Organisations and Development

Lecture notes for Development Economics I
(option course, MPhil in Economics)

Simon Quinn*

Michaelmas 2018

*Department of Economics, Centre for the Study of African Economies and St Antony's College, University of Oxford: simon.quinn@economics.ox.ac.uk. Without implicating them in any shortcomings, I would like to thank Marcel Fafchamps, Clément Imbert, Rocco Macchiavello, Daniela Scur and Kate Vyborny for their discussions on this topic.

Why study organisations?

(Or, 'But haven't we already discussed firms'?)

- Gibbons, R. and Henderson, R. (2013): "What Do Managers Do? Exploring Persistent Performance Differences among Seemingly Similar Enterprises," in *Handbook of Organizational Economics* (Gibbons, R. and Roberts, J., eds.), Princeton University Press.

What do managers do? Implicitly, we considered this question throughout our previous module on firms. In our previous module, managers accumulated capital, and decided how to combine capital and labour, made decisions on technology adoption, and supervised workers (subject to a limited 'span of control'). Yet all of these decisions were decisions taken by the manager *acting as the firm*. Our previous module said almost nothing about the incentives that might exist *within* a firm — or, indeed, within other forms of organisation.

Such organisational features are, however, important for understanding a wide range of economic phenomena. This is of particular interest in developing countries — where both private and public sectors seem often seem afflicted by a disproportionate number of 'dysfunctional organisations'. In this module, we aim to delve inside organisations — both private firms and public bureaucracies — to explore problems of communication, design and incentives. We will cover the following topics...

Michaelmas			
Week 3	Lecture 1	**Organisations, Rules and Management**	Monday, 11am – 1pm Seminar Room C
Week 3	Lecture 2	**Organisational Hierarchy**	Thursday, 11am – 1pm Seminar Room C
Week 4	Lecture 3	**Incentives within Organisations**	Thursday, 11am – 1pm Seminar Room C
Week 5	Lecture 4	**Relational Contracts**	Monday, 1pm – 3pm Seminar Room C

1 Lecture 1: Organisations, Rules and Management

- $\boxed{\star}$ Bloom, N., Eifert, B., Mahajan, A., McKenzie, D., and Roberts, J. (2013): "Does Management Matter? Evidence from India," *The Quarterly Journal of Economics*, 128(1), 1–51.
- $\boxed{\star}$ Bloom, N., Mahajan, A., McKenzie, D., and Roberts, J. (2018): "Do Management Interventions Last? Evidence from India," *World Bank Policy Research Working Paper*, 8339.
- $\boxed{\star}$ Bloom, N., Lemos, R., Sadun, R., Scur, D. and Van Reenan, J. (2014): "The New Empirical Economics of Management" *Journal of the European Economic Association*, 12(4), 835–876.
- $\boxed{\star}$ Ellison, G., and Holden, R. (2014): "A Theory of Rule Development," *Journal of Law, Economics, and Organization*, 30(4), 649–682.
- Adler, S. (2018). "Post No Evil," Radiolab (podcast): `www.wnycstudios.org/story/post-no-evil`.
- Bandiera, O., Hansen, S., Prat, A., and Sadun, R. (2017): "CEO Behavior and Firm Performance," *Working Paper*.
- Bloom, N., Sadun, R. and Van Reenen, J. (2016): "Management as a Technology?" *Working paper*.
- Bloom, N. and Van Reenan, J. (2010): "Why Do Management Practices Differ Across Firms and Countries?" *Journal of Economic Perspectives*, 24(1), 203–224.
- Callander, S. and Matouschek, N. (2016): "The Risk of Failure: Trial and Error Learning and Long-Run Performance," *Working paper*.
- Chassang, S. (2010): "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts," *The American Economic Review*, 100(1), 448–465.
- Garicano, L. and Rayo, L. (2016): "Why Organizations Fail: Models and Cases," *Journal of Economic Literature*, 54(1), 137–192.
- Rivkin, J.W. (2000): "Imitation of Complex Strategies," *Management Science*, 46(6), 824–844.
- McKenzie, D. and Woodruff, C. (2016): "Business Practices in Small Firms in Developing Countries," *Management Science*, 63(9), 2967–2981.

It is impossible to talk sensibly about *organisations* without also discussing *organisation*. What do managers do? Well, tautologically, they manage; they *organise*. There are several things that you might do to organise a group of people. For one thing, you might implement a *hierarchy* — in which different people take different roles, and are answerable to different middle managers in different ways. This is the topic of Lecture 2. You might need to manage *incentives* — deciding how to monitor the others in the group, and how to reward or penalise them on the basis of what you observe. This is the topic of Lecture 3. You may choose to share information and insights with other organisations; we will discuss this in Lecture 4. In this first lecture, we start with a more basic notion: management involves *instruction*, in which a manager directs subordinates to act in particular ways in particular circumstances. Put simply, management involves *rules*.

In this lecture, we consider an elegant model of 'rule by analogy' — in which a manager uses past organisational experience to teach a subordinate how to respond to future challenges. This model — due to Ellison and Holden (2014) — highlights several important features of rule-making, and provides a useful basis for thinking about management practices in organisations. In the second part of the lecture, we consider recent empirical evidence — both experimental and descriptive — that tests the role of such management practices for organisational performance.

For a fascinating illustration of the basic idea in Ellison and Holden (2014), I recommend that you listen to 'Post No Evil' (Adler, 2018) — a podcast episode about the evolution of Facebook's 'Community Standards' policy.

> So from then on as they run into problems, those rules just constantly get updated, with constant amendment — yeah, constant amendments. New problem? New rule. Another new problem? Updated rule. In fact, at this point, they're amending these rules up to twenty times a month. . .
>
> You know, it sounds a lot like common law. So, common law is this system dating back to early England where individual judges would make ruling which would sort of be a law but then that law would be amended or evolved by other judges. So the body of law was sort of constantly fleshed out in face of new facts. Literally every time this team at Facebook would come up with a rule that they thought was airtight — ka-plop! — something would show up that they weren't prepared for, that the rule hadn't accounted for. . .
>
> *This is a utilitarian document it's not about being right one hundred percent of the time — it's about being able to execute effectively. In other words we're not trying to be perfect here. And we're not even necessarily trying to be one hundred percent just or fair We're just trying to make something that works.*

1.1 A simple model of rules as analogies

We begin by considering Ellison and Holden's (2014) model of rule development. Consider two players: a Principal ('she') and an Agent ('he'); in this context, the former represents some kind of manager, and the latter is some kind of subordinate. The principal represents the interests of an organisation as a whole; the agent is responsible for taking decisions on behalf of the organisation, according to rules that the Principal has set for him.

The Principal and Agent interact in discrete time. In each period t, the following occurs:

(i) The Principal issues a rule to the Agent (r_t); the rule instructs the Agent that, in some circumstance, he should either take action $a_t = 1$ or $a_t = -1$.

(ii) The Agent observes a state (ω_t) and chooses one of three available actions: $a_t = -1$ or $a_t = 1$ (if the state is covered by an existing rule, either from period t or earlier), or $a_t = 0$ (if the state is not covered by any existing rule).[1]

(iii) The Principal receives a payoff, depending both on the state and the action: $a_t \cdot f(\omega_t)$.

The game continues for T periods; the Principal's payoff (indeed, more generally, the payoff of the organisation) is $V = \sum_{t=1}^{T} \delta^t \cdot a_t \cdot f(\omega_t)$. The Agent does not have a payoff function; this may seem strange, but remember that the Agent's role is simply to implement rules. As Ellison and Holden explain, the model is 'essentially making [the Agent] a robot who literally follows instructions' (page 10).

Crucially, all of the rules issued must take a particular form: they must be rules of analogy. Essentially, the Principal is constrained to telling the Agent, "if the state is *similar enough* to *this* particular state you've seen before, you should take this particular action; if the state is *similar enough* to *that* particular state you've seen before, you should take that other particular action", and so on. In this way, the Ellison-Holden model jointly captures two components: (i) what the authors call 'bounded communication' ('the Principal is unable to immediately communicate the mapping from states of the world to optimal actions') and (ii) what the authors call a 'recognition problem' (meaning that '[t]he agent must have experienced a state [himself] (or a nearby state) before the Principal can promulgate a rule relating to it').

Formally, this requires a 'commonly understood distance function' (p.7); Ellison and Holden denote this using '$||||$'. The Principal's rules must take the form $r_t = (\omega_t, d, a)$,

[1] Ellison and Holden consider just two actions ($a_t = -1$ or $a_t = 1$) and imagine the Agent tossing a fair coin in the case that no rule applies. This ends up being isomorphic to the case we consider, but I think the three-action example is simpler for teaching purposes.

interpreted to mean 'take action a if $||\omega_t - \omega|| < d$'.[2] This also requires us to assume something about how rules accumulate over time; in effect, a 'meta-rule' that determines how individual rules take precedence.[3] In this lecture, we will consider a stark and simple approach: older rules always prevail over newer rules. (Of course, we could consider many other approaches; in their paper, Ellison and Holden term our approach 'no overwriting', but also discuss 'incremental overwriting' and 'vast overwriting'.)

We will limit our attention — as Ellison and Holden do for most of their paper — to the illustrative case in which the state space is a circle with a circumference of 1. As the authors explain,

> Mathematically, this corresponds to setting $\Omega = [0, 1]$ and defining distances by $||x - y|| = \min\{|x - y|, 1 - |x - y|\}$.

Figure 1.1 illustrates; it shows values of ω on a circle, with an illustrative distance measure applied.

Figure 1.1: **The state space**

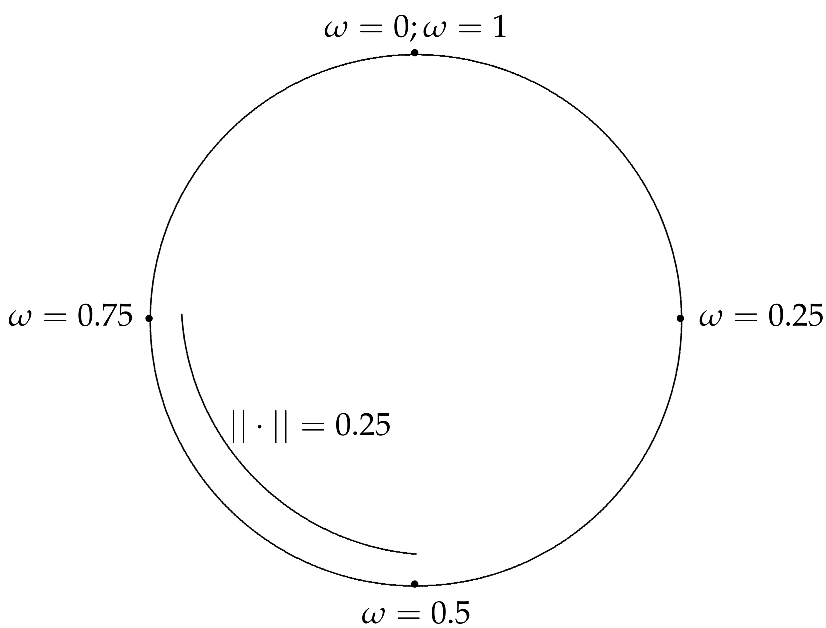

Further, we need to make an assumption on the distribution of ω. For simplicity, we will assume that ω is drawn from a Uniform distribution on $[0, 1]$: $\omega_t \sim U(0, 1)$.

[2] Ellison and Holden also consider what they term 'p' — namely, a notion of priority — but this is redundant in our context, because we only consider their 'no overwriting' case.

[3] The famous Oxford legal philosopher H.L.A. Hart described such 'meta-rules' as 'rules of recognition'.

For tractability, we will also consider a simple parametric form for the payoff function, f:

$$f(\omega_t) = \sin(2\pi \cdot \omega_t). \tag{1.1}$$

Figure 1.2 illustrates. Notice that $f(\omega_t) = f(\omega_t - 1)$; *i.e.* f is a periodic function with period 1. This will be convenient for applying our distance metric.

Figure 1.2: **The payoff function**

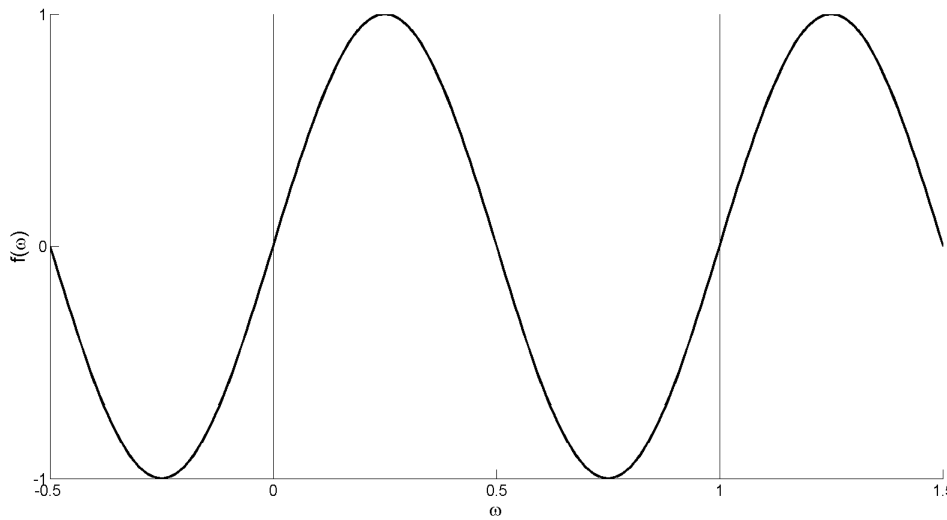

1.1.1 A two-period version

Let's start by considering a simple two-period model. Specifically, let's consider the case where $T = 2$ and $\delta = 1$.[4] To recall, this requires just *one* decision by the Principal: namely, having observed ω_1, the Principal needs to formulate a rule to pass to the Agent for period 2.

How should the Principal decide? To answer this, we need to write the Principal's expected utility, as a function of the rule r_1. Denote d as the distance used in rule r_1, and a as the instruction specified in the rule. Straightforwardly, the Principal's expected

[4] Of course, we could choose any value $\delta \in (0, 1]$ and we would reach the same result in this example.

payoff is:

$$V = \begin{cases} \displaystyle\int_{\omega_t - d}^{\omega_t + d} f(x)\, dx, & \text{if } a = 1; \\ \\ \displaystyle -\int_{\omega_t - d}^{\omega_t + d} f(x)\, dx, & \text{if } a = -1. \end{cases} \tag{1.2}$$

Let's start by assuming that $\int_{\omega_t - d}^{\omega_t + d} f(x)\, dx > 0$ — in which case the Principal will choose $a = 1$. (Of course, if this turns out not to be the case, the Principal can just solve the same optimisation problem, but choosing $a = -1$.) In our case, this corresponds to assuming — initially, at least — that $\omega_t \in [0, 0.5]$. In this case, the Principal faces the following decision:

$$\max_{d \in [0, 0.5]} \int_{\omega_t - d}^{\omega_t + d} f(x)\, dx = \int_{\omega_t - d}^{\omega_t + d} \sin(2\pi \cdot x)\, dx. \tag{1.3}$$

The first-order condition is:[5]

$$\sin\left[2\pi \cdot (\omega_t + d)\right] + \sin\left[2\pi \cdot (\omega_t - d)\right] = 0. \tag{1.4}$$

Notice that this first-order condition holds for *any* value of $\omega \in [0, 1]$; if $\omega > 0.5$, the Principal uses this same first-order condition to set d, but then chooses $a = -1$ rather than $a = 1$.

'Excess breadth': This is immediately a very interesting result. It tells us that the Principal optimally chooses d such that the *average* payoff at the two boundary cases (that is, $\sin\left[2\pi \cdot (\omega_t + d)\right]$ and $\sin\left[2\pi \cdot (\omega_t - d)\right]$) is zero. Except for the degenerate case $\omega_t = 0$ (which, after all, happens with probability zero), this means that *one of* the boundary cases will have a positive payoff and the other will be negative. Ellison and Holden describe this as a rule having 'excess breadth': 'rules are designed to have "excess breadth" in the sense that they are intentionally made to cover a broader domain than the domain on which they produce only correct decisions'.

Solving d: With the exception of the degenerate cases $\omega_t = 0$, $\omega_t = 0.5$ and $\omega_t = 1$, we can use the first-order condition to solve d uniquely. The first-order condition implies $\sin\left[2\pi \cdot (\omega_t + d)\right] = -\sin\left[2\pi \cdot (\omega_t - d)\right]$. But notice that $\sin(x) = -\sin(x + \pi)$. From this, it follows straightforwardly that $2\pi \cdot (\omega_t + d) = 2\pi \cdot (\omega_t - d) + \pi$, and $d = 0.25$. That is, irrespective of ω (aside from three degenerate cases), the Principal optimally sets a rule to cover *half* the state space (*i.e.* from $\omega_t - 0.25$ to $\omega_t + 0.25$).

[5] You *do* remember the Leibniz integral rule, don't you. . . ?

The value of the rule: So what, then, is the Principal's indirect utility? With the result $d = 0.25$, this is straightforward:

$$V = \begin{cases} \displaystyle\int_{\omega_t - 0.25}^{\omega_t + 0.25} \sin(2\pi \cdot x)\, dx, & \text{if } \omega \in [0, 0.5]; \\ \displaystyle -\int_{\omega_t - 0.25}^{\omega_t + 0.25} \sin(2\pi \cdot x)\, dx, & \text{if } \omega \in (0.5, 1]. \end{cases} \tag{1.5}$$

You should be able to check that this evaluates to:

$$V = \frac{1}{2\pi} \cdot \left| \cos\left[2\pi \cdot (\omega_t - 0.25)\right] - \cos\left[2\pi \cdot (\omega_t + 0.25)\right] \right| \tag{1.6}$$

Figure 1.3 illustrates; it shows how the indirect utility (that is, the value of the problem) varies with ω_1.

Figure 1.3: **Indirect utility as a function of** ω_t

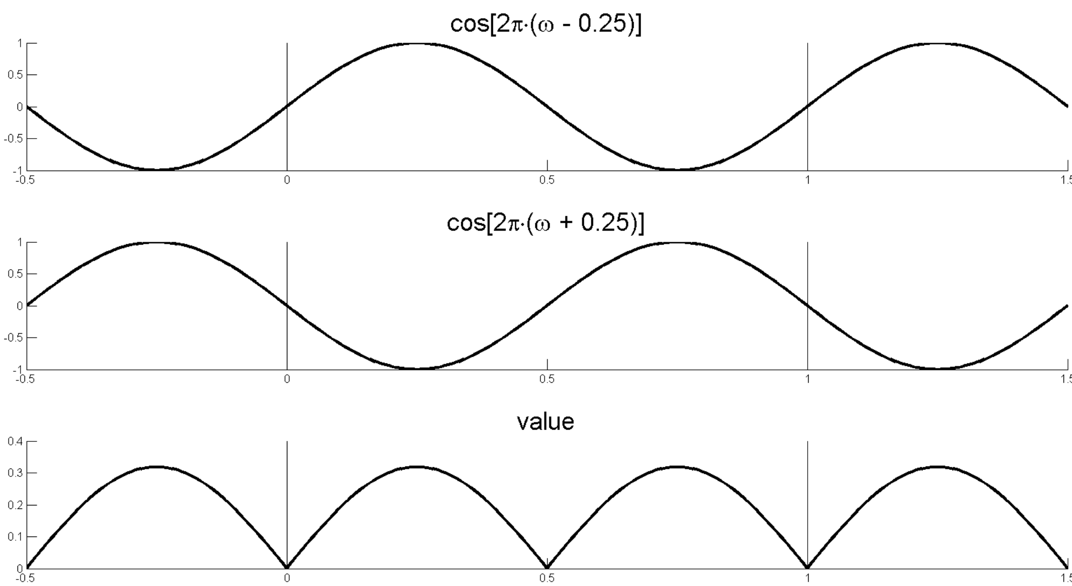

This, too, is a very insightful result: it shows that the organisation's performance depends crucially on its history of organisational challenges, because this history is necessary in order to form useful rules. As Ellison and Holden explain (page 13):

> In our model, historical accidents can be a source of performance differences: "lucky" early realizations of the state can lead to a very effective rule book being established, delivering highly efficient outcomes. Conversely, bad draws early on can have persistent effects which cannot be overcome.

'But isn't this all just driven by the assumption of rule symmetry'? Before we proceed to a more complicated case, it is worth pausing to consider the assumption of rule symmetry. Clearly, in this simple one-dimensional case, much of the interest inefficiencies that we have considered could be overcome if only the Principal could issue an asymmetric rule (*i.e.* if the Principal could demand that a rule apply from $\omega_t - \underline{d}$ to $\omega_t + \overline{d}$). When you read Ellison and Holden's paper, you should pay particular attention to section 2.4 on this issue. In short, we should think about the symmetry restriction as an elegant means to capture communication problems in a one-dimensional space, rather than as a substantive description of how managers must issue decisions in the real world.

1.1.2 Extending to three periods

The two-period setup is very useful for illustrating many of the key features of the model. But what about a longer timescale? Let's extend the model to the case where $T = 3$ and $\delta = 1$. Intuitively, it should be clear how this extension changes our earlier results. Under the assumption of 'no overwriting', the principal's choice of rule for period 2 must also apply to period 3. The broader the rule that is set in period 2, the less flexibility the principal will have in period 3 to change the rule book. In short, the shadow of the future leads the principal to formulate narrower rules that she otherwise would.

Ellison and Holden formalise this idea in their Proposition 2. They show that, in a three-period model, the optimal rule is always narrower in comparison to the optimal myopic rule (that is, the optimal rule for $T = 2$).[6] To illustrate this principal, consider Figure 1.4. This figure reports the results from a numerical exercise, in which we solve the optimal rule $d_2(\omega_1)$ while allowing for $T = 3$ and $\delta = 1$. This is equivalent to Ellison and Holden's Figure 4 (using our different parameterisation for f).

In my view, this result is both intuitive and fascinating. In particular, there seem to be many real-world scenarios in which decision-makers face trade-offs between the precedents they create today and the decisions they might like to take tomorrow. Ellison and Holden suggest that the result might provide one justification for 'starting small' (page 18):

> There is option value in developing rules and consequently rules can be under-inclusive, particularly early on (recall Proposition 2)... Early in its life, a firm may decline opportunities that would be profitable in the short run because accepting the opportunity would establish a less-than-ideal precedent.

[6] There is one important exception to this: namely, the case in which the myopic rule has no 'excess breadth'. In our parameterisation, this occurs for the special cases $\omega_1 = 0.25$ and $\omega_1 = 0.75$.

Figure 1.4: **Optimal rule choice and indirect utility for** $T = 3$

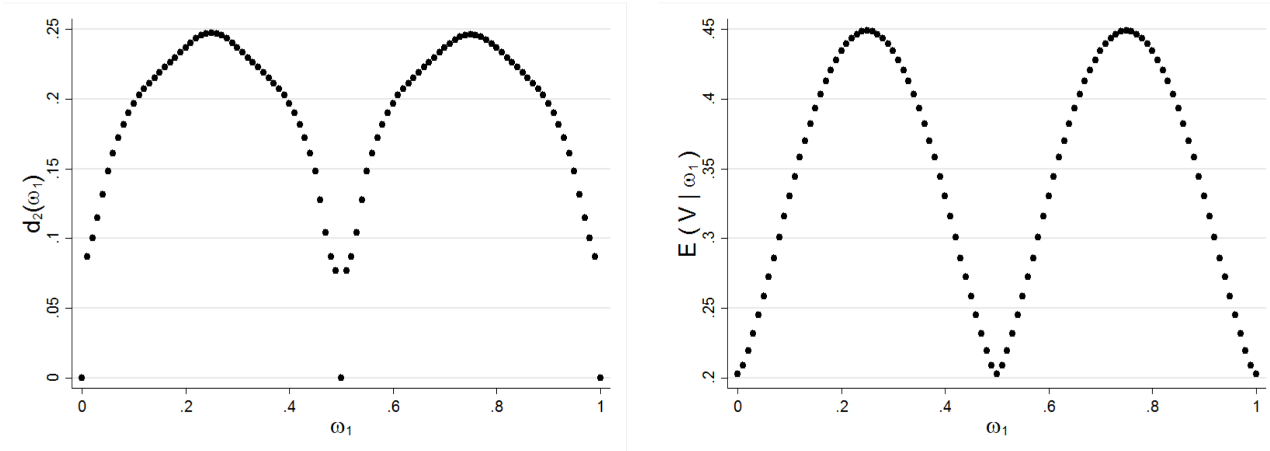

Similarly, the idea has clear resonance in understanding judicial decision-making. An old legal saying warns that 'hard cases make bad law'.[7] And so it is in Figure 1.4; decisions closest to the indifference points (that is, decisions near $\omega_1 = 0$ and $\omega_1 = 0.5$) should optimally generate narrower rules, lest they generate 'bad law' for future decisions.

1.1.3 An infinite horizon

Finally, let's consider an infinite horizon version of the model. For simplicity, we will consider the purely myopic case ($\delta = 0$); that is, we will consider a principal who has no concern about the future, but who nonetheless is constrained by her previous decisions through the 'no overwriting' restriction.

Following Ellison and Holden, we can understand the asymptotic behaviour of such a rule book through simulation. Specifically, we can run the following algorithm:

(i) Randomly draw ω_t.

(ii) Use ω_t to update the rule book, in order to optimise payoffs in the current period.

(iii) Return to step 1 until the rule book is 'complete' (*i.e.* until every possible case $\omega \in [0, 1]$ is covered by the rule book).

(iv) Calculate the expected value of the rule book (that is, the expected payoff for a single period, using the rules set down in the rule book).

[7] See, for example, the judgment of Justice Oliver Wendall Holmes Jnr in *Northern Securities Co. v United States* 193 US 197.

We can then repeat this exercise many times, and find the distribution of this expected payoff. (In their Proposition 5, Ellison and Holden show that this expected payoff converges in probability to a limiting random variable, π_∞.) Figure 1.5 shows the resulting empirical CDF; this is analogous to Figure 5 in Ellison and Holden (except that they authors show a PDF — and, of course, we use our alternative parameterisation for the function f). As Ellison and Holden note, there is substantial mass to the right of the graph; in many cases, the principal settles upon a rule book that is reasonably close to the maximum value attainable. However, this is far from universal; as the authors also note (page 23):

> ... there is a substantial amount of mass at much lower efficiency levels. This highlights that our model predicts substantial path dependence. A series of early "bad" realizations — in the sense that they lead the firm to promulgate substantially inefficient rules — have a persistent effect. Extremely inefficient rule books arise only rarely in this specification — they only occur when the Principal receives multiple unlucky draws.

Figure 1.5: **Empirical CDF of π_∞ under myopia**

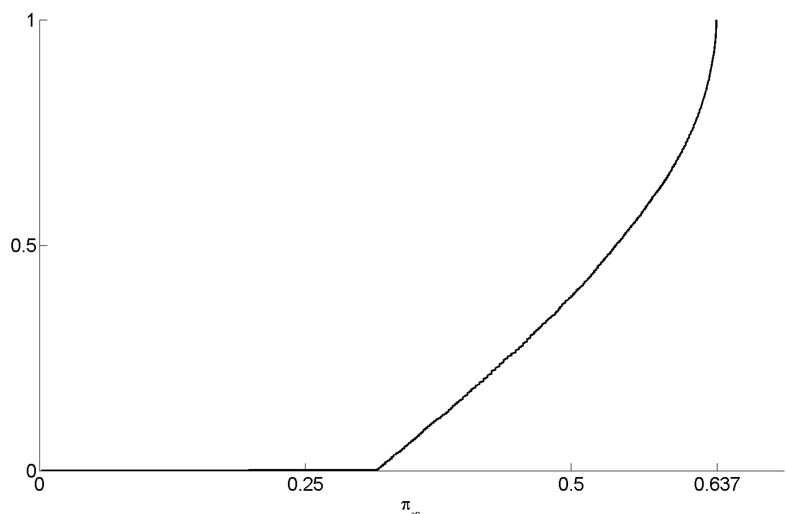

1.2 Empirical method: Measuring management practices

Over the past decade, economists have made substantial progress in understanding the role of management practices in organisational performance. This work has been led by Nick Bloom, John van Reenan, and co-authors at the World Management Survey. To conclude this lecture, we will discuss two recent empirical papers on management practices: one a survey article covering recent work, and the other a seminal experimental intervention.

You may find it useful to consider the following questions as you read the survey article by Bloom *et al* (2014).

(i) Bloom *et al* argue (page 843):

> There are a huge number of case studies discussing the importance of management, mostly focusing on CEOs of top corporations. Much can be learned from case studies in the formulation of hypotheses and the understanding of theories and mechanisms. They are wonderful tools for teaching, but they are poor tools for hypothesis testing.

Why are case studies a poor tool for hypothesis testing? How does the World Management Survey seek to overcome those weaknesses?

(ii) What is the response rate in the World Management Survey? Should we worry about this?

(iii) Consider the list of management practices in Appendix Table A1. What are the three areas on which the survey focuses? Which management practices fall into each area?

(iv) In your opinion, which of these map most neatly into the Ellison-Holden framework? Which map most awkwardly?

(v) What can the authors say about the separate effects of different management practices?

(vi) The World Management Survey has separate questionnaires for manufacturing firms, retailers, hospitals and schools. You can find these questions (indeed, you can 'benchmark your organisation') at http://worldmanagementsurvey.org/. How do the questionnaires differ? Do you agree with the way that questions are varied between different types of organisation? Why or why not?

1.3 Empirical analysis: An experiment to change management

You may find it useful to consider the following questions as you read the experimental results of Bloom *et al* (2013) and Bloom *et al* (2018).

(i) The authors motivate their experiment as follows (page 3):

> A growing literature measures many such [management] practices and finds large variations across establishments and a strong association between these practices and higher productivity and profitability. However, such correlations may be potentially misleading. For example, profitable firms may simply find it easier to adopt better management practices.

Explain this quote, particularly in light of the Lucas (1978) framework that we considered in our earlier lecture on 'Firm Size and the 'Missing Middle''.

(ii) How much did the intervention cost per firm? How much was it worth to the treated firms?

(iii) The intervention raised productivity by 17% in the first year. How?

(iv) What did control plants receive? What implications, if any, does this have for our interpretation of estimated treatment effects?

(v) Why would the number of male family members have such high explanatory power for firm size?

(vi) What do the authors mean by a 'permutation test'? How would such a test be implemented in this context? What is the advantage of such a test over a traditional t-test?

(vii) How many firms were contacted? How many expressed interest in receiving free consulting services? Why did the authors then limit attention to those firms expressing interest? What implications does this have for interpreting the results?

(viii) Which areas did the consultants identify to focus on? How closely — if at all — do these areas compare to the spirit of the Ellison-Holden model? How closely — if at all — do these areas compare to those used for the World Management Survey?

(ix) How was the randomisation implemented? (That is, did the authors stratify or re-randomise?) Why?

(x) The authors describe equation 1 as an 'intention to treat (ITT)' specification. What do they mean by this? What is the implication of this for interpreting the authors' estimates?

(xi) Do management interventions last?

2 Lecture 2: Organisational Hierarchy

- $\boxed{\star}$ Bloom, N., Sadun, R., and Van Reenen, J. (2012): "The Organization of Firms Across Countries," *The Quarterly Journal of Economics*, 127(4), 1663-1705.
- $\boxed{\star}$ Garicano, L. (2000): "Hierarchies and the Organization of Knowledge in Production," *Journal of Political Economy*, 108(5), 874-904.
- $\boxed{\star}$ Hayek, F.A. (1945): "The Use of Knowledge in Society," *The American Economic Review*, 35(4), 519-530.
- $\boxed{\star}$ Rasul, I. and Rogger, D. (2017): "Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service," *Economic Journal*, 128(608): 413-446.
- Rasul, I., Rogger, D. and Williams, M. (2018): "Management and Bureaucratic Effectiveness: Evidence from the Ghanaian Civil Service," *Working Paper*.
- Fafchamps, M., and Söderbom, M. (2006): "Wages and Labor Management in African Manufacturing," *Journal of Human Resources*, 41(2), 356-379.
- Honig, D. (2018): "Navigation by Judgment: Why and When Top Down Management of Foreign Aid Doesn't Work," *Oxford University Press*.
- Lemos, R. and Scur, D. (2018): "All in the Family? CEO Choice and Firm Organization," *Working Paper*.
- Macchiavello, R., Menzel, A., Rabbani, A., and Woodruff, C. (2016): "Challenges of Change: An Experiment Training Women to Manage in the Bangladeshi Garment Sector," *Working paper*.

2.1 'The Use of Knowledge in Society'

> If we can agree that the economic problem of society is mainly one of rapid adaptation to changes in the particular circumstances of time and place, it would seem to follow that the ultimate decisions must be left to the people who are familiar with these circumstances, who know directly of the relevant changes and of the resources immediately available to meet them. We cannot expect that this problem will be solved by first communicating all this knowledge to a central board which, after integrating *all* knowledge, issues its orders. We must solve it by some form of decentralization.

So wrote the Austrian economist Friedrich Hayek in 1945 (page 524), in what has become one of the most famous economics papers of all time.[8] In many respects, Hayek's

[8] See, for example, Arrow, Bernheim, Feldstein, McFadden, Poterba and Solow (2011). In October 2018, Hayek's paper has 16,026 citations on Google Scholar.

concern was to argue against central planning of an entire economy; for example, Hayek devotes a large share of his paper to discussing the importance of the decentralised price system. However, Hayek's fundamental point about the nature of knowledge is profound in itself.[9] If we think that important information is dispersed between a large number of separate individuals — as Hayek puts it, 'a body of very important but unorganised knowledge which cannot possibly be called scientific in the sense of knowledge of general rules' (page 521) — then every organisation will face a profound challenge of deciding its ideal hierarchical structure. Who should know what? What kinds of decisions should be made at each level of the organisation? How should lines of communication work within an organisation? As Hayek emphasises, this is very different to the usual kind of allocation problem with which economists are familiar. This is the problem that we will discuss in this lecture.

2.2 'Hierarchies and the Organization of Knowledge in Production'

In this paper, Garicano provides a framework for thinking about optimal organisational design, in a context in which it is costly to acquire information, and costly to acquire help from other parts of the organisation. Importantly, Garicano does not consider incentive issues. In one respect, this may seem unsatisfying — however, as Garicano emphasises (page 875), literature that focuses on incentive problems must generally treat 'hierarchical organization forms [as] assumed rather than obtained from the theory'. Following Garicano, we will focus on hierarchy in this lecture; we will consider issues of incentives in Lecture 3.

2.2.1 A single production worker

Let's start by considering a single production worker. This worker faces a variety of production challenges, which we will describe using the continuous random variable $Z \geq 0$ (such that we will denote an individual realisation of Z as z). For simplicity, we will assume throughout that Z has an exponential distribution; that is,

$$f(z) = \lambda \cdot \exp(-\lambda \cdot z) \tag{2.1}$$

$$F(z) = 1 - \exp(-\lambda \cdot z). \tag{2.2}$$

As Garicano explains (p.886), λ "uniquely determines the characteristics of the production environment: a higher λ is always preferred since it implied a more "predictable" environment."

Assume that the worker can solve problems $z \in [0, a]$. If $Z \in [0, a]$, the worker produces output of 1; otherwise, the worker produces output of 0. Before beginning production, the worker has a choice: *what kinds of problems should I know how to solve?* We allow the

[9] Hayek expanded on these ideas in 'Rules and Order', published in 1973 (and reprinted as part of his *Law, Legislation and Liberty*).

worker to pay cost ca to master problems $z \in [0, a]$ (where we assume that $c < \lambda$).

The solution to this problem is straightforward. The worker's net payoff is:

$$\mathbb{E}(y) = \Pr(Z < a) - ca = 1 - \exp(-\lambda \cdot a) - ca. \tag{2.3}$$

You should verify that the optimal choice of a is:

$$a^* = \frac{-1}{\lambda} \cdot \ln\left(\frac{c}{\lambda}\right). \tag{2.4}$$

This should be intuitively straightforward: as Garicano explains (p.878), "the worker learns those problems that are common enough to justify their learning costs and ignores the rest". The comparative statics are simple and intuitive: if $c \geq \lambda$, the worker learns nothing (and produces nothing); as c decreases, the worker learns more tasks; and in the limit as $c \to 0$, the worker learns everything. Indeed (remembering that $c < \lambda$), it is straightforward to show that the proportion of problems that the worker *cannot* solve is:

$$1 - \Pr(Z < a) = \frac{c}{\lambda}. \tag{2.5}$$

2.2.2 Introducing the organisation

We now introduce the notion of an organisation. Formally, Garicano describes an organisation as:

(i) A partition of workers into L classes (such that the proportion of workers in class i is β_i);

(ii) For each class i, a set of problems that the class can solve;

(iii) For each class i, a list of other classes that members of class i may ask for solutions; and

(iv) For each class i, a proportional allocation of time, either to helping other classes (t_i^h) or to producing (t_i^p).

We preserve the earlier assumption that each worker must pay a cost proportional to the measure of the problems that it can solve; thus, for example, if a worker can solve problems $z \in [\underline{a}, \overline{a}]$, the cost is $c \cdot (\overline{a} - \underline{a})$. We assume that a worker faces a 'helping cost', h, incurred by being asked by another worker for the answer to a particular problem (*i.e.* "regardless of whether the solution is known or not": p.879). As Garicano explains (p.875):

> The starting point is the observation that production requires physical resources and knowledge about how to combine them. If communication is

> available, workers do not need to acquire all the knowledge necessary to produce. Instead, they may acquire only the most relevant knowledge and, when confronted with a problem they cannot solve, ask someone else. The organization must then decide who must learn what and whom each worker should ask when confronted with an unknown problem.

For simplicity, we assume that the organisation is very large (*i.e.* infinitely large, in the limit), with an infinite number of classes. There is no leisure in the model — so, trivially, note that $t_i^h + t_i^p = 1$.

Garicano shows that any optimal organisational hierarchy has the following characteristics (p.880):

(i) "Workers specialize either in production or in solving problems. Only one class specializes in production."

(ii) "Knowledge acquired by different classes does not overlap."

(iii) "Production workers learn to solve the most common problems; problem solvers learn the exceptions. Moreover, the higher up in the list of production workers a problem solver is, the more unusual the problems she is able to solve. Information in the form of solutions to problems always flows in the same direction, from the highest to the lowest level, since this minimizes communication costs."

(iv) "The organization has a pyramidal structure, with each layer a smaller size than the previous one."

I leave it to you to follow Garicano's discussion and reasoning on this points (*i.e.* his Propositions 1 to 4); our focus will lie instead on the optimal design problem *subject to* these four characteristics holding. To understand the intuition behind this model and its solution, you may find it useful to play with the accompanying Matlab code, `RunGaricano.m`.

2.2.3 The organisational design problem

Given the four characteristics outlined earlier, the organisation's design problem becomes the following:

$$\max_{\{z_0, z_1, \ldots\}, \{\beta_0, \beta_1, \ldots\}} F\left(\sum_{i=0}^{\infty} z_i\right) \cdot \beta_0 - \sum_{i=0}^{\infty} c\beta_i z_i \tag{2.6}$$

That is, the organisation maximises the output of workers (the first term), less the cost of investing in problem-solving capacity (the second term).

This is subject to several constraints. First, we know that $\sum_{i=0}^{\infty} \beta_i = 1$. Second, for each separate class of problem-solvers (that is, each $i > 0$), it must be the case that the class is fully occupied in trying to solve problems that lower levels were unable to solve. That is, for every class of managers, it must be that the total time spent is the total time required to answer all of the problems passed up the hierarchy. Per worker, the probability that a problem is sufficiently complicated to reach level i is:

$$1 - F\left(\sum_{j=0}^{i-1} z_j\right). \tag{2.7}$$

The number of problems (if you like, the *mass* of problems) is simply this number multiplied by the mass of production workers:

$$\left[1 - F\left(\sum_{j=0}^{i-1} z_j\right)\right] \cdot \beta_0. \tag{2.8}$$

Finally, each problem takes time h to solve; so the total time required to solve problems at level i must then be:

$$\left[1 - F\left(\sum_{j=0}^{i-1} z_j\right)\right] \cdot \beta_0 \cdot h. \tag{2.9}$$

Therefore, for *each* separate class $i > 0$, we have the following time constraint:[10]

$$\left[1 - F\left(\sum_{j=0}^{i-1} z_j\right)\right] \cdot \beta_0 \cdot h = \beta_i, \tag{2.10}$$

$$\Rightarrow \exp\left(-\lambda \cdot \sum_{j=0}^{i-1} z_j\right) \cdot \beta_0 \cdot h = \beta_i. \tag{2.11}$$

Now, this looks complicated, because it is written as an optimisation *both* over (i) knowledge acquired and (ii) the size of each layer. However, we can simplify this to substitute for the β_i values, to create an unconstrained optimisation problem in the values for z.

First, let's use equation 2.11 to substitute into the objective function for 'β_i'; our objective function then becomes:

$$F\left(\sum_{i=0}^{\infty} z_i\right) \cdot \beta_0 - c\beta_0 z_0 - \sum_{i=1}^{\infty} chz_i \exp\left(-\lambda \cdot \sum_{j=0}^{i-1} z_j\right) \beta_0 \tag{2.12}$$

$$= \beta_0 \cdot \left[F\left(\sum_{i=0}^{\infty} z_i\right) - cz_0 - \sum_{i=1}^{\infty} chz_i \exp\left(-\lambda \cdot \sum_{j=0}^{i-1} z_j\right)\right] \tag{2.13}$$

[10] This is Garicano's equation 7. Note that there is a typo in the sentence preceding this equation in the paper; Garicano gives the correct formation just after his equation 5.

Next, let's substitute for β_0. By summing equation 2.11 across all $i > 0$ and adding β_0 to both sides, we obtain:

$$1 = \beta_0 + \sum_{i=1}^{\infty} \exp\left(-\lambda \cdot \sum_{j=0}^{i-1} z_j\right) \cdot h\beta_0 \tag{2.14}$$

$$\therefore \beta_0 = \frac{1}{1 + h\sum_{i=1}^{\infty} \exp\left(-\lambda \cdot \sum_{j=0}^{i-1} z_j\right)} = \frac{1}{1 + h\sum_{i=0}^{\infty} \exp\left(-\lambda \cdot \sum_{j=0}^{i} z_j\right)} \tag{2.15}$$

Therefore, our organisation maximises the following objective function:

$$\frac{F\left(\sum_{i=0}^{\infty} z_i\right) - cz_0 - \sum_{i=1}^{\infty} chz_i \exp\left(-\lambda \cdot \sum_{j=0}^{i-1} z_j\right)}{1 + h\sum_{i=0}^{\infty} \exp\left(-\lambda \cdot \sum_{j=0}^{i} z_j\right)}. \tag{2.16}$$

We need to find the first-order conditions. Note that z_0 enters in a different way than for $z_1, z_2, \ldots$; we therefore need to express one first order condition with respect to z_0 and one for z_k for $k > 0$. You should be able to differentiate this to obtain the first order conditions.[11] First, for z_0, you should obtain the following:

$$f\left(\sum_{i=0}^{\infty} z_i\right) \underbrace{-c}_{\text{worker learning}} + \underbrace{\lambda ch \sum_{i=0}^{\infty} z_{i+1} \exp\left(-\lambda \sum_{j=0}^{i} z_j\right)}_{\text{higher learning}} + \underbrace{y^* h\lambda \sum_{i=0}^{\infty} \exp\left(-\lambda \cdot \sum_{j=0}^{i} z_j\right)}_{\text{increased production time}} = 0. \tag{2.17}$$

For z_k for $k > 0$, you should obtain:

$$f\left(\sum_{i=0}^{\infty} z_i\right) \underbrace{-ch \exp\left(-\lambda \sum_{i=0}^{k-1} z_i\right)}_{\text{learning by } k} + \underbrace{\lambda ch \sum_{i=k}^{\infty} z_{i+1} \exp\left(-\lambda \sum_{j=0}^{i} z_j\right)}_{\text{higher learning}} + \underbrace{y^* h\lambda \sum_{i=k}^{\infty} \exp\left(-\lambda \cdot \sum_{j=0}^{i} z_j\right)}_{\text{increased production time}} = 0. \tag{2.18}$$

At this point, it is important to note a very useful property of the exponential distribution; namely, it is '*memoryless*'. Formally, this means that, for any non-negative real numbers s and t,

$$\Pr(Z > t + s \mid X > t) = \Pr(Z > s). \tag{2.19}$$

In the context of our problem, this implies that the firm should optimally choose $z_k = z_{k+1}$ for $k > 0$. As Garicano explains (page 886), "The value of the extra layer is given by the conditional probability that the problem solution is found in that layer given that it was not in the previous layers...". This is constant across all layers $k > 0$ (see, for example, Garicano's footnote 6).

[11] A little help...Consider a general problem of the form $y^* = \max u \cdot v^{-1}$. By the quotient rule, we know that the first derivative gives $(u'v - uv')/v^2 = 0$, implying $(u'v - uv') = 0$. We can divide by v, to get $u' - y^* \cdot v' = 0$. This is the basic form of the results considered here.

2.2.4 Solving the organisational design problem

Therefore, we actually need only to solve for two objects of interest: the knowledge of production workers (which we will term z_w) and the knowledge of each level of problem solver (which we will term z_s). Let's rewrite the two first-order conditions (where, for the second condition, we will set $k = 1$):[12]

$$-c + \lambda ch \cdot z_s \cdot \exp(-\lambda z_w) + \lambda ch \sum_{i=1}^{\infty} z_s \cdot \exp(-\lambda z_w - \lambda \cdot i \cdot z_s)$$
$$+ y^* h\lambda \cdot \exp(-\lambda z_w) + y^* h\lambda \sum_{i=1}^{\infty} \exp(-\lambda z_w - \lambda \cdot i \cdot z_s) = 0 \tag{2.20}$$
$$-ch \exp(-\lambda z_w) + \lambda chz_s \exp(-\lambda z_w) \sum_{i=1}^{\infty} \exp(-\lambda \cdot z_s \cdot i)$$
$$+ y^* h\lambda \exp(-\lambda z_w) \cdot \sum_{i=1}^{\infty} \exp(-\lambda z_s \cdot i) = 0. \tag{2.21}$$

Similarly, we can also write explicitly the value of the organisation's objective:

$$y^* = \frac{1 - cz_w - chz_s \sum_{i=1}^{\infty} \exp(-\lambda z_w + \lambda z_s - \lambda z_s i)}{1 + h \exp(-\lambda z_w) \cdot (1 + \sum_{i=1}^{\infty} \exp(-\lambda z_s \cdot i))}. \tag{2.22}$$

From these three equations, you should be able to solve for z_w^* and z_s^*; I show in the appendix to this chapter how we can do this. We obtain the following solution:

$$z_s^* = \frac{1}{\lambda} \cdot \ln\left(\frac{\lambda}{c} - \ln h\right) \tag{2.23}$$
$$z_w^* = z_s^* + \frac{1}{\lambda} \cdot \ln h \tag{2.24}$$

From these solutions, we can solve for the 'span of control' (s) — that is, the ratio of problem-solvers at level $i > 0$ to problem-solvers at level $i + 1$. Substituting from equation 2.11, we obtain:

$$s \equiv \frac{\exp(-\lambda z_w - \lambda(i-1)z_s) \cdot h\beta_0}{\exp(-\lambda z_w - \lambda i z_s) \cdot h\beta_0} = \exp(\lambda z_s) = \frac{\lambda}{c} - \ln h. \tag{2.25}$$

So, what does this all mean? You should check that you understand the comparative statics with respect to λ, c and h in equations 2.23, 2.24 and 2.25.

To see intuitively what these comparative statics mean, let's visualise some solutions for particular parameter values. To do this, we need to figure out (i) the proportion of

12 What happened to the first term in each first-order condition? Why?

workers at each level (β_i), and (ii) the proportion of problems solved 'at or below' each level. For the former, we can return to equation 2.15:

$$\beta_0 = \frac{1}{1 + h\sum_{i=0}^{\infty} \exp\left(-\lambda \cdot \sum_{j=0}^{i} z_j\right)} = \frac{1}{1 + h\exp(-\lambda z_w) + h\exp(-\lambda z_w)\sum_{i=1}^{\infty}\exp(-\lambda i z_s)} \tag{2.26}$$

$$= \frac{1}{1 + h\exp(-\lambda z_w) + \dfrac{h\exp(-\lambda z_w)}{\exp(\lambda z_s) - 1}}. \tag{2.27}$$

Equation 2.11 then tells us how to solve β_i for each $i > 0$:

$$\beta_i = \exp\left(-\lambda \cdot \sum_{j=0}^{i-1} z_j\right) \cdot h\beta_0 \tag{2.28}$$

$$= \exp\left(-\lambda z_w - \lambda(i-1) \cdot z_s\right) \cdot h\beta_0. \tag{2.29}$$

Finally, of course, the proportion of problems solved by production workers is:

$$1 - \exp(-\lambda \cdot z_w), \tag{2.30}$$

and the proportion of problems solved by production workers and the next i layers of problem solvers is:

$$1 - \exp(-\lambda \cdot z_w - \lambda \cdot i \cdot z_s). \tag{2.31}$$

To understand this solution intuitively, I encourage you to open Matlab and play with the provided graphical user interface ('`RunGaricano.m`'). What is the effect of changing h, c and λ? Under what conditions will the organisational hierarchy be flattest? When will it be steepest? Under what conditions do we lose the interior solution? Why?

2.3 Empirical analysis: Firm structure and trust

How do these predictions apply to the real world? To consider this, we will discuss the empirical results of Bloom, Sadun and Van Reenan (2012).

(i) Bloom *et al* present a very stylised extension of the Garicano model, explicitly incorporating a role for 'trust'. But suppose that we want to use the original Garicano model to make predictions on trust. How might we capture a notion of trust in the parameters h, c and λ? What predictions would we obtain about trust and hierarchy? How, if at all, do these predictions differ from the model extension presented by Bloom *et al*?

(ii) Bloom *et al* test for the effect of trust on firm size. What role does firm size play in the Garicano model?

(iii) How do the authors measure decentralisation? (Note that 'Online Appendix Table A1' is available on page 13 of `https://people.stanford.edu/nbloom/sites/default/files/orgappendix.pdf`.)

(iv) What was the response rate? Should we worry about this?

(v) How do the authors measure trust? At what geographical level?

(vi) The authors find a relationship between trust and decentralisation. At what level does this relationship exist — across countries, within countries, or both?

(vii) The authors explain (page 1664): 'To probe whether this effect is causal, we exploit the fact that some of our data are drawn from multinational subsidiaries.' How do the authors do this? What identifying assumption is necessary in order to give such results a causal interpretation?

(viii) Is this a paper about development? Why or why not?

2.4 Empirical analysis: Hierarchy in the Nigerian civil service

We began this module by noting that principles of organisations and development are more general than merely the internal organisation of firms — however important firms clearly are to the topic. To conclude our second lecture, we will discuss the seminal recent work of Rasul and Rogger (2018) — who study the role of management practices in a civil service, rather than in a firm.

(i) The authors motivate their paper by saying: "We thus provide among the first large-scale descriptive evidence on whether the management practices bureaucrats operate under, correlate to the quantity and quality of public services delivered." To what extent is this a test of a causal relationship?

(ii) The authors use data on 'over 4700 public sector projects that began in 2006/7' (page 2); they describe these as mostly 'small-scale rural infrastructure project[s]' (page 3). How should we think about the paper's results generalising to other kinds of project?

(iii) The authors focus on two dimensions of management in particular. What are they?

(iv) How do the authors measure project completion?

(v) How do the authors measure bureaucratic management practices?

(vi) What do the authors find on the role of 'autonomy'? How closely does this compare to Garicano's notion of hierarchy?

(vii) What do the authors find on the role of 'incentives/monitoring'? Is there a good explanation for this result?

(viii) Is this evidence of 'suboptimal management practices'? If so, why might such practices persist?

2.5 Appendix: Solving the organisational design problem

First, we need to remember the following rule about the sum of an exponential series (for $a > 0$):

$$\sum_{i=1}^{\infty} \exp(-a \cdot i) = \frac{1}{\exp(a) - 1}. \tag{2.32}$$

From equation 2.20, we obtain:

$$-c \cdot (1 - \exp(-\lambda z_s)) + \lambda ch \cdot z_s \cdot \exp\left(-\lambda z_w\right) + y^* h\lambda \cdot \exp\left(-\lambda z_w\right) = 0 \tag{2.33}$$

$$\therefore -\frac{\exp(\lambda z_w) \cdot (1 - \exp(-\lambda z_s))}{\lambda h} + z_s + \frac{y^*}{c} = 0 \tag{2.34}$$

From equation 2.21:

$$-ch \exp(-\lambda z_w) + \lambda chz_s \exp(-\lambda z_w) \cdot \frac{1}{\exp(\lambda z_s) - 1} + y^* h\lambda \exp(-\lambda z_w) \cdot \frac{1}{\exp(\lambda z_s) - 1} = 0 \tag{2.35}$$

$$\therefore -c + \lambda cz_s \cdot \frac{1}{\exp(\lambda z_s) - 1} + y^* \lambda \cdot \frac{1}{\exp(\lambda z_s) - 1} = 0 \tag{2.36}$$

$$\therefore \frac{1 - \exp(\lambda z_s)}{\lambda} + z_s + \frac{y^*}{c} = 0. \tag{2.37}$$

Substituting, we obtain:

$$-\frac{\exp(\lambda z_w) \cdot (1 - \exp(-\lambda z_s))}{\lambda h} = \frac{1 - \exp(\lambda z_s)}{\lambda} \tag{2.38}$$

$$\therefore \exp(\lambda z_w) \cdot (1 - \exp(\lambda z_s)) = h \cdot (1 - \exp(\lambda z_s)) \cdot \exp(\lambda z_s) \tag{2.39}$$

$$\therefore \exp(\lambda z_w) = h \cdot \exp(\lambda z_s) \tag{2.40}$$

$$\therefore z_w = z_s + \frac{1}{\lambda} \cdot \ln h \tag{2.41}$$

We can now substitute back into equation 2.22:

$$y^* = \frac{1 - cz_w - chz_s \cdot \dfrac{\exp\left(-\lambda z_w + \lambda z_s\right)}{\exp(\lambda z_s) - 1}}{1 + h\exp(-\lambda z_w) \cdot \left(\dfrac{1}{1 - \exp(-\lambda z_s)}\right)} \tag{2.42}$$

$$= \frac{1 - cz_w - chz_s \cdot \dfrac{\exp\left(-\lambda z_w + \lambda z_s\right)}{\exp(\lambda z_s) - 1}}{1 + h\exp(-\lambda z_w + \lambda z_s) \cdot \left(\dfrac{1}{\exp(\lambda z_s) - 1}\right)} \tag{2.43}$$

$$= \frac{\exp(\lambda z_s) - 1 - cz_w \cdot \exp(\lambda z_s) + cz_w - chz_s \cdot \exp\left(-\lambda z_w + \lambda z_s\right)}{\exp(\lambda z_s) - 1 + h\exp\left(-\lambda z_w + \lambda z_s\right)} \tag{2.44}$$

$$= \frac{1 - \exp(-\lambda z_s) - cz_w + cz_w\exp(-\lambda z_s) - chz_s \cdot \exp\left(-\lambda z_w\right)}{1 - \exp(-\lambda z_s) + h\exp\left(-\lambda z_w\right)}. \tag{2.45}$$

Equation 2.41 implies that the denominator to this last expression is 1. Therefore, we can write:

$$y^* = 1 - \exp(-\lambda z_s) - cz_w + cz_w\exp(-\lambda z_s) - chz_s \cdot \exp\left(-\lambda z_w\right) \tag{2.46}$$

Now, substituting into equation 2.37, we obtain:

$$\frac{1 - \exp(\lambda z_s)}{\lambda} + z_s + \frac{1 - \exp(-\lambda z_s)}{c} - z_w + z_w\exp(-\lambda z_s) - hz_s \cdot \exp\left(-\lambda z_w\right) = 0 \tag{2.47}$$

So let's substitute in for z_w using equation 2.41:

$$\frac{1 - \exp(\lambda z_s)}{\lambda} + \frac{1 - \exp(-\lambda z_s)}{c} - \frac{1}{\lambda} \cdot \ln h + \frac{1}{\lambda} \cdot \ln h \exp(-\lambda z_s) = 0 \tag{2.48}$$

$$\therefore c \cdot (1 - \exp(\lambda z_s)) + \lambda \cdot (1 - \exp(-\lambda z_s)) - c\ln h \cdot (1 - \exp(-\lambda z_s)) = 0 \tag{2.49}$$

$$\therefore c \cdot \exp(\lambda z_s) \cdot (1 - \exp(\lambda z_s)) - \lambda \cdot (1 - \exp(\lambda z_s)) + c\ln h \cdot (1 - \exp(\lambda z_s)) = 0 \tag{2.50}$$

$$\therefore c \cdot \exp(\lambda z_s) - \lambda + c\ln h = 0 \tag{2.51}$$

$$\therefore \exp(\lambda z_s) = \frac{\lambda}{c} - \ln h \tag{2.52}$$

So we're done; it follows straightforwardly that:

$$z_s^* = \frac{1}{\lambda} \cdot \ln\left(\frac{\lambda}{c} - \ln h\right) \tag{2.53}$$

$$\therefore z_w^* = \frac{1}{\lambda} \cdot \ln\left(\frac{\lambda}{c} - \ln h\right) + \frac{1}{\lambda} \cdot \ln h \tag{2.54}$$

$$= \frac{1}{\lambda} \cdot \ln\left(\frac{h\lambda}{c} - h\ln h\right). \quad \blacksquare \tag{2.55}$$

3 Lecture 3: Incentives within Organisations

- $\boxed{\star}$ Atkin, D., Chaudhry, A., Chaudhry, S., Khandelwal, A., and Verhoogen, E. (2017): "Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan," *Quarterly Journal of Economics*, 132(3), 1101–1164.

- $\boxed{\star}$ Bertrand, M., Burgess, R., Chawla, A. and Xu, G. (2018): "The Glittering Prizes: Career Incentives and Bureaucrat Performance," *Working paper*.

- $\boxed{\star}$ Callen, M., Gulzar, S., Hasanaian, A. and Khan, M.Y. and Rezaee, A. (2018): "Data and Policy Decisions: Experimental Evidence from Pakistan," *Working paper*.

- Breza, E., Kaur, S., and Shamdasani, Y. (2018): "The Morale Effects of Pay Inequality," *Quarterly Journal of Economics*, 133(2), 611-663.

- Crawford, V. and Sobel, J. (1982): "Strategic Information Transmission," *Econometrica*, 50(6), 1431-1451.

- Dal Bó, E., Finan, F., Li, N. and Schechter, L. (2018): "Government Decentralisation under Changing State Capacity: Experimental Evidence from Paraguay", NBER Working Paper 24879.

- Gibbons, R. and Murphy, K. (1992): "Optimal Incentive Contracts in the Presence of Career Concerns: Theory and Evidence," *Journal of Political Economy*, 100(3), 468–505.

- Habyarimana, J., Khemani, S. and Scot, T. (2018): "Political Selection and Bureaucratic Productivity," *Working paper*.

- Hjort, J. (2014): "Ethnic Divisions and Production in Firms," *The Quarterly Journal of Economics*, 129(4), 1899-1946.

- Khan, A.Q., Khwaja, A.I., and Olken, B.A. (2016): "Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors," *Quarterly Journal of Economics*, 131(1), 219–271.

- Khan, A.Q., Khwaja, A.I., and Olken, B.A. (2018): "Making Moves Matter: Experimental Evidence on Incentivizing Bureaucrats through Performance-Based Transfers," *American Economic Review* (forthcoming).

- Sobel, J. (2013): "Giving and Receiving Advice," in *Advances in Economics and Econometrics*, edited by Acemoglu, D., Arellano, M. and Dekel, E..

We have now spent two lectures discussing organisations – and yet we have not really discussed incentives. In various ways, the models in our past two lectures both considered issues of organisational design under communication costs — whether the 'bounded communication' of the Ellison-Holden model or the 'helping costs' of Garicano. Yet the alignment of incentives is clearly a serious challenge for any organisation. Different actors at different levels of an organisation are likely to have different objectives — and the tension between those objectives can cause serious problems for the performance of the organisation as a whole.

In this lecture, we will consider the challenge of incentivising performance in organisations. Primarily, we will focus on the recent work of Atkin *et al* in Pakistan; we start with a stylised model of incentives and cheap talk from that paper, then discuss the authors' experimental implementation and results. We finish by considering other recent empirical work on incentives within organisations in developing countries.

3.1 A model of technology adoption with cheap talk

The adoption of new technologies provides an ideal context to understand incentive problems: as we saw in Lecture 3 of 'Firms and Development', new technologies can be very beneficial, at least to some adopters. But new technologies can also be disruptive to many individuals within organisations: disturbing their established work patterns, changing their relationship with others in the organisation and, in extreme cases, even making their positions redundant.

Atkin *et al* provides an extremely creative example of an experiment to test technology diffusion. We will discuss these authors' specific empirical results later in the lecture. To motivate our theoretical discussion, we being simply by noting the most important stylised facts from the experiment. First, the authors introduced a new die for cutting pentagons for the manufacture of footballs. Figure 3.1 repeats Figure 1 and Figure A.8 from the original paper; it shows the key innovation.[13] As the authors report (page 1):

> A conservative estimate is that the new die reduces rexine cost per pentagon by 6.76 perccent and total costs by approximately 1 percent — a modest reduction but a non-trivial one in an industry where mean profit margins are 8 percent.

[13] The authors clarify — on page 8 — that 'A two-pentagon variant of our design can be made using the specifications in the blueprint (with the two leftmost and two rightmost pentagons in the blueprint. . . cut separately). This version is easier to maneuver with one hand and can be used with the same cutting rhyhtm as the traditional two-pentagon die. It is the version that has proven more popular with firms.' That is, this is a two-pentagon off-set die, rather than the two-pentagon design shown in Atkin *et al*'s Figure 1.

Figure 3.1: **The new technology**

Atkin *et al*'s Figure 1 | Atkin *et al*'s Figure A.8

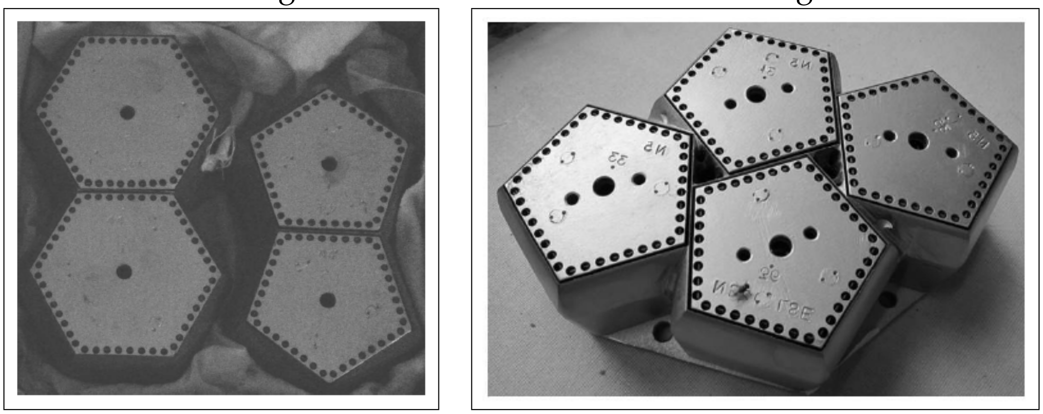

Second, to their surprise, the authors found very limited adoption; of the 35 firms who received the new die in May 2012, only five had adopted by August 2013 (and one firm from the group of 79 firms that received nothing). When asked why they had not adopted, 18 firms responded; of these, 10 listed as their most important reason: 'The cutters are unwilling to work with the offset die' (Table VI).

Thus, a paper that was intended to be about diffusion of adoption became a paper about misalignment of incentives within firms.[14] The authors therefore followed the initial technology-drop experiment with an experiment designed to improve the alignment of incentives within firms. We will discuss this experiment later in the lecture; before we do, we will consider in detail the authors' theoretical model of incentive misalignment. This model is designed to provide a framework for answering two key questions (page 19):

> First, given that owners should be aware of that workers have an incentive to discourage adoption of the offset dies, why are they influenced by what the workers say? Second, why do owners not simply offer a different labor contract, to give workers an incentive to support the adoption [of] the cost-saving technology?

To consider these questions, Atkin *et al* present an elegant model of incentive conflicts and technological change (primarily in their 'Appendix B'). In this model, we consider a single-shot interaction between a principal ('she', representing the organisation), and an

[14] As the authors explain (page 2): "We expected the technology to be adopted quickly by the tech-drop firms, and we planned to focus on spillovers to the cash-drop and no-drop firms; we are pursuing this line of enquiry in a companion project".

agent ('he': the production worker). The principal's payoff is:

$$\pi = p \cdot q - w(q) - c \cdot q, \tag{3.1}$$

where q represents output, c is the constant marginal cost of production, and $w(q)$ is a wage paid to the agent.

The agent produces $q = s \cdot e$, where e represents effort and s is the 'speed of the technology'. We assume that the agent's utility takes a quadratic form:

$$U = w(q) - 0.5e^2. \tag{3.2}$$

The principal cannot contract on e, so instead must condition the wage payment upon output, q. For simplicity, we assume that the principal is constrained to a linear contract structure:

$$w(q) = \beta \cdot q. \tag{3.3}$$

In their paper, Atkin *et al* actually consider a wage of the form $w(q) = \alpha + \beta \cdot q$. They then impose $\alpha \geq 0$. Atkin *et al* describe this as an assumption 'that the agent has limited liability'; they justify the assumption empirically 'given that no worker in our setting pays an owner to work in the factory' (page 1, appendix B).[15] Having assumed that $\alpha \geq 0$, the authors then proceed to show in every case that $\alpha = 0$. Therefore, for simplicity, we drop α from the outset.

The model concerns four separate types of production technologies. The technologies are characterised by their material cost (c) and their speed (s). Notice that, for a given output quantity q, the material cost directly affects the payoff of the *principal*, whereas the speed directly affects the payoff of the *agent*.

The status quo technology is referred to as 'θ_0', having cost c_0 and speed s_0. The authors then consider the introduction of a new technology, which might take one of three types:

(i) *'Type θ_1', having cost $c_1 = c_0$ and $s_1 < s_0$.* This technology has no effect on cost, but is slower than the status quo. Both principal and agent can agree that this technology is inferior.

(ii) *'Type θ_3', having cost $c_3 < c_0$ and $s_3 > s_0$.* This technology has lower cost than the status quo, and is faster; both principal and agent can agree that this technology is superior.

[15] Personally, I am not convinced that this is the right description or a very persuasive justification; there is no variability in q in this model, so if $q >> 0$ in equilibrium, the agent faces no risk of 'liability', and would never have to pay to work in the factory.

(iii) *'Type θ_2', having cost $c_2 < c_0$ and $s_2 < s_0$.* This technology has *lower cost* than the status quo, *but is slower*. The existence of this technology sets up a tension between principal and agent: the principal may wish to adopt,[16] whereas the agent would prefer not to do so.

We will use θ to denote the type of technology (that is, $\theta = 1$, $\theta = 2$ or $\theta = 3$).

The game unfolds as follows:

(i) The principal offers a contract to the agent.

(ii) Nature reveals the type of the new technology to the agent, but not to the principal. Nature draws this from a categorial distribution with probabilities ρ_1, ρ_2 and ρ_3 (such that $\rho_1 + \rho_2 + \rho_3 = 1$); these probabilities are common knowledge among all players.

(iii) The agent can then send a message to the principal about the technology.

(iv) The principal chooses whether or not the firm adopts the new technology.

(v) The agent chooses an effort level.

(vi) Output and the technology type are observed and payoffs are realised.

Before we go further, it is worth noting several points. First, notice that the agent chooses e to maximise:

$$U = \beta s e - 0.5e^2. \tag{3.4}$$

From this, it follows straightforwardly that (i) for a given β, the agent optimally chooses $e = \beta s$, and (ii) the agent's indirect utility is $0.5\beta^2 \cdot s^2$. Thus, as Atkin *et al* explain, "conditional on β the agent prefers faster technologies".

Given this choice, the principal's payoff from adopting technology θ_i (with cost c_i and speed s_i) is:[17]

$$\pi_i(\beta) = \beta s_i^2 \cdot (p - \beta - c_i)\,. \tag{3.5}$$

3.1.1 Benchmark 1: The principal is fully informed

Let's start — as Atkin *et al* do — by considering a case in which the principal observes the technology type directly. In that case, the principal's problem is simply to choose β to maximise $\pi_i(\beta)$. You should verify that, trivially, the principal will choose:

$$\beta_i^* = \frac{p - c_i}{2}. \tag{3.6}$$

[16] More on this shortly!

[17] We can ignore the last term in Atkin *et al*'s equation B3; the authors are allowing here for a fixed 'switching cost' of adoption, but nothing depends on this for our purposes.

You should check that you understand the intuition here: why does the principal pay more if price is higher, or cost lower?

So, should the principal adopt the technology, or not? Substituting equation 3.6 into equation 3.5, you should find the following profit:

$$\pi_i(\beta_i^*) = \frac{(p - c_i)^2 \cdot s_i^2}{4}. \tag{3.7}$$

Trivially, the principal will adopt technology 3 (since $c_3 < c_0$ and $s_3 > s_0$), and will not adopt technology 1 (since $c_1 = c_0$ and $s_1 < s_0$). What about technology 2? In principle, this is ambiguous: we know that $c_2 < c_0$ and $s_2 < s_0$, but it is still possible that, even with full information, the agent and principal could agree that technology 2 should not be adopted. We will come back to this shortly.

3.1.2 Benchmark 2: The principal is completely uninformed

Second, let's consider the alternative benchmark — in which the principal is completely uninformed, and in which the agent is unable to send any messages. The principal's expected payoff here is:

$$\tilde{\pi}(\beta) = \sum_{i=1}^{3} \rho_i \cdot \beta s_i^2 \cdot (p - \beta - c_i). \tag{3.8}$$

You should verify that the principal optimally chooses:

$$\tilde{\beta} = \frac{\sum_{i=1}^{3} \rho_i s_i^2 (p - c_i)}{2 \sum_{i=1}^{3} \rho_i s_i^2} = \frac{\sum_{i=1}^{3} \rho_i s_i^2 \cdot \beta_i^*}{\sum_{i=1}^{3} \rho_i s_i^2}. \tag{3.9}$$

So, does the uninformed principal adopt, or not? As in the previous section, the answer is ambiguous — but, as in the previous section, we want to restrict our model to generate an interesting tension. Let's turn to such restrictions now.

3.1.3 Some parameter restrictions

It is important to restrict the model to generate interesting tensions between the principal and agent; after all, there is a little in this context to learn from a model in which the principal and agent always agree on the optimal technology adoption decision.

Atkin *et al* therefore impose three parameter restrictions…

Parameter Restriction 1 $\pi_2(\beta_0) > \pi_0(\beta_0)$
This imposes that 'the principal finds technology 2 more profitable than the status quo, even under the status quo incentive contract'. Notice that $\pi_2(\beta)$ is maximised at $\beta = \beta_2$. Therefore, this parameter restriction also implies that $\pi_2(\beta_2) > \pi_0(\beta_0)$; that is, 'if the principal knows that the true technology is technology 2, she will adopt'.

Parameter Restriction 2 $\pi_3(\beta_2) > \pi_0(\beta_2)$
In effect, we require that 'if the principal were forced to use the incentive for technology 2, she would still prefer to adopt technology 3 over the status quo'.

Initially, this looks like a strange restriction. However, this is useful for several reasons — one being that (as Atkin et al *show in their 'Remark 2'), this implies that:*

$$\pi_3(\beta_0) > \pi_0(\beta_0). \tag{3.10}$$

Parameter Restriction 3 $\pi_0(\beta_0) > \tilde{\pi}(\tilde{\beta})$
This restriction has a straightforward interpretation: 'a principal who is completely uninformed will prefer not to adopt'.

As Atkin *et al* show, Parameter Restriction 2 and Parameter Restriction 3 jointly imply that:

$$\pi_0(\beta_0) \geq \frac{\rho_1}{\rho_1 + \rho_2} \cdot \pi_1(\beta_0) + \frac{\rho_2}{\rho_1 + \rho_2} \cdot \pi_2(\beta_0). \tag{3.11}$$

Figure 3.2 reproduces the authors' Figure B.1; it shows an illustrative set of functions that obey these parameter restrictions.

Figure 3.2: **'Profit functions for different technology types'** (Atkin *et al*'s Figure B.1)

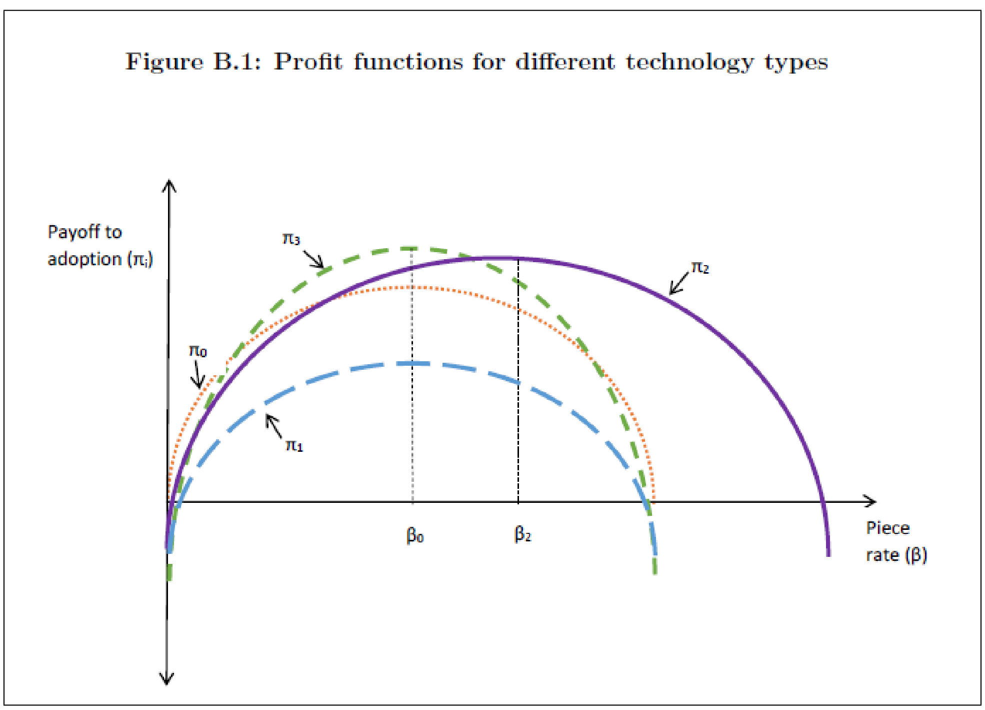

3.1.4 Introducing cheap-talk

Now let's consider the full model — in which the agent is allowed to send a message. In this context, our interaction becomes a 'cheap talk game', of the kind famously studied by Crawford and Sobel (1982). In particular, we seek a Perfect Bayesian Equilibrium. In this game, a Perfect Bayesian Equilibrium is a combination of (i) a strategy for the agent, mapping from the technology type to the message space (that is, $m^*(\theta)$), (ii) a response for the principal, mapping from the message space to an adoption decision (that is, $a^*(m)$) and (iii) a set of posterior beliefs for the principal about the probability that the technology is either type 1, 2, or 3, conditional on the message received (that is, $p(1 \mid m)$, $p(2 \mid m)$ and $p(3 \mid m)$). We require that:

(i) The principal acts as if $\Pr(\theta = k \mid m) = p(k \mid m)$. For any message m that occurs with positive probability on the equilibrium path, the principal must form $p(k \mid m)$ using Bayes' Rule. Otherwise, the principal sets $p(k|m) = \rho_k$; *i.e.* the principal simply uses the prior belief as the posterior.[18]

(ii) The principal chooses $a^*(m)$ to maximise her utility, given her beliefs.

(iii) The agent chooses $m^*(\theta)$ to maximise his utility, given the principal's strategy $a^*(m)$.

Atkin *et al* allow for mixed strategies — both in the agent sending messages and in the principal responding. For simplicity, we will limit ourselves to pure strategies; fortunately, none of the intuition of the model is lost by doing so. With this restriction, the principal either chooses to adopt the new technology or chooses not to do so. Therefore, without further loss of generality, we can restrict the set of agent communication to just two messages:[19]

(i) 'The technology is good' (or any 'equivalent message'), or

(ii) 'The technology is bad' (or any 'equivalent message').

Existence of a 'babbling equilibrium': Suppose, initially, that the principal ignores whatever message the agent sends. In that case, a Perfect Bayesian Equilibrium exists in which the agent simply reports 'the technology is good', irrespective of the technology. Let's check the intuition...

[18] Note that Atkin *et al* take a slightly different approach to messages off the equilibrium path: they assume (in footnote 9 of Appendix B) that 'if the principal receives a message that is off the equilibrium path... she takes one of the actions induced on the equilibrium path for that β'.

[19] Atkin *et al* discuss this in some detail. As they explain, " we describe two messages m and m' as *equivalent* in a given subgame equilibrium if they induce the same action". It is worth noting here the simplicity of the principal's decision; namely, the principal is taking a single binary decision. Therefore, if we were to allow the principal to mix, we could fully characterise the set of messages by allowing only messages of the form 'you should adopt with probability p'.

(i) In such a case, the principal uses $p(1 \mid m) = \rho_1$, $p(2 \mid m) = \rho_2$ and $p(3 \mid m) = \rho_3$.

(ii) By Parameter Restriction 3, the principal refuses to adopt, regardless of the message.

(iii) Therefore, the agent is indifferent between any message sent (and, therefore, can do no better than to say, 'the technology is good').

(iv) Finally, the principal is correctly using Bayes' Rule: whether the principal says 'the technology is good' (*i.e.* a message on the equilibrium path) or any other message (*i.e.* off the equilibrium path), the principal still correctly infers $p(k \mid m) = \rho_k$.

Of course, by the same logic, it is *also* an equilibrium for the agent to report 'the technology is bad' in every case — or, for that matter, for the agent to report *any message*.

This is known as a '*babbling equilibrium*' — in which the agent's message is uninformative of the unknown type. It is a characteristic of cheap-talk games that such a babbling equilibrium always exists (see, for example, Sobel (2013)). To me, the existence of such an equilibrium is profoundly important for understanding the intuition behind cheap-talk games. Usually, when we study Perfect Bayesian Equilibrium and communication, we start by considering *costly* communication — that is, a context in which an agent can send a costly signal (for example, by completing a tertiary degree), in order to distinguish herself from other agents. In cheap-talk games — as the name suggests — the communication is not costly... or, at least, is not *inherently* costly. In such games, the cost of communication arises endogenously, *depending on how the principal interprets the message*. The babbling equilibrium emphasises this, by showing that communication can always 'break down' — in the sense that the principal chooses not to listen to the agent, and the agent therefore can say anything that he pleases.

An informative equilibrium: More interesting, of course, is the case in which communication can be informative. In addition to the babbling equilibrium, there is an equilibrium in which:

(i) The principal sets $\beta^* = 0.5(p - c_0)$.

(ii) The agent:

   (a) Reports 'the technology is bad' if the technology is type θ_1 or type θ_2, and

   (b) Reports 'the technology is good' if the technology is type θ_3.

(iii) The principal adopts if the agent reports 'the technology is good', and does not adopt if the agent reports 'the technology is bad'.

This is Proposition 1 of Atkin *et al*. Atkin *et al* are very careful in proving this proposition; we will take a more intuitive approach. Intuitively, we can reason as follows...

(i) Start by fixing $\beta = \beta_0$ (that is, $\beta = 0.5(p - c_0)$)...

(a) The principal forms her beliefs using Bayes' Rule...

i. If the agent reports 'the technology is good', the principal sets $p(1) = p(2) = 0$ and $p(3) = 1$; and

ii. If the agent reports 'the technology is bad', the principal sets $p(1) = \rho_1 \cdot (\rho_1 + \rho_2)^{-1}$, $p(2) = \rho_2 \cdot (\rho_1 + \rho_2)^{-1}$ and $p(3) = 0$.

iii. If the agent were to deviate to any message other than these two (or equivalent messages), the principal would set $p(k) = \rho_k$.

(b) Given these beliefs, the principal should adopt if the agent reports 'the technology is good'. This is because, as we discussed earlier,

$$\pi_3(\beta_0) \geq \pi_0(\beta_0). \tag{3.12}$$

(c) Conversely, given her beliefs, the principal should not adopt if the agent reports 'the technology is bad'. This is because, we we discussed earlier,

$$\pi_0(\beta_0) \geq \frac{\rho_1}{\rho_1 + \rho_2} \cdot \pi_1(\beta_0) + \frac{\rho_2}{\rho_1 + \rho_2} \cdot \pi_2(\beta_0). \tag{3.13}$$

(d) Given the principal's response, the agent should report 'the technology is bad' if the technology is type θ_1 or type θ_2 and should report 'the technology is good' if the technology is type θ_3. (This follows straightforwardly from our earlier assumptions about the agent's preferences.)

(ii) Now, consider whether the principal would have any incentive to deviate from $\beta = \beta_0$. This requires checking four different cases on possible values for β — depending on (i) whether, for a given β, the principal wants to adopt if the technology is type 3, and (ii) whether, for the same given β, the principal wants to adopt if the technology is drawn from $\theta \in \{\theta_1, \theta_2\}$ (with proportions according to ρ_1 and ρ_2). We will not pause to consider all of these cases; you should check that you understand the reasoning of Atkin *et al* on this step (in section B.2.1.3).

Finally — and to check your intuition about cheap-talk games — suppose that the agent reports 'the technology is good' if the technology is type θ_1 or θ_2, and reports 'the technology is bad' if the technology is type θ_3. How should the principal respond? Can this behaviour support a Perfect Bayesian Equilibrium?

3.2 Allowing for conditional contracts

The previous section illustrates nicely an important organisational problem: namely, reports from employees contain *some* information — and, therefore, form a useful basis for action — but employees may nonetheless succeed in systematically 'persuading' management to take actions that are not in the interests of the organisation as a whole. In this simple model, the problem arises because, if the technology is 'type 2', the principal would like to adopt but the agent would not.

Suppose, instead, that the principal can commit to pay an *additional* per-unit reward in the event that the cost is c_2; that is, such that $w(q) = (\beta + \gamma_2) \cdot q$ if the cost is c_2. What would the principal like to pay? We know from the 'benchmark 1' case earlier that, if the principal knows with certainty that the technology is θ_2, she would ideally like to set $\beta + \gamma_2 = 0.5(p - c_2)$. If the technology is θ_1 or θ_3, she would ideally like to continue to set $\beta = 0.5(p - c_0)$, as before. Therefore, she would ideally like to set γ_2 to 'make up the difference'; that is, $\gamma_2 = 0.5(c_0 - c_2)$.

But what effect does this have on the agent's reporting? Well, from our earlier discussion, we know that, if the principal adopts under technology θ_2, the agent's utility will now be $0.5(\beta + \gamma_2)^2 \cdot s_2^2$. If the principal does not adopt, the agent's utility will be $0.5\beta^2 \cdot s_0^2$. In order for the agent to want the principal to adopt technology θ_2, it must be that:

$$(\beta + \gamma_2)^2 \cdot s_2^2 \geq \beta^2 \cdot s_0^2 \tag{3.14}$$

$$\Leftrightarrow \left(\frac{p - c_2}{2}\right)^2 \cdot s_2^2 \geq \left(\frac{p - c_0}{2}\right)^2 \cdot s_0^2. \tag{3.15}$$

From equation 3.7, we know that this condition is equivalent to saying $\pi_2(\beta_2) \geq \pi_0(\beta_0^*)$ — something that follows immediately from Parameter Restriction 1.

This is a nice result; it shows that, if the principal can pay G to invest in a monitoring technology to check whether the adopted technology is θ_2, then it will be in the agent's interest to encourage adoption of θ_2. Under what circumstances, then, will such investment make sense? Without such a technology, the principal's payoff is:

$$\rho_1 \cdot \pi_0(\beta_0) + \rho_2 \cdot \pi_0(\beta_0) + \rho_3 \cdot \pi_3(\beta_0). \tag{3.16}$$

With the technology, the payoff is:

$$\rho_1 \cdot \pi_0(\beta_0) + \rho_2 \cdot \pi_2(\beta_2) + \rho_3 \cdot \pi_3(\beta_0) - G. \tag{3.17}$$

Trivially, the principal should invest in the technology if:

$$G \leq \rho_2 \cdot \left[\pi_2(\beta_2) - \pi_0(\beta_0)\right]. \tag{3.18}$$

In their section B.3, Atkin *et al* show that, if this condition holds, there is a Perfect Bayesian Equilibrium in which:

(i) The principal sets $\beta^* = 0.5(p - c_0)$ and $\gamma_2 = 0.5(c_0 - c_2)$.

(ii) The agent:

   (a) Reports 'the technology is bad' if the technology is type θ_1, and

(b) Reports 'the technology is good' if the technology is type θ_2 or θ_3.

(iii) The principal adopts if the agent reports 'the technology is good', and does not adopt if the agent reports 'the technology is bad'.

That is, a conditional contract should allow the principal to shift the agent's incentives such that the agent reports in a way that is consistent with the principal's desired adoption behaviour.

3.3 The second experiment: Introducing conditional contracts

Prompted by their puzzling result on non-adoption, Atkin *et al* ran a second experiment, run just among the 31 active 'tech drop' firms.[20] I leave it to you to consider how the authors stratified this randomisation. As for the new treatments, the authors explain (page 23):

> To firms in Group B we gave a reminder about the offset die and the new cutting pattern, and explicitly informed them about the two-pentagon variant of the offset die (which, as noted above, had proven more popular than the four-pentagon offset die we originally distributed.) We also offered to do a new demonstration with their cutters. To each firm in Group A, we gave the same refresher, the same information about the two-pentagon variant, and the same offer of a new demonstration. In addition, we explained the misalignment of incentives to the owner and offered to pay one cutter and one printer lump-sum bonuses roughly equivalent to their monthly incomes — 15,000 Rs (US$150) and 12,000 Rs (US$120), respectively — on the condition that within one month they demonstrate competence in using the new technology in front of the owner. If the owner agreed to the intervention, we explained the intervention to one cutter and one printer chosen by the owner, paid them 1/3 of the incentive payment on the spot, and scheduled a time to return to test their performance using the die.

I leave it to you to consider the authors' discussion of their results. What do we learn from this paper about technology adoption? What do we learn about alignment of incentives within organisations? What does it matter — if at all — that the technology here was designed by an outside research team, who needed actively to encourage adoption?

In considering these questions, you may also find it interesting to look at the recent work of Hjort (2014), considering a flower-packaging firm in Kenya. As Hjort summarises (on page 1942):

[20] Once again — as in Lecture 1 of this module — we see a novel experiment with a very small 'N'. Once again, it is worth understanding why the authors would use a 'permutation test' in this context.

When contentious presidential election results led to political conflict and violent clashes between the two ethnic groups represented in the sample in early 2008, a dramatic, differential decrease in the output of mixed teams followed. The reason appears to be that workers' taste for discrimination against non-coethnic co-workers increased. Six weeks into the conflict period, the plant implemented a new pay system in which biased upstream workers were unable to increase the relative pay of favored downstream workers by distorting relative supply. As a result, horizontal misallocation of flowers was eliminated and total output in teams in which the two downstream workers were of different ethnic groups increased.

3.4 Empirical analysis: Career concerns as organisational incentives

So far, we have considered incentives in the sense of communications and contract structure. However, incentives within organisations may take many other forms. In particular, literature in organisational economics has often emphasised the role of career concerns in providing incentives (as a starting point for a large literature, see, for example, Gibbons and Murphy (1992)). We will therefore discuss the recent empirical work of Bertrand *et al* (2018) on a different kind of organisation — namely, a professional bureaucracy.

(i) The authors argue that '[a] few key features distinguish professional bureaucratic organizations from other organizations: selection through competitive examinations, a virtual absence of discretionary firing (and hence limited exit), seniority-based progression rules and a fixed retirement age' (page 2). To what extent, if any, can we generalise between government bureaucracies and other kinds of organisation in developing countries?

(ii) What was the 'Northcote-Trevelyan (1854) report'? (For additional credit, come to the lecture via Jowett Walk…)

(iii) The authors conduct a survey to determine the effectiveness of civil servants. How do the authors choose the respondents for this survey? What underlying population is this sample intended to represent?

(iv) Why might some officers enter the IAS earlier in life and others later? What do the authors say about this heterogeneity? Why might it matter?

(v) The authors' survey data is collected on a "5 point integer scale, where 1 is the lowest and 5 the highest performance" (page 8). How does this affect the authors' choice of estimator?

(vi) The authors argue (page 13) that "Compared to individuals who enter younger, those entering older and less likely to reach teh top may be less motivated to do well on the job." Why should prospects of progressing to the top echelons matter for motivation?

(vii) The authors find significant negative effects of age at entry on various measures of bureaucratic effectiveness. Discuss the magnitude of these estimates; in a practical sense, should we consider that age at entry has a large effect on bureaucratic effectiveness?

(viii) The authors include state fixed effects in equation 6. What key assumption is necessary in order to justify this specification? How reasonable is that assumption in this context?

3.5 Empirical analysis: Monitoring and reporting on public sector absence

To conclude this lecture, we will jump from senior bureaucrats in India to front-line public sector employees in Pakistan; this is the focus of the recent work by Callen *et al* (2018).

(i) This paper considers medical staff. How broadly should we generalise the results? (*i.e.* What kinds of public representatives should we think about when considering these results?)

(ii) What proportion of public-sector doctors are absent from their workplace during normal working hours in Punjab, Pakistan? (Do a quick Google search; can you find the equivalent statistic for the UK National Health Service?)

(iii) How was the randomisation implemented? How does this affect the authors' approach to inference?

(iv) Describe the authors' main result on the probability of a facility being inspected. How much can the authors say about multi-tasking problems for health inspectors?

(v) Explain the intuition of the identification strategy in equation 3. Suppose that the officials reviewing the 'dashboard' paid no attention to it; what value would we expect for β_1? Alternatively, suppose that the same officials paid extremely close attention; what value would we then expect for β_1?

Before concluding, it is worth noting that Punjab (Pakistan) has recently been a fertile ground for some very creative interventions to test the role of incentives in organisations. If you are interested in reading further on the topic, I strongly recommend the two references from Khan *et al* ('Tax Farming Redux' and 'Making Moves Matter').

4 Lecture 4: Relational Contracts

- $\boxed{\star}$ Cai, J. and Szeidl, A. (2018): "Interfirm Relationships and Business Performance," *Quarterly Journal of Economics*, 133(1), 1229–1292.
- $\boxed{\star}$ Fafchamps, M. and Quinn, S. (2016): "Networks and Manufacturing Firms in Africa: Results from a Randomized Field Experiment," *The World Bank Economic Review*, `https://doi.org/10.1093/wber/lhw057`.
- $\boxed{\star}$ Stein, J. (2008): "Conversations among Competitors," *American Economic Review*, 98(5), 2150–2162.
- Breza, E. and Liberman, A. (2017): "Financial Contracting and Organizational Form: Evidence from the Regulation of Trade Credit," *Journal of Finance*, 72(1), 291–323.
- Bubb, R., Kaur, S. and Mullainathan, S. (2018): "The Limits of Neighborly Exchange," *Working paper*.
- Macchiavello, R., and Morjaria, A. (2015): "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports," *The American Economic Review*, 105(9), 2911-2945. (Online appendix available at `https://assets.aeaweb.org/assets/production/articles-attachments/aer/app/10509/20120141_app.pdf`.)
- Macchiavello, R. and Miquel-Florensa, J. (2018): "Vertical Integration and Relational Contracts: Evidence from the Cost Rica Coffee Chain," *Working paper*.
- McMillan, J. and Woodruff, C. (1999): "Interfirm Relationships and Informal Credit in Vietnam," *Quarterly Journal of Economics*, 114(4), 1285-1320.

The three previous lectures all considered issues *internal* to organisations in developing countries — namely, their rules and management practices, their hierarchies, and their handling of internal incentive problems. We end this module by considering organisations' relations with the outside world. Specifically, we will discuss the way that trust can build between organisations, in order to facilitate contracting in situations where formal contractual enforcement is not available; that is, we discuss *relational contracts*. We first consider a theoretical model in which competing firms have the opportunity to exchange ideas that are useful for their production. In the second half of the lecture, we will discuss a recent field experiment by Cai and Szeidl (2018), generating opportunitie for Chinese manufacturing firms to share business ideas.

4.1 Model: ‘Conversations among Competitors’

To begin, we consider Stein’s elegant 2008 model of ‘conversations among competitors’. In my view, this model is useful for capturing several key themes in underestanding relational contracts between organisations. First, we imagine a context in which players cannot make binding contracts — and, therefore, cooperation must be justified purely by the prospect of future cooperation: the ‘shadow of the future’. Second, the model seeks to capture the trade-off between cooperating and competing the in the generation of new ideas — as Stein explains (p.2151), it provides “a theory of incentive-compatible information exchange among players who. . . are in competition with one another”. Third, the model provides an interesting counterpoint to the ‘sender-receiver’ framework that we considered in the previous lecture. As Stein explains (p.2151):

> . . . information flows in both directions — from player A to player B and vice versa — during the course of a conversation, as players quite literally take turns bouncing ideas off of one another. This differs from the classic framework of [Crawford and Sobel (1982)]. Like I do, Crawford and Sobel pose the question of whether “cheap talk” can be credible in a situation where the two parties involved have partially conflicting interests. But in their model, one party is always the better-informed “sender”, and the other is always the less-well-informed “receiver,” so there is no scope for two-way communication.

4.1.1 Two-player model

We begin by considering a two-player model (where we denote players as Player A (‘she’) and Player B (‘he’)). Figure 4.1 illustrates the basic structure. Player A acts in odd-numbered periods, and Player B acts in even-numbered periods. Each player serves a market of mass 1, of which a proportion $\theta < 0.5$ overlaps with the other player.

Figure 4.1: **Game structure: Two players**

Player A *moves at* $t \in \{1, 3, \ldots\}$

Player B *moves at* $t \in \{2, 4, \ldots\}$

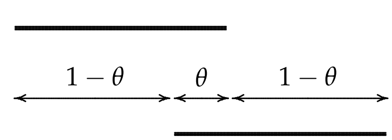

We assume that, in period 1, there is probability p that Player A will have a good idea. If Player A has a good idea, she must decide whether to pass this idea to Player B. If (and only if) she decides to pass the idea, Player B then has a probability p, in period 2, of having a good idea that builds upon Player A’s idea. If Player B has a good idea, he must decide whether to pass this idea back to Player A. We allow this process to continue until either (i) a player fails to have a good idea, or (ii) a player has a good idea but decides not to pass it on. (Note that this is a simplification of the structure in Stein (2008); Stein also allows for players to “pass along a bogus (i.e. informationally useless) version of the idea”. We ignore this possibility for simplicity.)

Once this exchange of ideas has ended, production and selling occurs. We assume that, for Player i, the cost of manufacturing is $1 - h(n_i)$, where n_i is the number of ideas that Player i has access to (that is, the total of (i) ideas that the player came up with, and (ii) ideas that were shared with that player). We assume that, in the part of the market where Player i is the monopolist, the price is 1 — therefore, the monopolist profit per consumer is $1 - (1 - h(n_i)) = h(n_i)$. In the part of the market where two players' markets overlap, we assume a Bertrand competition structure — such that the price is the larger of the two competing marginal costs. Therefore, under competition, the player with higher costs earns a profit of zero, and the player with lower costs has profit equal to the difference in costs. (Of course, if the players have identical costs, this implies that both earn a profit of zero under competition.)

You should check that you agree that, for Player A and Player B respectively, these assumptions imply the following payoff functions:

$$U_A = (1 - \theta) \cdot h(n_A) + \theta \cdot \max\{0, h(n_A) - h(n_B)\}; \tag{4.1}$$
$$U_B = (1 - \theta) \cdot h(n_B) + \theta \cdot \max\{0, h(n_B) - h(n_A)\}. \tag{4.2}$$

So... under what conditions will the players choose to cooperate? For simplicity, we will immediately make an additional assumption of functional form:

$$h(n) = 1 - \beta^n, \tag{4.3}$$

where $\beta \in (0, 1)$. (Note that Stein imposes this functional form as a ‘parametric example’ to illustrate his more general results; we will use this functional form throughout.)

Imagining cooperation: Suppose that the players decide to cooperate: that is, suppose that each player passes on any good idea that she or he has. In that case, we know that, when production and sales occur, profits are zero in the competitive region of the market. (Why?) Further, we know that the game will continue until one player fails to come up with a good idea. Suppose that, in period t, a player has a good idea and decides to share it (so the game has produced $n_A = n_B = t$ ideas). What will the expected payoff be? Well, it depends upon when the game stops. Since — by assumption, for now — the game only stops because a player fails to have a good idea, we can say that the payoff will be...

$$\begin{array}{lll}
(1-\theta)\cdot h(t) & \text{with probability} & (1-p); \\
(1-\theta)\cdot h(t+1) & \text{with probability} & p\cdot(1-p); \\
(1-\theta)\cdot h(t+2) & \text{with probability} & p^2\cdot(1-p); \\
 & \vdots & \\
(1-\theta)\cdot h(t+i) & \text{with probability} & p^i\cdot(1-p).
\end{array}$$

Therefore — taking the infinite sum — we can say:

$$\begin{aligned}
\mathbb{E}_t\left[U(\texttt{continue})\right] &= (1-\theta)\cdot\sum_{i=0}^{\infty} p^i\cdot(1-p)\cdot h(t+i) && (4.4)\\
&= (1-\theta)\cdot\sum_{i=0}^{\infty} p^i\cdot(1-p)\cdot\left(1-\beta^{t+i}\right) && (4.5)\\
&= (1-\theta)\cdot\left[(1-p)\cdot\sum_{i=0}^{\infty} p^i-(1-p)\cdot\beta^t\cdot\sum_{i=0}^{\infty}(p\beta)^i\right] && (4.6)\\
&= (1-\theta)\cdot\left[(1-p)\cdot\frac{1}{1-p}-(1-p)\cdot\beta^t\cdot\frac{1}{1-p\beta}\right] && (4.7)\\
&= 1-\theta-(1-\theta)\cdot(1-p)\cdot\beta^t\cdot\frac{1}{1-p\beta}. && (4.8)
\end{aligned}$$

Imagining defection: Suppose that, in period t, a player has a good idea and then decides not to share it. You should check that you agree that the player's payoff is:

$$\begin{aligned}
U(\texttt{stop}) &= (1-\theta)\cdot h(t)+\theta\cdot[h(t)-h(t-1)] && (4.9)\\
&= h(t)-\theta\cdot h(t-1) && (4.10)\\
&= 1-\beta^t-\theta\cdot\left(1-\beta^{t-1}\right) && (4.11)\\
&= 1-\theta-\beta^t+\theta\beta^{t-1} && (4.12)
\end{aligned}$$

Conditions for cooperation: So, under what conditions will the players cooperate? Comparing the previous utility specifications, we obtain:

$$\begin{aligned}
\mathbb{E}_t\left[U(\texttt{continue})\right] &\geq U(\texttt{stop}) && (4.13)\\
1-\theta-(1-\theta)\cdot(1-p)\cdot\beta^t\cdot\frac{1}{1-p\beta} &\geq 1-\theta-\beta^t+\theta\beta^{t-1} && (4.14)\\
(1-\theta)\cdot(1-p)\cdot\beta\cdot\frac{1}{1-p\beta} &\leq \beta-\theta && (4.15)\\
(1-\theta)\cdot(1-p)\cdot\beta &\leq \beta-\theta-p\beta^2+p\beta\theta && (4.16)\\
\beta-p\beta-\theta\beta+p\beta\theta &\leq \beta-\theta-p\beta^2+p\beta\theta && (4.17)\\
p\beta &\geq \theta. && (4.18)
\end{aligned}$$

You should check that you understand the intuition for this condition. Why do p and β enter symmetrically here? (And, indeed, why do they enter multiplicatively?) Why does θ enter? (And should we worry that θ also relates to total market size?)

4.1.2 An N-player model

Following Stein, we now extend the model to N players. To do this, we maintain a sequential structure — that is, we assume that Player A and Player B exchange ideas; when these players have ended their conversation, Player B and Player C may start a new conversation, building upon the ideas that have been generated between Player A and Player B (and so on). Figure 4.2 illustrates the new structure. We assume that production and selling occurs after *all* conversations have ended; this is important, because it implies (for example) that if Player B gains a new idea from a conversation with Player C, this idea can be used to out-compete Player A in the θ market share that Player A and Player B have in common. (Note that the overlap between players is still θ — so that, with the exception of Player A, each player is now a monopolist only for a share $1-2\theta$ of her or his market.)

Figure 4.2: **Game structure: Multiple players**

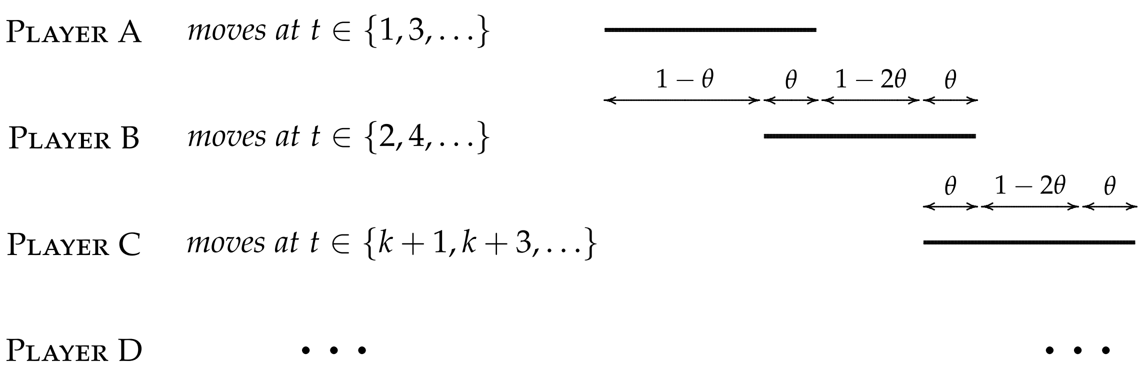

This extension captures an interesting trade-off; as Stein explains (p.2157):

> As before, the potential advantage of [B sharing with C] is that player C will, with probability p, be able to come up with a further refinement $\mathcal{X}_{k+1}$, and that the conversation between B and C will continue on for several more stages from there.The disadvantage to B of starting a conversation with C is that by giving C access to idea $\mathcal{X}_k$, B effectively gives C *all* of the first k signals s_1 through s_k — i.e. B repeats everything of value that was learned during the A-B conversation. The implicit assumption here is that the refined idea $\mathcal{X}_k$ also embodies all of the cumulative knowledge in the previous-stage ideas

$\mathcal{X}_1$ through $\mathcal{X}_{k-1}$... when A first speaks to B, he gives away a brand-new idea that is equivalent to just one signal, in the hopes of getting one more signal in return. By contrast, when B first speaks to C, he gives away a more fully developed idea that is equivalent to k signals, in the hopes of getting one more signal in return.

You should check that, for Player B (and, symmetrically, every player from C onwards) the utility function now looks like this:

$$U_B = (1 - 2\theta) \cdot h(n_B) + \theta \cdot \max\{0, h(n_B) - h(n_A)\} + \theta \cdot \max\{0, h(n_B) - h(n_C)\}. \quad (4.19)$$

Note two important features of this payoff structure (see page 2157):

(i) "[C]onditional on completing a conversation with player C, player B does not care if C then goes ahead and improves his information set via a subsequent conversation with player D. This is because once B and C have access to the same number of signals, B can never earn a profit from those customers they have in common; this does not change if C eventually becomes even better informed than B."

(ii) "[O]nce a conversation [between B and C] is started, C will always have a stronger incentive to continue it than B. This is because C has the option value associated with the possibility of moving on to a further conversation with D."

We now assume that $p\beta \geq \theta$. This ensures that Player A and Player B will exchange ideas until they can do so no longer, just as in the two-player version of the game. (You may wonder whether Player B, with his new utility function, would still find it optimal to behave in this way; note that, precisely as Stein explains, B will have a stronger incentive than Player A in this N-player model.) To understand the communication incentives between Player B and Player C, we need to think about how such a conversation would progress if it were started — and, with that result in hand, we can then ask when such a conversation might start.

If B and C initiate a conversation, will they want to continue it? Suppose that the conversation between A and B has ended, having generated k signals. Suppose that B and C then initiate a conversation; will they want to continue it as long as possible? As just noted, "C will always have a stronger incentive to continue it than B"; therefore, we can answer this question by focussing on Player B's incentives. Suppose that the conversation between B and C continues until n_B ideas have been generated. In that case, player B's payoff is:

$$U_B = (1 - 2\theta) \cdot h(n_B) + \theta \cdot [h(n_B) - h(k)] \quad (4.20)$$
$$= (1 - \theta) \cdot h(n_B) - \theta \cdot h(k). \quad (4.21)$$

Suppose, instead, in period $t > k+1$, player B has a good idea and decides not to share it. (Note that, because $t > k+1$, B has at least exchanged the first idea with C; therefore, if B stops the conversation, C's costs will be $1 - h(t-1)$, rather than $1 - h(0) = 1$.) In that case, player B's payoff will be:

$$U(\texttt{stop}) = (1-2\theta) \cdot h(t) + \theta \cdot [h(t) - h(k)] + \theta \cdot [h(t) - h(t-1)] \tag{4.22}$$
$$= h(t) - \theta \cdot h(t-1) - \theta \cdot h(k). \tag{4.23}$$

If we were feeling keen, we could work out the expected payoff to player B from continuing under this new payoff structure, and compare it to the new payoff from stopping. But, happily, we don't need to. Hopefully, you see why. Compare equation 4.21 with the in-period payoff to each player in the two-player game (namely, $(1-\theta) \cdot h(t)$), and compare equation 4.23 with equation 4.10: the new payoffs are the same as those in the two-player case, less the final term — which is fixed and, crucially, is the same term in both equations. Therefore, assuming that Player B and Player C had the possibility of starting a conversation — that is, given $p\beta \geq \theta$ — it follows that, *if* they start, they will want to continue. But do they want to start?

Will Player B want to start a conversation with Player C? Replacing n_A with k, we can simplify equation 4.19 to:

$$U_B = (1-2\theta) \cdot h(n_B) + \theta \cdot [h(n_B) - h(k)] + \theta \cdot \max\{0, h(n_B) - h(n_C)\} \tag{4.24}$$
$$= (1-\theta) \cdot h(n_B) + \theta \cdot \max\{0, h(n_B) - h(n_C)\} - \theta \cdot h(k). \tag{4.25}$$

Note that this is identical to U_B in the two-player game — except that, in this N-player game, we subtract the additional term $\theta \cdot h(k) = \theta \cdot \left(1 - \beta^k\right)$. Therefore, in this N-player game, we can follow equation 4.8 and say:

$$\mathbb{E}_t\left[U(\texttt{continue})\right] = 1 - \theta - (1-\theta) \cdot (1-p) \cdot \beta^t \cdot \frac{1}{1-p\beta} - \theta \cdot \left(1 - \beta^k\right). \tag{4.26}$$

If Player B chooses not to start the conversation with Player C, this implies $n_A = n_B = k$ and $n_C = 0$, so (substituting in to equation 4.19), Player B receives utility of:

$$U_B(\texttt{withdraw}) = (1-2\theta) \cdot h(k) + \theta \cdot h(k) \tag{4.27}$$
$$= (1-\theta) \cdot h(k) \tag{4.28}$$
$$= (1-\theta) \cdot \left(1 - \beta^k\right) \tag{4.29}$$

Therefore, Player B will choose to start the new conversation if:

$$\mathbb{E}_t\left[U(\texttt{continue})\right] \geq U_B(\texttt{withdraw}) \tag{4.30}$$

$$\Leftrightarrow 1-\theta-(1-\theta)\cdot(1-p)\cdot\beta^k\cdot\frac{1}{1-p\beta}-\theta\cdot\left(1-\beta^k\right) \geq (1-\theta)\cdot\left(1-\beta^k\right) \tag{4.31}$$

$$1-\theta-(1-\theta)\cdot(1-p)\cdot\beta^k\cdot\frac{1}{1-p\beta} \geq 1-\beta^k \tag{4.32}$$

$$(1-\theta)\cdot\left[1-(1-p)\cdot\beta^k\cdot\frac{1}{1-p\beta}\right] \geq 1-\beta^k \tag{4.33}$$

$$\Leftrightarrow \frac{1-\beta^k\cdot(1-p)/(1-p\beta)}{1-\beta^k} \geq \frac{1}{1-\theta}. \tag{4.34}$$

4.1.3 What do we learn from this model?

To interpret this result, we will consider a few comparative statics; these essentially follow closely some of the comparative statics discussed by Stein. Following Stein, we have the following definitions:

(i) k: As discussed, this is "the duration of the initial A-B conversation" (p.2159);

(ii) $I(k^*)$: This is defined as the largest integer such that equation 4.34 holds (it can be shown that the left-hand side of equation 4.34 is decreasing in k). In a sense, $I(k^*)$ measures the most sophisticated idea that is capable of further development.

(iii) c: This is defined as "the number of conversations that occur in a given play of the game *subsequent* to the initial A-B conversation" (p.2159, emphasis in original).

As Stein explain in the Online Appendix,[21] we can think of c as "the random number of failures in a series of Bernoulli (p) trials before there are $I(k^*)$ successes". Therefore, c has a Negative Binomial distribution, and we can say that $\mathbb{E}(c) = I^*(k)\cdot p/(1-p)$. We can use these results for numerical description of comparative statics of $I(k^*)$ and $\mathbb{E}(c)$. We will do this with respect to two model parameters: θ (degree of competition) and p (which, as Stein explains (p.2160), can be thought of "as a proxy for the talent of the players in a given network — more talented players are less likely to draw a blank at any point in a conversation".[22]

[21] This is available at `https://assets.aeaweb.org/assets/production/articles-attachments/aer/data/dec08/20071005_app.pdf`, but is *not* a required reading.

[22] Except where otherwise stated, the following graphs use the same parameters that Stein used for numerical simulation: namely, $\beta = 0.95$, $\theta = 0.1$, and $p = 0.5$.

Varying competition: We begin by varying the degree of competition, θ. Figure 4.3 shows how $I(k^*)$ varies with respect to θ (where I graph for $\theta \in [0.01, 1]$). The vertical dotted line shows $\theta = p\beta$; communication is supported to the left of this line. As one might expect, competition reduces the extent of conversation — in the sense that, as competition increases, players are less willing to discuss sophisticated ideas. Note that, because we use $p = 0.5$, $I(k^*) = E(c)$ for all values of θ in this particular case — so this is also a graph showing how $\mathbb{E}(c)$ varies with θ. Competition not only impedes conversation; it reduces the expected number of conversations.

Figure 4.3: **Comparative statics for $I(k^*)$ with respect to θ**

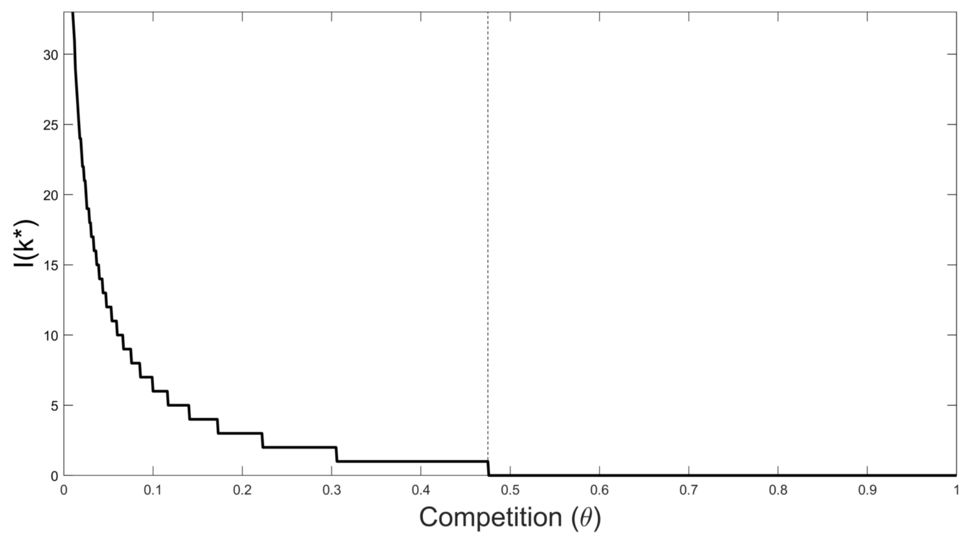

Varying player ability: What about player ability? Here, the results are quite different. In Figure 4.4, we consider how the willingness to start a conversation, $I(k^*)$, changes with p. (Again, the vertical dotted line shows the case $p\beta = \theta$; communication is feasible to the right of this line.) $I(k^*)$ is a weakly increasing function of p: if players are more able, there is a higher expected payoff from collaborating with them.

Finally, in Figure 4.5, we consider the expected number of conversations. This has a very strange shape: even though the upper bound on $E(c)$ is decreasing in p, the function itself is highly non-linear. Why? What competing intuitions might explain this result? I leave this for you to consider in reading Stein's work.

Figure 4.4: **Comparative statics for $I(k^*)$ with respect to p**

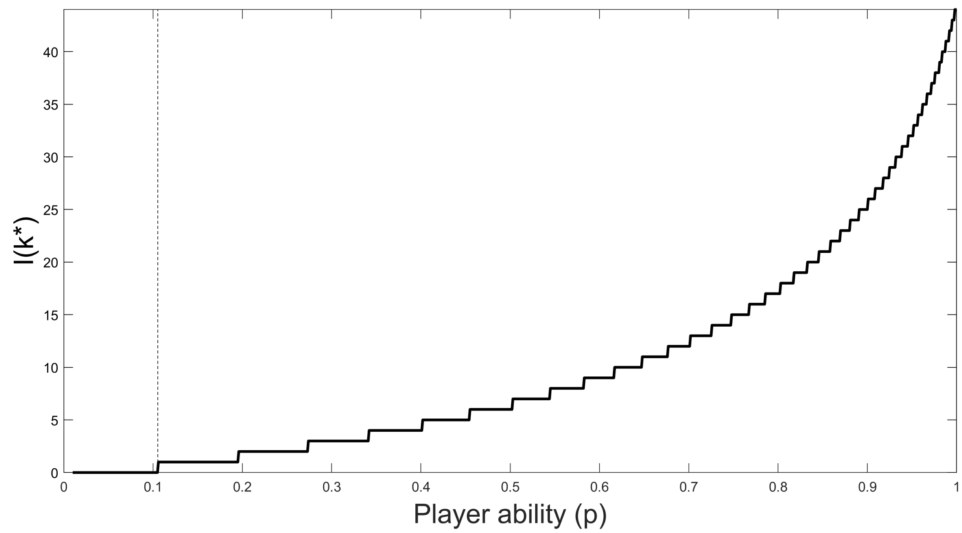

Figure 4.5: **Comparative statics for $E(c)$ with respect to p**

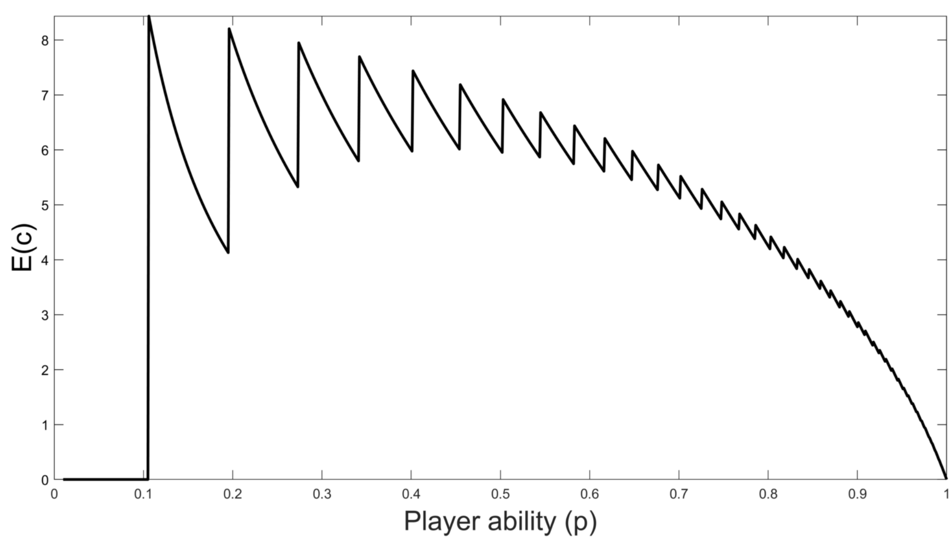

4.2 Cai and Szedil (2018)

We will conclude our discussion of organisations and development by considering Cai and Szeidl's 2018 paper about inter-firm relationships in China. In doing so, we will draw comparisons with the other required reading for this lecture: Fafchamps and Quinn's 2016 paper on similar themes in three sub-Saharan countries.

(i) How many firms expressed interest? How many were assigned to meetings?

(ii) Cai and Szeidl use quite a different firm sampling strategy to Fafchamps and Quinn. In your view, what are the most important differences in sampling? (At the very least, you should think, for example, about firm size and firm age. . .)

(iii) How did the treatment in Cai and Szeidl (2018) differ from the treatment in Fafchamps and Quinn (2016)?

(iv) How did the authors create variation in the composition of firm groups? What did the authors learn from this?

(v) Explain how the authors seeded specific information into the groups. What information did the authors inject? What did the authors learn from this?

(vi) The authors randomly varied the proportion of managers in each meeting group receiving the information. Why?

(vii) What is 'exclusion bias'? Why might it matter for a network experiment such as this?

(viii) To what extent do Cai and Szeidl's results align with the theoretical framework of Stein (2008)?