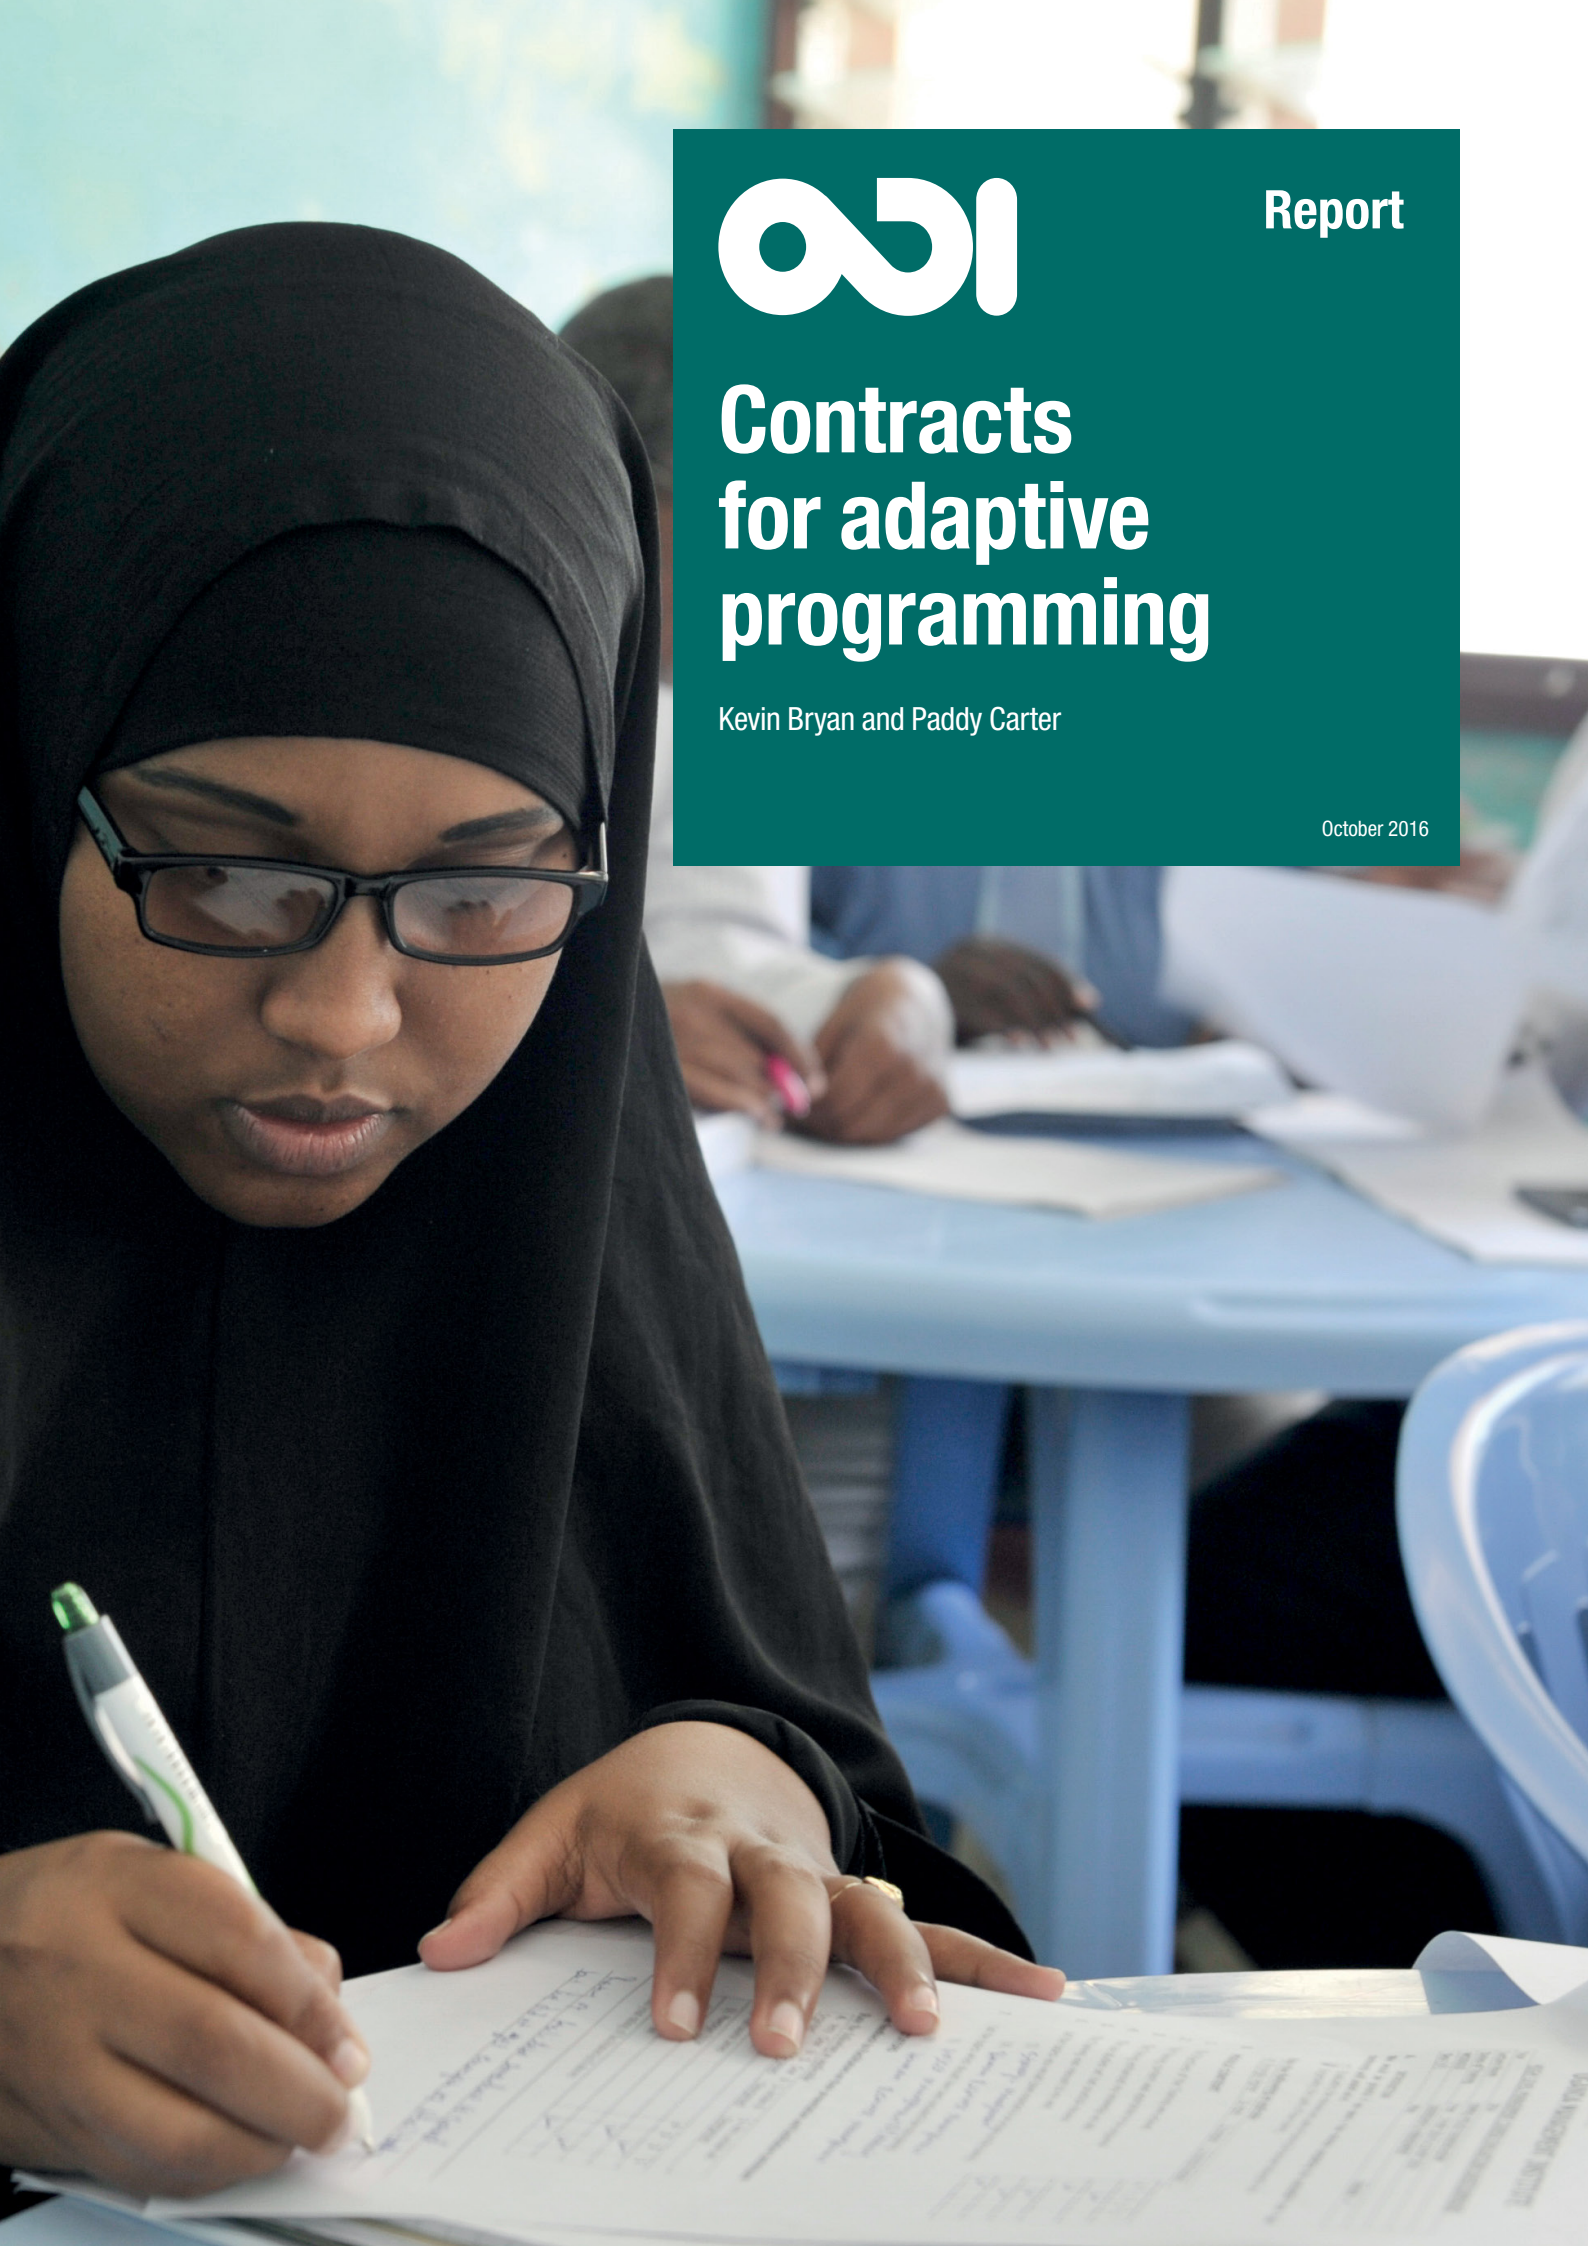

Report

Contracts for adaptive programming

Kevin Bryan and Paddy Carter

October 2016

Overseas Development Institute
203 Blackfriars Road
London SE1 8NJ

Tel. +44 (0) 20 7922 0300
Fax. +44 (0) 20 7922 0399
E-mail: info@odi.org.uk

www.odi.org
www.odi.org/facebook
www.odi.org/twitter

Cover photo: AU UNIST PHOTO/David Mutua 2014. A female senior civil servant during training at General Kaahiye Police Academy in Mogadishu, Somalia, by consultants from the Ugandan Management Unit, 24th May 2014.

Acknowledgements

The authors are grateful to Myra Bernadi, Maitreesh Ghatak, Romilly Greenhill, Marina Halac, Leni Wild and Liam Wren-Lewis for helpful comments and suggestions. Errors and omissions are the responsibilities of the authors. Funding from the Bill and Melinda Gates Foundation is also gratefully acknowledged. The usual disclaimer applies.

Contents

Boxes

Executive summary

The idea that development actors should experiment, learn and adapt is hard to disagree with. But how can this approach, referred to here as adaptive programming, be delivered through contractual arrangements with implementation partners?

This paper surveys a branch of economics known as contract theory, with the objective of drawing attention to some of the challenges that emerge when viewing adaptive programming from a contracting perspective. We ask: what insights does contract theory offer to practitioners of adaptive programming?

The premise of contract theory is that, because the aid agencies or governments who commission the services of non-governmental organisations (NGOs) and consultants often cannot know what actually happens during experiments, service providers may not genuinely experiment or may conceal the results if it is in their interest to do so. The essence of adaptive programming is to learn and change. When contracts are subject to revision, service providers will worry about aid agencies using the results of experimentation against their interests. Contract theory is concerned with designing contracts that overcome this incentive to conceal information, and induce genuine experimentation and truthful disclosure. This paper sketches some of the solutions.

Contract theory also presents a warning for proponents of adaptive programming. Despite the great potential to improve on the status quo of development practice, which has so often produced disappointing results, more open-ended contracts can introduce more scope for manipulation. Under fixed contracts, service providers face little incentive to try to manipulate the terms of their contract tomorrow via their actions today; adaptive contracts allow aid agencies to take advantage of learning, but service providers may also take advantage and shift the contract to more favourable terms. The promise of adaptive programming could be subverted when poorly designed contracts are used with service providers whose first concern is their convenience and profitability, or the continuity of their own business.

Academic contract theory is abstract and not tailored to the context of development interventions. The motivations that economists assume of parties to a contract, and the nature of the problems that they face, may not always map very obviously onto the motives and problems of development practitioners. The intention of this paper is to introduce development practitioners to contract theory and the incentives created by adaptive contracts, and to identify areas where future theoretical work could better address the problems faced by aid agencies. There is also a need to complement theory with empirical research, to understand the nature of contracts currently in use by development actors and to gather evidence on their comparative performance.

Box 1: What is an adaptive contract?

Many practitioners and scholars now advocate a new approach to international development cooperation: tackling locally-defined problems through a process of experimentation, learning and adaptation.

From contracting perspective, the contract for a traditional development project is static and simply specifies how much will be paid for delivering pre-specified actions or outputs. But best laid plans often go awry. Introducing the flexibility to adjust plans during implementation implies not fully pinning down objectives and methods in advance. From a contracting perspective, that can mean not completely specifying in the initial contract how much will be paid for doing what (although, as this paper will show, such uncertainty creates problems to which the solution can be a contract that makes limited commitments).

For the purpose of this paper we define an adaptive contract as one that encourages experimentation, learning and adaptation. Although this set of contracts includes payment-by-results (PbR) and other forms of static contract designed to encourage experimentation, this paper also concerns itself with dynamic contracts which do not pin down how what is required (and what will be paid) in later periods may respond to what occurred in earlier periods.

1. Contracts and contracting problems

A contract is any arrangement between two parties to procure goods or services, ranging from a tightly defined contract to deliver specific goods and services, to a grant that imposes very few constraints on the grantee. Contracts matter when the respective parties have different objectives, when implementing agents have access to private information, and when outcomes are not fully under their control.

The two parties to a contract are known as the principal – who wants something done – and the agent – who is paid to do it. In this paper we have in mind contracts between aid agencies and services providers like NGOs or consulting firms, although the relationship between donors and recipient governments can also be seen through the lens of contract theory, as can that between head office and field workers.

The critical assumption of principal-agent theory is that the principal and the agent do not share the same goals (a standard assumption of economists is that agents are self-interested). In the context of development practice, however, principals and agents might more often share similar goals, meaning that contracting problems may be less severe. But whilst it is easy for development contractors to claim intrinsic motivation, it would be rash to assume that their objectives are perfectly aligned with aid agency goals. A simple example of private information which agents may possess but face incentives not to reveal is that they are ineffective and the donor would do better to terminate their employment. We might expect few contractors to volunteer that information.

2. Complete contracts, payment-by-results and messy reality

In an ideal setting, the best contract is complete and optimal, and simply pays the agent for doing whatever it is that the principal wants done in all circumstances. This is reminiscent of a contracting arrangement known as payment by results (PbR) in development circles. But whilst a PbR contract might be complete (because it defines payments in all circumstances) it typically will not be optimal; there will usually be circumstances in which the principal would have been better served by offering different terms or by asking the agent to do something else (see Box 4). Complete and optimal contracts are rarely feasible.

Contract theory is occupied with the question of how incomplete contracts perform in less than ideal settings. PbR might be the best choice in less than ideal settings; much has already been written on the circumstances in which PbR is likely to perform well. This paper also considers settings where PbR is not suitable, perhaps because there may be no measurable 'result' that adequately captures what the principal is trying to achieve, or because desired results cannot be specified in advance. We focus on contracts which are subject to revision in the light of new information.

In the commercial world, corporations make surprisingly little use of contracts that tightly tie payments to specific actions, preferring more vaguely-worded contracts that leave plenty of room for interpretation. This reveals that the conditions under which tightly-specified formal contracts perform well are rare. Two problems in particular limit the use of these ideal contracts: hidden actions and hidden information. When agents are not wholly responsible for outcomes, contracts where payments are based on outcomes cannot induce the desired actions without paying a premium for risk, which can make such contracts costly. A particular problem arises when a task involves both observable and unobservable elements. In these cases, basing payments on what can be measured will distort behaviour away from what the principal really wants (a problem that may often be relevant to PbR).

Ideally all the information relevant to a contract will be verifiable, so that disputes can be settled in court or by another third-party. Otherwise, contract theorists call less formal contracts relational, when rather than being enforced by the courts, the contract is given force by the parties' desire to maintain a relationship in the future. Adaptive contracts are more likely to be couched in terms of high-level objectives, making them less easy to enforce in court, and hence more likely to be relational. The performance of relational contracts rests on clarity and credibility, both of which are relevant in the context of adaptive programming. Credibility might be a challenge; the agent must trust that the principal genuinely wants experimentation, and will reward the admission of failure in a political context where donors are under pressure to be accountable and demonstrate results.

3. Contracts for experimentation and the ratchet effect

The big problem with adaptive contracts is that agents may be reluctant to experiment and truthfully reveal results due to what theorists refer to as the 'ratchet effect'. This term originally referred to the idea that if an agent reveals that a task is easier to perform than thought, the principal will 'ratchet up' what is required in the future. But there are other reasons why agents may either be reluctant to genuinely experiment or divulge the results honestly, such as the risk of putting oneself out of a job. Ratchet effect worries can be 'solved' if the principal can commit to not using information generated against the interests of the agent, however this limits the usefulness of adaptive contracts because the principal may want to alter the contract based on what has been learned. So the desire for adaptation can mean that the principal is unwilling to commit, and agents will know that the actions they take today, or information they reveal, may affect the nature of the contractual relationship tomorrow.

4. Contract design principles for effective experimentation

Adaptive contracts that hope to benefit from learning and experimentation must address the problem that principals do not know everything that the agent knows, cannot observe everything that the agent does, and must also

limit any incentives for agents to manipulate the process of adaptation in their favour.

Contract theory cannot completely solve these and other problems. in the sense of delivering the same results as would be possible in an ideal setting, with perfect information. The goal of contract theory is to find the contract that does the best possible job, in an imperfect world. Successful adaptive programming in challenging circumstances will require clever contractual design. The theories reviewed in this paper suggest the following principles:

1. A contract that makes a partial commitment can mitigate the ratchet effect. The commitment is only partial, because after a specified period of time the principal can terminate the arrangement. The principal fully delegates the choice of activity to the agent but, once time is up, if the results do not look good then the contract reverts to a status-quo project. This limits the downside to the aid agency from self-interested experimentation by contractors, and is suitable when results-based bonuses are not appropriate.
2. When it is possible to base payments on results, early successes should not be heavily rewarded (or at all) in cases where genuine experimentation is likely to deliver fewer good outcomes at first, but increases the chances of good outcomes in the long run.
3. It may be possible to offer a menu of contracts so that agents with different levels of ability select different contracts, where those less likely to succeed take low-risk, low-reward contracts, and higher ability agents choose risky and potentially more lucrative contracts, which impose greater penalties upon failure.

The contracting solutions described in this paper arise from how contract theorists conceive the environment that principals and agents operate in, what determines their payoffs, and what different models assume about what can be observed.

It is not always clear how well the theory corresponds to the reality of development practice, and many open questions remain about how to apply the tools developed by contract theorists to problems faced by development practitioners. The gap between theoretic insight and pragmatic contract design also needs to be closed. From an empirical perspective, much could be gained by studying the performance of contracts used in the real world by development agencies and contractors doing adaptive programming.

1. Introduction

'The best laid schemes o' mice an' men
Gang aft a-gley'
Robert Burns, 'Tae a Moose (To a mouse)', 1785

The further one travels from relatively straightforward interventions like conducting vaccination programmes towards engagement with complex problems like general health system strengthening, the harder it is to find foreign aid success stories.[1]

In response to the disappointing track record of many traditional development projects, many practitioners and scholars now advocate a different approach to international development cooperation, consisting of tackling locally-defined problems through a process of experimentation, learning and adaptation. This approach to development goes by different names – here we call it adaptive programming.[2]

A review of 44 health sector projects pursued by the World Bank and Global Fund provides some insight into the difference between success and failure:

> *Successful projects pursued locally defined, specific problems in a demonstrable and continuous fashion... the projects were initiated as responses to locally defined problems, baseline indicators of these problems were measured in the early stages of the project, project activities were directly determined as solutions to these problems, and progress in solving problems was routinely evaluated and considered in adjusting project content.*
> (Andrews et al., 2015)

The idea that development agencies should respond to local needs, learn and adapt, is hard to disagree with. But how exactly can this approach be delivered through contractual arrangements with implementation partners? A ubiquitous feature of contemporary development practice is that governments and donors often subcontract the implementation of development projects to NGOs and private firms. Even when projects are implemented directly by donor agency staff, contracting problems still arise through the objectives given to staff in the field and in how their performance is evaluated and rewarded.[3]

Traditional development projects, structured around the delivery of pre-determined outputs like the numbers of teachers trained or clinics constructed, involve relatively straightforward contracts. The problem, of course, is that the contract can be fulfilled without having much development impact – for example teachers can attend training courses without educational outcomes improving. Adaptive programming is a response to that problem. Rather than implementing a fixed course of action, adaptive programming entails experimenting with different alternatives and changing the course of action over time in response to the results that have been observed. If the goal is to improve education outcomes, adaptive programming might start with discovering what problems local actors are interested in solving and, if teaching ability is identified, experimenting with different ways of raising standards.

We define an adaptive contract as one designed to encourage experimentation, learning and adaption. Although this broad definition includes fixed 'payment-by-results' (PbR) contracts, which implicitly encourage experimentation, this paper also considers contracts which are more explicitly adaptive in that the terms of the contract are subject to revision following the results of experimentation.

Despite widespread support in principle, in practice, adaptive programming is more honoured in the breach than in observance. One potential explanation for the slow uptake of adaptation by donors could be that they are reluctant to write such contracts when large sums of money are involved, because they prefer the comfort of stating in advance how much money will be spent on what, and

1. Andrews (2013) recounts the failures of development intervention in the public finance and governance sector, and the reasons for persistent failure; Wild et al. (2015) focus on the social sectors and argue that adaptive programming holds the key to better results. USAID (2016) provides an overview of the evidence that adaptive approaches are more effective than traditional approaches.

2. The term 'problem-drive iterative adaptation' is introduced in Andrews et al. (2013). A good introduction to the principles of adaptive programming can be found at doingdevelopmentdifferently.com. For analysis of adaptive approaches used by the UK's Department for International Development (DFID), see Derbyshire and Donovan (2016); Pritchett et al. (2013) propose methods that aid agencies can use to try out alternatives and then adapt projects based on the results.

3. Contract theory also addresses the question of when things should be done in-house or contracted out.

where it is relatively easy to say whether the contract has been fulfilled (despite notionally successful fulfilment rarely coinciding with development success). But even if donors are willing and able to write contracts that encourage adaptation, this may not yield the hoped-for results because underlying contracting problems may still be present.

This paper introduces development practitioners to recent results from *contract theory,* which explicitly discuss how to incentivise experimentation, to encourage reflection about the relevance of its insights for development practice, and to draw attention to the sometimes-perverse incentives created by adaptive contracts. Whilst some of the lessons of contract theory might be of immediate use to development practitioners, we hope that this paper will also prompt practitioners to identify areas where contract theory fails to adequately address their problems, and point to directions for future research.

The paper begins with an introduction to contracts and contract problems, before reprising the basics of contract theory to explain why and when contracts matter. It then looks at the status quo of development practice from a contracting perspective, before presenting some principles for adaptive contracting, drawn from the contract theory literature. The final sections provide some discussion and conclusions.

Box 2: The agent-principal problem

In contract theory the parties to a contract are known as *principals* – those who want something done – and *agents* – those who are contracted to do it.

Contract theorists usually approach principal-agent problems from the principal's perspective, where the problem is how to induce the right behaviour by the agent. This paper takes the same approach. From the point of view of development practitioners, however, the problem might feel more like one of the agent having to induce the right behaviour from the principal – what could be called the 'agent-principal problem'. For example, donors may pay lip service to adaptive programming, but in practice do not allow project spending to vary from a pre-determined budget.

In this paper we consider a two-tier structure where the donor or aid agency is the principal and an NGO or private firm is the implementation agent. In reality, there is often a three-tier structure, where the donors can also be seen as the agents of politicians and the general public (who are the ultimate principals). Seen in this light, it is obvious that donors also have private information which they do not want to truthfully report to their own principals; they face incentives to look successful and provide easily understood stories about what they do, otherwise their funding may cease (Tirole, 1986). This can make adaptive programming difficult, because it might be misunderstood as failure and an admission that a project is ineffective. From the perspective of the development practitioners – contractors and NGOs, who we treat as agents in this paper – they may thus be faced with the problem of working with donors who face incentives not to genuinely experiment. Worse, the principals (the public and donors) might not even really want the agents to reveal the truth. Niehaus (2014) suggests that even altruistic politicians and members of the public do not actually want to learn about development, because doing so undermines the 'warm glow' feeling of doing good. The reasons why donors may be resistant to adaptive programming are beyond the scope of the present paper.

2. Contract theory

Contracts are a pervasive feature of economic life, both within firms (employment contracts) and between customers and suppliers. Contract theory – the pioneers of which have won a clutch of Nobel prizes – is part of the trend in economics away from studying simple exchange under perfect information, towards understanding the market failures that arise in more realistic settings.[4] Contract theory takes place in an environment of imperfect information. So, for example, you can hire somebody to work for you, but you often cannot observe what they actually do, nor even know what they ought to be doing.

Contracting matters when the respective parties have different objectives, when some information about which actions should be taken is unknown, or when what happens during implementation cannot be fully observed. These features of the economic environment require clever design of contracts to limit inefficiencies that either harm performance, raise costs, or both.

Much academic work on contract theory is concerned with exploring all the ways in which contracting problems can arise, and tracing their economic consequences. Theorists have devised contracts that might overcome some of these problems, but contracting problems can rarely be solved entirely – it is usually a matter of doing as well as possible in the circumstances.

The solutions found by contract theorists can be hard to translate into practice – for example, an economist might calculate the optional duration of an experiment as a function of a particular model's parameters. This is of limited help to a donor trying to decide whether to fix contract terms for six-months or three years. However, a model might provide some insight into when a longer fixed term is more likely to be appropriate. This paper aims to draw out 'high level' insights such as this from abstract theory.

It would be a mistake to ignore highly-stylised models because they are 'unrealistic'. Contract theory could have told us that giving bankers incentives based on short-run profits could lead to disaster, or that basing teachers' pay on exam results might produce students who are good at passing exams but not well-educated in the wider sense that society desires. The possibility of perverse incentives under poorly designed adaptive contracts should not be dismissed because the models motivating that concern are abstract.

Box 3: Contract theory and development

In the wider context of development, contract theory is most often applied to the relationship between donors and recipient governments, and between bilateral donors and multilaterals. Explicit conditionality in aid relationships (i.e. tying aid disbursements to recipient actions) is now less common, but the donor-recipient relationship can still be seen as an implicit contract, in that there remain some actions that the donor wants the recipients to undertake, and some actions that would cause the donor to suspend aid, even if these are not explicitly specified. Many of the same contracting problems apply: recipients are likely to have objectives that diverge from those of the donor; recipients have access to private information and can take actions that are hidden from the donor (see Williamson, 2002; and Hegen, 2015).

2.1 The status-quo: conventional contracts for development

A conventional development contract is for performing specific tasks decided in advance.[5] One of the movements behind the push for adaptive programming is called 'Doing Development Differently'. If one thing characterises 'doing development conventionally', it is deciding what to do in advance (perhaps with little input from the intended beneficiaries), and then ploughing on regardless.

4. Contract theory emerged from early advances in the economics of information, associated with economists such as Ken Arrow (who won the Nobel in 1972). The 2016 Prize in Economic Sciences in Memory of Alfred Nobel was won by Oliver Hart and Bengt Holmström, for creating theoretical tools to understand the performance of incomplete contracts in more complicated realistic settings. Other relevant Nobel winners include James Mirrlees and William Vickrey, awarded the prize in 1996 for work on incentives under asymmetric information, George Akerlof, Michael Spence and Joseph Stiglitz in 2001, again for work on asymmetric information, Leonid Hurwicz, Eric Maskin and Roger Myerson who shared the prize in 2007, Oliver Williamson who shared the prize in 2009, and Jean Tirole, who won in 2014.

5. The distinction between contracting for pre-specified actions and adaptive programming is not actually very sharp, because adaptation can potentially be achieved via a series of short-term inflexible contracts, where each new contract changes in response to what was learnt during the previous one.

Worse, traditional development contracts are typically based on delivering inputs (such as materials or labour) or outputs (such as building schools or raising attendance) as opposed to the intended outcome (such as giving children the education they need). From the perspective of contract theory, this has long been recognised as a bad idea. The management scholar Steven Kerr famously called it 'The folly of rewarding A, while hoping for B' (Kerr, 1975). This observation does not necessarily imply basing contracts on B is a good idea either, however, which might not be feasible, particularly when outcomes are difficult to measure. Rather, it just means that basing payments on something that is not really what you want to achieve, is unwise.

Forty years later, the idea of basing aid contracts on B – an idea commonly known as Payment by Results (PbR) is still thought of as a controversial innovation in development circles. Contract theory provides plenty of reasons why it may only be appropriate to make payments heavily dependent on results in rare circumstances, and these should be borne in mind in later sections of this paper, which propose forms of result-based contracts.

2.2 When contracts matter

Contracts are not always important. If both parties want all the same things, the principal can just leave the agent to get on with things. In most circumstances, however, the principal would not get what they want without a contract in place. In these cases, a contract is both a problem to be solved (what is the right contract?) and the solution to the underlying problems that would otherwise lead things astray. The misalignment of incentives, along with the inability to completely monitor the actions and local knowledge of agents, means that agents need *inducement* if they are to perform tasks exactly as a fully-informed principal would want them to do. Contracting will be particularly important in settings with the following characteristics:

When the principal and agent have different objectives

To caricature, agents want to get paid as much and do as little work as possible, while the principal wants to pay as little and get as much output as possible. More realistically, even hard-working agents may have objectives that do not align perfectly with those of the principal – for example a consultant may want to give work to local contacts, to maintain good relations for the sake of future projects for other clients, when it would be in the principal's interest not to. Thus the problem arises of writing a contract that gives agents the right incentives, from the principal's perspective.

In the context of development, it might be reasonable to think that principals and agents have closely aligned objectives more often than not, in comparison to purely commercial settings. But shared objectives do not imply completely aligned incentives. Both parties may happily sign up to having the same altruistic development objectives, but as realisation dawns that under a given contract certain tasks are easier or cheaper, or certain outcomes may affect their international reputation, these factors may permeate the agent's behaviour to a greater degree than they are willing to admit.[6] When incentives are not a problem, the optimal contract would simply be to give the agent a sum of money and allow them to use it at their discretion. The fact that contracts are rarely of this form, reveals that perfect alignment of objectives between principal and agent is also rare.

When principal and agent have different information

Employees working in the field know things that their paymasters back at headquarters do not. One way of looking at contract design is as a mechanism to elicit truthful information from agents about what they have done and about what they know. There is an asymmetry here: agents will only readily reveal private information when it helps them, not harms them. As a general rule, contract theory suggests that agents will be able to earn 'information rents'. That is, the optimal contract will wind up rewarding agents for revealing their private information. In the absence of such rewards, agents can be expected to use their informational advantage to the detriment of the principal's objectives. These information rents are particularly salient in adaptive or exploratory contexts, where the results of the experiment are often known privately by the agent.

When outcomes are not under the control of the agent

A basic precept of contract theory is that any contract must satisfy the agent's 'participation constraint', meaning that if the contract terms are not sufficiently favourable, the agent will not sign. If an agent is contracted to deliver an outcome that they only have partial control over, they are likely to want sufficient compensation to cover the risk of not getting paid for reasons outside their control. Hence 'high powered incentive contracts' – where payments vary greatly with performance – can be more expensive (NAO, 2013). As a result, rather than tie payments tightly to results, it may be more efficient to offer agents a flat fee when they have little control over outcomes, even if this reduces their incentives to exert effort.

6. Huysentruyt (2011) analyses bidding and contract award patterns at DFID across NGO and private for-profit contractors. She finds that NGOs have an advantage in bids when quality investments (that are not measurable for contracting purposes) are more important, and that contract renegotiations is more costly with private contractors.

Box 4: Payment by Results (PbR)

A simple but potentially seismic shift in contracts for development is to base payments on the ultimate intended outcome, rather than on the inputs that you hope will deliver them. For example, payments could be based on the proportion of children that pass basic literacy and numeracy tests, rather than on building schools and training teachers. This idea could be applied to contracts with implementation partners or form the basis of disbursements of aid to recipient governments. The link to adaptive programming is that the recipient government or subcontractor is not bound by any predetermined method of delivery, and is free to experiment, learn and adapt their approach to find the best way of delivering the outcome. Pioneers of this approach include the Global Partnership for Output Based Aid, hosted by the World Bank, which also introduced a Program for Results instrument in 2002, and DFID, which has conducted a number of PbR trials. The Center for Globalisation and Development (CGD) was heavily involved in developing these concepts under the banner 'Cash on Delivery Aid'.

Results-based aid does not overcome all contracting problems and is subject to many of the same criticisms that have been directed towards performance-based management more broadly.[7] These include the possibility that results-based aid will distort activity towards outcomes that are more easily measured at the expense of those which are harder to observe, and may be susceptible to gaming. Depending on design, PbR may also transfer risks to implementing agencies that they may not be willing to bear; development impact bonds are a potential response to this problem, which transfer risks to investors willing to bear them (CGD, 2013). PbR will only be appropriate in certain circumstances and contract theory can help identify these. For recent analysis of experiences with PbR, see Perakis and Savedoff (2015).[8]

A PbR contract of the form 'get paid X for every Y' will not be complete and optimal because there will be some states of the world where the agent would not do what the principal most wants. Suppose a contract pays $10 for every girl who passes a literacy exam – if the average cost of getting a girl through the exam is only revealed to the agent during implementation (and never to the principal), it could turn out that the cost exceeds $10 so a self-interested agent would make zero effort to minimise their losses. In this event, perhaps the optimal contract would have paid $12. Or perhaps a drought hits, and the principal would then prefer the agent to stop trying to promote female education and instead switch to providing free meals for all school children. This would require contract renegotiation. In the example of unknown costs, a complete and optimal contract would be impossible because the contract can never be based on information known only to the agent. In the example of a drought, one could say that the contract is complete but not optimal (because the contract applies in all states of the world, but is not optimal when a drought hits), or is incomplete if the contract attempts to condition payments on states of the world, but omits to mention droughts.

PbR can be seen as a static adaptive contract that only leaves one dimension of the contract open-ended: it specifies a goal, not how to achieve it. PbR is only appropriate when there is a readily observable outcome, suitable for basing contracted payments upon, however. Donors may want to contract for experimentation in other cases, perhaps because there is no suitable result to condition payments upon, or because experimentation is necessary to establish what results to target.

Rather than add to the substantial body of work on PbR, this paper takes a more general approach, relevant in cases where experimentation and adaptation is desirable but the conditions are not right for PbR. Many of the problems addressed in this paper arise when agents do not know in advance how much they will be paid for doing what. Some of the papers surveyed later in this paper resemble dynamic PbR contracts that vary payments over time.

It may be tempting to turn experimentation and learning into a 'result' in itself for contacting purposes: the agent is paid for delivering learning. But the reasons for being sceptical of PbR in general would also apply here. Learning is not an end in itself, and is also hard to verify and quantify – agents are always going to be able to claim that they have learnt something.

7. See Van Theil and Leeuw (2002) for a critique of performance-based management in the public sector more generally, and Shutt (2016) for a critique of 'the results agenda' in development.

8. See Clist (2016), Clist and Verschoor (2014) and Clist and Dercon (2014).

3. The building blocks of contract theory

In this section we introduce some fundamental concepts and results of contract theory, as well as some context, as a precursor to section four which covers contracts that are explicitly concerned with experimentation.

3.1 Complete contracts are always best, rarely possible

A *complete, formal and optimal* contract specifies the responsibilities and rewards of the agent for every possible contingency, with no grey areas, with the contract written such that the agent performs precisely the principal's most preferred action after each contingency (Tirole, 1999). A well-known result from contract theory is that if all information relevant to a contract is verifiable and known in advance, then the optimal contract simply states what is to be done by both parties in every circumstance, and pays just enough to induce the agent to perform those actions. In practice, however, most contracts are incomplete (or complete but not optimal). In reality, ideal contracts are usually impossible either because some of the information needed is unverifiable, or because some of the desired actions cannot be known in advance.

Verifiability matters because a formal contract is ultimately enforced by the courts, which must be able to verify the information needed to determine whether the contract has been breached. Information may be observable but not verifiable, so for example the principal may recognise low-quality work, but be unable to demonstrate breach of contract in court.

The inability to state in advance exactly what should be done can arise for two reasons. First, the number of potential contingencies in the world can be too large to enumerate. If a third party begins development work on a similar issue in the same region, how should our actions alter? If the local leader is replaced by one who is hostile to the contracted project, what should be done? Some, but not all, of these contingencies will appear in any given contract, but there will almost always be some circumstance where the parties enter the 'incomplete part of the contract' and must still decide what to do. Second, some factors which the principal would like to make a condition of the contract are not observable by them. For instance, an aid agency purchasing equipment would like to get the best price, but cannot observe the minimum price each vendor will accept.

A complete and optimal adaptive contract would need to specify how payments would change in response to *every* relevant piece of knowledge that could be learned.[9] The goal of learning and adapting means the complete contract is particularly challenging to write in advance: the principal will struggle to anticipate all possible outcomes of experimentation and hence, what should be done and at what price.

3.2 Hidden information and hidden actions limit the use of complete contracts

There are two types of knowledge that agents possess, and that principals would like to contract on if they could, namely, hidden information and hidden actions.

In the case of hidden information, only the agent may know the true costs of different methods for achieving a specific goal, for example. If the principal knew these costs, they would select the most efficient method and pay the agent the corresponding sum. Not knowing these costs, the principal may offer a fixed-fee contract. In turn, the agent will not choose the most efficient method, in terms of output per dollar, if the costs are higher than their fee.

Hidden information arises quite naturally in the context of adaptive contracting. Consider a simple experiment where a local agent is tasked with delivering AIDS drugs to village health clinics, and is asked to experiment with methods for delivering the drugs more efficiently. Suppose delivery by car maintains relationships between the agent and a delivery service they use in other contexts. Delivery by bicycle requires forming a new relationship with cyclists, and provides no side benefit for the local agent. If the result of the experiment (i.e. the true efficiency of the two alternatives) is unobservable by the principal, the agent has an incentive to misreport the result to favour delivery by car. A more serious concern would be

9. A sensible approach may be to write contracts that specify procedures for decision-making in the light of new information. But that would still not be a complete contract in the sense used by contract theorists, because the procedures would inevitably omit something.

an experiment where the results suggest that the agent themselves should be replaced by a different contractor: surely even a very altruistic local partner will hesitate to truthfully reveal information that makes them redundant. To induce truthfulness, an additional incentive for the agent may be needed. The bigger point is that astute principals understand that contract-relevant information is an input that needs to be procured like any other, and subsequently attempts to procure that information cheaply.

Hidden action is when what one party *does* is not totally visible to the other. For example, a salesperson may not exert themselves, but then claim that they simply had bad luck. Hidden action is not a problem for incentive-heavy contracts if there is a perfect relationship between hidden action and visible output. If a salesperson *certainly* sells one widget for every hour of hard work, a contract based on widget sales can induce the optimal level of hard work, which can be perfectly inferred from the total number of sales. A problem occurs when the link between the agent's action and the visible output is subject to some randomness. Many agents, for example small local partners to a development agency, are hesitant to take on risk which might bankrupt them, and hence will want compensation for accepting contracts which specify payment in terms of partially-random output. There is therefore a trade-off between the benefit of providing strong incentives for effort, and the cost of compensating for risk.

A well-known case of hidden action occurs when production depends on the actions of teams. A team of agents is asked to produce some output, but it is unknown to the principal who exactly did what (Holstrom, 1982). Every individual realises that if they shirk their responsibilities, it is hard for them to be caught: how could the principal know who didn't pull their weight? Hence, agents may attempt to free ride on the efforts of others in their team. Contract theory suggests using smaller or less complex teams in situations where individual effort is difficult to monitor, as the benefits from mitigating the incentive to free ride outweigh the risk of using teams smaller than the theoretical optimum.

An even more worrying type of hidden action in a development context is *multitasking* (Holstrom and Milgrom, 1991). The multitasking problem involves situations where agents perform more than one task, and where effort is easier to observe for some tasks than for others. An example is of a teacher who can spend time teaching mathematics or mentoring students. The teacher's efforts to teach maths can be inferred indirectly (with some error) by testing students, but student mentoring cannot be measured at all. A bonus for high test results will cause the teacher to spend relatively more time teaching maths and less time mentoring, which in actual fact might not be what the education ministry wants. The general lesson is that if many things matter but only a few can be observed and rewarded in a contract, then doing so can distort behaviour in an unhelpful way. This problem is likely to resonate with critics of PbR, who question whether these contracts can sometimes divert effort from important but less-observable activities.

The above problems with results-based payments may mean that some of the adaptive contracts described later in this paper, which make use of results-based payments, may only be suitable in particular circumstances. It may be possible to attach incentives to outcomes that are highly correlated with other desired yet hard-to-monitor efforts, in which case results-based payments will induce the desired behaviour. If that is impossible, it may sometimes be possible to divide responsibilities so that one agent can be given results-based incentives to produce the observable output, and the other given a fixed salary to perform the hard-to-monitor task. If unobservable efforts matter, it may be necessary to restrict use of strong results-based incentives altogether.

In summary, the problems of hidden information and hidden action limit the usefulness of results-based contracts. Solutions can involve costly efforts to acquire information or better contract design to limit the scope for agents to act against the interests of the principal.

These lessons become particularly salient in the context of adaptive contracts. Hidden information, in terms of what is learned from an experiment, and hidden action, in terms of whether the agent is genuinely experimenting in the right way, both imply that strong incentives for running experiments may not generate the desired effect.

3.3 Relational contracts as an alternative to formal contracts

Formal contracts are enforced by courts, but for many development projects it would be difficult for a court to determine whether a contract has been honoured or breached. An alternative is to use *relational contracts* (see Baker et al., 2002; Levin, 2003; and Gibbons and Henderson, 2011 amongst many others).

A relational contract, for the purpose of this paper, is simply a contract where the agent's cooperation is induced by the 'shadow of the future' rather than by a formal, court-enforceable contract. A principal lays out a desired set of actions for the agent, with the understanding that performance will be rewarded in the future in some way. The agent desires to cooperate in order to garner subsequent business from the principal, and the principal does not renege on promises because it wants the agent's cooperation in the future. A very common form of relational contract is a promotion decision within a firm: rarely are there explicit criteria by which an employee is promoted, but firms have an incentive to reward effort (for example, if a firm consistently promotes politically connected individuals rather than the best performers, employees may begin to shirk their responsibilities). Relational contracts in the private sector govern some very substantive relationships: Coca-Cola and McDonald's have

famously relied primarily on a handshake agreement to act in each other's interests (*New York Times*, 2014).

In a development context, contractors may often be more motivated by the desire to get hired again, than by the threat of legal censure, and so principals need to think about what makes a contract operate successfully on that basis, rather than how to demonstrate breach of contract in court. Although they may not see things in these terms, practitioners who believe that their formal contract has little relevance to their work are probably working 'relationally' on the basis of mutual understanding and credible expectations about future behaviour from all parties. Of course, many contexts will be something of a hybrid, with some sort of formal contract that sets out requirements in broad terms, perhaps a formal auditing arrangement, complemented by a less formal holistic evaluation that depends on non-verifiable aspects of performance.[10]

Two major factors – *clarity* and *credibility* – affect the performance of relational contracts. Clarity means that both parties need to understand what is expected of them. If the principal is going to evaluate an agent subjectively, it is critical that both sides understand what will cause the other party to stop cooperating in the future. Formal contracts require only that agents know what they are supposed to do, and the consequences are clear. Relational contracts further require that both parties understand what will happen if the desired action is not taken or the desired outcome is not achieved. For example, if an anti-corruption campaign does not deliver as strong results as had been seen in the past, but the contractor can credibly show that they have spent as much time as was believed necessary on the campaign, ought the contractor be retained next year? Ought a bonus be paid?

For adaptive programming, clarity requires that principal and agent are on the same page about what experimentation implies, what actions by the agent are desired and will be viewed favourably by the principal, and what will happen if experimentation suggests the agent's specific skills are no longer needed, etc. It may even be desirable to hire a relationship broker as part of an investment in mutual understanding (Pellini and Nixon, 2016).

And even clarity is not always enough. Since relational contracts are those which cannot be enforced in a court, either side may renege on the contract if the short-term benefit from doing so exceeds the long-term harm from ending a fruitful relationship. A contractor who is on the verge of bankruptcy may cut costs or run experiments in a shoddy way, since the harm of damaging their relationship with the principal is not as bad as the benefit of avoiding bankruptcy. Promises made under relational contracts may also not be credible when agents work with multiple principals – perhaps when things go badly 'head office' may take charge of proceedings and overrule promises to tolerate failure made by field offices. Agents may also worry about how other potential clients will perceive failure on a project, even if the principal on that project has given them incentives to be honest when things go wrong.

Without the comfort of formal contacts, trust is a big issue. When General Motors set up the joint venture NUMMI with Toyota in the 1980s in an attempt to implement Toyota's 'Lean Manufacturing', it was fantastically productive, with the exact same workers and machines that had been producing substandard GM cars now producing highly-reliable, cost-effective vehicles (Helper and Henders, 2014). NUMMI was heavily reliant on relational contracts, with a famous example being the 'andon cord'. Any line worker who noticed a flaw could pull a cord, stopping the entire assembly line at a cost of roughly $10,000 per minute to the firm. Managers at NUMMI wanted the cord pulled if there were actual flaws, but did not want the cord pulled by workers who simply wanted a break. Clarity was required from both managers and line workers about when exactly the cord should be pulled, what the punishment would be in terms of career progression for pulling the cord when it ought not have been pulled, and what the reward would be for spotting important flaws and pulling the cord. Workers needed to trust that the line would not be unduly sped up, or their jobs replaced by machines, as they gave the firm information about how to run the line more effectively.

Though these mutual expectations were clearly and credibly laid out at the NUMMI plant, GM had a great deal of difficulty in expanding this success to other plants outside the joint venture. Workers who had spent their career being told never to stop the line, and who had a combative relationship with management, needed to be convinced that there would be no retribution if they pulled the andon cord by mistake after wrongly suspecting a defect. Managers needed to believe that NUMMI-style processes would be retained into the future, as they would not want to link their career to a 'management fad' that may later be seen as a failure. Within NUMMI, managers and workers were able to establish a relationship based on trust, but in established GM plants the desired behaviour and promised rewards were not regarded as credible.

10. Banerjee and Dulflo (2000) is a very interesting analysis of relational contracting in the Indian software industry, which might resemble a development setting in so far as 'at the outset … neither the client or firm entirely understands what they are setting out to build' and that in the event of a dispute 'the description of the project in the contract is so complex that courts cannot be relied upon to adjudicate correctly between competing claims'. They propose a model (and provide evidence for it) in which firms with good reputations will pay for overruns if it is their fault, but when both parties are not sure of each other's reputation they take refuge in either fixed cost or cost-plus contracts that protect one side from cost overruns.

Credibility can also be challenging for relational contracts that involve experimentation. Halac (2012) studied this formally, showing that if parties to a relationship face uncertainty about the value each side places on continuing the relationship, then a period of past cooperation (which in a sense 'confirms' the relationship value) is necessary to generate cooperation in the future. When agents experiment, they recalculate the value of the relationship on the basis of the results of the experiment. This recalculation makes it difficult for the agent to know that the principal still values their relationship in a way that will lead to future cooperation. This line of reasoning means that under some adaptive contracts agents may place more weight on short-run rewards, having less faith in the future.

Despite relational contracts presenting some challenges, the potential for integrating hard-to-anticipate and hard-to-observe factors into a working relationship means that they are often the most ppropriate way of operating development partnerships. Sometimes formal procurement procedures can work against establishing partnerships between principal and agent based on repeated interaction – in some jurisdictions regulations may specify that contract awards must be made on the merit of the application without advantage gained by historical performance, although in practice incumbents often have many advantages.

4. Principles for adaptive contracting

The preceding sections have considered broad general principles in contract design. Here we survey contract theory that attempts to solve problems in situations where the principal would like the agent to experiment and learn. The papers surveyed are the closest that contract theory comes to explicitly studying adaptive programming, although the link between the applied development problem and the original abstract theory can sometimes be hard to see.

If there is no asymmetric information – no hidden information or potential for hidden action – then the optimal contract is simple: i) write a contract to perform an experiment and then write another contract making use of the results; or ii) write a long contract that covers any possible information that might be revealed later on. The interesting cases are where information revelation over time *interacts* with hidden information or actions of the agent (Freixas et al., 1985; Laffont and Tirole, 1988).

4.1. The ratchet effect makes experimentation difficult

Actions taken by the agent early on can reveal information that will be useful further down the line, hence the attraction of contracting arrangements that may be revised in the light of new information. From the agent's perspective, a contract that is subject to renegotiation is worrying, however. Experimentation early on could reveal information that the agent would prefer the principal does not know, such as whether the agent is worth hiring at all, and therefore the agent may hesitate to experiment or report the results truthfully. This pattern is known as the 'ratchet effect' due to the agent's worries that expectations will be 'ratcheted up' if the principal learns something about the agent early on and has the ability to rewrite parts of the contract to take advantage of that information (Weitzman, 1980). Consider a street sweeper who is paid a bonus for finishing the street sweeping quickly. If the sweeper works for only one day, the bonus may effectively incentivise him to finish the job quickly. However, the sweeper also realises that if they finish the job very quickly today, then tomorrow's contract may simply say 'do the sweeping as quick as you did on the first day or you are fired'. Understanding this, the sweeper will forgo the short-run bonus for the sake of reducing long-term contract requirements.

Note that the ratchet effect happens because the principal has not committed to a long-term contract – in contrast to a PbR contract, which we consider as falling under the broader concept of contracts for adaptive programming because it encourages learning and experimentation, although the contract itself is fixed. If the agent faces a short-term PbR contract knowing that the terms of subsequent contracts will be revised in the light of experience, then the ratchet effect will apply.[11]

There is an important distinction to make between a fixed contract for experimentation, and an adaptive contract. A fixed contract asks for experimentation but also specifies what will happen afterwards and is not subject to revision; an adaptive contract seeks to take advantage of what is learned and leaves future contracting arrangements to be determined. It is this open-ended nature that gives rise to the 'ratchet effect'.[12] An agent on a fixed-term contract is effectively operating under an adaptive contract if they are hoping for that contract to be renewed, and there is scope for renegotiation on renewal.

If the information revealed by experimentation is not relevant to an agent's income or costs, then adaptive contracts can look like a particular form of complete contract: simply pay the agent for experimenting initially, and then pay the agent for implementing the results of that experiment later on. In development contexts it may often be the case that experimentation and implementation

11. In some respect the division between fixed and adaptive contracts is artificial – the future is never entirely written and even fixed contracts are only fixed for a finite time, leaving the agents to worry about the future. So it's a sliding scale: in some cases, potential future revisions are a pressing concern, in others a distant prospect.

12. Of course a badly written fixed contract for experimentation could still be problematic if, for example, it says 'admit failure and we will stop paying you' then agents will still face incentives to conceal information. The point is that that a fixed contract can be written so that agents have no reason not to be truthful, an adaptive contract by its nature will always give agents something to worry about, because the contract is yet to be written.

cannot easily be separated, and that experimentation may yield information that could limit the scale or scope of a project, to the detriment of the agent. This suggests that an important challenge when writing contracts for adaptive programming is giving agents an incentive to truthfully reveal negative information (at an acceptable cost).

4.2. Spreading experimentation and learning across projects is not easy

Why might a development agency need to learn over time on a given project anyway? An alternative type of adaptive contract might simply use results learned from experimentation in one setting to slowly and steadily improve the operation of programmes in other settings. In this case, the ratchet effect is less of a worry since information learned will only affect other agents in other settings.

This idea of learning and diffusing best practices seems sensible, but a large body of literature on so-called 'persistent productivity differences' shows that some firms have costs that are two or three times higher than other firms in the exact same industry (see Melitz, 2003; and Syverson, 2004; 2011). Why don't these underperforming firms simply replicate what the efficient firms are doing? Or, analogously in the development context, why don't development agencies concentrate on learning best practice and stop worrying about adaptive experimentation on a contract-by-contract basis?

One answer comes from the 'rugged landscape' model (Levinthal, 1997). The essential idea is that rather than performance as a function of action being smooth with a single optimum, the 'performance function' is filled with many minor peaks and valleys. Locally, taking actions more similar to a known high-performing action in general may not improve outcomes. Intuitively, the optimal way to structure a programme may depend on a dozen complementary tasks, including which workers are hired, the political environment of the host country and so on. Getting eleven out of twelve components right but making a mistake on the twelfth may in a rugged landscape result in a worse outcome than a radically different programme. Hence, lessons learned even in similar settings may be difficult to apply and a process of minor trial and error is required. The importance of local context is a familiar idea in development – in fact it is something of a cliché that every performance evaluation report blames disappointments on a lack of understanding of local context.

Even abstracting away from the incentive concerns in the previous two sections, if two programmes involve very different actions but have very similar performance, in a rugged landscape it will be not be worthwhile to try to imitate the slightly better performing programme, since the maximal potential gain is small but the maximal potential harm of slight mistakes in imitation is very large (Callender and Matouschek, 2014).

These ideas suggest that the strategy of dividing responsibilities between agents in order to mitigate against the ratchet effect might not be viable, because learning is not transferable. Another problem in the context of development is that the pool of agents is quite small, so the idea that they won't worry about the implications for future agents is less plausible.

4.3. Delegation can mitigate the ratchet effect

If spreading experimentation across different projects in different places is not possible, or if experimentation requires learning-by-doing (Arrow, 1971) and so cannot be separated from implementation, then an adaptive contract will have to induce experimentation by agents who will also then be asked to perform whatever actions are chosen as a result. Contracts can still be structured so that agents experiment and take decisions on what to do *without* that decision negatively affecting their own payoff, however.

One class of contracts with that property involves *delegation*. Under delegation the principal fully commits to give decision-making authority over how to use certain resources to the agent under a set of predetermined conditions. Payments to the agent are fixed, but the agent has discretion over tasks performed. A simple version is 'static delegation' (Holmstrom, 1982) where the principal must decide whether to give power over a one-time decision to an agent who possesses private information of interest to the principal, but objectives that do not fit perfectly with the principal. The trade-off here is that under delegation the agent will take advantage of their private information and maximise their own payoff, but without delegation there is no way to exploit the agent's private information. The best contract in this simple setting involves the principal giving the agent a list of permissible actions from which to freely choose, with a larger list in cases where the principal and agent have similar objectives, and a smaller list when preferences differ wildly.

In a dynamic setting, such as the case where experimentation is potentially informative, useful contracts can involve *partial commitment via delegation* (Bennedsen and Schultz, 2005). Suppose there is some status-quo method of delivery that is acceptable to the principal but that could be improved on via experimentation. Assume both the principal and the agent are uncertain about the gains from experimentation, and that the potential experiments are too complex to be contracted upon. This leaves three options: contract for delivery of the status quo, delegate decision-making authority completely to the agent, or delegate to the agent for an initial period and then renegotiate in later periods once the initial outcomes are observed. A contract that specifies the status-quo policy cannot benefit from learning. A contract that pays the

agent a fixed fee and delegates decision-making authority forever can potentially be abused: the agent will choose to experiment on policies with the highest private upside for the agent. A better contract gives the agent complete freedom over what policy to pursue early on, at the cost of potentially opportunistic behaviour by the agent (which could actually be worse than the status quo from the principal's perspective), and will continue granting that freedom if the agent delivers improvements, but will revert to the status quo otherwise. Under this contract the agent still will not reveal information that harms them, but will at least be willing to experiment in the knowledge that the results will not be used against them. In effect the principal is hoping the agent will discover something mutually beneficial (or at least which achieves the donor's development goals without being worse for the agent). The price the principal pays is a period of time during which the agent may be working against their interests.

Bennedsen and Schultz (2005) give the example of the Copenhagen bus system reforms in the 1990s. Contractors were initially given very limited instructions on where to locate stops, what routes to run, and so on. Although the bus contractors were rightly aware of the ratchet effect dangers of their learning being used against them in the future under an adaptive contract, the commitment to use either the experimental bus system or a status quo in later periods meant that the contractor's worse case from experimenting was simple reversion to the status quo.

A model by Guo (2016) shows that delegation for a fixed period to an experimenter is, in some settings, the optimal way to induce experimentation. Consider a setting where the only choice is whether to continue experimenting with a new policy or not. Experiments come in two varieties: good or bad. If good, the experiment will deliver success with some probability in any given unit of time, but if bad, never will. Let delegation mean that the principal supplies resources (the agent just invests their time) and leaves the decision as to how to use them – what experiment to run – to the agent. As experimentation proceeds, the agent and principal's belief about the probability the experiment is 'good' falls unless success is observed. Since the agent doesn't pay for the experiment, they would like to continue running the experiment for longer than the principal would want. So in this model the ratchet effect is not that the agent is worried about demands being ratcheted up, but that they are worried about revealing information that would result in the experiment getting shut down sooner than they would like.

Assume that the principal knows nothing about the initial potential success of the policy, but that the agent correctly knows the initial *probability* that the policy will be successful. For instance, a local contractor may correctly believe that an experimental health delivery system will improve outcomes with 40% probability, and proposes to an aid agency to experiment to learn whether the modified health system actually does improve outcomes. The aid agency, being less informed about the project and the context which will determine its success, cannot put a probability on success, but can observe outcomes when an experiment is run. This is a problem of hidden information, not hidden action, since we will assume the principal can observe whether the agent is actually experimenting at any time.

As the experiment proceeds without generating success, both principal and agent become more pessimistic in their belief that the experiment is worthwhile. The contracting problem is how to allow for experimentation without continuing to pay for experiments with a low probability of success.

Guo (2016) shows that the optimal policy in this setting is to delegate the decision as to whether to experiment to the agent, but with a firm deadline. If you see evidence of success before a given amount of time has passed, then let the agent continue to experiment. If not, terminate the experiment. Before the deadline, the agent should have complete power to end the experiment early. This policy ends up stopping very promising experiments too early (those where the agent's belief that the experiment will prove worthwhile is very high) but all other experiments too late, relative to the stopping times that the principal would choose if, like the agent, they knew the true probability that the project is good. A perennial complaint of implementing agents is that donors claim to want experimentation but only allow tiny variation from budget. Guo's model shows that it is always best either to delegate full experimentation on the part of the agent, or none at all.

The reason even optimal delegation contracts generate inefficiency is that an agent could always tell the principal, 'we know this is a high probability project, just give us more time!' but because the agent cares more about experimentation than the principal, the principal cannot trust such claims.[13] If the principal knew the true success probabilities they would vary the contract terms accordingly, but not knowing, they have to set the same terms for all projects. Stopping some genuinely high-probability projects too soon is the price that the principal pays for not funding other low-probability projects for too long.

In the more complex, though realistic, case where even 'bad' projects can show signs of success during

13. Under the opposite assumption, that the awwgent finds experimentation less rewarding than the principal, the agent can credibly communicate that an experiment should be continued, there is no risk of them recommending continuation for longer than the principal would want, if the principal knew the true probability.

experimentation, but where 'good' projects are more likely to succeed, then the optimal contract is delegation with a sliding deadline: every time news of success is observed, the deadline is extended. This type of sliding deadline is often seen in practice. Google employees have historically been allowed to spend some of their time on experimental prototype projects. If the projects work out, they are given more time to work; if they do not show early success, the worker is told to move on. Legislatures often perform policy experiments, but many of these have sunset provisions: if the policy shows a lot of success by the time of the sunset, it can be renewed, else it is cancelled.[14]

A practical question is how to calculate the length of the initial deadline. This will depend on the specifics of a given project, and cannot be given a general answer. In this particular model what determines the right duration is rather straightforward: the higher the payoff from success, relative to the costs of running the experiment, the longer the deadline will be. The interesting result is the simplicity of the optimal mechanism.

Delegation contracts, such as those just examined, are most useful when it is impossible to commit to the nature of the contract following experimentation, and where the incentives of the agent to shirk responsibilities are not terribly strong. The first condition is required because delegation is naturally a (partial) solution to the ratchet effect, and the second condition is required since misaligned incentives will guarantee that the agent chooses a project harmful to the principal in the initial period of the delegation contract.

4.4 Conditioning payments on successful experiments

Delegation contracts naturally minimise worries regarding the ratchet effect, since by their very nature they limit both the maximal harm to the principal (since delegation can be clawed back) and the ability of the principal to use information generated by the agent in order to harm them. Furthermore, they are the only class of contracts available in situations where bonuses based on outcomes are not feasible (or sensible). Otherwise, more complex contractual arrangements for experimentation may prove useful. For example, contracts can be written to pay bonuses conditional on how experiments proceed, or demand claw backs from failed experimentation, and so on.[15]

It is probably worth emphasising that in the papers that follow, where payments are conditional on results, the results are pre-defined. One of the main motivations behind adaptive programming is to permit the discovery of what problems need addressing most, and what results are feasible. But if, for example, targeted results are only defined after an initial period of experimentation, then the agent will have a clear incentive to manage down expectations during that period and make contracted results easier to achieve. Bhaskar (2014) looks at a case where both principal and agent must learn about the difficulty of a task in a two-period model, and explore solutions where results payments are used in the first period to counteract the incentive to pretend the task is harder than it really is. This solution does not translate well into cases where the task itself must be defined through experimentation, however.

Manso (2011) considers optimal experimentation contracts with bonuses conditional on results. The agent can perform a task using existing best practice, or can experiment, or can shirk altogether. Experimentation is costly for the agent and the returns are uncertain. Because experimentation imposes costs on the agent, they will only experiment if the contract (indirectly) compensates them for doing so. Unlike in Guo (2016), there is hidden action: the principal cannot be sure genuine experimentation has taken place, so cannot directly base compensation on experimentation having occurred. The contractual problem is to ensure the agent has the incentive to experiment when it is worthwhile, and never to shirk responsibility completely. This differs from the delegation contracts considered, where experimentation could be observed and the problem was to ensure that the agent did not waste funder money pursuing unpromising experiments.

Imagine that the status quo produces a policy success 50% of the time, and the policy always fails if the agent shirks. Assume the experiment is initially successful only 40% of the time, but if it succeeds this year, then next year it will succeed 80% of the time. Experiments, almost by definition, are programmes where efficacy is unclear initially, but potential is great if they show initial promise. If the experiment fails this year, we know it is best just to revert to the status quo. Whether a policy shows success

14. This simple sliding deadline policy is not as trivial as it may seem. Even with delegation, one might imagine that an optimal contract would adjust the resources available to the agent over time, rather than permit the agent full control over resources until experimentation is cut off altogether. These cut-off rules arise naturally in a mathematical class of problems called the 'multi-armed bandit' which arises frequently in experimentation contexts. Essentially, imagine a slot machine with a 'safe' arm that always spits out one coin when pulled, and various 'risky' arms whose payoff functions are not known. It can never be worthwhile to, at a given time, pull both the risky and safe arms: either the expected number of coins you get from a risky arm plus the value of what you learn by pulling it is more valuable than the one coin from the safe arm, or it isn't. Hence delegation of experimentation is either fully worthwhile or not worthwhile at all.

15. The ratchet effect is a problem when the contract is so open-ended that the principal may use learning against the interests of the agents. In this section we consider contracts that are fixed in the sense that the contract is based on results of experimentation, but not adaptive in the sense that the contract can be completely rewritten.

is observable to both principal and agent, but the lack of policy success can occur either because the policy is in fact a bad one or because the agent shirked. The principal would therefore like to experiment this year, and then continue with the new experimental policy next year if it shows initial success, or if not use the status-quo best practice.

If initial experiments are more likely to fail than existing best practice, then failure indicates that the agent either experimented or shirked. Future success is more likely if an experiment was performed, hence to encourage experimentation the agent should be compensated for failures early on and then for strings of successes later (Manso, 2011). Paying for early failure alone will simply cause the agent to shirk rather than experiment – the promise of future payment for success is necessary to induce experimentation.

Payment by results in both years, where an agent is simply compensated for every success, overpays for early successes and pushes agents towards pursuing the status-quo policy, which initially has a greater chance of success, rather than experimenting.[16] Instead, the agent can be encouraged to pursue experiments early on by the promise of future payment for results. This resembles a PbR contract that provides some unconditional upfront financing, perhaps by using a development impact bond structure, but then payments conditional on results in later periods.

The ability of the principal to credibly commit to future payments for success is crucial – and avoids the ratchet effect. As a general rule, being able to commit to future rewards allows the principal to incentivise action today. Without the principal tying its hands, the agent will not believe that the principal will generously reward results once the costly work of discovering how to achieve them has already been done.

In some cases, the challenge for an aid agency is incentivising agents to experiment when effort cannot be observed *and* when the ability of the agent to experiment cannot be observed. For instance, some contractors may be very efficient in the speed with which they evaluate and implement experiments. Again, assume both the principal and the agent learn over time whether an experimental project appears successful or not, but assume that the principal wants to ensure that genuine experiments are conducted, and also that they are done as cheaply as possible given the constraint that the skills of the contracting partner are unknown.[17]

An optimal contract needs to keep inducing effort from agents of different skill levels but also terminate experiments that look like they will never succeed. Halac et al. (2016) show how to achieve this optimally using a menu of bonus contracts, where the bonus is based on the time at which the success occurs. As time passes without success, the agent becomes more pessimistic about the probability of success, and so less inclined to exert effort in order to capture a bonus. This can be counteracted by a bonus that increases over time, however if bonuses in later periods are too high, the agent may postpone efforts in order to manipulate the timing of success and obtain a higher reward. In the Halac model there is no ratchet effect – an experiment is performed over time, but the principal cannot then use the results against the interests of the agent.

The relationship between optimal contract length and agent type is complicated: on the one hand, it makes sense to induce more effort for longer from high-skilled experimenters because they are more likely to succeed at any point. On the other hand, they also discover that experiments are 'bad' more quickly, therefore the efforts of the agent will fall off with good reason and so it can be more efficient to have them terminate sooner than a low-skilled. Which of these effects dominates depends on the skill level: as an agent's ability increases, for a while this implies the optimal contract length increases too, but then at higher skill levels, the optimal contract length starts to fall. Whether the contract for low or high types is longer depends on where the two types are on that curve.

Either way, an optimal menu of contracts must be designed so that the high skilled agents prefer to accept the contract that is intended for them rather than pretend to be low skilled, and vice versa. The basic idea is that low-skilled agents will prefer relatively low-risk, low-reward contracts, whilst high-skilled agents will choose potentially more lucrative contracts, where failure is costlier, because they are more confident in their chances of success.

But contract length also matters; the optimal menu of contracts implies terminating the contract of the low – skilled agents inefficiently early, which makes those contracts unattractive to high-skilled types. This is the price that the principal pays for not being able to observe the agent's type nor whether a genuine experiment is conducted. The principal cannot avoid over paying high-skilled agents, relative to what they could pay if types were observable (what economists call an 'information rent'), and by shortening the length of experimentation

16. This feature of the model may not sit well with practitioners who think that status-quo projects are extremely likely to fail and experimentation will immediately raise the probability of success. What squares the circle is that even successful experimentation surely involves some period of learning about the precise nature of execution, during which time news of success may be in short supply.

17. This feature of the model may not sit well with practitioners who think that status-quo projects are extremely likely to fail and experimentation will immediately raise the probability of success. What squares the circle is that even successful experimentation surely involves some period of learning about the precise nature of execution, during which time news of success may be in short supply.

by low-skilled agents the principal ensures that pretending to be low skilled is not worthwhile.

In some cases (such as when the difference in skills is great) the principal will find it optimal to offer a single contract that only high-skilled agents would accept. Otherwise, when a menu is optimal, would the principal prefer to see the job go to a high or low-skilled agent? Under the optimal menu of contracts, both types are worth hiring but the high-type deliver more value. However, the Halac model considers a one-off situation in which the principal faces one agent of unknown ability; it does not cater for the possibility of searching over time. If waiting for high types were possible but costly, the magnitude of the difference in skills that would justify excluding low-types altogether would shrink, but in cases when offering a menu is worthwhile, it would still sometimes be better to hire a low type than wait for a high type.

The authors also discuss whether the principal is better off when there are more high-skilled agents in the population. Paradoxically, they show that principals may have to pay higher information rents when high-skilled agents are more common, and ultimately end up worse off.

4.5 Parallel experimentation

It is sometimes suggested that multiple experiments should be carried out simultaneously, with the learning shared. In the development context, it might be more natural to expect this to be done by different teams working for the same contracting agent, whereas the theoretical treatments tend to consider experimentation by competing agents. There may still be some competition within organisations, however, in which case some of the lessons may still apply.

Optimally, each experimenter would completely share information about their successes and failures, so that the best programme can be rapidly reached without going down a series of dead-ends. However, each experimenter derives some private payoff from a successful new development policy (the credit of being first with the idea) and may not want to share information fully. How should experiments be structured when many experiments are run at once?

Akcigit and Liu (2016) consider a situation in which everyone can see when an experiment is proving successful, but only the experimenter can see if their project is failing. The experiments overlap, in the sense that more than one agent is experimenting with the same idea, so agents can learn from each other. They note that there are two inefficiencies that may arise: i) either some agents will work too long on projects that are known by others to be dead-ends, or ii) some agents will switch too early from promising experiments to safer well-known policies because they begin to believe that other agents working on similar experiments must have learned the experiment is a dead-end (no news is bad news). By paying agents for proving when they have had failures, in addition to when they have successes, a principal can restore efficiency: in essence, to know that an experiment has failed requires the agent to exert effort, so optimal experimentation in this context simply involves extending payment by results to paying for failures as well. This result is similar to that of Manso (2011) that was discussed earlier, but relies on differing agent incentives to innovate intersecting, rather than the problem of motivating a single agent.

On a practical note, both the individual incentives to do costly experimentation, and the incentives to share information about dead ends, suggest a system that tolerates failure early on (when experiments are likely to induce failure), but not later on (when genuine experiments are more likely to have been terminated and hence where continuing failure is more likely a sign of low effort). Manso (2011) discusses a story of the Bank of America's exploratory research team, which was initially given a target failure rate of 30%. The first year, only 10% of projects failed, suggesting the experiments were not sufficiently radical. Headquarters, as a result, upped the expected failure rate to 40%, with bonuses lowered if not enough projects failed! The broad lessons are that tolerance and payment for failed experimentation early on in a new setting can be necessary both to make information sharing from experimentation worthwhile, and to encourage experimenters to pursue high-risk projects rather than the status quo.

With multiple experimenters, only a few of which will be given long-term contracts on the basis of their experimentation, Gross (2016) shows that it is critical not to choose either too few or too many experimenting teams. Why? Imagine contractors can choose how radical their particular experiment is (where being more radical means having a potentially higher payoff). If there is very little competition, then one has a good chance of winning with relatively safe bets. If there is too much competition, then the probability of winning becomes so low that the costs of experimentation mean it is not worth the effort. Somewhere in between teams will understand that the best chance of winning comes from taking a risk to find an unusually successful new method. Gross looks at online contests for logo design and sees that if, upon initial submission, there is another similar and highly-rated logo, the designer will radically revamp the design, but if there is either no strong competitor *or* many competitors, the designer will make only cheap and simple modifications. This suggests that if a development agency wants to incentivise contractors to 'compete' for future grants on the basis of policy experiments, care needs to be taken to ensure this competition is neither too limited in participant number nor too intense.

Box 5: Using contracts to screen agents

Not all agents are equal – some may have more experience, better ideas, may be more motivated by development goals whilst others are more motivated by making money for minimal effort, others may have ideological biases. Contract theory often involves using 'menus of contracts' so that agents sort themselves into different contracts according to their hidden information. In theory, these menus perform better than contracts which offer set terms to all agents. A simple example is to offer a flat fee contract that only an agent who is confident of their ability to produce an output at low cost would choose, whilst others would choose a cost-plus contract. Wren-Lewis (2016) describes these ideas at more length in the context of financing for health.

Systematic evidence is lacking for the use of screening contracts in practice. Examples include publishers who often ask writers to choose different payments depending on the time taken to deliver a book, and agricultural landlords offering different combinations of contract duration and type (share cropping or fixed rent) to different kinds of farmer. The feasibility of screening in highly complex settings where agents may have limited knowledge of their own abilities is open to question.

In the context of contracts for adaptive programming, principals would benefit from screening agents for their degree of intrinsic developmental motivation. Perhaps agents who are more motivated by development goals could be offered a more open-ended contract that gives them more discretion, but such a contract might be expected to deliver poor results if accepted by a self-interested agent. So how can the principal prevent the wrong kind of agent accepting the contract? It might be possible to design a menu of contracts so that self-interested agents do not choose the more open-ended option (perhaps because they are able to make more money under the alternative with more tightly-defined parameters). This can get very complicated, however, if, for example, some more self-interested agents also have higher ability.

The screening literature is vast, and a result that appears in many models is 'no distortion at the top'. Essentially, in the presence of hidden information, the optimal menu of contracts is designed so that the agent whose hidden information is most valuable to the principal does the 'optimal' thing, whereas agents with less valuable information or traits must be distorted (are not as efficient as they would be if there was no need to sort agents). A simple example can illustrate the principle at work (realism is not the objective here). Consider selling 2 hectares of land. Half of the potential buyers value each hectare at 10, the other half value the first hectare at 10 and the second hectare at only 4. A seller who is able to discriminate between buyers would set a price of 20 for high types and 14 for low types, and expect to make (0.5*20)+(0.5*14)=17 per lot. This is the first-best, full-information benchmark.

But what if the buyer's valuation is unknown to the seller? If the price is set at 20, half the time the potential customer will not purchase the land, so the seller will only expect to earn (0.5*20)=10. If the price is set at 14, a sale will always be made, earning 14. Even better, though, is to use a menu of contracts: if the buyer wants both hectares, the price is 20, else the price is 10 but only a single hectare may be bought. Under this contract, the high-value buyer still buys both hectares and the low-value buyer only buys a single hectare. This contract earns the seller (0.5*20)+(0.5*10)=15, more than a single fixed-price contract would, but still less than the perfect information benchmark.

This menu of contracts sees the high-type always buying the full parcel ('no distortion at the top') while the low-type only gets one hectare ('distortion at the bottom'). The only way to get the low type to buy both hectares would be to offer such a low price that the high type would 'pretend' to be a low-type.

Halac et al. (2016) integrate the problem of screening agent types into an experimental setting, where the agents vary by ability. A handful of papers examine screening agents for levels of intrinsic motivation – as opposed to self-interest – but not in the context of experimentation. In one stream of this literature which looks at recruitment, bonuses or monetary incentives have the potential to reduce intrinsic motivation, because they signal that a particular job is burdensome or unfulfilling (Benabou and Tirole, 2003). This form of 'crowding out' occurs in settings where the principal is more informed than the agent about important characteristics of a task, however, so might be less relevant in the context of development (Ashraf et al., 2015). A second stream considers whether the desire of publicly-minded organisations to recruit 'mission-oriented' agents should affect contracts. The answer is that when publicly-minded organisations compete for workers with the private sector, and observe neither the worker's willingness to work for less money on publicly-minded projects nor the worker's quality, the optimal contracts sort all the publicly-minded workers into the publicly-minded jobs, but then use traditional 'no distortion at the top' contracts to induce effort. That is, agents may work at a discount if they are intrinsically motivated, but once that discount is accounted for, the problem of how to incentivise effort given hidden information remains, and its solution involves the standard trade-offs (Barigozzi and Burani, 2016).

5. The gaps between theory and practice

Inside development organisations, the challenges of adaptive programming can look quite different from the standpoint of contract theory. For this reason, there remain open theoretical questions in the design of adaptive contracts.

Contract theorists see barriers to experimentation and truthful disclosure arising from the incentives faced by agents. They tend to assume that if these problems are solved, then disclosing and using information would be a straightforward matter. In contrast, practitioners are more concerned with the pragmatic challenge of ascertaining what information is required and how to obtain it, and with getting the feedback loop working so that activities can be adapted during implementation, with the consent of the principal. From this point of view, even if everybody's incentives are aligned, generating useful information and responding to it is still challenging.

The most obstructive barriers to adaptive programming can be the practicalities of operating within a cumbersome bureaucracy. Here, principals are trying to manage budgets in a predictable fashion and provide lines of accountability and demonstrable results to their ultimate paymasters (i.e. politicians and the public). Accountability is not a word used by contract theorists, but it dominates the discussion in development circles. De Renzio (2016) argues that foreign aid is undermined by the pressure for donors to be accountable for short-term results. There are other barriers to change too. Bain (2016) writes about the changes to managerial culture within the World Bank that will be necessary for adaptive programming to flourish; the difficulties of securing money and time for discovery before rushing to start work and disperse funds; the fact that reporting systems have mostly been designed to catch and penalise failures, rather than create feedback loops that put new knowledge into action; the difficulties of deciding what to monitor; and the fact that staff are rewarded for getting projects done (even if they do little good) and fear of getting punished for following adaptive programming principles (being experimental, acknowledging and learning from failure).

This paper has introduced some ideas from contract theory and its application to experimentation. The problems identified by contract theory flow naturally from the assumptions made about the environment that principals and agents operate in, and what determines their pay-offs. It is not clear how well the theory corresponds to the reality of development practice, however, and our intention in this paper is to stimulate discussion.

The motivation behind adaptive programming is often not that there is technical knowledge that needs to be discovered – such as the most cost-effective way of doing something – but rather about the need to discover how to negotiate paths around political and institutional barriers. Problems such as the 'ratchet effect' are easier to understand when experimentation produces technical information which the principal can use against the interests of the agent in the future, in the sense of making them work harder, or accept lower profits.

It is harder to see how this idea applies when the purpose of experimentation is to solve political economy problems. At a high level of abstraction, finding ways of working politically could be viewed as a form of technical problem that boils down to discovering what can be done at what cost. Costs might not feel like the relevant problem to the practitioner who will be paid the same regardless, but wants to be given the freedom to find their way through a maze rather than being asked to continually go down a dead end.

Another motivation behind adaptive programming is the need to discover what results development interventions ought to target in a given context – which can be conceptualised as 'searching design space' (Pritchett et al., 2013) to locate the activity with the highest payoff. Some of the papers surveyed here involve discovering whether an experiment is promising or not, but this might not do full justice to the problem of designing contracts, when figuring out what to attempt is as important as figuring out how to do it.

Contract theorists see adaptive contracts as introducing new problems, but theorists tend to start from simple benchmarks and work outwards towards more difficult problems. Practitioners start from a status quo which has plenty of problems and want to move in the direction of solving them. Few theorists would think of starting where practitioners often find themselves: being asked to implement a predesigned project with little chance of success. A contract theorist might not call payment-by-results an adaptive contract, because it might not occur to a theorist that the status quo would be a contract that pins down inputs and actions, as opposed to specifying the

required product and leaving production decisions to the agent.

Contracts that introduce scope for adaptation introduce new problems, but a problematic adaptive contract may still perform far better than an 'unproblematic' contract to do the wrong thing in the wrong place. Of course the relative merits of adaptive versus traditional contracts depend on context, and where the status quo approach is already reasonably successful, the risks and rewards of introducing scope for agents to potentially manipulate results in their favour may look less favourable.

Conclusion

Development agencies often implement projects by hiring contractors. Contracts matter, because agents will rarely have exactly the same objectives as the principal, and because there are information problems that require clever contract design to overcome.

Adaptive programming has the potential to deliver better outcomes than traditional pre-designed projects, because it introduces scope to search for novel ways of achieving development goals, directed by local partners, and to learn from experience and adapt the project as it goes along. But just as payment-by-results is not a panacea, adaptive programming brings some new problems of its own.

The theoretical literature on experimentation in contracts is rapidly growing. Our survey has generated six broad lessons to keep in mind for practitioners hoping to use adaptive contracts.

1. **Uncertainty makes relational contracts more difficult to sustain.** Most development contracts are likely to be hard to enforce in court and will be 'relational', meaning that they are given force by desire for repeat business. Information from experiments can change the value of the continued relationship and introduces uncertainty. When the relationship is fairly stable, this gives agents confidence about the future and they are more likely to exert effort to avoid the relationship breaking down. But if the future is too uncertain, the incentive to think about short-term payoffs, and possibly shirk or cut corners, can increase.
2. **The 'ratchet effect' means agents may conceal information.** The big problem with adaptive contracts is that agents may be reluctant to experiment and truthfully reveal the results due to the 'ratchet effect'. This problem arises when the future contractual relationship is not pinned down, as will often be the case when the principal wishes to retain the option of using the results of experimentation to change how much will be paid for doing what. Ratchet effect worries can be 'solved' if the principal can commit to not using information generated against the interests of the agent, but this limits the usefulness of adaptive contracts. Adaptive contracts must find the right balance between protecting agents from having information used against them, so that they feel free to experiment, and retaining some ability to use the results to change the contractual relationship.
3. **Contracts that make a partial commitment can mitigate the ratchet effect.** When transfers like cash bonus payments are not feasible or desirable, one way around the ratchet effect is partial commitment through delegation: contractors are given discretion over the use of resources, but after some period of time if the results look bad, experimentation is terminated and the contract reverts to the status quo. When project quality is not known and even 'bad' projects, which the aid agency would rather not support, may sometimes show signs of success, a sliding deadline can be used where the deadline is extended after each sign of success, so bad projects will still be terminated eventually.
4. **Contracts that commit to future results-based payments induce effort.** When payments on results are feasible and appropriate, the schedule of payments needs to induce genuine experimentation early on and keep inducing effort until, if no success is observed, the contract is terminated. This requires that successes should not be heavily rewarded early on (because it can discourage risky experiments that are initially less likely to be successful but raise the chances of long-run success), but also should not be excessively back-loaded (otherwise agents will intentionally delay success). Payments for early negative results can be optimal when encouraging experimentation, if done in conjunction with strong payments for future success. It should be remembered that contract theory also predicts that in some circumstances results-based payments will perform poorly, however.
5. **A menu of contracts can sort between different types of agent.** When neither the effort that contractors exert when experimenting, nor their ability can be observed – and it is appropriate to base payments on results – then it may be possible to design a menu of contracts where low-ability agents choose lower-reward, lower-risk contracts and high-ability agents choose potentially more lucrative contracts that impose a greater penalty on failure. Discriminating between agent types can be better value for money than a single contract.
6. **The number of competing teams is important when running experiments in parallel.** When experiments are run in parallel, there needs to be an incentive to reveal failures (otherwise other teams will continue chasing dead-ends), and there needs to be neither too few nor too many teams. With too few teams, contractors will all work on relatively safe experiments rather than pursue riskier high-return experiments. With too many

teams, the chances of 'winning' will fall too low and too few resources will be devoted to experimentation.

There is more work to be done from a theoretical perspective, taking the tools developed by contract theorists and applying them to problems that more explicitly resemble those faced by development practitioners. The gap between theoretic insight and pragmatic contract design also needs to be closed. From an empirical perspective, much could be gained by studying the performance of contracts used in the real world by development agencies and contractors doing adaptive programming.

References

Akcigit, Ufuk, and Qingmin Liu. (2016) 'The role of information in innovation and competition.' Journal of the European Economic Association. 14(4): 828–870

Andrews, M. (2013) *The limits of institutional reform in development: changing rules for realistic solutions*. Cambridge: Cambridge University Press.

Andrews, M., Pritchett, L., and Woolcock, M. (2013) 'Escaping capability traps through problem driven iterative adaptation (PDIA)', *World Development* 51: 234–244.

Andrews, M and Pritchett, L and Woolcock, M. (2015) 'Doing problem driven work'. CID Working Paper No. 307.

Arrow, K. J. (1971) 'The economic implications of learning by doing', in F. H. Hahn, *Readings in the theory of growth* (pp. 131–149). Palgrave Macmillan UK.

Ashraf, N., Bandiera, O. and Lee, S.S. (2015) 'Do-gooders and go-getters: career incentives, selection and performance in public service delivery'. Harvard Business School Working Paper. Cambridge, MA: Harvard.

Baker, G., Gibbons, R. and Murphy, K. (2002) `Relational contracts and the theory of the firm', *Quarterly Journal of Economics* 117(1): 39–84.

Bain, K. (2016, forthcoming) 'The Implications for large aid organizations of Doing Development Differently – updating the plumbing to fit the architecture'. London: ODI.

Banerjee, A. V., and Duflo, E. (2000) 'Reputation effects and the limits of contracting: a study of the Indian software industry', *Quarterly Journal of Economics* 115(3): 989–1017.

Barigozzi, F., and Burani, N. (2016) 'Screening workers for ability and motivation', *Oxford Economic Papers* 68(2): 627–650.

Benabou, R., and Tirole, J. (2003) 'Intrinsic and extrinsic motivation', *Review of Economic Studies* 70(3): 489–520.

Bennedsen, M. and Schultz, C. (2005) 'Adaptive contracting: a trial-and-error approach to outsourcing', *Economic Theory* 25.

Bhaskar, V. The Ratchet Effect Re-examined: A Learning Perspective. No. 9956. CEPR Discussion Papers, 2014.

Callander, S. and Matouschek, N. (2014). Optimal search in rugged landscapes. Working paper, graduate school of business, stanford university 2014

CGD (2013) 'Investing in Social Outcomes: Development Impact Bonds' Development Impact Bond Working Group.

Clist, P. and Dercon, S (2014) '12 Principles for Payment By Results (PbR) In International Development', DFID Research for Development. London: DFID.

Clist, P. and Verschoor, A. (2014) 'The conceptual basis of payment by results'. School of International Development, University of East Anglia.

Clist, P. {2016} 'Payment by Results in Development Aid: All that Glitters is not Gold' The World Bank Research Observer

De Renzio, P. (2016) 'Accountability dilemmas in foreign aid'. London: ODI.

Derbyshire, D. and Donovan, E. (2016) 'Adaptive programming in practice: shared lessons from the DFID-funded LASER and SAVI programmes' (http://www.laserdev.org/).

Freixas, X., Guesnerie, R., and Tirole, J. (1985) 'Planning under incomplete information and the ratchet effect', *Review of Economic Studies* 52(2): 173–191.

Gibbons, R and Henderson, R. (2011) `Relational contracts and organizational capabilities', *Organizational Science* 23(5): 1350–1364.

Gross, Daniel P. Creativity Under Fire: The Effects of Competition on Creative Production. No. 16–109. Harvard Business School, 2016.

Gou, Y. (2016) 'Dynamic delegation of experimentation), forthcoming at *American Economic Review*.

Halac, M. (2012) 'Relational contracts: the value of relationships', *American Economic Review* 102 (2).

Halac, M., Kartik, N., and Liu, Q. (2016, forthcoming) 'Optimal contracts for experimentation', *Review of Economic Studies*.

Hegen, R. J. (2015) 'Dancing to the donors' tune? Policy-choice in aid-dependent countries', *Scandinavian Journal of Economics* 117(1): 126–163.

Helper, S. and Henderson, R. (2014) 'Management practices, relational contracts and the decline of General Motors', Working Paper 2014–062. Cambridge, MA: Harvard Business School.

Holmstrom, B. (1982) 'Moral hazard in teams', *Bell Journal of Economics* 13(2): 324–340.

Holmstrom, B. and Milgrom, P. (1991) 'Multitask principal-agent analyses: incentive contracts, asset ownership, and job design', JLEO 7: 24–52.
Huysentruyt, M. (2011) *Development aid by contract: outsourcing and contractor identity.* London: London School of Economics.
Kerr, S. (1975) 'The folly of rewarding A, while hoping for B', *Academy of Management Journal* 18: 769–783.
Laffont, J. J., and Tirole, J. (1988) 'The dynamics of incentive contracts', *Econometrica: Journal of the Econometric Society* 56(5): 1153–1175.
Levin, J. (2003) `Relational incentive contracts', *American Economics Review* 93(3): 835–857.
Levinthal, D. (1997) 'Adaptation on rugged landscapes', *Management Science* 43(7): 934–950.
Manso, G. (2011) 'Motivating innovation', *Journal of Finance* 66: 1823–1860.
Melitz, M. (2013) 'The impact of trade on intra-industry reallocation and aggregate industry productivity', *Econometrica* 71(6).
NAO (2015) 'Outcome-based payment schemes: government's use of payment by results'. London: National Audit Office.
New York Times, The (2014) 'Coke and MacDonald's: growing together since 1955', 15 May, David Gelles (http://www.nytimes.com/2014/05/16/business/coke-and-mcdonalds-working-hand-in-hand-since-1955.html?_r=1).
Niehaus, P. (2014) 'A theory of good intentions'. San Diego, CA: University of California and Cambridge, MA: NBER.
Pellini, A. and Nixon, N. (2016) 'Are partnership agreements a way forwards for Doing Development Differently?' (http://doingdevelopmentdifferently.com/are-partnership-agreements-a-way-forward-for-doing-development-differently/).
Perakis, R.and Savedoff, W. (2015) 'Does results-based aid change anything? Pecuniary interests, attention, accountability and discretion in four case studies'. CGD Policy Paper 052. Washington, DC: CGD.
Pritchett, L., Samji, S., and Hammer, J. S. (2013) 'It's all about MeE: using structured experiential learning ('e') to crawl the design space'. Center for Global Development Working Paper (322). Washington, DC: CGD.
Shutt, C. (2016) 'Towards an alternative development management paradigm? Rapport 2016:07. Stockholm: Expertgruppen för biståndsanalys (EBA).
Syverson, C. (2011) 'What determines productivity', *Journal of Economic Literature* 49(2): 326–365.
Syverson, C. (2004) 'Market structure and productivity: a concrete example', *Journal of Political Economy* 112(6): 1181–1122.
Tirole, J. (1999). 'Incomplete contracts: where do we stand?', *Econometrica* 67(4): 741–781.
Tirole, J. (1986) 'Hierarchies and bureaucracies: on the role of collusion in organizations', *Journal of Law, Economics, & Organization* 2(2): 181–214.
USAID (2016) 'Evidence base for collaborating, learning, and adapting: Summary of the Literature Review'. Washington, DC: USAID.
Van Thiel, S. and Leeuw, F. (2002) 'The performance paradox in the public sector', *Public Performance & Management Review* 25 (3): 267–281.
Weitzman, M. (1980) 'The "Ratchet Principle"' and performance incentives', *Bell Journal of Economics* 11(1): 302–308.
Wild, L., Booth, D., Cummings, C., Foresti, M. and Wales, J. (2015) 'Adapting development: improving services to the poor'. London: ODI.
Williamson, O. (2002) 'The lens of contract: applications to economic development and reform'. Forum series on *the institutional economics approach to aid effectiveness*. Washington, DC: USAID and College Park, MA: IRIS.
Wren-Lewis, L. (2016) 'Designing contracts for the Global Fund: lessons from the theory of incentives'. Center for Globalization and Development Working Paper 425. Washington, DC: CGD.

ODI is the UK's leading independent think tank on international development and humanitarian issues.

All ODI Reports are available from www.odi.org

Cover photo: AU UNIST PHOTO/ David Mutua 2014. A female senior civil servant during training at General Kaahiye Police Academy in Mogadishu, Somalia, by consultants from the Ugandan Management Unit, 24th May 2014.

Overseas Development Institute
203 Blackfriars Road
London SE1 8NJ
Tel +44 (0) 20 7922 0300
Fax +44 (0) 20 7922 0399

odi.org