Working Paper

Hainz, Christa and Hendrik Hakenes (2012), “The politician and his banker – How to efficiently grant state aid,” *Journal of Public Economics*, 96, issue 1, p. 218-225,

Lach, Saul, Zvika Neeman and Mark Schankerman (2018), “Government Financing of R&D: A Mechanism Design Approach,” CEPR Discussion Paper 12199 (revised January 2020), under revision for the *American Economic Journal: Microeconomics*



WWW.CGDEV.ORG

This work is made available under the terms of the Creative Commons Attribution-NonCommercial 4.0 license.

PADDY CARTER is director of research and policy at CDC Group.

MARK PLANT is a senior fellow at the Center for Global Development.