
INTRODUCTION TO GAME THEORY

1 Discrete Games

Lecturer: Simon Quinn

Reading:

- Jehle and Reny, section 7.1, 7.2 and 7.3.1 to 7.3.5.
- Osborne, *An Introduction to Game Theory*, chapters 1 to 6.
- Varian, chapter 15.

1.1 The concept of a ‘game’

The word ‘game’ means many things to many different people. This is true even within economics, where game theory has become central to a very large body of literature across a very wide range of topics. In his book *The Undercover Economist*, Tim Harford gives the following definition (p.160):

A game, to a game theorist, is any activity in which your prediction of what another person will do affects what you decide to do. Such games include poker, nuclear war, love or bidding for thin air [*i.e.* 3G spectrum] in an auction.

This is a useful definition, particularly for a popular audience. However, for the purpose of this course, we will use a narrower approach. We will think about a game not as an *actual* interaction in the real world, but as a mathematical object used to *model* such interactions. I think this is a useful distinction, because it forces us to separate two distinct challenges: (i) the *logical* challenge of solving a game mathematically, and (ii) the *practical* challenge of using that solution to think about human interactions. In my experience, game theory can quickly become unnecessarily complicated if we try to combine these two challenges. On the one hand, we can easily find ourselves missing logical aspects of the game, because they do not conform to our intuitive ideas about human behaviour. On the other hand, we may start deluding ourselves that our model is intended as a literal representation of the world — and therefore start making unrealistic predictions about human behaviour merely because a model suggests that such outcomes are a theoretical possibility. In introducing their famous game theory textbook, Osborne and Rubinstein (1994, page 1) say that “The art of applying an abstract model to a real-life situation should be the subject of another tome.” We will certainly try to apply game-theoretic insights in these lectures — but we will try always to keep this ‘art’ distinct from the art of game theory itself.

We will therefore use Osborne and Rubinstein’s definition (1994, page 2, emphasis in original):

A game is a description of strategic interaction that includes the constraints on the actions that the players *can* take and the players' interests, but does not specify the actions that the players *do* take. A *solution* is a systematic description of the outcomes that may emerge in a family of games. Game theory suggests reasonable solutions for classes of games and examines their properties.

Time for an example.

1.2 The Prisoner's Dilemma

[Player 1] and [Player 2] have been arrested for robbing the Hibernia Savings Bank and placed in separate isolation cells. Both care much more about their personal freedom than about the welfare of their accomplice. A clever prosecutor makes the following offer to each. "You may choose to confess or remain silent. If you confess and your accomplice remains silent I will drop all charges against you and use your testimony to ensure that your accomplice does serious time. Likewise, if your accomplice confesses while you remain silent, they will go free while you do the time. If you both confess I get two convictions, but I'll see to it that you both get early parole. If you both remain silent, I'll have to settle for token sentences on firearms possession charges. If you wish to confess, you must leave a note with the jailer before my return tomorrow morning."

Stanford Encyclopedia of Philosophy, Prisoner's Dilemma, October 2007

This is an elegant description of the 'Prisoner's Dilemma' — arguably the most famous of all the games ever studied. As the *Stanford Encyclopedia of Philosophy* explains:

Puzzles with the structure of the prisoner's dilemma were devised and discussed by Merrill Flood and Melvin Dresher in 1950, as part of the Rand Corporation's investigations into game theory (which Rand pursued because of possible applications to global nuclear strategy). The title "prisoner's dilemma" and the version with prison sentences as payoffs are due to Albert Tucker, who wanted to make Flood and Dresher's ideas more accessible to an audience of Stanford psychologists. Although Flood and Dresher didn't themselves rush to publicize their ideas in external journal articles, the puzzle attracted widespread attention in a variety of disciplines. Christian Donninger reports that "more than a thousand articles" about it were published in the sixties and seventies.

Philosophers may be interested in the Prisoner's Dilemma as a tool for thinking about the nature of human morality, and the Rand Corporation were concerned to avoid nuclear annihilation. We will use it for a more modest task: understanding the basic principles of game theory. Figure 1.2 shows one way of thinking about this problem; each cell shows a pair of payoffs, as a function of the combination of players' actions.

Figure 1.1: ‘Prisoner’s Dilemma’ in Normal Form

		PLAYER 2	
		Cooperate	Defect
PLAYER 1	Cooperate	2, 2	0, 3
	Defect	3, 0	1, 1

1.3 Nash Equilibrium

So how should we solve this game? To do this, we need a *solution concept* — a criterion for deciding the outcomes of the game. We also need a systematic method for finding all of the solutions to a particular game. In this section, I suggest a series of questions that, I think, we should ask about *all* of the games that we study. (These questions are based loosely on section 2.8 in Osborne (2003).)

Who are the players? This step is usually trivial — but it is also critically important. If a game is an abstract mathematical object (as I argue we should consider it to be) then it is fundamentally important that we define precisely the constituent optimising parts — that is, the players. Osborne and Rubinstein (1994, page 2) say that “A player may be interpreted as an individual or as a group of individuals making a decision.” In our Prisoner’s Dilemma, we have two players, conveniently labelled ‘PLAYER 1’ and ‘PLAYER 2’.

What actions can they take, and when? Every player has a set of available actions. It may be that different players have different actions available (indeed, it may even be the case that one player’s action changes the set of actions available to another player). In our Prisoner’s Dilemma, each player must decide between two actions: ‘Cooperate’ and ‘Defect’.¹ Both players must decide simultaneously — that is, each player must decide its action *before* knowing what actions its opponent is taking.

What payoff does each player get for all combinations of player actions? This is illustrated in Figure 1.2.

For each player, what is the set of best-responses to all combinations of other player actions? There are different ways of finding this set for different kinds of games. For example, when we consider the Cournot game in the next lecture, we will take the first derivative of profit with respect to quantity, because we treat quantity as a continuous variable. For now, we will simply inspect each cell; the best responses are underlined.

¹ Notice that we are *not* allowing either player to make this choice by, say, tossing a weighted coin. We will talk about this possibility later, when we consider ‘mixed strategies’. For now, we are confining ourselves to ‘pure strategies’.

Figure 1.2: ‘Prisoner’s Dilemma’ in Normal Form: Best Responses

		PLAYER 2	
		Cooperate	Defect
PLAYER 1	Cooperate	2, 2	0, 3
	Defect	3, 0	1, 1

For each player, the *set* of best responses can be described by a ‘best response function’. This is a mapping that tells us, for *every* combination of opponents’ actions, the entire *set* of best responses for a player.² This is sometimes termed a player’s ‘strategy’; for example, in our Prisoner’s Dilemma, each player has a unique optimal strategy: “If the other player chooses Cooperate, I should choose Defect; if the other player chooses Defect, I should also choose Defect.”

What is the solution concept? A ‘solution concept’ is a criterion that tells us how to solve the game. For example, in our Prisoner’s Dilemma, we could have a solution concept that says, “The game is solved if the aggregate utility is maximised.” In that case, the solution would be (Cooperate, Cooperate). Unfortunately, this is not a very good solution concept in this context, because there is no good reason to think that actions in this kind of interaction will aim to maximise aggregate utility (for example, the moves are not determined by a ‘social planner’). Instead, we will use a more traditional (and more famous) solution concept, which reflects the assumption that each player acts to maximise its individual welfare: *Nash equilibrium*.

Definition 1 (NASH EQUILIBRIUM) A *Nash equilibrium* is an intersection of best-response functions. That is, an outcome is a *Nash equilibrium* if and only if no player can deviate to another action to improve its payoff, holding fixed all of the other players’ actions.

What are the solutions? Having answered all of the previous questions, the answer to this last question should be straightforward. In our Prisoner’s Dilemma, for example, there is a unique solution: Player 1 chooses Defect, and so does Player 2.

Let’s illustrate these principles with another game. Figure 1.3 shows the (financial) payoffs from a game on *Golden Balls*, a (sadly defunct) British TV show. You should check that you can answer all of the previous questions for this game. This is a good point to introduce another useful concept: *strictly dominated actions*. In the Prisoner’s Dilemma, Defect was a strictly dominating strategy for each player; Cooperate was strictly dominated. Is the same true of Steal and Steal in Golden Balls?

² Some people would prefer to call this a ‘correspondence’, preserving the word ‘function’ for a mapping that has a unique outcome. We will use the word ‘function’ for simplicity (as, for example, do Osborne and Rubinstein (1994)).

Figure 1.3: ‘Golden Balls’ in Normal Form

		IBRAHIM	
		Split	Steal
NICK	Split	£ 6800, £ 6800	£0, £13600
	Steal	£13600, £0	£0, £0

Definition 2 (STRICT DOMINATION) For player i , an action a_i **strictly dominates** an action b_i if player i receives a strictly higher payoff from a_i than from b_i , irrespective of the other players’ actions. In that situation, we would say that action b_i is **strictly dominated**.

Finally, it is worth noting a new dominance concept, first introduced by Li (2017): *obviously dominant strategies*.³

Definition 3 (OBVIOUSLY DOMINANT STRATEGY) For player i , a strategy s is **obviously dominant** if, at the time when the player must choose, the best possible outcome from any other strategy is no better than the worst possible outcome for strategy s .

Is Defect an obviously dominant strategy in the Prisoner’s Dilemma? You will have the chance to explore the concept of obvious dominance further in the class exercises.

1.4 Subgame Perfect Equilibrium

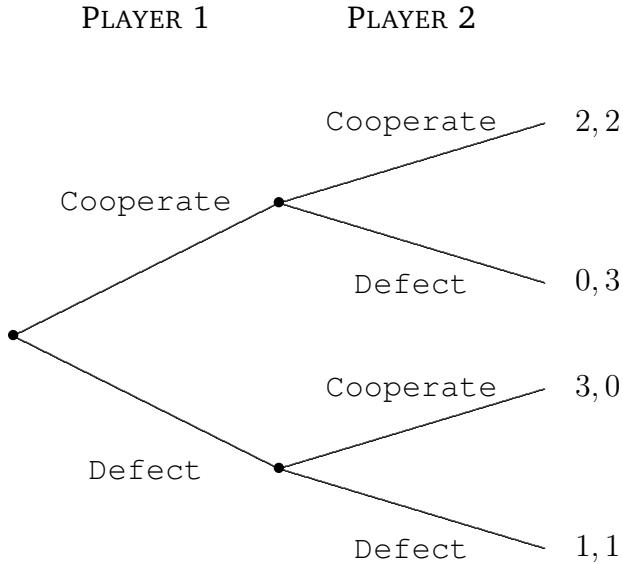
Hopefully, we have a good way of thinking about strategic interactions in which players act simultaneously. But many strategic interactions are not simultaneous. The Royal Navy, for example, maintains four nuclear-armed submarines — at least one of which is always at sea — so that Britain can launch a nuclear strike *after* an opponent chooses to do so.⁴ For the same reason, the United States Air Force kept at least one ‘Looking Glass’ command and control plane airborne at all times from February 1961 to July 1990. Clearly, we need a way of thinking about sequential strategic decisions. This will require a new kind of game diagram, and a new solution concept.

Suppose that we change the rules of our Prisoner’s Dilemma, so that we now require Player 1 to act before Player 2. (Of course, this implies that Player 2 will know Player 1’s action before deciding how to respond; if Player 1 acts first but Player 2 does not know what decision is taken, we may as well model the decisions as simultaneous.) We can illustrate this new version of the game using Figure 1.4. We describe this kind of diagram as an ‘extensive form’ game. As we noted earlier, Defect is a strictly dominating strategy for both players in the Prisoner’s Dilemma, so there is no reason that changing the order of play should affect our prediction: we still anticipate that each player will choose Defect.

³ The definition given here is a slight simplification of Li’s definition; however, it will serve well for our purposes. You should read Li’s excellent paper if you would like to know more.

⁴ This is known in Britain as ‘Continuous at Sea Deterrence’.

Figure 1.4: ‘Prisoner’s Dilemma’ in Extensive Form



In other situations, though, the order of play may matter a lot. Suppose that a wealthy NGO approaches two neighbouring communities in a developing economy, and asks each community to choose whether it wants a school or a hospital. The NGO commits that it will provide either one school or one hospital — but it will do so only if the two communities can agree.⁵ Suppose that the two communities have different priorities: community 1 would prefer a school over a hospital, and community 2 would prefer a hospital over a school. Both communities would prefer *either* a school or a hospital over receiving nothing.

Figure 1.5 shows a simple illustration of this game. (Note that, for better or for worse, this structure has traditionally been known as the ‘Battle of the Sexes’.) You should check that you can step through the questions we discussed earlier, and find all of the (pure strategy) Nash Equilibria.

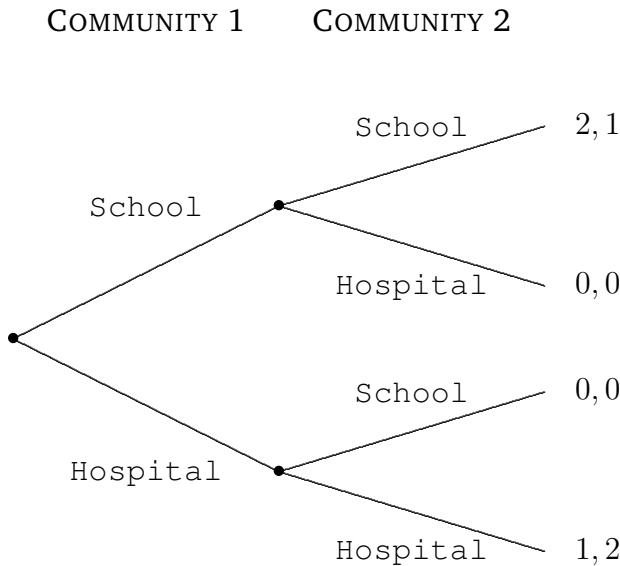
Figure 1.5: ‘Battle of the Sexes’ in Normal Form

		COMMUNITY 2	
		School	Hospital
COMMUNITY 1	School	2, 1	0, 0
	Hospital	0, 0	1, 2

⁵ The NGO explains that its donors are keen to see evidence of genuine community engagement — just not so much engagement that the communities might hold differing opinions.

Now suppose that the NGO decides to approach Community 1 before Community 2, and announces Community 1's decision before Community 2 responds. Figure 1.6 shows this 'Battle of the Sexes' game in extensive form. How should we think of the outcome in this game? We haven't changed our definition of Nash equilibrium, so we clearly still have two Nash Equilibria in pure strategies: (School, School) and (Hospital, Hospital). But one of these equilibria — (Hospital, Hospital) — now seems extremely unlikely as an outcome of the game. After all, why would Community 1 choose Hospital? Community 1 should surely reason that, if it chooses School, Community 2 will make the same choice. We need a new solution concept that extends Nash equilibrium to sequential decisions.

Figure 1.6: 'Battle of the Sexes' in Extensive Form



Definition 4 (SUBGAME) *For our purposes, a **subgame** is a game that begins at a node in an extensive form game.*

Under this definition, there are three subgames in our current example: a subgame after Community 1 chooses School, a subgame after Community 1 chooses Hospital, and a subgame before Community 1 makes its choice (i.e. the entire game). (The first two of these are sometimes described as '*proper subgames*'). We can now define our new solution concept.

Definition 5 (SUBGAME PERFECT EQUILIBRIUM) *A **subgame perfect equilibrium** is a combination of strategies that induces a Nash equilibrium in every subgame.*

Note that this implies that every subgame perfect equilibrium is a Nash equilibrium — but that not every Nash equilibrium is subgame perfect.

Most sequential games that we study are *finite* games — that is, they have a fixed end, and all of the players know at the beginning of the game when this end will be. (In our case, for example, both Community 1 and Community 2 know that the NGO will consult each of them, and then leave.) For finite sequential games, we can find the Subgame Perfect Equilibria by *backward induction*: start at the final period and work backwards, assuming that each player acts at each point to maximise its payoff. Using this method, you should obtain (School, School) as the unique subgame perfect equilibrium for the game in Figure 1.6.

2 Continuous Games

Lecturer: Simon Quinn

Reading:

- Jehle and Reny, section 4.2.
- Varian, chapters 14, 15 and 16.
- Osborne, *An Introduction to Game Theory*, chapters 3 and 14.

The previous lecture introduced key concepts in game theory. To do this, we focussed on *discrete games* — that is, games in which each player was confined to a ‘countable set’ of possible actions. In this lecture, we will extend these concepts to consider games in which players may choose their actions from a ‘continuous’ (i.e. ‘uncountable’) set. This will be particularly useful for modelling imperfect competition between firms — but will also have a broader range of applications, including as a foundation for several topics later in the course.

2.1 Mixed strategies

We have noted several times that our previous results were restricted to ‘pure strategies’. That is, we forced each player to choose one action and one action only: we did not allow players to ‘mix’ between actions by, say, tossing a coin. In many situations, it is entirely reasonable *not* to consider mixed strategies; after all, we can hardly imagine President Biden allowing a coin toss to decide whether or not to launch a nuclear strike, nor a community doing the same with its choice of projects. But in other situations, randomisation seems much more reasonable. Consider a striker stepping up to take a penalty kick. Suppose that she has just two choices: Kick left and Kick right. Suppose that her opponent, the goalkeeper, can either choose Dive left or Dive right. (To be clear, we will use ‘Dive left’ to refer to the *striker’s* left, not the *goalkeeper’s*!) Suppose that, if the goalkeeper chooses the correct direction to dive, she saves the shot; otherwise, the striker scores. This game is represented in Figure 2.1. (Note that this kind of structure is usually termed a ‘matching pennies’ game.)

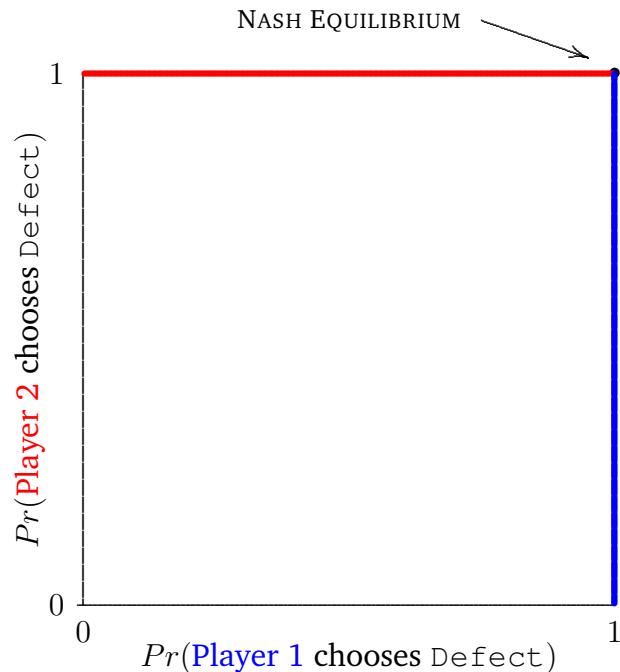
Figure 2.1: ‘Matching Pennies’ in Normal Form

		GOALKEEPER	
		Dive left	Dive right
STRIKER	Kick left	0, 1	1, 0
	Kick right	1, 0	0, 1

Hopefully, it’s obvious that this game has no Nash equilibrium in pure strategies. The reason is clear: for any cell in our diagram, one or other of the players wants to change its

action. Fortunately, we can still make reasonable predictions about the outcome if we allow each player to ‘mix’; indeed, ‘mixing’ seems like a good intuitive description of the way that many goalkeepers and strikers would actually approach this kind of game in the real world.⁶ We’ve discussed, in very loose terms, the possibility of ‘tossing a weighted coin’. Let’s now be more formal. By *mixing*, we mean that each player will be allowed to choose a probability weight for each action. In the Prisoner’s Dilemma, for example, we would allow Player 1 to assign a probability p_1 to the action `Defect`, so that the probability of action `Cooperate` is $1 - p_1$. Similarly, Player 2 would assign a probability p_2 to `Defect`, leaving $1 - p_2$ as the weight on `Cooperate`. In mixed strategies, a Nash equilibrium therefore comprises a combination of *probability weights* — not a combination of actions — such that each player is maximising its *expected utility*. (For this reason, the concepts that we covered in Lecture 8 are fundamentally important for understanding mixed strategies.)

Figure 2.2: ‘Prisoner’s Dilemma’: Best Response Functions



⁶ See, for example, Azar and Bar-Eli (2011).

Let's return to the Prisoner's Dilemma, and let's assume that Player 2 chooses Defect with probability p_2 . Player 1's expected utility from Cooperate is $2 \cdot (1 - p_2) + 0 \cdot p_2 = 2 - 2p_2$; expected utility from Defect is $3 \cdot (1 - p_2) + 1 \cdot p_2 = 3 - 2p_2$. The utility from Defect is higher, irrespective of p_2 — so Player 1 should choose $p_1 = 1$ for any value of p_2 . Symmetrically, Player 2 should choose $p_2 = 1$ for any value of p_1 . Figure 2.2 shows the result. Admittedly, this is not a very interesting diagram — thanks to the fact that the Prisoners Dilemma has strictly dominating strategies — but Figure 2.2 shows how we can take our earlier result and understand it in a context of mixing.

Things become more interesting when we apply the same principles to our goalkeeper and striker. Suppose that Striker chooses Kick left with probability p_s (so Kick right has probability $1 - p_s$). Suppose that Goalkeeper has to choose the probability of Dive left, which we will denote p_g . The optimal choice of p_g is straightforward. If $p_s < 0.5$, Goalkeeper should choose $p_g = 0$. If $p_s > 0.5$, Goalkeeper should choose $p_g = 1$. And if $p_s = 0.5$, Goalkeeper will be indifferent between Dive left and Dive right, so can choose any value $p_g \in [0, 1]$. Figure 2.3 illustrates; note that Striker's best response function follows the same kind of logic.

Having carefully solved for the individual best responses, it is simple to find the Nash equilibrium: we just superimpose the two figures. Figure 2.4 illustrates. Notice that the Nash equilibrium in this case can be described by a simple rule of thumb: *keep your opponent indifferent!*

You should think about how these principles would extend to other contexts. Would it ever be possible to have Nash equilibria in pure strategies *and* in mixed strategies? And what if, for example, Striker is naturally more effective in kicking to one side than the other? And, finally, (when) are we guaranteed that a Nash equilibrium will exist?

Figure 2.3: ‘Matching Pennies’: Best Response Functions

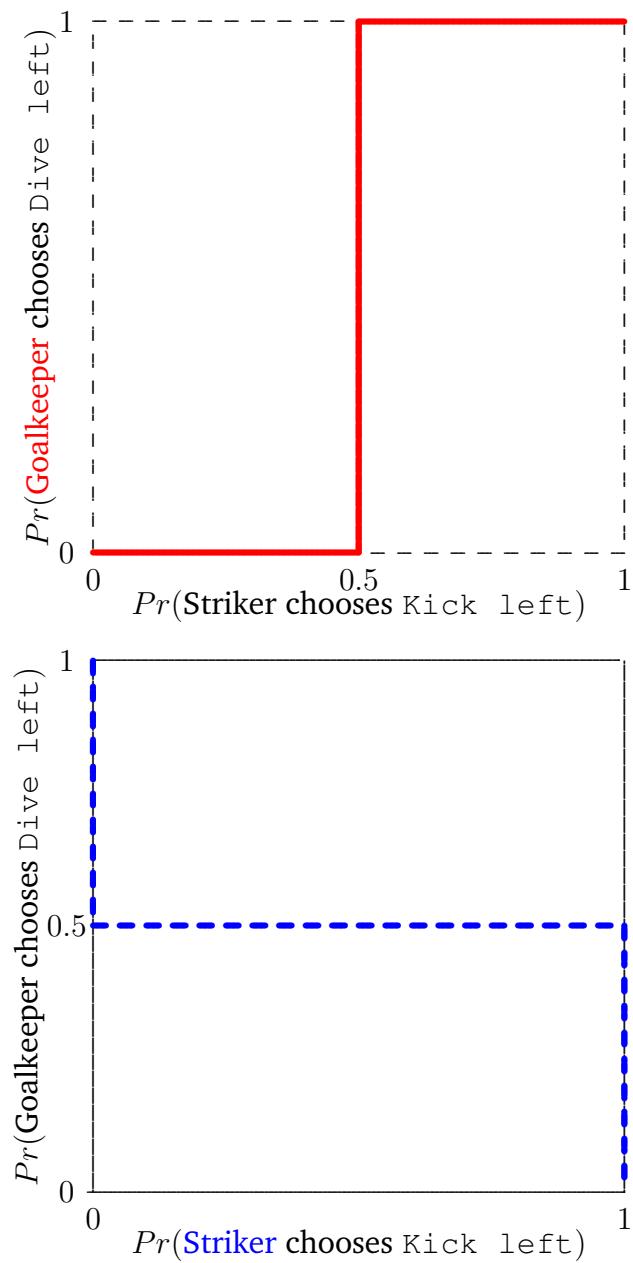
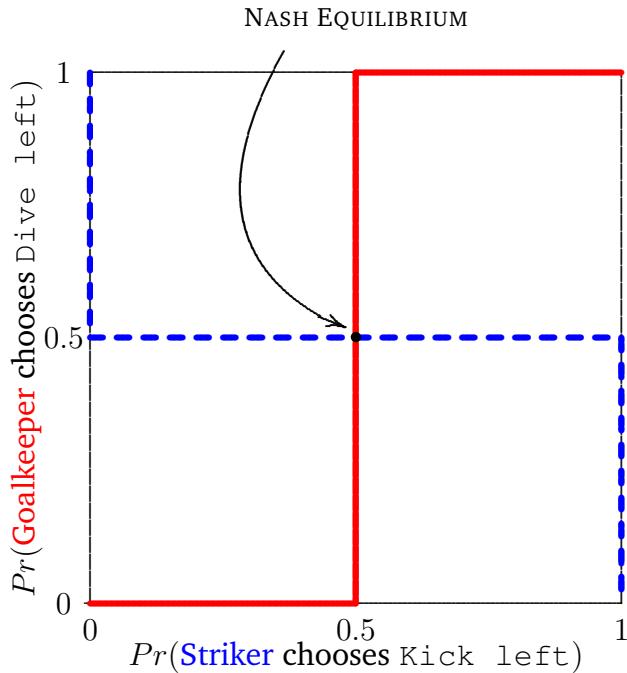


Figure 2.4: ‘Matching Pennies’: Nash Equilibrium in Mixed Strategies



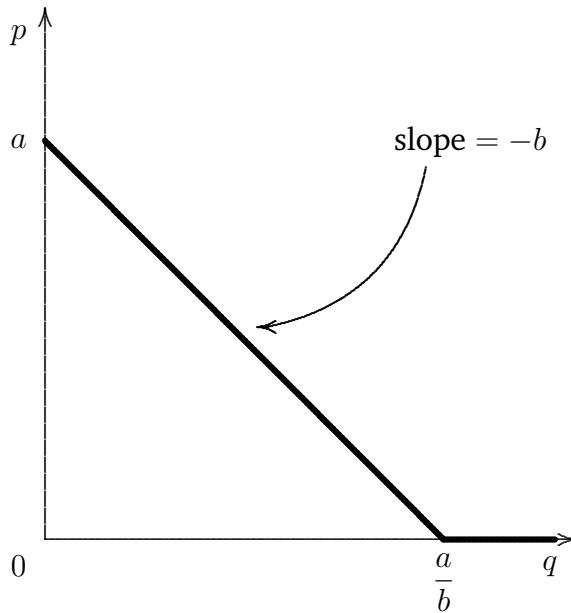
2.2 The Cournot model

2.2.1 The benchmark: A monopolist

We will begin by considering a single firm, producing a single good. The firm faces a very simple demand curve (also illustrated in Figure 2.5):

$$p(q) = \begin{cases} a - b \cdot q & \text{for } q < a \cdot b^{-1}; \\ 0 & \text{for } q \geq a \cdot b^{-1}. \end{cases} \quad (2.1)$$

Figure 2.5: Inverse demand function



For simplicity, we will assume that the firm has constant marginal cost, $c < a$. It is straightforward, then, to solve the ‘monopolists’ problem’:

$$\pi(q) = p(q) \cdot q - c \cdot q \quad (2.2)$$

$$= (a - b \cdot q) \cdot q - c \cdot q \quad (2.3)$$

$$= (a - c) \cdot q - b \cdot q^2 \quad (2.4)$$

$$\frac{\partial \pi(q)}{\partial q} \Big|_{q=q^*} = (a - c) - 2bq = 0 \quad (2.5)$$

$$\therefore q^* = \frac{a - c}{2b} \quad (2.6)$$

$$\therefore p(q^*) = a - \frac{a - c}{2} = \frac{a + c}{2} \quad (2.7)$$

$$\therefore \pi(q^*) = \frac{(a - c)^2}{4b}. \quad (2.8)$$

In a competitive market, of course, firms would be forced to set $p = c$, implying a total market production of $q = (a - c)/b$ (with zero profits). Thus, compared to a market with an infinite number of firms, a single firm produces *half* as much, and charges a price that is higher by $(a - c)/2$.

2.2.2 A Cournot duopoly game

What happens if we introduce a second firm? For simplicity, we will assume that total market demand remains the same, and that this new firm has an identical cost structure to the incumbent. We will analyse this duopoly model using the concepts (and the key guiding questions) that we covered in the previous lecture.

Who are the players? We have two players: Firm 1 and Firm 2. Generically, we can refer to these players as i and j .

What actions can they take, and when? Each player i chooses a quantity to produce, $q_i \geq 0$. Actions are taken simultaneously.

What payoff does each player get for all combinations of player actions? Total quantity supplied is $q = q_i + q_j$. For some combination of quantities (q_i, q_j) , the payoff (profit) to firm i is therefore:

$$\pi_i(q_i, q_j) = [a - b \cdot (q_i + q_j)] \cdot q_i - c \cdot q_i \quad (2.9)$$

$$= (a - c - b \cdot q_j) \cdot q_i - b \cdot q_i^2. \quad (2.10)$$

For each player, what is the set of best-responses to all combinations of other player actions? Suppose that firm j produces q_j . Then firm i 's best response is solved as follows:

$$\frac{\partial \pi_i(q_i, q_j)}{\partial q_i} \Big|_{q_i=q_i^*} = a - c - b \cdot q_j - 2b \cdot q_i = 0 \quad (2.11)$$

$$\therefore q_i^*(q_j) = \frac{a - c - b \cdot q_j}{2b} = \frac{a - c}{2b} - \frac{q_j}{2}. \quad (2.12)$$

What is the solution concept? Nash equilibrium.

What are the solutions? For player i , we have:

$$q_i^*(q_j) = \frac{a - c}{2b} - \frac{q_j}{2}. \quad (2.13)$$

By symmetry, we can therefore also write:

$$q_j^*(q_i) = \frac{a - c}{2b} - \frac{q_i}{2}. \quad (2.14)$$

We now simply find the intersection:

$$q_i^*(q_j^*) = \frac{a - c}{2b} - \frac{q_j^*}{2} \quad (2.15)$$

$$\therefore 4q_i^*(q_j^*) = 2\left(\frac{a - c}{b}\right) - 2q_j^* \quad (2.16)$$

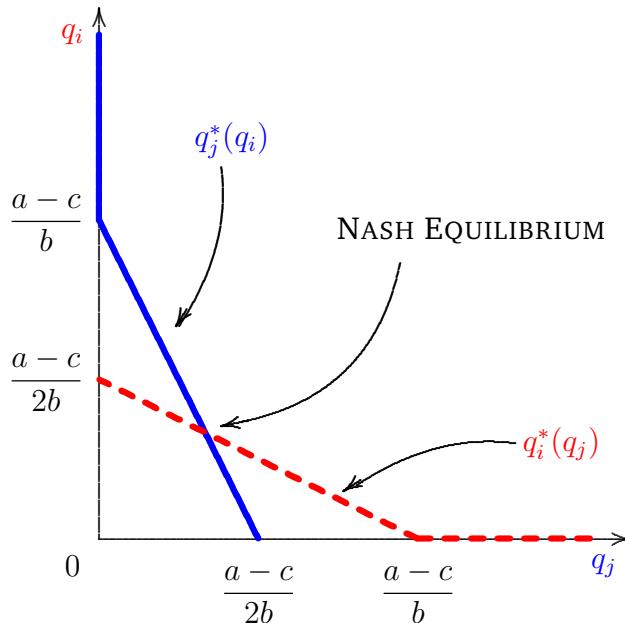
$$2q_j^*(q_i^*) = \left(\frac{a - c}{b}\right) - q_i^* \quad (2.17)$$

$$\therefore 4q_j^*(q_i^*) = 2\left(\frac{a - c}{b}\right) - \left(\frac{a - c}{b}\right) + q_i^* \quad (2.18)$$

$$\therefore q_i^*(q_j^*) = q_j^*(q_i^*) = \left(\frac{a - c}{3b}\right). \quad (2.19)$$

Figure 2.6 illustrates.

Figure 2.6: Cournot game: Best response functions



It is straightforward to show that the equilibrium price is now $(a + 2c)/3$, and that the total profit earned between both firms is $(2/b) \cdot [(a - c)/3]^2$. Therefore, the introduction of a second player has (i) shifted the price closer to marginal cost, and (ii) decreased total firm profits.

2.2.3 An N -player oligopoly model

Now suppose that we take the same structure, but allow for N identical firms. Let's solve the model again...

Who are the players? We have N players, indexed by $i \in \{1, \dots, N\}$.

What actions can they take, and when? Each player i chooses a quantity to produce, $q_i \geq 0$. Actions are taken simultaneously.

What payoff does each player get for all combinations of player actions? For some combination of quantities (q_1, \dots, q_N) , the payoff (profit) to firm i is therefore:

$$\pi_i(q_i, q_j) = \left(a - b \cdot \sum_{j=1}^N q_j \right) \cdot q_i - c \cdot q_i \quad (2.20)$$

$$= \left(a - c - b \cdot \sum_{j \neq i}^N q_j \right) \cdot q_i - b \cdot q_i^2. \quad (2.21)$$

For each player, what is the set of best-responses to all combinations of other player actions? Suppose that firm j produces q_j . Then firm i 's best response is solved as follows:

$$\frac{\partial \pi_i(q_i, q_j)}{\partial q_i} \Big|_{q_i=q_i^*} = a - c - b \cdot \sum_{j \neq i}^N q_j - 2b \cdot q_i = 0 \quad (2.22)$$

$$\therefore q_i^*(q_1, \dots, q_{i-1}, q_{i+1}, \dots, q_N) = \frac{a - c - b \cdot \sum_{j \neq i}^N q_j}{2b} = \frac{a - c}{2b} - \frac{\sum_{j \neq i}^N q_j}{2}. \quad (2.23)$$

What is the solution concept? Nash equilibrium.

What are the solutions? For player i , we have:

$$q_i^* = \frac{a - c}{2b} - \frac{\sum_{j \neq i}^N q_j}{2} \quad (2.24)$$

$$\therefore 2q_i^* = \frac{a - c}{b} - \sum_{j \neq i}^N q_j \quad (2.25)$$

$$\therefore q_i^* = \frac{a - c}{b} - \sum_{j=1}^N q_j = \frac{a - c}{b} - Q, \quad (2.26)$$

where we now use Q to denote the total output produced across all N firms.

By symmetry, we can therefore write:

$$Q = \sum_{i=1}^N \left(\frac{a-c}{b} - Q \right) = N \cdot \left(\frac{a-c}{b} \right) - NQ \quad (2.27)$$

$$\therefore Q = \left(\frac{N}{N+1} \right) \cdot \left(\frac{a-c}{b} \right) \quad (2.28)$$

$$\therefore q_i^* = \frac{a-c}{b} - \left(\frac{N}{N+1} \right) \cdot \left(\frac{a-c}{b} \right) \quad (2.29)$$

$$= \left(\frac{1}{N+1} \right) \cdot \left(\frac{a-c}{b} \right). \quad (2.30)$$

You can then show that:

$$p^* = \frac{a + N \cdot c}{1 + N} \quad (2.31)$$

$$\sum_{i=1}^N \pi_i = \frac{N}{b} \cdot \left(\frac{a-c}{1+N} \right)^2. \quad (2.32)$$

You should verify that increasing N monotonically (i) increases total quantity produced, (ii) decreases the quantity produced by any one firm, (iii) decreases price, and (iv) decreases total profits. You should verify that, in the limit $N \rightarrow \infty$, we recover the competitive market outcome.

2.3 A Stackelberg game

In the previous model, we assumed that all players choose q simultaneously. In some contexts, however, we may think that one player is required to announce its value of q before the other player(s) respond. This is known as a ‘Stackelberg game’. It is an important and interesting extension on the basic ideas that we have developed in this lecture so far, and you will consider it in the class problems for this lecture.

2.4 Price competition: A Bertrand game

Let’s return to think of two firms. However, instead of assuming that the firms compete on *quantity* — as in the Cournot model — we will now think of the firms as competing on *price*. To do this, we will make a strong simplifying assumption that whichever firm charges the lower price can sell to the entire market. This will generate a discontinuity in the firm’s payoff function — which, in turn, will leave us with some very different predictions to the Cournot model. This kind of model is known as a ‘Bertrand game’.

Let’s return to our standard game theoretic questions.

Who are the players? We have two players: Firm 1 and Firm 2.

What actions can they take, and when? Each player i chooses a price to charge, $p_i \geq 0$. Actions are taken simultaneously. Production costs are incurred if and only if the contract is won.

What payoff does each player get for all combinations of player actions? Firm 1 has marginal costs c_1 ; firm 2 has marginal costs $c_2 > c_1$. We assume that there is a fixed total market demand, which we normalise to 1, and that this is not responsive at all to price. The firm charging the lower price takes the entire demand. If firms charge the same price, the entire demand goes to firm 1 (having the lower marginal cost). Therefore, the payoff for firm 1 is:

$$\pi_1(p_1, p_2) = \begin{cases} p_1 - c_1 & \text{if } p_1 \leq p_2; \\ 0 & \text{if } p_1 > p_2. \end{cases} \quad (2.33)$$

The payoff for firm 2 is:

$$\pi_2(p_1, p_2) = \begin{cases} p_2 - c_2 & \text{if } p_2 < p_1; \\ 0 & \text{if } p_1 \geq p_2. \end{cases} \quad (2.34)$$

For each player, what is the set of best-responses to all combinations of other player actions? Start by considering firm 1. First, suppose that $p_2 > c_1$. Notice that $\pi_1(p_1, p_2)$ is strictly increasing in p_1 for $p_1 < p_2$. Therefore, firm 1 should choose $p_1 = p_2$. Second, suppose that $p_2 \leq c_1$. In that case, firm 1 can (i) make $\pi_1 \leq 0$ by choosing $p_1 \leq p_2 \leq c_1$ or (ii) make $\pi_1 = 0$ by choosing $p_1 > p_2$ (in effect, by choosing to exit).

Therefore, firm 1's best-response function is:

$$p_1^*(p_2) = \begin{cases} p_2 & \text{if } p_2 > c_1 \\ [p_2, \infty) & \text{if } p_2 = c_1 \\ (p_2, \infty) & \text{if } p_2 < c_1 \end{cases} \quad (2.35)$$

Now consider firm 2. First, suppose that $p_1 > c_2$. In this case, firm 2 should choose $p_2 = p_1 - \varepsilon$ (where ε is the smallest currency unit available). Second, suppose that $p_1 \leq c_2$. In that case, firm 2 should choose $p_2 \geq p_1$ (i.e. exiting). Therefore, firm 2's best-response function is:

$$p_2^*(p_1) = \begin{cases} p_1 - \varepsilon & \text{if } p_1 > c_2 \\ [p_1, \infty) & \text{if } p_1 \leq c_2 \end{cases} \quad (2.36)$$

Figures 2.7 and 2.8 illustrate.

Figure 2.7: Bertrand Game: Best response function for player 1

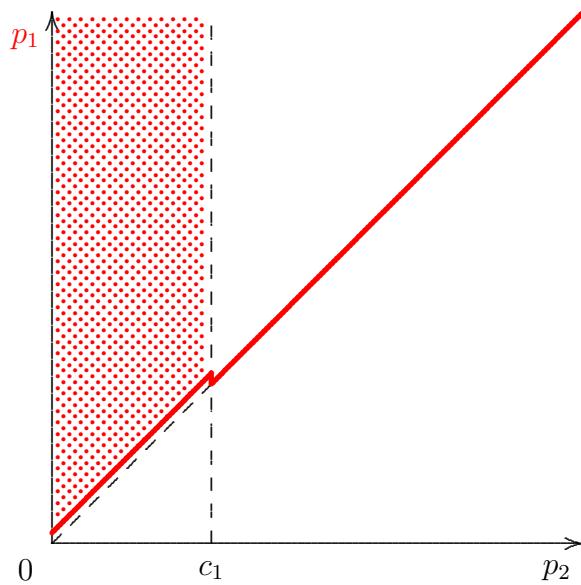
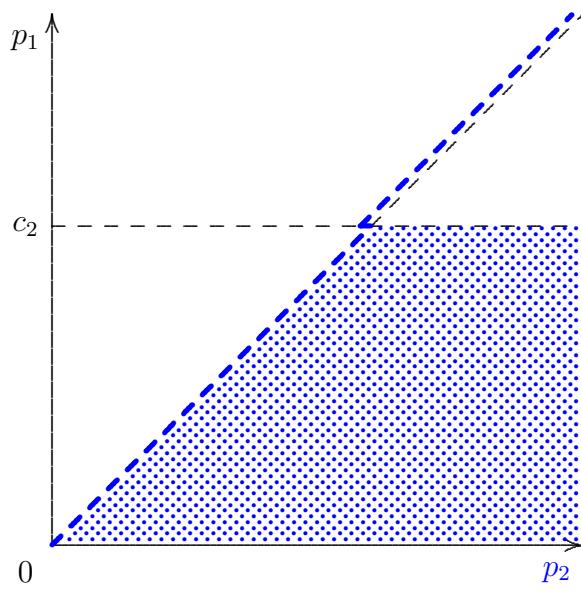


Figure 2.8: Bertrand Game: Best response function for player 2



What is the solution concept? Nash equilibrium.

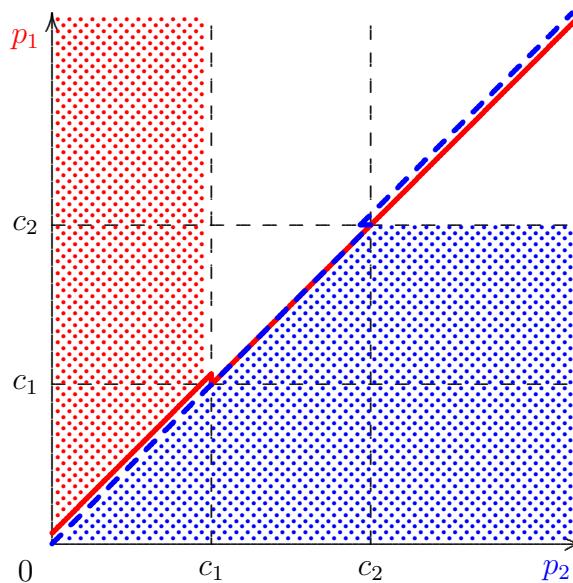
What are the solutions? Figure 2.9 overlays the two best response functions. It shows that there are multiple Nash equilibria: $p_1^* = p_2^* \in [c_1, c_2]$.

Varian (page 292) considers a very similar version of this game, and says:

Thus a Nash equilibrium in this game is for firm 1 to set $p_1 = c_2$ and to produce [1 unit] of output, while firm 2 sets $p_2 \geq c_2$ and produces zero.

Do you agree?

Figure 2.9: Bertrand Game: Best responses



2.5 Class exercises: Introduction to Game Theory

1. In his paper, Li (2017) says that “obviously dominant strategies can be recognized as dominant without contingent reasoning” — by which he means “to reason state-by-state about hypothetical scenarios”. Explain the relationship between obvious dominance and contingent reasoning. To illustrate, revisit the following example games; which of them (if any) has an obviously dominant strategy?
 - (a) Figure 1.2 ('Prisoner's Dilemma' in Normal Form)
 - (b) Figure 1.3 ('Golden Balls' in Normal Form)
 - (c) Figure 1.4 ('Prisoner's Dilemma' in Extensive Form)
 - (d) Figure 1.5 ('Battle of the Sexes' in Normal Form)
 - (e) Figure 1.6 ('Battle of the Sexes' in Extensive Form)
 - (f) Section 2.2.2 (A Cournot duopoly game)
 - (g) Section 2.4 (Price Competition: A Bertrand game)
2. Consider again the 'battle of the sexes' game between two communities in deciding whether to ask a donor for a school or a hospital. Consider just the simultaneous-move version of the game. Find the mixed strategy Nash equilibria, showing the best response functions in mixed strategies.
3. Consider again the 'matching pennies' game between Striker and Goalkeeper. In that example, we assumed that Striker scores every time she kicks at the goal, unless Goalkeeper dives the same way as the kick. Assume that this is still the case when Striker chooses Kick right. However, now assume that, if she chooses Kick left, Striker has a 50% probability of missing the goal completely (*i.e.* irrespective of what Goalkeeper does).
 - (a) Write the normal form for this new version of the game. (You may assume that Goalkeeper obtains just as much utility from seeing Striker miss the goal as from saving the shot; *i.e.* Goalkeeper cares only about whether the goal is scored.)
 - (b) Find all pure strategy Nash equilibria.
 - (c) Find all mixed-strategy Nash equilibria, showing (graphically) the best response functions.

4. Find all mixed-strategy Nash equilibria for the following games, showing (graphically) the best response functions.

		PLAYER 2	
		Left	Right
		Up	0, 0 1, 1
(a)		Down	0, 0 1, 1

		PLAYER 2	
		Left	Right
		Up	0, 1 1, 0
(b)		Down	0, 0 1, 1

		PLAYER 2	
		Left	Right
		Up	0, 0 0, 0
(c)		Down	0, 0 0, 0

5. Consider again the two-player Cournot duopoly game. Assume now that the two firms have asymmetric costs, so that $c_1 > c_2$. Find the pure strategy Nash equilibrium. (Check that, for the case $c_1 = c_2$, this result collapses to the result we obtained in the lecture.)
6. Consider again the two-player Cournot duopoly game. Maintain the earlier assumption that both firms have the same marginal costs: $c_1 = c_2 = c$. However, assume now that Firm 1 chooses q_1 first, and that this is known to Firm 2 before it chooses q_2 . Find the subgame perfect equilibrium. How much profit does each firm make in equilibrium, and how does this compare to the Cournot duopoly game? How much profit is made in total, and how does that compare?

CONTRACTS

3 Hidden Actions

Lecturer: Simon Quinn

Reading:

- Jehle and Reny, section 8.2.2.
- Varian, chapter 25.
- Feltham and Xie (1994), ‘Performance Measure Congruity and Diversity in Multi-Task Principal/Agent Relations’, *The Accounting Review*, 69(3), pp. 429 – 435.
- Gibbons (2010), ‘Inside Organisations: Pricing, Politics, and Path Dependence’, *Annual Review of Economics*, 2, pp. 337 – 365

Note: The notes for our lectures on hidden action, signalling and bargaining all draw heavily on Dr Meg Meyer’s lecture notes on ‘Bargaining, Contracts, and Theories of the Firm’.⁷ I’m most appreciative to Dr Meyer for the opportunity to draw on this resource in preparing these notes!

3.1 The problem of hidden actions

In the real world, many strategic interactions involve ‘hidden action’: actions that are partially observed, but not fully observed, by other players. This generates a problem known as ‘*moral hazard*’, where one party (the ‘agent’) has better information about its actions than its contracting partner (the ‘principal’). As Holmström (1979, p.74) explained:

The source of this moral hazard or incentive problem is an asymmetry of information among individuals that results because individual actions cannot be observed and hence contracted upon. A natural remedy to the problem is to invest resources into monitoring of actions and use this information in the contract. In simple situations complete monitoring may be possible, in which case a first-best solution (entailing optimal risk sharing) can be achieved by employing a forcing contract that penalizes dysfunctional behavior. Generally, however, full observation of actions is either impossible or prohibitively costly. In such situations interest centers around the use of imperfect estimators of actions in contracting. Casual observation indicates that imperfect information is extensively used in practice to alleviate moral hazard, for instance in the supervision of employees or in various forms of managerial accounting.

⁷ These notes are available at <http://www.nuffield.ox.ac.uk/teaching/Economics/Bargaining/bargainingindex.htm>.

Critically, we will assume *both* (i) that the agent is better informed than the principal and (ii) that the agent is more risk averse. Risk aversion will be an important concept for thinking about contracting under hidden action. This is because, as Holmström (1979, p.79) explained, “The deviation from perfect risk sharing implies that the agent is forced to carry excess responsibility for the outcome and this points to the implicit costs involved in contracting under imperfect information.”

3.2 A tractable model of sharecropping

In this lecture, we will develop – and then extend – a model of contracting under moral hazard. To fix ideas, we will focus initially on a simple case of ‘sharecropping’: a situation where a landowner allows a tenant farmer to work a plot of land in return for a share of the crops produced. We will model this interaction as a game between the owner of a plot of land (whom we will refer to as `Farm`) and a tenant labourer (whom we will refer to as `Worker`). Implicitly, we will imagine that our worker is drawn from a large pool of available labourers. We will imagine that the farm can offer the worker a ‘two-part contract’, involving a lump sum (γ) and a per-unit fee (δ).

Who are the players? We have two players: `Farm` (the principal) and `Worker` (the agent).

What actions can they take, and when? `Farm` commits to a linear piece-rate contract in plot output x :

$$w = \gamma + \delta \cdot x. \quad (3.1)$$

`Worker` sees the offer, (γ, δ) , and chooses whether or not to work on the plot; if working, `Worker` then chooses effort, $a > 0$ (measured in currency terms). Output (x) is realised, such that:

$$x = a + \varepsilon \quad (3.2)$$

$$\varepsilon \sim \mathcal{N}(0, \sigma^2). \quad (3.3)$$

`Farm` observes x , but never observes a or ε separately. `Farm` then pays w on the terms agreed.

What payoff does each player get for all combinations of player actions? The two players have different utility functions.

Assumption 3.1 (PRINCIPAL’S UTILITY) `Farm` sells output at price p , and pays w if the `Worker` has agreed to the contract. Therefore, profit is:

$$\pi = \begin{cases} p \cdot x - w & \text{if } a > 0; \\ p \cdot \varepsilon & \text{if } a = 0. \end{cases} \quad (3.4)$$

We assume that `Farm` payoff is simply profit; this implies that `Farm` is risk-neutral.

Assumption 3.2 (AGENT'S UTILITY) *Effort has quadratic cost $0.5a^2$. Worker has exponential utility over net wage $w - 0.5a^2$, with coefficient of absolute risk aversion r :*

$$u(w, a; r) = -\exp[-r \cdot (w - 0.5a^2)]. \quad (3.5)$$

Worker also has an outside option, having certainty equivalent \bar{u} , which we will normalise to zero.

Note that, if *Worker's* net wage has a distribution $(w - a) \sim \mathcal{N}(\mu, \sigma^2)$, the certainty equivalent is:

$$\bar{u} = \mu - \frac{1}{2} \cdot r\sigma^2. \quad (3.6)$$

This is an extremely useful result: it means that we can solve *Worker's* problem by considering the certainty equivalent \bar{u} , rather than needing explicitly to integrate over the possible values for w . We will come back to this soon.

What is the solution concept? The solution concept is subgame perfect equilibrium.

For each player, what is the set of best-responses to all combinations of other player actions? Consequently, what is the solution? We can solve this using backward induction. Remember that *Farm* offers the contract and then *Worker* decides whether or not to accept (and, if so, how much effort to exert); therefore, we should start by solving for the best-response of *Worker*:

Step 1: Solve for *Worker's* optimal effort, conditional on accepting the contract: If *Worker* accepts the contract ($a > 0$), *Worker* will solve:

$$\max_{a>0} u(w, a; r). \quad (3.7)$$

Note that $w = \gamma + \delta x = \gamma + \delta a + \delta \varepsilon$, and $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. This implies that:

$$w \sim \mathcal{N}(\gamma + \delta a, \delta^2 \sigma^2) \quad (3.8)$$

$$\therefore w - 0.5a^2 \sim \mathcal{N}(\gamma + \delta a - 0.5a^2, \delta^2 \sigma^2) \quad (3.9)$$

Therefore, the certainty equivalent is:

$$\bar{u}(a, \gamma, \delta, \sigma) = \gamma + \delta a - 0.5a^2 - 0.5r\delta^2\sigma^2. \quad (3.10)$$

It follows that, if *Worker* accepts the contract, *Worker* will solve:

$$\max_{a>0} \gamma + \delta a - 0.5a^2 - 0.5r\delta^2\sigma^2 \quad (3.11)$$

$$\therefore a^*(\delta, \sigma) = \delta. \quad (3.12)$$

Step 2: Solve for Worker's decision to accept or reject the contract: The worker will accept the contract if and only if:

$$\bar{u} \geq 0 \quad (3.13)$$

$$\Leftrightarrow \gamma + \delta a^* - 0.5a^{*2} - 0.5r\delta^2\sigma^2 \geq 0 \quad (3.14)$$

$$\Leftrightarrow \gamma + 0.5\delta^2 \cdot (1 - r\sigma^2) \geq 0. \quad (3.15)$$

Therefore, we can summarise Worker's best response as:

$$a^*(\delta, \sigma, \gamma) = \begin{cases} 0 & \text{if } \gamma + 0.5\delta^2 \cdot (1 - r\sigma^2) < 0; \\ \delta & \text{if } \gamma + 0.5\delta^2 \cdot (1 - r\sigma^2) \geq 0. \end{cases} \quad (3.16)$$

Step 3: Solve for Farm's optimal piece-rate contract: Suppose Farm anticipates some Worker response function $a(\delta, \sigma, \gamma)$. We need to consider the case $a(\delta, \sigma, \gamma) = 0$ and the case $a(\delta, \sigma, \gamma) > 0$.

If $a(\delta, \sigma, \gamma) = 0$, Farm pays no wage ($w = 0$), so earns expected profit of:

$$\pi = p \cdot \mathbb{E}(\varepsilon) = 0. \quad (3.17)$$

If $a(\delta, \sigma, \gamma) > 0$, Farm pays $w = \gamma + \delta \cdot a(\delta, \sigma, \gamma)$, so earns expected profit of:

$$\pi = p \cdot a(\delta, \sigma, \gamma) - (\gamma + \delta \cdot a(\delta, \sigma, \gamma)) \quad (3.18)$$

$$= (p - \delta) \cdot a(\delta, \sigma, \gamma) - \gamma. \quad (3.19)$$

For now, we will assume that $(p - \delta) \cdot a(\delta, \sigma, \gamma) - \gamma > 0$ (that is, we will assume that Farm always wants Worker to agree the contract), but we will need to come back and check this.

We can therefore solve Farm's problem as a 'nested' optimisation; that is, by taking Worker's optimisation problem as a constraint:

$$(\gamma^*, \delta^*) = \arg \max_{(\gamma, \delta) \in \mathbb{R}^2} (p - \delta) \cdot a^* - \gamma \quad (3.20)$$

subject to

$$\gamma + 0.5\delta^2 \cdot (1 - r\sigma^2) \geq 0 \quad (3.21)$$

$$a^* = \delta. \quad (3.22)$$

Equation 3.21 is referred to as the 'individual rationality constraint' (in effect, the 'participation constraint'); equation 3.22 is referred to as the 'incentive compatibility constraint' (in effect, the 'effort constraint').

We can solve this in two parts. First, notice that π is strictly decreasing in γ ; therefore Farm wants to set γ as low as possible (subject to the constraints being met). γ only enters

the individual rationality constraint. Therefore, Farm sets γ so that equation 3.21 holds with equality:

$$\gamma^* = 0.5\delta^2 \cdot (r\sigma^2 - 1) \quad (3.23)$$

Second, having solved this, we can (at last!) rewrite Farm's optimisation problem as an unconstrained maximisation problem over a single variable:

$$\max_{\delta} \pi(\delta) = (p - \delta) \cdot \delta - 0.5\delta^2 \cdot (r\sigma^2 - 1) \quad (3.24)$$

$$\pi(\delta) = p \cdot \delta - 0.5\delta^2 \cdot (1 + r\sigma^2) \quad (3.25)$$

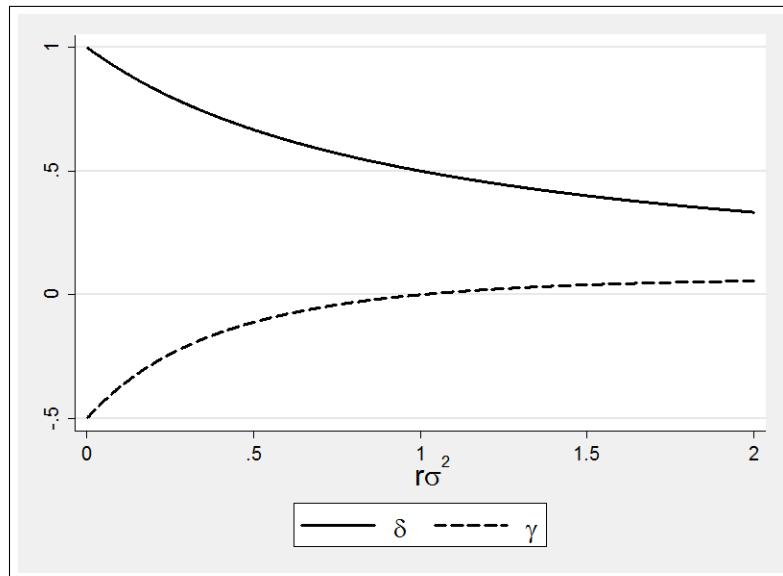
$$\therefore \delta^* = \frac{p}{1 + r\sigma^2} \quad (3.26)$$

$$\therefore \gamma^* = 0.5p^2 \cdot \left[\frac{r\sigma^2 - 1}{(r\sigma^2 + 1)^2} \right] \quad (3.27)$$

$$\therefore \pi^* = 0.5 \cdot \left(\frac{p^2}{1 + r\sigma^2} \right). \quad (3.28)$$

3.3 Understanding the solution

Figure 3.1: (δ^*, γ^*) versus $r\sigma^2$ (for $p = 1$)



Equations 3.26, 3.27 and 3.28 are very interesting, for several reasons. First, note that r and σ^2 only ever enter multiplicatively: it is their combined effect that matters. Second, note that, as $r\sigma^2$ increases, the optimal contract places ever more weight on the lump-sum

payment (γ^*), and less weight on the per-unit payment (δ^*); Figure 3.1 illustrates, for the case $p = 1$.

This emphasises again the critical importance of risk aversion under hidden action: if the agent is more risk averse, the principal optimally chooses to shift less risk onto the agent. In order to act effectively as an incentive contract, the wage contract is (implicitly) also acting as an *insurance contract*. Indeed, in the limiting case $r\sigma^2 \rightarrow \infty$, the per-unit payment tends towards zero – but this also means that the agent’s effort tends towards zero, along with the principal’s expected profit. (Incidentally, we should note at this point that equation 3.28 shows that Farm profit from contracting is always positive; this confirms that, as we assumed earlier, Farm will always prefer to contract than not to contract.)

Third, consider the limiting case as Worker tends to being risk neutral (that is, the limit as $r \rightarrow 0$). In that case, we have $\delta^* = p$, $\gamma^* = -0.5p^2$ and $\pi^* = 0.5p^2$. This is a very interesting case – because Farm is choosing to pass on *all* of the marginal revenue from each sale (i.e. p). What’s more, Farm is setting $\gamma^* < 0$; that is, the Farm is *taking* money from Worker to produce on the plot! We can describe this as ‘selling the firm’ (or, if you like, ‘selling the farm’): because there is nothing to be gained from implicit insurance, Farm optimises by fully incentivising Worker, and making all of its profit from the transfer of the asset.

3.4 Extending the model to think about multi-tasking

In the previous game, Worker was required to decide how much effort to exert on a *single* task (namely, working on the farm), where that effort was measured imprecisely through its effect on output. In the real world, however, many workers need to decide how to divide their time between several alternative tasks – and organisations need to design contracts to incentivise an appropriate balance of different responsibilities. These different responsibilities often differ in how precisely they can be monitored; as an academic, for example, I am required to divide my time between research (whose outputs can be precisely measured, through academic publications), teaching (whose outputs can be measured less precisely, through student feedback), and administration (which is very difficult to measure well).

We will now extend our earlier model, to allow for these kind of multi-tasking considerations. In doing so, we will follow closely the model of Feltham and Xie (1994), which itself follows the work of Holmström and Milgrom (1991). Specifically, we will keep the same structure of the game – including the solution concept of Subgame Perfect Equilibrium – and will make several modifications.

Introducing a second task: We will now allow Worker to exert effort on either (or both) of two tasks; we will denote this effort as a_1 (for task 1) and a_2 (for task 2). Continuing our agricultural example, we can imagine perhaps that the farm has now expanded

to become a tourist attraction,⁸ and that `Worker` must decide how much effort to spend working the fields (task 1) and how much to spend selling souvenirs to tourists (task 2).⁹

We assume that costs for `Worker` are again quadratic in effort; specifically, we allow costs to be quadratic in the effort for each task separately:

$$C = 0.5a_1^2 + 0.5a_2^2. \quad (3.29)$$

Further, we keep the assumption that `Worker` has exponential utility: $u(w) = -\exp(-rw)$.

Introducing a new performance measure: The famous management theorist Peter Drucker is quoted as having said that ‘what gets measured gets managed’. Previously, we assumed that `Farm` could observe only total output. We now extend this idea, to allow `Farm` instead to measure some variable y , which is a weighted sum of `Worker` effort on task 1 and task 2 (and which, again, is measured with noise):

$$\begin{aligned} y &= \beta_1 \cdot a_1 + \beta_2 \cdot a_2 + \varepsilon_y; \\ \varepsilon_y &\sim \mathcal{N}(0, \sigma_y^2). \end{aligned}$$

As before, we impose that `Farm` offers a linear piece-rate contract; this contract is written in terms of the performance measure (y):

$$w = \gamma + \delta \cdot y.$$

Output for `Farm`: Finally, we assume that the output for `Farm` (that is, before paying `Worker`) is simply the sum of the effort on task 1 and task 2, plus a random shock; we denote this as x :

$$\begin{aligned} x &= a_1 + a_2 + \varepsilon_x; \\ \varepsilon_x &\sim \mathcal{N}(0, \sigma_x^2) \end{aligned} \quad (3.30)$$

(Of course, we could also apply weights to a_1 and a_2 in equation 3.30 – but I don’t think this adds anything useful to the model.)

Model solution: We can solve this model using the same series of steps as we used in the first half of this lecture...

⁸ See, for example, <https://www.biggpineapple.com.au/>.

⁹ We might think that working in the fields is much more exhausting than working in the gift shop – which, one hopes, would be air-conditioned. The cost function in equation 3.29 treats task 1 and task 2 equivalently; this is something that we could change, but we will ignore this issue for simplicity now.

Step 1: Solve for Worker's optimal effort, conditional on accepting the contract:

$$\max_{a_1, a_2 \geq 0} \gamma + \delta\beta_1 a_1 + \delta\beta_2 a_2 - 0.5a_1^2 - 0.5a_2^2 - 0.5r\delta^2\sigma_y^2.$$

By differentiating, we obtain:

$$\begin{aligned} a_1^* &= \delta\beta_1; \\ a_2^* &= \delta\beta_2. \end{aligned}$$

Step 2: Solve for Worker's decision to accept or reject the contract: If Worker accepts the contract, (s)he will therefore get a certainty equivalent of:

$$\begin{aligned} &\gamma + \delta^2\beta_1^2 + \delta^2\beta_2^2 - 0.5\delta^2\beta_1^2 - 0.5\delta^2\beta_2^2 - 0.5r\delta^2\sigma_y^2 \\ &= \gamma + 0.5\delta^2 \cdot (\beta_1^2 + \beta_2^2 - r\sigma_y^2) \end{aligned}$$

Therefore, Worker should accept the contract if this expression is at least as good as the outside option (having certainty equivalent of zero).

Step 3: Solve for Farm's optimal piece-rate contract: Again, we start by solving for γ , determined so as to keep Worker indifferent between accepting and not accepting the contract:

$$\gamma^* = 0.5\delta^2 \cdot (r\sigma_y^2 - \beta_1^2 - \beta_2^2).$$

Substituting, we can write the expected profit for Firm as:

$$\begin{aligned} &p \cdot (a_1 + a_2) - \delta \cdot (\beta_1 a_1 + \beta_2 a_2) - 0.5\delta^2 \cdot (r\sigma_y^2 - \beta_1^2 - \beta_2^2) \\ &= p\delta \cdot (\beta_1 + \beta_2) - 0.5\delta^2 \cdot (\beta_1^2 + \beta_2^2 + r\sigma_y^2). \end{aligned}$$

Differentiating, we obtain:

$$\delta^* = p \left(\frac{\beta_1 + \beta_2}{\beta_1^2 + \beta_2^2 + r\sigma_y^2} \right),$$

such that the expected profit for Firm is:

$$\begin{aligned} \pi^* &= p\delta \cdot (\beta_1 + \beta_2) - 0.5\delta^2 \cdot (\beta_1^2 + \beta_2^2 + r\sigma_y^2) \\ &= 0.5p^2 \left[\frac{(\beta_1 + \beta_2)^2}{\beta_1^2 + \beta_2^2 + r\sigma_y^2} \right]. \end{aligned}$$

3.5 Understanding the solution

You should check that you understand the intuition for this solution. What drives `Worker` effort (a_1^* and a_2^*), and why? How does σ_y^2 affect γ^* and δ^* , and why? (This question should be very familiar from the single-task model that we considered in the first half of the lecture.) And whatever happened to σ_x^2 ?

Two aspects of the solution are particularly worth noting...

Two reasons for welfare loss: You should be able to show that the maximum possible expected profit for `Firm` occurs when $r\sigma_y^2 = 0$, and $\beta_1 = \beta_2$; in that case, expected profit is:¹⁰

$$\pi = p^2.$$

Therefore, `Farm`'s loss of profit, relative to this first-best case, is:

$$p^2 - 0.5p^2 \left[\frac{(\beta_1 + \beta_2)^2}{\beta_1^2 + \beta_2^2 + r\sigma_y^2} \right] = p^2 \cdot \left(\frac{\overbrace{0.5(\beta_1 - \beta_2)^2}^{\text{loss due to misalignment}} + \overbrace{r\sigma_y^2}^{\text{loss due to noise}}}{\beta_1^2 + \beta_2^2 + r\sigma_y^2} \right). \quad (3.31)$$

As Feltham and Xie (1994) explain:

The key point here is that if the gross payoff to the principal is not contractible information, then risk neutrality or a noiseless performance measure are not sufficient to achieve the first-best result. The performance measure must also be perfectly congruent.

Optimal piece-rates and misalignment of performance measures: Finally, let's consider the special case in which $r\sigma_y^2 = 0$; that is, let's consider the case in which `Worker` is risk-neutral. Let's also set $p = 1$ for algebraic simplicity. Then δ^* becomes:

$$\delta^* = \left(\frac{\beta_1 + \beta_2}{\beta_1^2 + \beta_2^2} \right)$$

Recall that, for any two vectors α and β , having angle θ between them, we know that:¹¹

$$\cos \theta = \frac{\alpha \cdot \beta}{\|\alpha\| \times \|\beta\|},$$

¹⁰ Why? What's the intuition for `Firm` preferring $\beta_1 = \beta_2$ in this model?

¹¹ 'Recall', in this context, means 'Personally, I should have remembered this, but had long ago forgotten it'...

where (for example) $\|\beta\|$ is the length ('norm') of the vector – so, if $\beta = (\beta_1, \beta_2)$, then $\|\beta\| = \sqrt{\beta_1^2 + \beta_2^2}$.

In our case, profit for Farm is simply the sum of effort on the two tasks (plus noise); we can therefore specify $\alpha = (1, 1)$. We can specify $\beta = (\beta_1, \beta_2)$. Therefore, it follows that:

$$\cos(\theta) = \frac{\beta_1 + \beta_2}{\sqrt{2} \cdot \sqrt{\beta_1^2 + \beta_2^2}}.$$

Therefore, we can say:

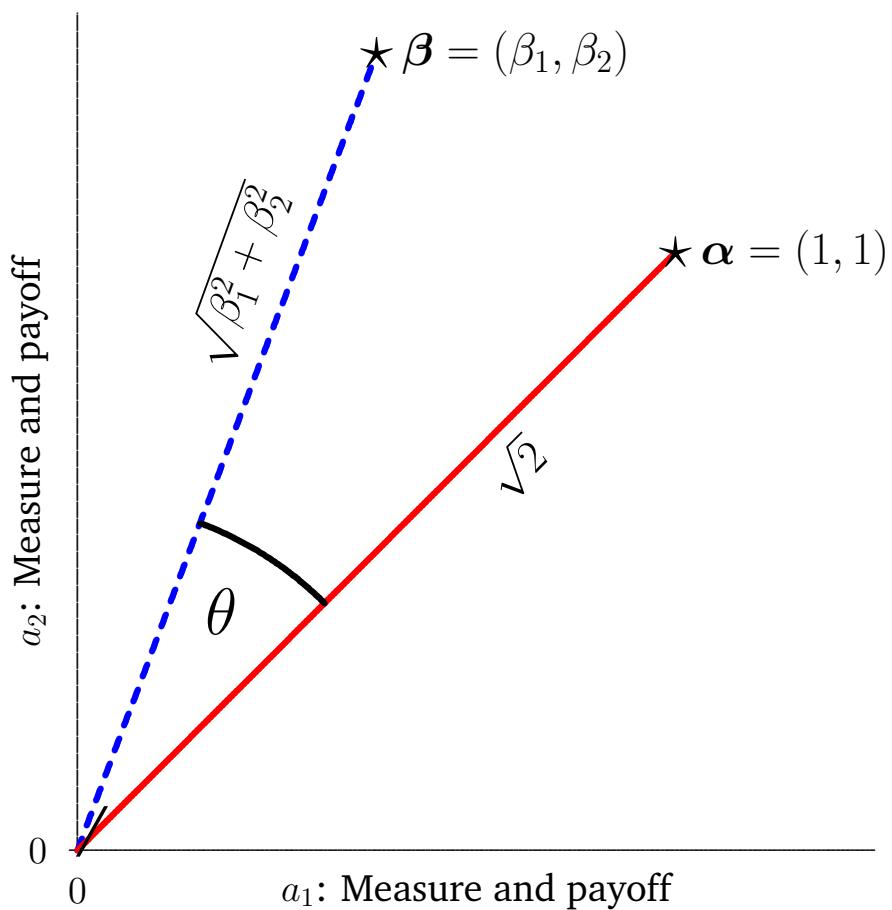
$$\delta^* = \frac{\|\alpha\|}{\|\beta\|} \cdot \cos(\theta).$$

As Gibbons (2010) explains (where I substitute our current notation):

There are two important features in δ^* : scaling and alignment, reflected by $\|\alpha\|/\|\beta\|$ and $\cos(\theta)$, respectively. Scaling is intuitive but uninteresting. For example, if β_1 and β_2 are both much larger than 1, then the efficient contract puts a small bonus rate on y , as shown by $\|\alpha\|/\|\beta\|$. Alignment, however, is the key to the model. As one example, if the α and β vectors lie almost on top of one another (regardless of their lengths), then the incentives created by paying on y are valuable for increasing x . As a second example, if the α and β vectors are almost orthogonal to each other, then the incentives created by paying on y are almost useless for increasing x . More generally, the efficient contract has a larger bonus rate δ^* when α and β are more closely aligned, as measured by $\cos(\theta)$.

Figure 3.2 illustrates.

Figure 3.2: Alignment of measures and payoffs under risk-neutrality



4 Signalling

Lecturer: Simon Quinn

Reading:

- Mas-Colell, Whinston and Green (1995), pages 450 – 460.
- Lecture 4 of Dr Meyer's course on 'Bargaining, Contracts and Theories of the Firm'.

4.1 Markets and development

There are many differences between developed economies and developing. Arguably, one of the most important distinctions lies in the quality of institutions of legal regulation and enforcement. In many different ways — and for many different reasons — it is generally much easier to guarantee the quality of goods and services sold in developed economies. Partly, this may be thanks to stricter regulatory standards; it may also be due to faster and more certain resolution of legal disputes.

For this reason, it is often the case in developing economies that one party in a transaction knows much more than the other about the quality of the good or service being sold. Of course, this is also true in developed economies — where this literature was born. But there are many reasons to think that this is *particularly* true in developing economies; for this reason, I think that such 'information asymmetries' are very important for understanding incentives in poor countries.

In this lecture, we will consider a simple model of '*signalling*', in which an agent may (or may not) manage to persuade a principal of something that only the agent directly observes. This is a natural extension to our previous lecture, on moral hazard. In that case, we considered a hidden *action* — namely, the effort spent by a worker on a farm. In this lecture, we consider a hidden *characteristic* — something that is not itself a choice variable, but whose asymmetric observability creates important incentives.

4.2 Unproductive education as a signal

To fix ideas, we will consider a simple model, in which a single agent (`Worker`) tries to signal his or her quality to a single employer (`Firm`). We will treat this single employer as being drawn from a large pool of potential employers; for this reason, we will assume that (i) the employer optimises, but (ii) in doing so, the employer earns zero profit. This is directly analogous to our sharecropping model, in which the employee was drawn from a large pool of other potential employees (and, therefore, received zero utility in equilibrium). (Some versions of this model explicitly introduce a second firm, and assume that the firms engage in a Bertrand-like competition to hire the worker. We could take this path

if we wanted to be very clear about the microfoundations of the zero-profit condition — but, for the sake of simplicity, I would prefer to stick to a single firm.)

Who are the players? We have three players: Nature, Worker and Firm.

What actions can they take, and when? Nature draws a quality for worker: $n \in \{H, L\}$. This is drawn randomly, such that $\Pr(n = H) = q$. Nature reveals this information to Worker, but not to Firm.

Worker then chooses a level of education, $e \geq 0$. Firm observes this level of education.

Firm then makes an offer to Worker, w . Worker decides whether to accept or reject.

What payoff does each player get for all combinations of player actions? Nature receives no payoff.

Worker has utility $U(w, e, n)$, such that:

$$U(w, e, n) = w - C(n, e); \quad (4.1)$$

$$\frac{\partial C(n, e)}{\partial e} > 0; \quad (4.2)$$

$$\frac{\partial^2 C(n, e)}{\partial e^2} > 0; \quad (4.3)$$

$$C(H, e) < C(L, e); \quad (4.4)$$

$$\frac{\partial C(H, e)}{\partial e} < \frac{\partial C(L, e)}{\partial e}. \quad (4.5)$$

That is, Worker prefers a higher wage to a lower wage, and prefers less education to more (with increasing marginal cost). Low-quality workers suffer higher cost of education than high-quality workers. And — critically — the *marginal cost* of education is always lower for high-quality workers than low-quality. This last condition (in equation 4.5) is sometimes known as the ‘*single-crossing property*’, and will be fundamental to our story. Figure 4.1 illustrates, by showing an indifference curve for each quality of Worker (where the two arrows indicate the direction in which, for each player, utility is increasing).

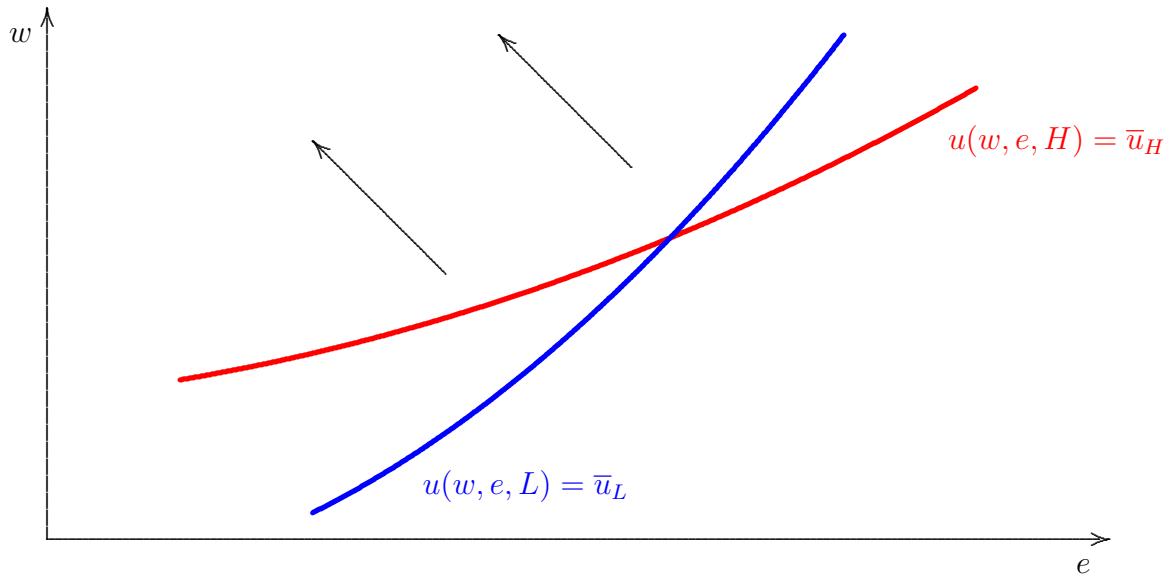
Firm receives profit, y , which is increasing in the quality of Worker and (weakly) increasing in Worker’s education:

$$y(H, e) > y(L, e); \quad (4.6)$$

$$\frac{\partial y(n, e)}{\partial e} \geq 0. \quad (4.7)$$

We will return to this shortly.

Figure 4.1: Indifference curves for Worker



What is the solution concept? The solution concept is Perfect Bayesian Equilibrium ('PBE'), where we restrict attention to pure strategies. In this context, a PBE is a combination of (i) a strategy for Worker, $e^*(n)$, (ii) a strategy for Firm, $w^*(e)$, and (iii) a belief for Firm, $\mu(e) \in [0, 1]$, such that:

1. Firm acts as if $\Pr(n = H | e) = \mu(e)$.
2. For each $e \geq 0$, $w^*(e)$ maximises the payoff of Firm, given its belief $\mu(e)$.
3. For each $n \in \{H, L\}$, $e^*(n)$ maximises the utility of Worker, given Firm's strategy $w^*(e)$.
4. For each $e \geq 0$ and each $n \in \{H, L\}$, if $\Pr(e^*(n) = e > 0)$, then $\mu(e)$ must be formed using Bayes' Rule and the strategy $e^*(n)$. Loosely, we can say that 'Firm forms its beliefs using Bayes' Rule wherever possible'. That is, we can say that, 'For any equilibrium that might occur, it must be that Firm's beliefs about n turn out to be correct. But Firm can hold *any* beliefs for outcomes that do not occur in equilibrium.'

For each player, what is the set of best-responses to all combinations of other player actions? What, therefore, are the solutions? This, in a sense, is where the story really begins — and this is where we will focus the remainder of the lecture. Specifically, the rest of the lecture will consider various permutations on the shape of y and μ . However, we will maintain throughout the assumption that (i) as noted above, Firm acts as if $\Pr(n = H | e) =$

$\mu(e)$, and (ii) Firm makes zero profit in expectation. Therefore, we can say immediately that:

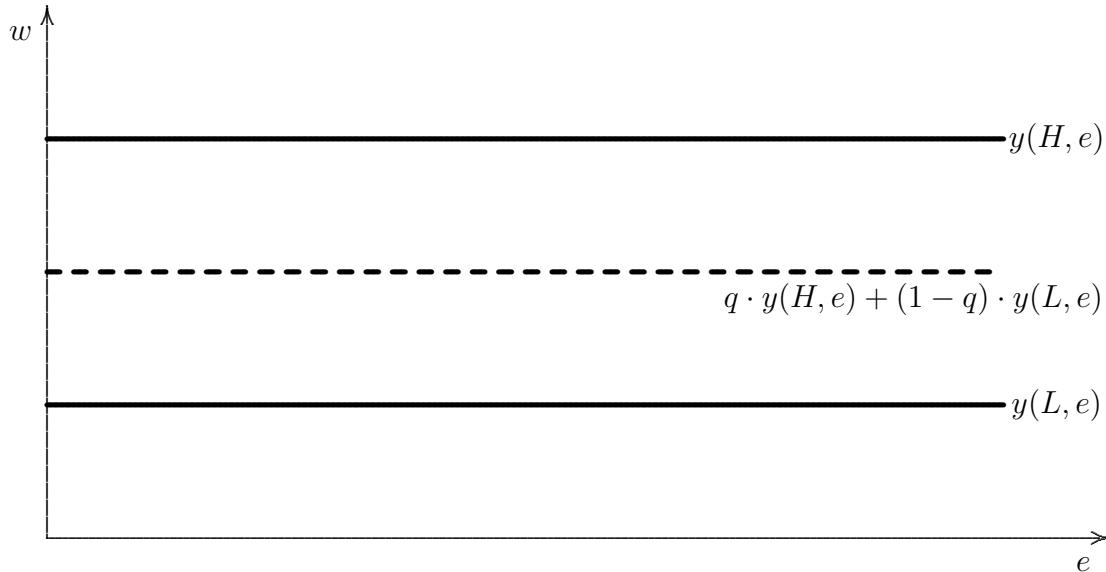
$$w^*(e) = \mu(e) \cdot y(H, e) + (1 - \mu(e)) \cdot y(L, e). \quad (4.8)$$

Of course, we could relax this assumption if, for example, we wanted to think about signalling under firm market power. But, as we will soon find, this ‘simple’ model already generates some very intricate solutions...

4.2.1 Two benchmark cases

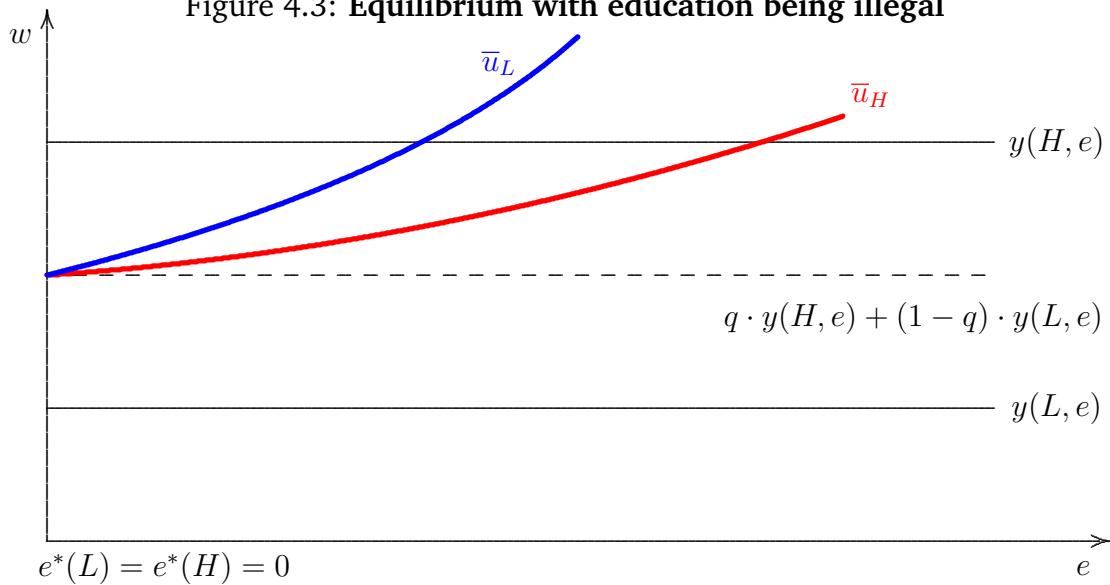
For most of this lecture, we will use an assumption that is very simple (and, indeed, very cynical): education is *useless* for firm productivity. That is, we assume that Firm prefers to hire Worker of type H than of type L — but does not otherwise care how much education that worker has. Figure 4.2 illustrates Firm profit for types $n = H$ and $n = L$. It also shows the expected Firm profit if Firm knows only that it is contracting with a pool of workers (of which proportion q have type H).

Figure 4.2: Firm profit under unproductive education



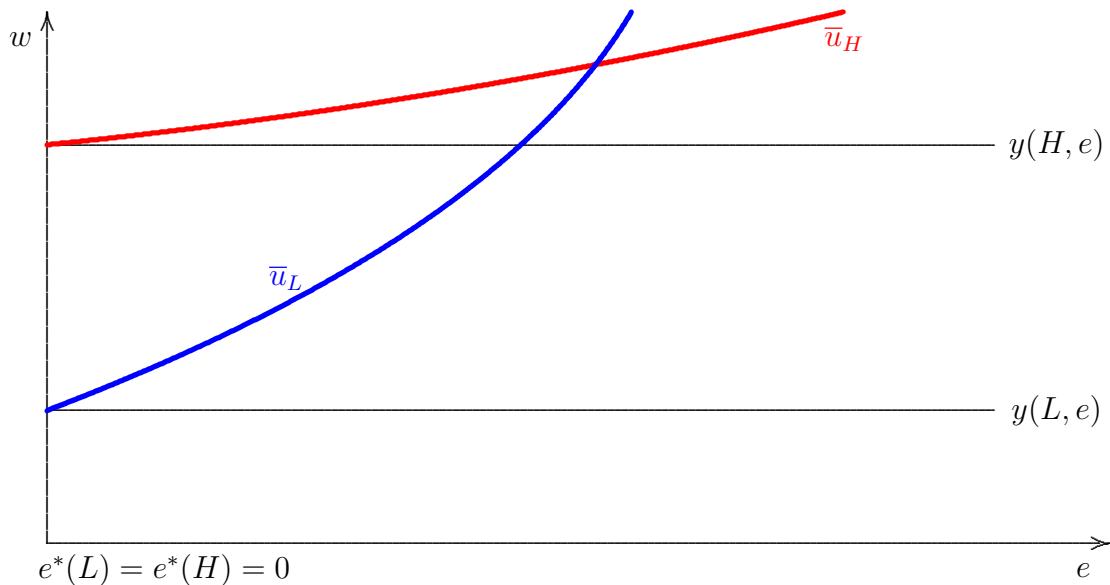
Now let’s consider a benchmark case in which education is illegal. Figure 4.3 illustrates the equilibrium: by assumption, we have $e^*(L) = e^*(H) = 0$. Since Firm has no basis for differentiating workers of different quality, its expected profit is $q \cdot y(H, e) + (1 - q) \cdot y(L, e)$. Since we assume that it must make zero profit in equilibrium, this is also the wage.

Figure 4.3: Equilibrium with education being illegal



Now let's consider another benchmark case, in which Firm can observe n (i.e. Worker quality) directly. (This, of course, is a deviation from our model; we will return to our model soon, but this is a useful benchmark to consider.) Let's also make education legal again.

Figure 4.4: Equilibrium with full information (and education being legal)



Hopefully, the intuition for this solution is simple: since education has no value to the firm, no worker gets any education — and each worker is paid the full value of production.

4.2.2 Separating equilibria

Let's now turn to our model — without the benchmark assumptions of either full information or illegal education. Figure 4.5 shows one possible PBE of this model. This is a '*separating equilibrium*' — meaning that type $n = L$ and type $n = H$ invest in different levels of education, and are therefore paid different wages. *Even though Worker type is not directly observed by Firm, it can be inferred perfectly in equilibrium, because different types fully 'separate'.*

So how, then, is this a PBE? First, notice that — for the reasons we discussed earlier — it does not make sense for type L to invest in *any* education, given that Firm will know its type. We can therefore immediately solve $e^*(L) = 0$, and $w^*(0) = y(L, 0)$. Second, notice that there is a level of education (which I denote \tilde{e}) such that Worker of type L is indifferent between (i) receiving no education and wage $y(L, 0)$ and (ii) receiving education \tilde{e} and a wage $y(H, \tilde{e})$. This is a critical point — because, for any education level $e > \tilde{e}$, Worker of type L would *prefer to get no education and to be paid the lower wage*. That is, *even if* a Worker of type L could somehow 'pretend' to be type H by choosing $e > \tilde{e}$, it would choose not to do so. It is therefore 'credible' for Worker of type H to choose $e = a$ — in a sense that is broadly analogous to the concept of credible actions in Subgame Perfect Equilibrium. We will return to this idea repeatedly in this lecture.

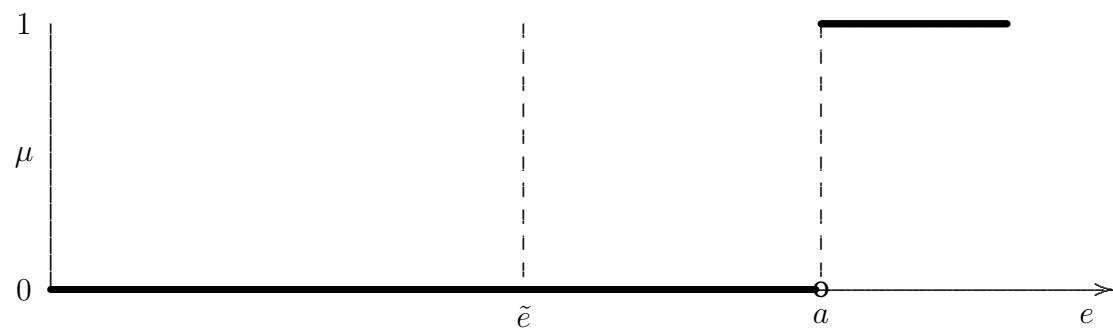
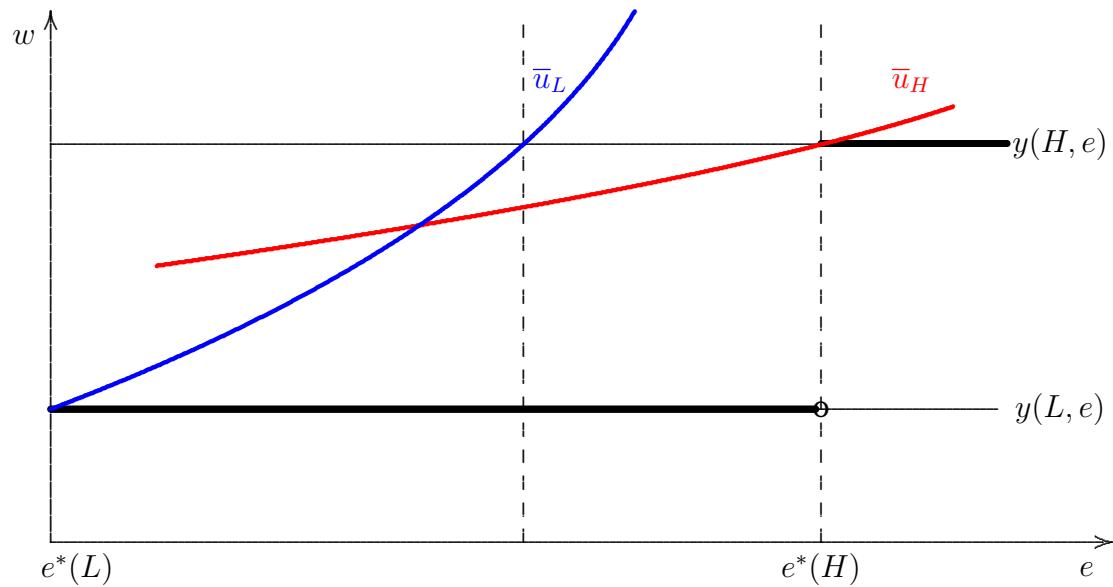
Third, suppose that Firm chooses some arbitrary education level $a > \tilde{e}$ and decides, "If Worker can get education $e \geq a$, Worker must be type H . Otherwise, Worker must be type L ." That is, suppose that Firm adopts the following function:

$$\mu(e) = \begin{cases} 0 & \text{if } e < a; \\ 1 & \text{if } e \geq a. \end{cases} \quad (4.9)$$

As Figure 4.5 shows, Worker of type H optimises by choosing $e(H) = a$. This is the minimum level of education that Worker of type H can get in order to obtain a wage $w = y(H, e)$. This is a PBE: both players are optimising, Firm is using $\mu(e)$ as the subjective conditional probability of observing type H , and this belief is correct *for all of the cases observed in equilibrium* (namely, $e = e^*(L)$ and $e = e^*(H)$).

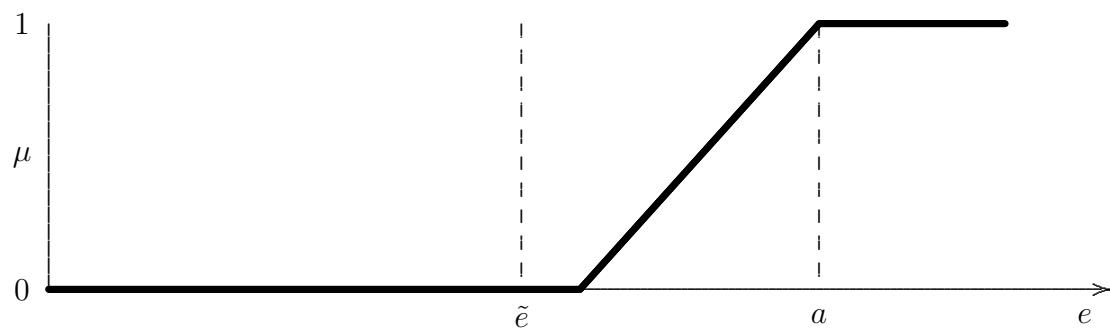
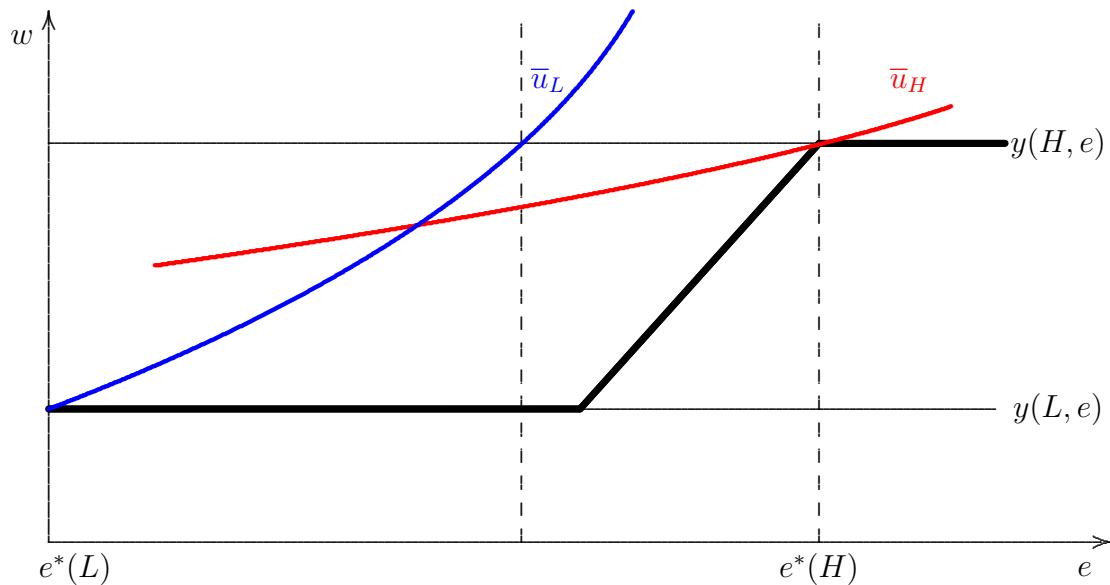
Check that you understand the intuition for this result: does this reasoning mean that a can take *any* value $a \geq e^*(H)$? Or is there somehow an upper bound on the values of a that will support a separating equilibrium of this kind?

Figure 4.5: A separating equilibrium with unproductive education



It should immediately be clear that the equilibrium in Figure 4.5 is not the unique equilibrium to this game. For one thing, there are lots of different values of a that can produce this kind of behaviour. There are also lots of different kinds of functions μ that will produce this behaviour. For example, Figure 4.6 shows an equilibrium with a different function μ , but with *identical* $e^*(L)$ and $e^*(H)$ to Figure 4.5. Since the function μ is a necessary part of the definition of a PBE, we should consider the result in Figure 4.6 to be a *different* equilibrium to that in Figure 4.5.

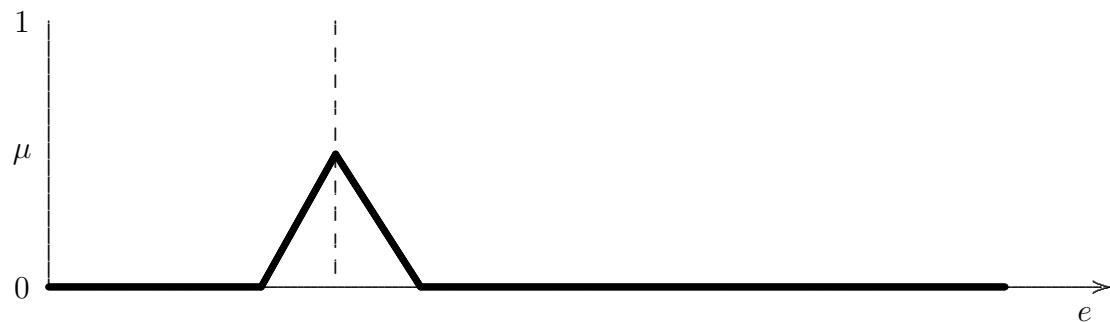
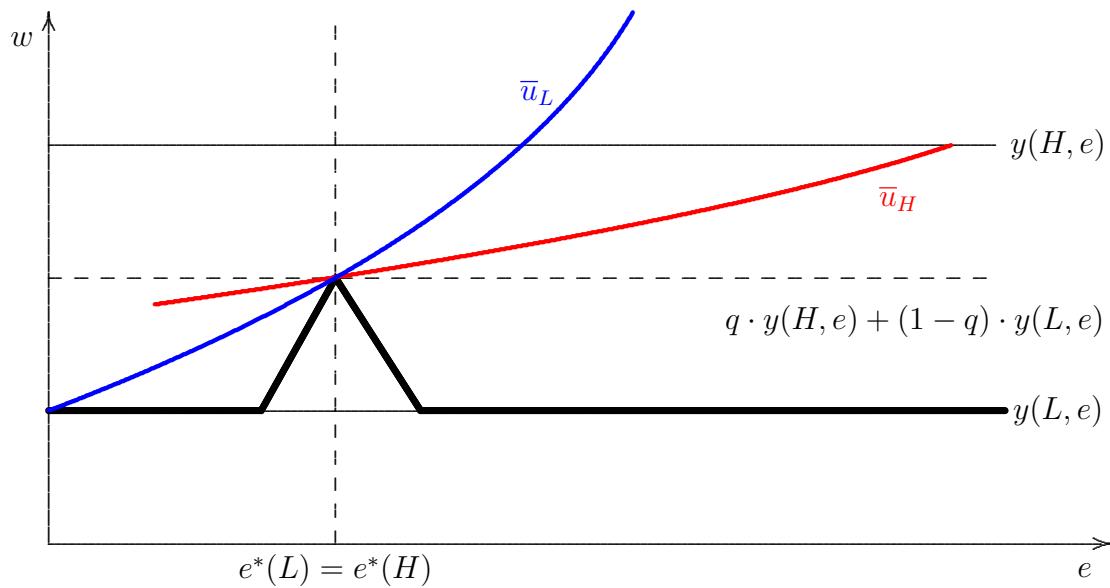
Figure 4.6: Another separating equilibrium with unproductive education



4.2.3 Pooling equilibria

I hinted earlier that there might be some value of a so large that even Worker of type H would prefer $e = 0$; that is, there might be some situation in which *both* types of Worker would rather be seen as ‘one of the pool’. Figure 4.7 shows one such equilibrium. Here, Firm believes (for reasons best known to itself) that workers with low education *and* workers with high education are of type L. Here, there is an equilibrium education level, $e^*(L) = e^*(H)$, representing an optimal response by both types of worker. This means that, *at that equilibrium point*, the true probability of type H is q . Firm therefore pays $w^*(e^*) = q \cdot y(H, e^*) + (1 - q) \cdot y(L, e^*)$ (which is what, in turn, induces the pooling behaviour by both type of Worker).

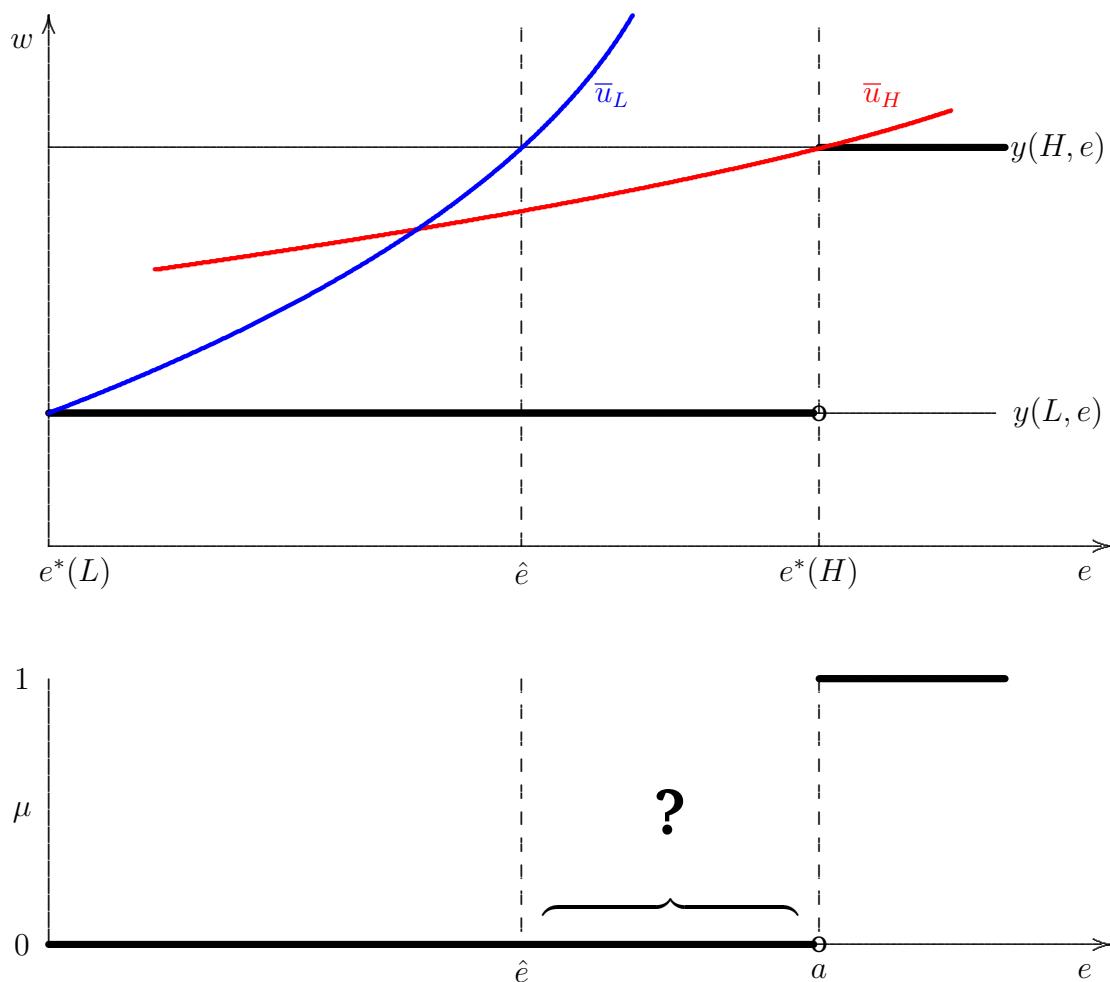
Figure 4.7: A pooling equilibrium with unproductive education



4.2.4 Restrictions on beliefs

Figure 4.7 is a very strange model of human behaviour. So too, perhaps, are the equilibria in Figure 4.5 and Figure 4.6. The reason should be clear: so far, the *only* restriction that we have placed on beliefs is that they must be correct *in equilibrium*. That is, we have equilibria supported by beliefs that could not themselves represent equilibrium outcomes. Or, to put it bluntly, our definition of PBE just doesn't give human decision-makers enough credit in the way that we form our beliefs.

Figure 4.8: A separating equilibrium with unproductive education



So how, then, should the model be extended? Let's look again at the separating equilibrium we considered earlier (now presented as Figure 4.8). Remember that this scenario allowed for multiple separating equilibria, because we allowed Firm to believe $\mu = 1$ for some arbitrary value $a \geq \hat{e}$. This, clearly, is a very strange belief. In particular, it is strange because it could *never* be an improvement for Worker of type L to choose $e > \hat{e}$, regardless

of Firm's strategy. That is, Worker of type L will always prefer $e = 0$ to $e > \hat{e}$ — yet Firm reasons that, if it observes $e \in [\hat{e}, a]$, it should be certain that Worker is low type (that is, $\mu(e) = 0 \forall e \in [\hat{e}, a]$).

There are many ‘restrictions on beliefs’ that can improve the model to rule out this kind of outcome. We will take what is, arguably, the simplest approach, by adding a new requirement to our equilibrium concept:

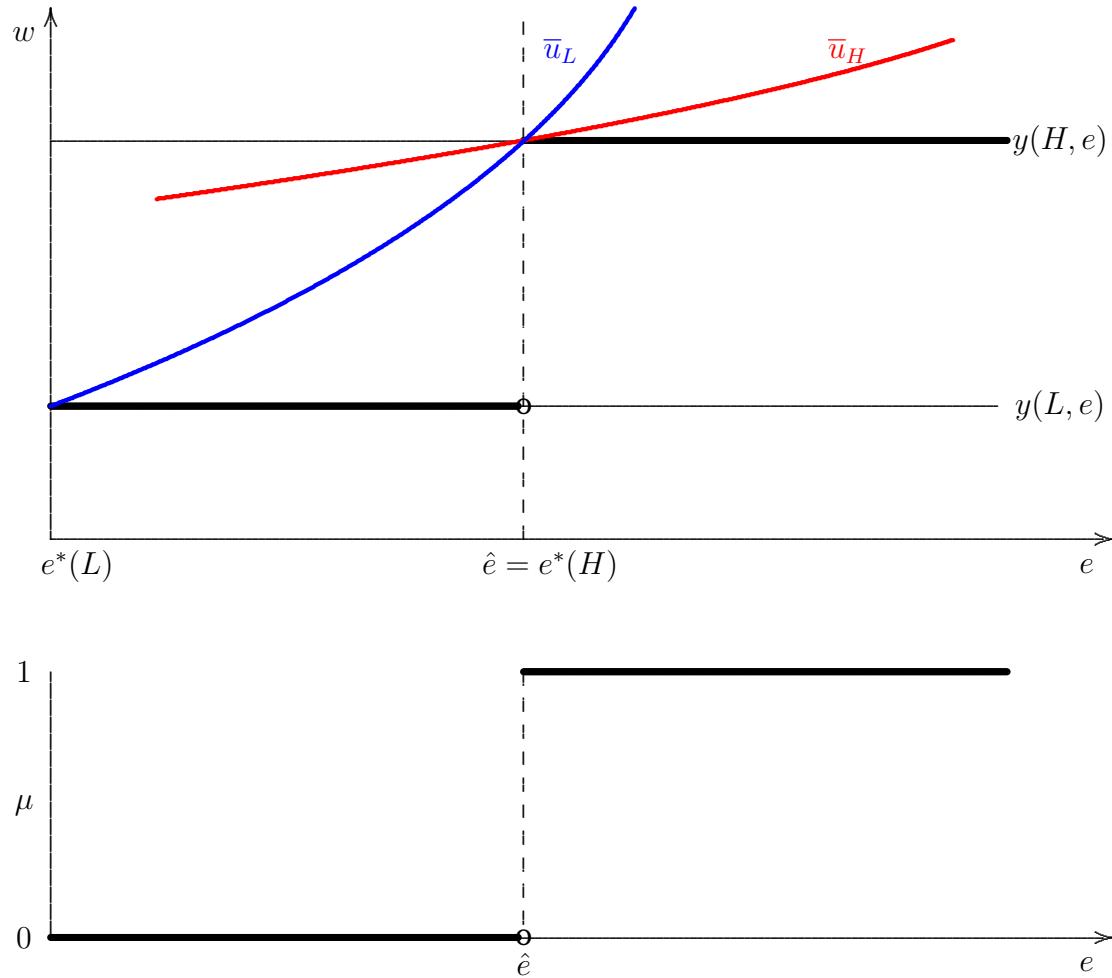
5. Suppose that Firm observes a deviation from equilibrium, to education level e . Suppose that an action could never result in a higher payoff than the equilibrium action for player j , but could result in a higher payoff to k for some beliefs of Firm. Then:

$$\mu = \begin{cases} 1 & \text{if } j = H \\ 0 & \text{if } j = L. \end{cases} \quad (4.10)$$

This is sometimes known as the ‘Cho-Kreps Intuitive Criterion’. Let’s apply the Intuitive Criterion to Figure 4.8. Under this additional assumption, Firm now reasons, “Suppose I observe Worker choosing some e between \hat{e} and a . How could any Worker possibly want to choose such e ? This e must decrease utility for type L, irrespective of how I respond. On the other hand, this e may increase utility for type H, if I assume that the worker is indeed of type H. Therefore, I should assume with certainty that this is type H.”

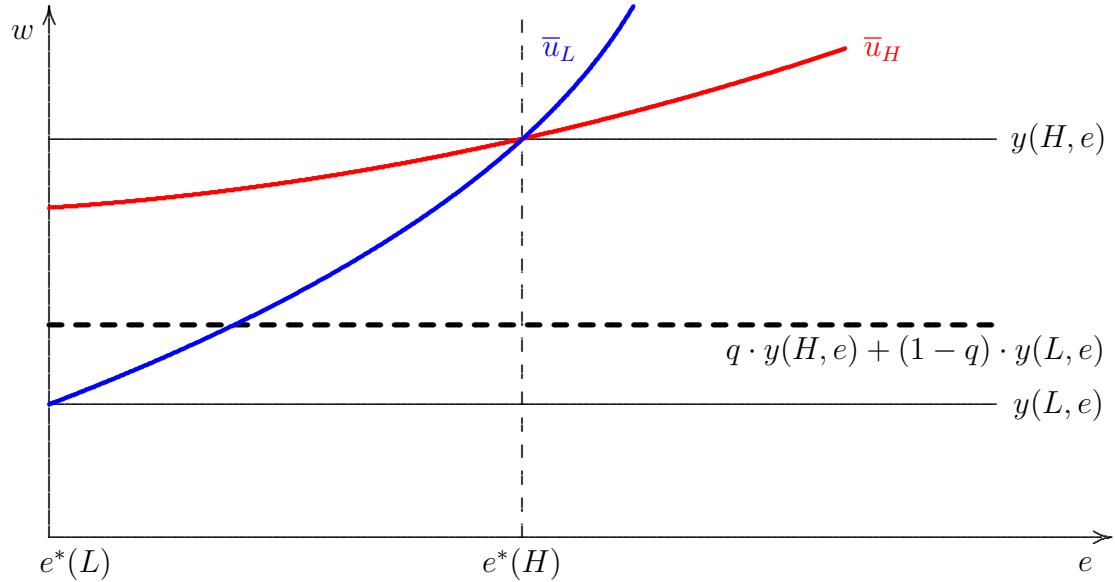
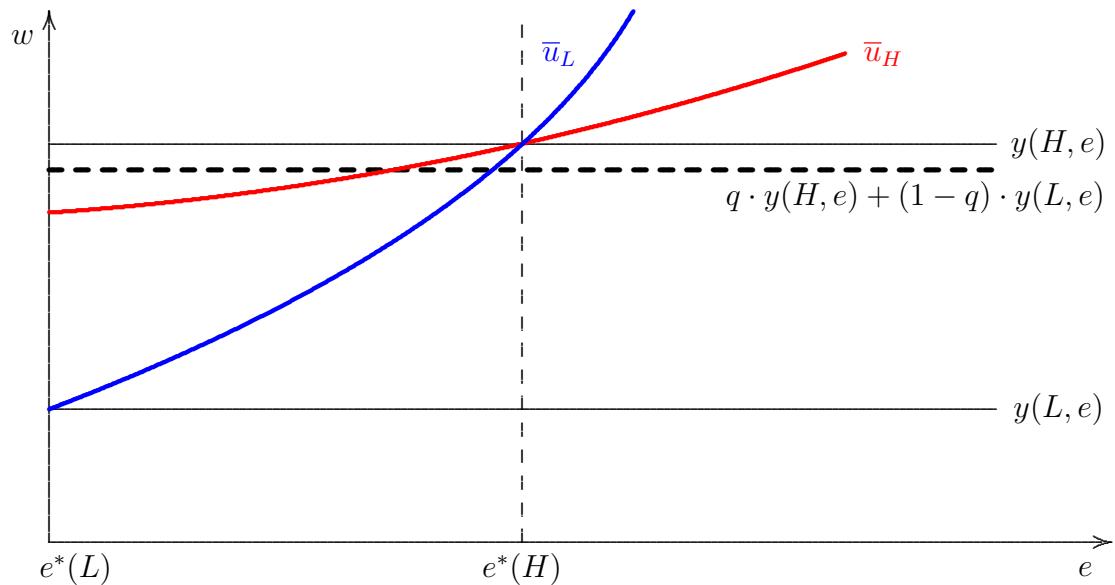
Figure 4.9 shows that, under this assumption, only one equilibrium now survives: where $e^*(L) = 0$ and $e^*(H) = \hat{e}$, and $\mu(e) = 1$ if $e = e^*(H)$ and $\mu(e) = 0$ if $e < e^*(L)$. (What value should $\mu(e)$ take if $e > \hat{e}$?) You should also confirm that no pooling equilibrium can survive the imposition of the Cho-Kreps Intuitive Criterion in this context.

Figure 4.9: A separating equilibrium under the Cho-Kreps Intuitive Criterion



4.2.5 Welfare and signalling

So is this separating equilibrium welfare enhancing for Worker, relative to the ‘illegal education’ benchmark? The answer is interesting: *it depends on q , the proportion of type H*. Figures 4.10 and 4.11 illustrate. In Figure 4.10, we see a case where q is low: it is relatively unusual to be type H, so type H prefers to be able to signal. (Of course, type L would prefer *not* to have signalling, so is worse off; the welfare consequences of signalling are therefore ambiguous.) However, consider Figure 4.11. In this case, q is large; it is very common to be type H, and very unusual to be type L. We now find that *both* type L and type H would prefer not to be allowed to signal! This makes intuitive sense: without signalling, both types could obtain a relatively high wage — whereas, under signalling, type H invests substantial resources to avoid being thought of as one of the (rare) low types.

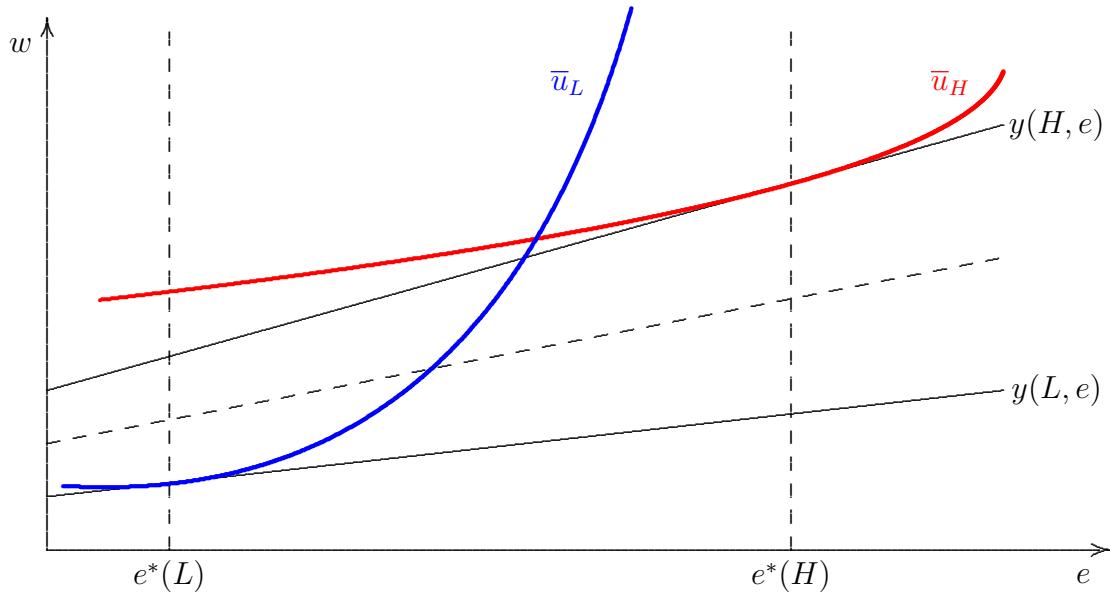
Figure 4.10: Signalling and welfare: Low q

 Figure 4.11: Signalling and welfare: High q


4.3 Productive education as a signal

So far, we have made the simplifying assumption that education has no intrinsic value for Firm. This has been very useful for introducing key concepts, but very few economists would accept this as a reasonable description of the real world. In this section, we will explore briefly how the earlier results can extend to a case where education is productive.

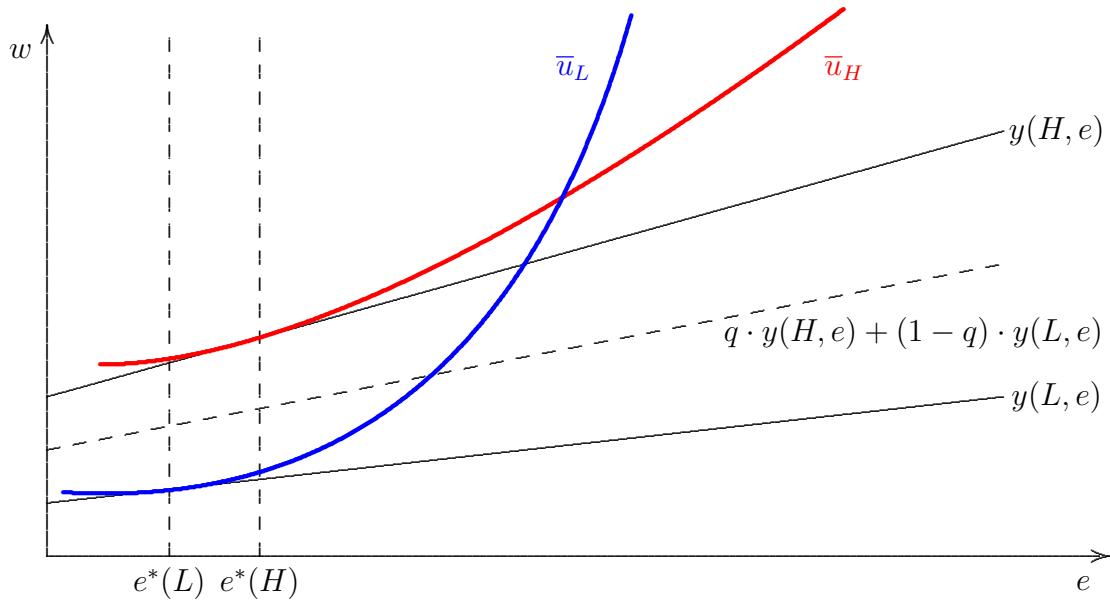
Figure 4.12 shows how we can treat education as productive: the isoquants $y(H, e)$ and $y(L, e)$ now slope upwards in (e, w) space. Figure 4.12 shows one other complication that can arise when education is productive: we can have ‘no-envy’ cases, in which type L and type H each chooses an optimal education level, and *neither* type would like to be confused for the other. In Figure 4.12, for example, type L optimally chooses e_L^* and type H optimally chooses e_H^* , *irrespective of any signalling value of education*. (It is straightforward to see how this result holds: for each type, the pair $(e, w^*(e))$ for the other type lies *below* the indifference curve at the type’s optimum.)

Figure 4.12: A full information outcome in a ‘no-envy’ case



We should rule out this kind of ‘no-envy’ case for the rest of our analysis — not because it is unreasonable (indeed, it may be a fair description of many real-world contexts), but because it removes entirely the asymmetric information problem. Instead, we will assume that the shape of the indifference curves is such that, in a full-information context, type L *would* envy the contract provided to type H. Figure 4.13 illustrates; note that the pair $(e_H^*, w^*(e))$ lies *above* the indifference curve for type L.

Figure 4.13: A full information outcome in an ‘envy’ case



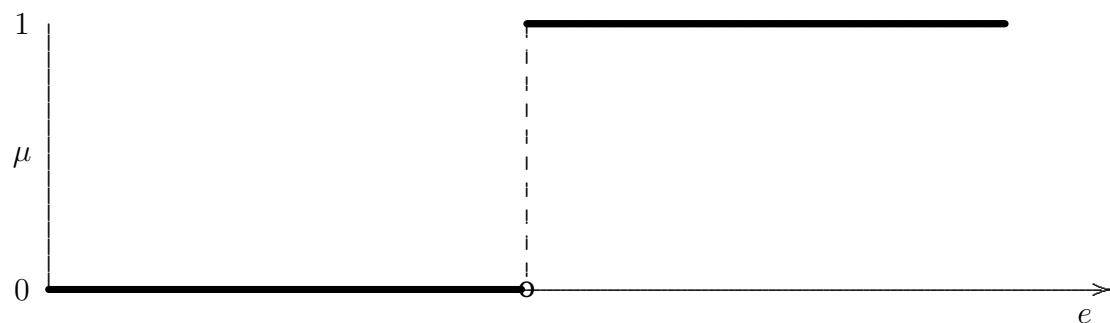
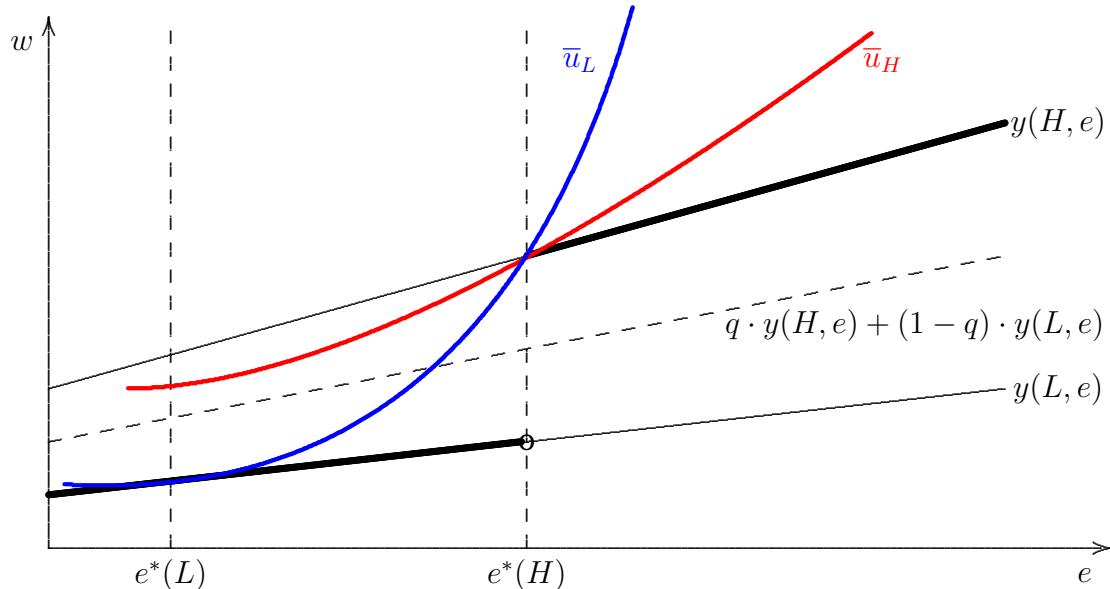
So how, then, should we solve the model in this case? For simplicity, we will maintain the Intuitive Criterion. In that case, the result from Figure 4.9 extends straightforwardly. Figure 4.14 shows how the solution structure there can extend to a case with productive education, with a unique outcome for $e^*(n)$ and $w^*(e)$.

4.4 Lessons for development

So what should we learn from all of this for development? I think there are several lessons (and I’m sure that you can draw several more). First, education can serve as a signal of unobservable ability, because it is a costly investment. However, this is almost certainly a ‘second best’ outcome compared to having a good technology for assessing and verifying employee quality. (After all, remember that the welfare analysis in section 4.2.5 assumed, for simplicity, that all of the costs of education are borne by the *worker*. In reality, we know that much of the cost of education is actually borne by the state — so not internalised by the worker. This is not an argument against public spending on education — but we should remember that the state may not have a compelling interest simply for funding private signals of ability.) Second, as we noted when discussing Akerlof’s ‘market for lemons’, it may even be the case that some markets fail completely because uninformed principals cannot adequately discern the quality of agents. In such cases, innovative policy interventions — for example, ‘active labour market policies’ like referral systems, internships, screening technologies and so on — may generate substantial welfare gains. Third, our theoretical results have interesting econometric implications. Suppose that we have a sample of data on education and earnings, and that the world is accurately described by Figure 4.9. In

that case, we *would* want to consider $w^*(e^*(H)) - w^*(e^*(L))$ as the ‘causal effect of education’ — because this is indeed the counter-factual for an individual deviating from $e^*(L)$ to $e^*(H)$. But this is a *private* return to education, rather than an increase in worker productivity being caused by education. Finally, we should remember that education is merely one illustration of the principles of costly signalling under information asymmetry. There are likely to be many other contexts in which costly choices can signal unobservable ability, including in developing economies. I leave them for you to consider.

Figure 4.14: A separating equilibrium in an ‘envy’ case with productive education



4.5 Class exercises: Hidden Actions and Signalling

- Suppose that there are two players, Farm and Worker. Each is risk neutral. Worker chooses to divide her time between two tasks; the effort on each task is a_1 and a_2 respectively, and the cost of effort is:

$$C = 0.5a_1^2 + 0.5a_2^2. \quad (4.11)$$

Farm is able to observe a noisy measure of the *total* effort on the two tasks:

$$y = a_1 + a_2 + \varepsilon_y; \quad (4.12)$$

$$\varepsilon_y \sim \mathcal{N}(0, \sigma_y^2). \quad (4.13)$$

Farm offers a linear piece-rate contract, in terms of the performance measure (y):

$$w = \gamma + \delta \cdot y. \quad (4.14)$$

The profit for Farm is a *weighted* sum of a_1 and a_2 , less the wage paid to Worker:

$$\pi = \alpha_1 \cdot a_1 + \alpha_2 \cdot a_2 - w. \quad (4.15)$$

The payoff for Worker is simply wage minus cost of effort:

$$U = w - C. \quad (4.16)$$

Worker has an outside option that guarantees utility of zero.

- Assume that Worker accepts the contract. Solve for a_1^* and a_2^* (that is, the effort that Worker spends on each task).
- Under what circumstances will Worker accept the contract?
- What contract will Farm offer? (That is, solve for γ^* and δ^*).
- Assume for the remainder of the problem that the length of the vector $\alpha = (\alpha_1, \alpha_2)$ is fixed at $\sqrt{2}$ (that is, so $\alpha_1^2 + \alpha_2^2 = 2$). What is the *maximum* value of δ^* that Farm might pay – and in what circumstance? What is the *minimum* value of δ^* that Farm might pay – and in what circumstance?
- What is the minimum *absolute* value of δ^* that Farm might pay (that is, $|\delta^*|$) – and in what circumstance will it do so?
- Explain the intuition for your result. Draw a graph for this problem that is similar to Figure 11.2 in the lecture notes; can we think of this solution in terms of ‘cos θ ’?
- Under what circumstances – if any – can the solution to this problem be described as ‘selling the farm’? Explain.

2. (*From the 2019 exam...*) Consider a small butcher shop in a developing country, trying to attract a large order from a luxury hotel. The hotel knows that there are two types of butcher ('high-quality' butchers, selling good meat, and 'low-quality' butchers, selling poor meat); it also knows that the proportion of high-quality butchers is q . The hotel cannot determine the quality of the meat before purchasing it — but can observe the effort that the butcher has spent in making the shop look good. Denote this effort by $e \geq 0$. The hotel will purchase all of the butcher's product. High-quality butchers have a utility function $U_H(p_H, e_H) = p_H - e_H$, where p_H is the total payment received by the high-quality butcher and e_H is the effort made. Low-quality workers have a utility function $U_L(p_L, e_L) = p_L - 2e_L$ (where p_L and e_L refer to the low-quality butcher). If the meat is high-quality, the hotel's revenues will increase by \$2,000; if the meat is low-quality, the hotel's revenues will not increase at all. Suppose that, when negotiating over the price, the butcher holds all of the bargaining power.

- (a) Consider the utility functions $U_H(p_H, e_H)$ and $U_L(p_L, e_L)$. Does the 'single-crossing property' hold? Give an intuitive explanation for why a high-quality butcher and a low-quality butcher might have these utility functions.
- (b) Explain what is meant here by the butcher holding all of the bargaining power. What price will a high-quality butcher receive in a separating equilibrium? What price will a low-quality butcher receive in a separating equilibrium? What price will the hotel pay in a pooling equilibrium? (You may assume that the hotel is risk-neutral and acts to maximise profit.)

Assume that, if $e \geq a$, the hotel will believe with certainty that the butcher is high-quality; otherwise, the hotel will believe with certainty that the butcher is low-quality.

- (c) Show that, for $a = 1100$, there is a Perfect Bayesian Equilibrium in which $e_H = 1100$ and $e_L = 0$.
- (d) What is the maximum value of a that can support a Perfect Bayesian Equilibrium? Explain.
- (e) What value(s) of a will support a Perfect Bayesian Equilibrium under the Cho-Kreps Intuitive Criterion? Explain.

Now assume that the hotel decides to ignore completely the presentation of the butcher shop; that is, it decides that the probability of a butcher being high-quality is q (regardless of the butcher's effort in making the shop look good).

- (f) Explain how this belief can support a pooling equilibrium. How does the butcher's utility differ between this case and the case of Perfect Bayesian Equilibrium with the Cho-Kreps Intuitive Criterion? (Explain both for the case of a low-quality butcher and for the case of a high-quality butcher.)

5 Auctions I

Lecturer: Simon Quinn

I plan to use this lecture slot to run a stylised auctions exercise, in preparation for the next lecture. (We will decide together whether we want to run this in person or on Zoom; to be discussed!)

6 Auctions II

Lecturer: Simon Quinn

Reading:

- Osborne, *An Introduction to Game Theory*, section 3.5 and section 9.6.

6.1 Auctions with perfect information

6.1.1 A first-price sealed bid auction

In this lecture, we will discuss auctions. There are three reasons to do this:

1. The auction is an important mechanism for buying and selling in many contexts — including developing countries.
2. Many other interesting economic situations can be treated as a kind of auction — for example, the basic structure of an auction problem can be useful for thinking about political lobbying.
3. More generally, auctions provide a useful application with which to consider a new solution concept: *Bayesian Nash Equilibrium*. This solution concept is useful for a very wide variety of contexts, in which one or more players is uncertain about some aspect of their opponents' payoff functions.

To begin, let's start by considering a simple example: an auction between two players, where each player bids simultaneously, and where the item is sold to the player with the higher bid at the price of that higher bid. We can refer to this auction as a 'first-price sealed bid auction'. This is a common way of auctioning many items; for example, we might think of firms as submitting sealed bids to purchase an asset from the government. Similarly, this is formally equivalent to a 'descending price' auction (sometimes known as a 'Dutch auction'). We will start by assuming that each player knows exactly the value that the other player attaches to the item being sold.

Who are the players? We have two players: Player 1 and Player 2.

What actions can they take, and when? Each player i chooses a bid, $b_i \geq 0$. These bids are chosen simultaneously.

What payoff does each player get for all combinations of player actions? Suppose that Player 1 values the item as v_1 , and that Player 2 values the item as v_2 ; these values are common knowledge between the players. For simplicity, let's assume that the auctioneer imposes an arbitrary 'tie-breaker' rule; if the bids are equal, Player 1 will be sold the

item. (We will also assume that $v_1 > v_2$. We don't *need* to assume that the tie-breaker rule favours the player with the higher valuation; however, this assumption is innocuous and is useful for the elegance of the solution.)

Therefore, we can write the payoff functions as follows:

$$\pi_1(b_1; b_2) = \begin{cases} v_1 - b_1 & \text{if } b_1 \geq b_2; \\ 0 & \text{if } b_2 > b_1. \end{cases} \quad (6.1)$$

$$\pi_2(b_2; b_1) = \begin{cases} v_2 - b_2 & \text{if } b_2 > b_1; \\ 0 & \text{if } b_1 \geq b_2. \end{cases} \quad (6.2)$$

For each player, what is the set of best-responses to all combinations of other player actions? I claim — following the reasoning when we considered Bertrand games — that the best-responses are illustrated by Figure 6.1 and Figure 6.2. (You should check that you agree... and should let me know if you do not!)

Figure 6.1: First-price sealed bid auction: Best response function for player 1

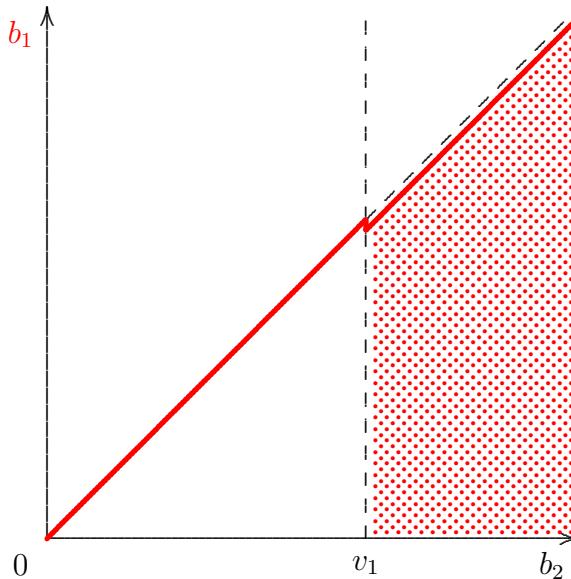
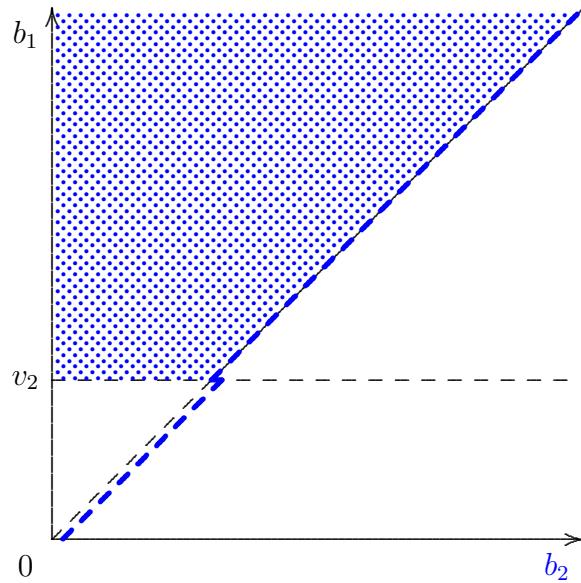


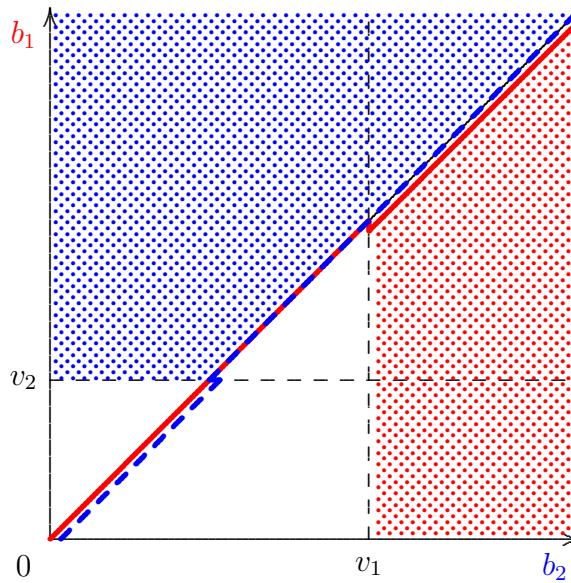
Figure 6.2: First-price sealed bid auction: Best response function for player 2



What is the solution concept? Nash equilibrium.

What are the solutions? I claim that the set of Nash equilibria is illustrated by Figure 6.3; that is, I claim that a Nash equilibrium is a pair (b_1, b_2) such that $b_1 = b_2$, and $b_1 \in [v_1, v_2]$. Note that, in all cases, Player 1 purchases the item.

Figure 6.3: First-price sealed bid auction: Nash equilibria



6.1.2 A second-price sealed bid auction

Of course, a first-price mechanism is not the only way to run an auction. In contrast, let's now consider a *second*-price sealed bid auction. We can think of this as an auction in which players submit sealed bids, and in which the item is won by the higher bidder — but where (s)he pays the bid of the *second*-highest bidder. Alternatively, this is equivalent to an ascending auction (sometimes known as an ‘English auction’).

Who are the players? What actions can they take, and when? The answers to these questions are the same as for the first-price auction.

What payoff does each player get for all combinations of player actions? Again, we assume valuations of v_1 and v_2 , and we will keep the tie-breaker assumption. However, with a second-price rule, we obtain the following payoffs:

$$\pi_1(b_1; b_2) = \begin{cases} v_1 - b_2 & \text{if } b_1 \geq b_2; \\ 0 & \text{if } b_2 > b_1. \end{cases}$$

$$\pi_2(b_2; b_1) = \begin{cases} v_2 - b_1 & \text{if } b_2 > b_1; \\ 0 & \text{if } b_1 \geq b_2. \end{cases}$$

For each player, what is the set of best-responses to all combinations of other player actions? I claim that the best-responses to the second-price auction are illustrated in Figure 6.4 and Figure 6.5. Do you agree?

Figure 6.4: Second-price sealed bid auction: Best response function for player 1

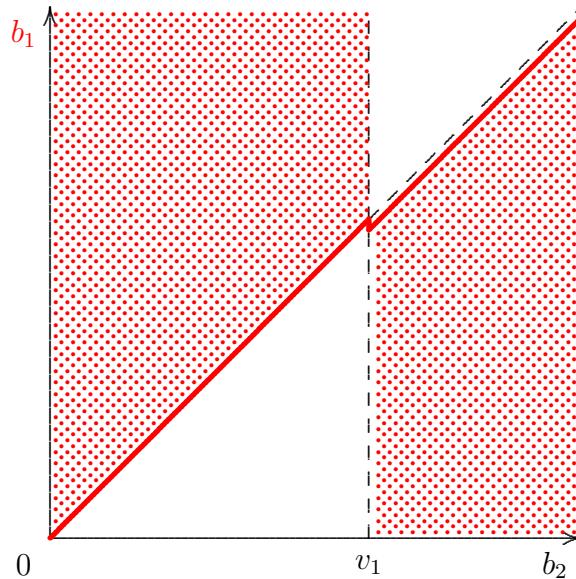
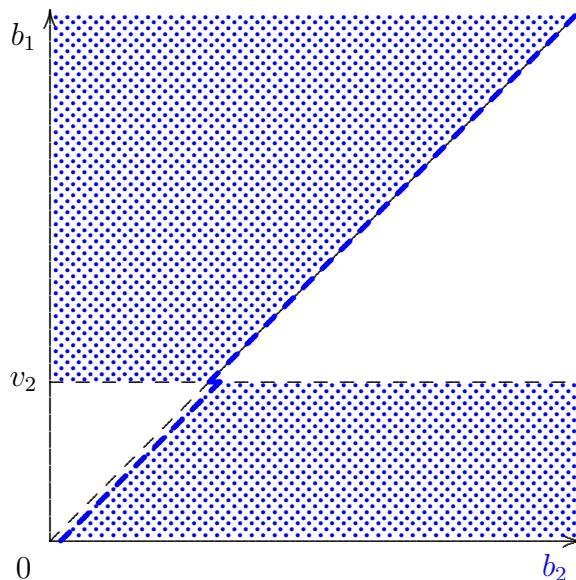


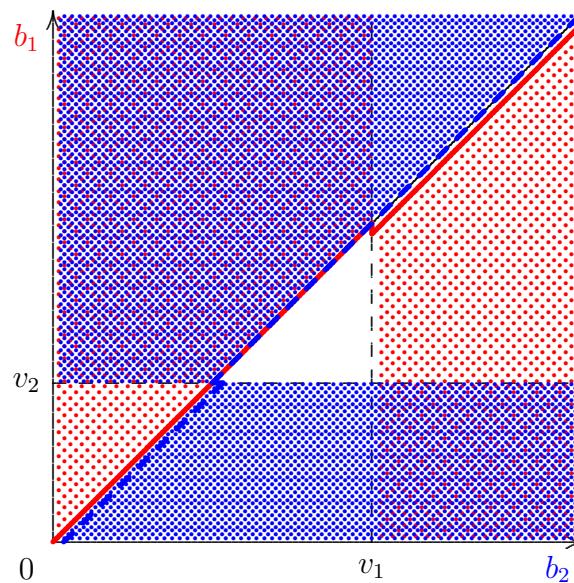
Figure 6.5: Second-price sealed bid auction: Best response function for player 2



What is the solution concept? Nash Equilibrium.

What are the solutions? I claim that the set of Nash equilibria are illustrated in Figure 6.6. This, it must be said, is a rather bewildering set of possible Nash equilibria. Note that, in many of the equilibria, Player 2 purchases the item — despite Player 2 valuing the item less than does Player 1.

Figure 6.6: Second-price sealed bid auction: Nash equilibria



6.2 Auctions with independent private values

In many respects, the two previous models seem like inadequate descriptions of auctions. This is true for several related reasons. First, in the vast majority of auctions, no bidders know exactly how the other bidders value the item. This poses a problem for the Nash Equilibrium solution concept. Second, the models — particularly the second-price model — do not seem to make very useful predictions.

Clearly, we would like a way to allow each player to be uncertain about its opponent's valuation. To do this, we need a new solution concept: *Bayesian Nash Equilibrium*. Varian (p.279, emphasis in original) gives very useful description of this concept:

A **Bayes-Nash equilibrium** of this game is then a set of strategies for each type of player that maximizes the expected value of each type of player, given the strategies pursued by the other players. This is essentially the same definition as in the definition of Nash equilibrium, except for the additional uncertainty involved about the type of the other player. Each player knows that the other player is chosen from a set of possible types, but doesn't know exactly which [type he or she] is playing. Note in order to have a complete description of an equilibrium we must have a list of strategies for *all* types of players, not just the actual types in a particular situation, since each individual player doesn't know the actual types of the other players and has to consider all possibilities.

Let's apply this solution concept to think about a first-price auction. To do this, it will be convenient to rely on a Uniform distribution; we will therefore take a moment to revise two key features of this distribution.

6.2.1 Two useful facts about the Uniform distribution

Suppose that a random variable, X , has a Uniform distribution. Then X has the following *cdf*:

$$\Pr(X \leq z) = \Pr(X < z) = \begin{cases} 0 & \text{if } z < 0; \\ z & \text{if } z \in [0, 1]; \\ 1 & \text{if } z > 1. \end{cases} \quad (6.3)$$

In this lecture, however, we will only ever worry about situations in which $z \in [0, 1]$. Assuming $z \in [0, 1]$, we can also say:

$$\mathbb{E}(X | X \leq z) = 0.5z. \quad (6.4)$$

6.2.2 A first-price sealed bid auction

Who are the players? What actions can they take, and when? What payoff does each player get for all combinations of player actions? The answer to these questions has not changed from the earlier first-price auction.

What characterises the player types? How are these types distributed? Player 1's type is defined by player 1's valuation, v_1 . Player 2's type is defined by player 2's valuation, v_2 . v_1 and v_2 each have Uniform distributions, and v_1 and v_2 are drawn independently of each other.

What does each player know? Each player knows all of the parameters of the game, and knows its own type. However, neither player knows the type of its opponent.

What is the solution concept? Bayesian Nash Equilibrium.

What is the solution? Player i learns v_i , but does not know v_j . Therefore, the player needs to form a *bid function*, which we will denote $b(v_i)$. We can think of $b(v_i)$ as a best-response — but, in a sense, we are now talking about Player i best-responding to its own type, v_i , conditional on some belief about the way that Player j will best-respond to v_j . In this context, the Bayesian Nash Equilibrium is an intersection of *bid functions*, rather than simply an intersection of *bids*.

For simplicity, we will assume throughout that bid functions are symmetric across players (*i.e.* this does not mean that the players make the same bid — just that, if they were to receive the same signal, they would make the same bid). Therefore, Player i 's expected payoff — given a realised value of v_i — is:

$$\mathbb{E} [\pi(b(v_i); v_i)] = \mathbb{E} (v_i - b(v_i) \mid i \text{ wins}) \cdot \Pr(i \text{ wins}) \quad (6.5)$$

$$= (v_i - b(v_i)) \cdot \Pr(b(v_j) \leq b(v_i)). \quad (6.6)$$

This poses quite a difficult problem: *solve for the function b , such that b maximises this expected payoff given v_i* . Indeed, this is known as a *functional equation*. If you like, you can pursue a general solution for a problem of this form¹² — but we will proceed by making a convenient guess...

Specifically, let's guess that the optimal bid function is linear: $b(v_i) = \kappa \cdot v_i$. If this is so, then player i faces the following problem:

$$\max_{s_i \geq 0} \mathbb{E} [\pi(s_i; v_i)] = \mathbb{E} (v_i - s_i \mid i \text{ wins}) \cdot \Pr(i \text{ wins}) \quad (6.7)$$

$$= (v_i - s_i) \cdot \Pr(\kappa \cdot v_j \leq s_i) \quad (6.8)$$

$$= (v_i - s_i) \cdot \frac{s_i}{\kappa}, \quad (6.9)$$

where the last line follows from the assumption that v_j has a Uniform distribution.

¹² For example, you could see Chapter 2 of *Paarsh and Hong* (2006), “An Introduction to the Structural Econometrics of Auction Data”. This provides a very interesting discussion, which is comfortably beyond the requirements of this course!

Before solving, you should check carefully that you understand how this problem has been constructed. Player i essentially reasons as follows:

I have just drawn my valuation, v_i . I assume that my opponent is using the bid function $b_j = \kappa \cdot v_j$. Given this belief, and given my valuation, I need to choose a submitted bid, s_i , to maximise my expected payoff.

(That is, we treat player i as choosing the submitted bid, s_i . We do *not* impose anywhere that player i must use a linear bid function; this is something that we need to check, once we find the optimal choice of s_i .) Notice that s_i appears twice in equation 6.9; why?

Once you understand this reasoning, you can solve player i 's problem straightforwardly by differentiating, to obtain $s_i = 0.5v_i$. Therefore, $b(v_i) = \kappa \cdot v_i$ does indeed constitute a Bayesian Nash Equilibrium, with $\kappa = 0.5$.

In some respects, $b(v_i) = 0.5v_i$ seems like an intuitive bidding strategy. In other respects, it may seem strange; for example, why not simply bid $b(v_i) = v_i$?

6.2.3 A second-price sealed bid auction

We can apply similar reasoning to the case of a second-price sealed bid auction with uncertainty...

Who are the players? What actions can they take, and when? What payoff does each player get for all combinations of player actions? The answer to these questions has not changed from the earlier second-price auction.

What characterises the player types? How are these types distributed? What does each player know? What is the solution concept? The answer to these questions has not changed from the preceding game.

What is the solution? Under a second-price mechanism, player i solves the following problem:

$$\max_{s_i \geq 0} \mathbb{E} [\pi(s_i; v_i)] = \mathbb{E} (v_i - \kappa \cdot v_j \mid i \text{ wins}) \cdot \Pr(i \text{ wins}) \quad (6.10)$$

$$= \mathbb{E} (v_i - \kappa \cdot v_j \mid \kappa \cdot v_j \leq s_i) \cdot \Pr(\kappa \cdot v_j \leq s_i) \quad (6.11)$$

$$= (v_i - 0.5s_i) \cdot \frac{s_i}{\kappa}. \quad (6.12)$$

Again, this can be solved straightforwardly by differentiating; you should confirm that we obtain $s_i = v_i$; *i.e.* $\kappa = 1$, and each player should bid his or her valuation. This is quite a different result than the first-price mechanism; why?

6.3 Auctions with common values

The assumption of independent private values seems appropriate for situations where value is subjective — for example, in buying artwork. However, there are many situations in which players receive separate signals relating to the same underlying value. For example, two firms may be bidding for oil exploration rights in a developing country; each firm may have a different initial assessment of what the rights will be worth but, ultimately, a commonly agreed value of the rights will be realised in due course.

We can model this using a ‘common values’ auction. We will consider a simple common-values auction, in which each player values its own signal with weight 1, and (additively) its opponent’s signal with weight γ ; that is, the realised payoff for player i will be $v_i + \gamma v_j$.

6.3.1 A first-price common values auction

To characterise and solve this game, let’s keep the same structure that we used for the earlier first-price auction (with uncertainty); however, as noted, we will assume a payoff $v_i + \gamma v_j$ for player i . Again, let’s assume that each players use a linear bid function, $b(v_i) = \kappa \cdot v_j$. Therefore, we can write player i ’s problem as follows:

$$\max_{s_i \geq 0} \mathbb{E} [\pi(s_i; v_i)] = \mathbb{E} (v_i + \gamma \cdot v_j - s_i \mid i \text{ wins}) \cdot \Pr(i \text{ wins}); \quad (6.13)$$

$$= \mathbb{E} (v_i + \gamma \cdot v_j - s_i \mid \kappa \cdot v_j \leq s_i) \cdot \Pr(\kappa \cdot v_j \leq s_i); \quad (6.14)$$

$$= \left(v_i + 0.5\gamma \frac{s_i}{\kappa} - s_i \right) \cdot \frac{s_i}{\kappa}. \quad (6.15)$$

You should find that this is maximised for:

$$v_i + \gamma \cdot \frac{s_i}{\kappa} - 2s_i = 0; \quad (6.16)$$

$$\kappa v_i + s_i \cdot (\gamma - 2\kappa) = 0; \quad (6.17)$$

$$\therefore s_i = \frac{\kappa v_i}{2\kappa - \gamma}. \quad (6.18)$$

Now, if $s_i = \kappa v_i$, we therefore obtain:

$$\kappa v_i = \frac{\kappa v_i}{2\kappa - \gamma} \quad (6.19)$$

$$\therefore \kappa = 0.5(1 + \gamma). \quad (6.20)$$

That is, there is a Bayesian Nash Equilibrium for this common values game characterised by a bid function of the form:

$$b(v_i) = 0.5(1 + \gamma) \cdot v_i. \quad (6.21)$$

You should check that you understand the intuition of this solution. (What happens for the special case $\gamma = 0$?)

6.3.2 How *not* to solve common values problems

Suppose that you are advising a firm that is deciding how to bid in a common values auction of the kind described here. Suppose that the firm management reasons as follows:

Our payoff is $v_i + \gamma v_j$. We know that v_j has a Uniform distribution, so its expectation is 0.5. Therefore, in expectation, our valuation is $v_i + 0.5\gamma$. With this simplification, we can think about this as an ‘independent private values’ auction in which each firm has the payoff $v_i + 0.5\gamma$.

You should be able to show that such an auction would have a Bayesian Nash Equilibrium in which each firm uses the following bidding function:

$$b(v_i) = 0.5\gamma + 0.5v_i. \quad (6.22)$$

Compare this to the Bayesian Nash Equilibrium of the common values auction: notice that, under this current reasoning, the firm will bid an additional amount of $0.5\gamma(1 - v_i)$. *Why the difference?*

6.4 Class exercises: Auctions

1. (*From the 2016 exam...*) Zoolandia is a developing country ruled by a long-standing autocrat, President Mugatu. President Mugatu must decide which of two mining companies — Hansel Exploration or Gretel Energy — is to receive a lucrative mining licence for the north of the country. President Mugatu invites Hansel and Gretel each to make a voluntary contribution to his re-election fund ($c_g \geq 0$ for Gretel and $c_h \geq 0$ for Hansel). President Mugatu promises each company that the licence will be awarded to the company making the larger donation; in the event that the two companies donate the same amount, he will assign the licence to Gretel. Gretel values the mining licence as v_g and that Hansel values it as v_h .

You should assume that Mugatu will keep his promise, and that both Gretel and Hansel know this. Further, you should assume that, unfortunately, President Mugatu will *not* refund any donations after they are made (*i.e.* even the firm that does not receive the licence will not be able to recover its donation). Assume that Gretel and Hansel make their contributions simultaneously, and do not collude in doing so. Assume that both Gretel and Hansel are risk-neutral.

- (a) Assume that $v_g > v_h > 0$, and that v_g and v_h are known both to Gretel and to Hansel. Write the payoff functions for Gretel and for Hansel and graph the best-response functions. How many pure-strategy Nash Equilibria does this game have?
- (b) President Mugatu is considering whether he should invite Gretel to donate first, so that he can publicise its contribution before Hansel is invited to donate. How would this change the solution concept and the reasoning in part (a)?

Now assume that each firm's valuation is drawn independently from a Uniform distribution between 0 and 1. Assume that each firm knows its own valuation before choosing its donation, but does not know the valuation of its competitor. Assume, as in part (a), that Hansel and Gretel choose their contributions simultaneously.

- (c) Show that there is a Bayesian Nash Equilibrium in which Gretel donates $d_g = 0.5 \cdot v_g^2$ and Hansel donates $d_h = 0.5 \cdot v_h^2$. What is the expected value of the total donation that President Mugatu is to receive?
- (d) Assume now that Mugatu changes his mind, and decides that he will refund any donation from a company not receiving the licence. Find a symmetric Bayesian Nash Equilibrium to describe Gretel's and Hansel's donation strategies in this game. What is the expected value of the total donation that President Mugatu is to receive? Is this greater or less than the total donation expected in part (c)? Explain the intuition for this result. (*Note:* Consider any two independent random variables, X and Y , each having a Uniform distribution between 0 and 1. Define Z as the larger of the two, *i.e.* $Z = \max(X, Y)$. You may rely on the fact that the expected value of Z is $2/3$.)

2. (*From the 2021 exam...*) Bluey Energy and Chloe Resources are each large mining companies, operating in the Republic of Oz. The Government of Oz is conducting a public auction for a new mining permit. The Government requires both Bluey and Chloe to bid, and no other bids are permitted. Bluey and Chloe are each allowed to bid only in round billions of US dollars, up to (and including) a maximum of US\$2 billion (that is, each may submit a bid of zero, of US\$1 billion, or of US\$2 billion). Bids must be submitted simultaneously. The Government of Oz commits that the permit will be sold to whichever company submits the larger bid; in the event that both companies submit identical bids, Bluey will be declared the winner. The Government commits to a first-price auction: the winner must pay the bid that it submitted, and the loser is not required to pay anything. Denote Bluey's valuation (in US\$ billion) by v_b ; denote Chloe's valuation (also in US\$ billion) by v_c . Denote Bluey's submitted bid (in US\$ billion) as s_b , and Chloe's submitted bid (also in US\$ billion) as s_c . (Thus, because of the Government's requirement to bid in round billions, we can say $s_b \in \{0, 1, 2\}$ and $s_c \in \{0, 1, 2\}$.)

Assume initially that Bluey values the permit at $v_b = 2.1$, and that Chloe values the permit at $v_c = 1.9$.

- (a) Show graphically Bluey's best response to Chloe's submitted bid: $s_b^*(s_c)$.
- (b) Show graphically Chloe's best response to Bluey's submitted bid: $s_c^*(s_b)$.
- (c) Describe the set of Nash Equilibria.

Now suppose that Bluey's valuation and Chloe's valuation are each drawn independently from a continuous uniform distribution: $v_b \sim U(0, 3)$ and $v_c \sim U(0, 3)$. Suppose that each company knows its valuation – but not the valuation of its competitor – at the time that it bids. (The requirement to bid in round billions remains.) Denote by p_k^b the probability that Bluey submits a bid of k ; denote by p_k^c the probability that Chloe submits a bid of k . Denote by $\pi_k^b(v_b)$ the expected payoff to Bluey from submitting a bid of k ; denote Chloe's equivalent payoff as $\pi_k^c(v_c)$.

- (d) Show that Bluey's expected payoffs are as follows:

$$\begin{aligned}\pi_0^b(v_b) &= p_0^c \cdot v_b; \\ \pi_1^b(v_b) &= (p_0^c + p_1^c) \cdot (v_b - 1); \\ \pi_2^b(v_b) &= (p_0^c + p_1^c + p_2^c) \cdot (v_b - 2).\end{aligned}$$

- (e) Write the expected payoffs for Chloe (that is, $\pi_0^c(v_c)$, $\pi_1^c(v_c)$ and $\pi_2^c(v_c)$).

Suppose that Bluey decides to bid $s_b = 1$ if $v_b \geq 1.5$ and $s_b = 0$ otherwise, and that Chloe decides to bid $s_c = 1$ if $v_c \geq 1$ and $s_c = 0$ otherwise.

- (f) Does this combination of bid functions form a Bayesian Nash Equilibrium?

BARGAINING

7 Bargaining: The strategic approach

Lecturer: Simon Quinn

Reading:

- Lecture 1 of Dr Meyer's course on 'Bargaining, Contracts and Theories of the Firm'.

In this final part of the course, we discuss bargaining. Bargaining is a central part of life in every economy, and the theory of bargaining is useful for thinking about a very wide range of scenarios — including, in particular, bargaining within firms and bargaining within households.

Broadly, there are two different approaches to modelling bargaining: the 'strategic approach', and the 'axiomatic approach'. We will consider the former approach in this lecture, and the latter approach in the final lecture. The strategic approach will feel familiar from the general approach to game theory that we have taken in our previous lectures; the axiomatic approach will probably feel quite different. Nash (1953) explained the difference as follows:

In the [strategic approach], the cooperative game is reduced to a non-cooperative game. To do this, one makes the players' steps of negotiation in the cooperative game become moves in the non-cooperative model. Of course, one cannot represent all possible bargaining devices as moves in the non-cooperative game. The negotiation process must be formalized and restricted, but in such a way that each participant is still able to utilize all the essential strengths of his position.

The second approach is by the axiomatic method. One states as axioms several properties that it would seem natural for the solution to have and then one discovers that the axioms actually determine the solution uniquely. The two approaches to the problem, via the negotiation model or via the axioms, are complementary; each helps to justify and clarify the other.

7.1 Dividing a single good under Nash Equilibrium

In this lecture, we will focus primarily on a finite-horizon 'alternating-offers' bargaining game. To motivate that model and to highlight some key features of bargaining problems, we will begin with a simpler model. Specifically, we will imagine two players bargaining over a single good — without loss of generality, a £1 coin. We imagine that this good is

infinitely divisible. We imagine a simple mechanism: the two players each simultaneously submit their proposed shares, and the coin is divided if (and only if) the two shares do not exceed 1; if the shares exceed 1, both players receive nothing. In the event that both players receive nothing from the bargain, we allow each player to receive some positive ‘outside option’.

The flavour of this game should feel very familiar from our previous lectures — particularly the last lecture, on auctions. Hopefully, you now feel very comfortable in analysing this kind of game...

Who are the players? We have two players: Player 1 and Player 2.

What actions can they take, and when? Each player i submits a proposed share, $s_i \in [0, 1]$. These proposals are submitted simultaneously.

What payoff does each player get for all combinations of player actions? Denote the currency value of the outside option for player i as d_i . Without loss of generality, we can allow each player’s utility to be linear in its payoff; therefore, we have:

$$\pi_i(s_i, s_j) = \begin{cases} s_i & \text{if } s_i + s_j \leq 1; \text{ and} \\ d_i & \text{otherwise.} \end{cases}$$

For each player, what is the set of best-responses to all combinations of other player actions? When we studied auctions, we asked, ‘does the player actually want to win the auction, or not?’. We now ask an analogous question: ‘does the player actually want to divide the £1, or would the player prefer to take the outside option?’. With this in mind, player i ’s best-response becomes:

$$s_i^*(s_j) = \begin{cases} 1 - s_j & \text{if } 1 - s_j \geq d_i; \\ (1 - s_j, 1] & \text{if } 1 - s_j \leq d_i. \end{cases} \quad (7.1)$$

Figures 7.1 and 7.2 illustrate.

Figure 7.1: Simple bargaining game: Best response function for player 1

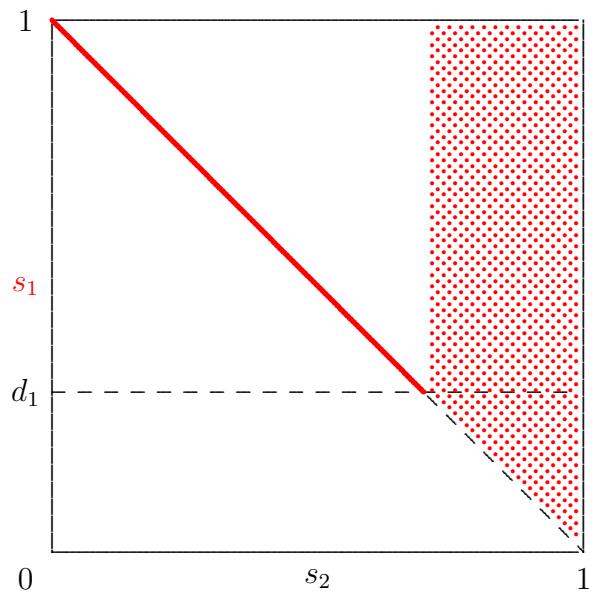
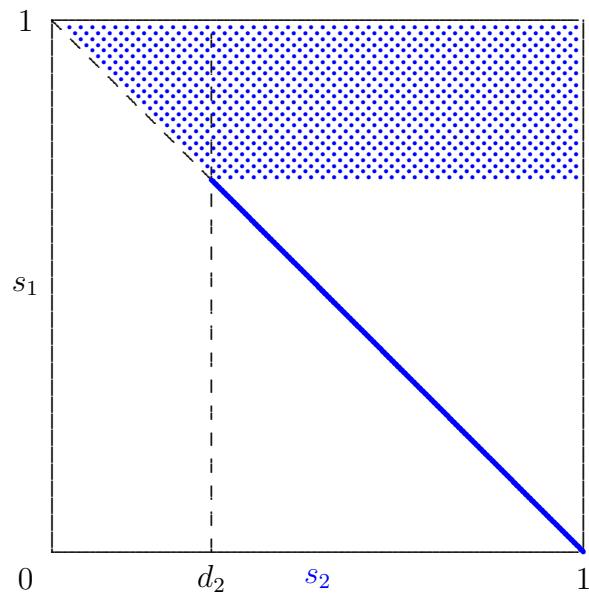


Figure 7.2: Simple bargaining game: Best response function for player 2



What is the solution concept? We will use Nash Equilibrium. (Note that there is no uncertainty here about the players' payoffs — therefore, unlike the previous lecture, we would gain nothing from using Bayesian Nash Equilibrium here.)

Figure 7.3: Simple bargaining game: Nash Equilibrium

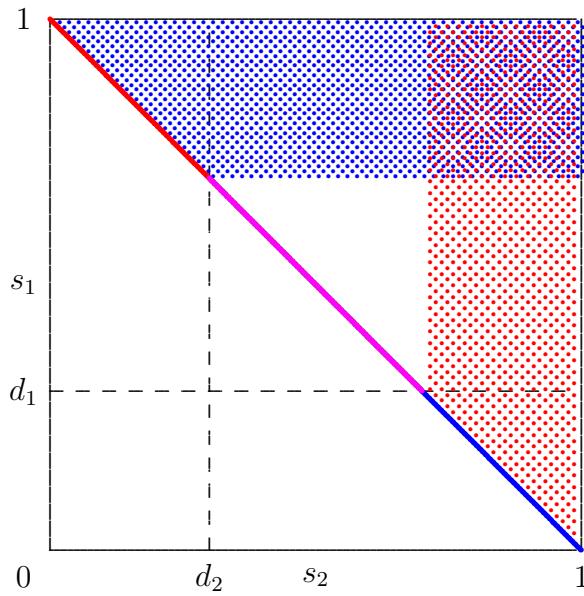


Figure 7.3 illustrates the result: the Nash Equilibria comprise (i) all of the points (s_1, s_2) such that $s_1 + s_2 = 1; s_1 \geq d_1; s_2 \geq d_2$ (that is, cases in which the £1 is divided), and (ii) all of the points such that $s_1 \geq 1 - d_2; s_2 \geq 1 - d_1$ (that is, cases in which the bargaining fails).

Just as in our study of auctions, Nash Equilibrium admits a very wide range of predictions — and, therefore, does not seem to provide a particularly useful solution concept here. First, consider the cases in which where bargaining fails. Of course, we all know that bargaining *does* sometimes fail in the real world. However, the explanation implicitly given by this model — namely, coordination failure between the players — hardly seems like a satisfactory description of real-world bargaining problems. Second, even among those cases in which the £1 is divided, Nash Equilibrium gives very little guidance as to what factors might drive the division. For this reason, models using the strategic approach to bargaining prefer to place much more structure on the problem...

7.2 A bargaining game with alternating offers

7.2.1 A finite horizon

Let's now consider a version of the famous bargaining game with alternating offers.¹³ The structure of the alternating-offer bargaining game is simple. As in the earlier case, two players seek to divide £1 between them. The players take turns in making 'take it or leave it' offers to each other. If an offer is accepted, the game ends. We begin by considering a finite-horizon version of the game. Players know in advance the maximum number of periods (T); if no agreement is reached by the end of the final period, the game ends, and both players receive nothing. We assume that Player 1 has a discount factor $\delta_1 < 1$; Player 2 has a discount factor $\delta_2 < 1$.

We index periods by $t \in \{1, \dots, T\}$. For ease of understanding, we will count periods *backwards* — so we will speak about the game starting in period T , and ending in period 1.¹⁴ In odd-numbered periods (again, counting *backwards*), we will allow Player 1 to offer to Player 2; in even-numbered periods, Player 2 offers to Player 1.¹⁵

We will use x to denote the amount proposed for Player 1. (Thus, in odd-numbered periods, we can imagine Player 1 saying to Player 2, "I propose to *take* x , leaving you $1 - x$ "; in even-numbered periods, we can imagine Player 2 saying, "I propose to *give* x , taking $1 - x$ for myself".)

We can analyse this game using the same set of questions as before...

Who are the players? We have two players: Player 1 and Player 2.

What actions can they take, and when? In any odd-numbered period, Player 1 offers $1 - x$ to Player 2. Player 2 may then accept or refuse. In any even-numbered period, Player 2 offers x to Player 1.

What payoff does each player get for all combinations of player actions? We assume the players' utility is linear in their payoffs, that each player uses exponential discounting, and that each player has an outside option of zero. There are two possible final outcomes to the game:

1. *Acceptance in some period $t \geq 1$:*

$$\begin{aligned}\pi_1(x, t) &= \delta_1^{T-t} \cdot x; \\ \pi_2(x, t) &= \delta_2^{T-t} \cdot (1 - x).\end{aligned}$$

¹³ Our discussion here is based on Ichiro Obara's excellent lecture notes (available online at <http://www.econ.ucla.edu/iobara/Bargaining201B.pdf>).

¹⁴ This feels weird, I know. But it will make the discussion easier, as you will soon see...

¹⁵ Therefore, if T is odd, Player 1 makes the first offer; if T is even, Player 2 makes the first offer.

2. *Refusal in the final period, $t = 1$:*

$$\begin{aligned}\pi_1(x, t) &= 0; \\ \pi_2(x, t) &= 0.\end{aligned}$$

What is the solution concept? As in the previous example, there are many Nash Equilibria to this game. Instead, we will exploit the sequential structure of the alternating-offers game, by using Subgame Perfect Equilibrium ('SPE').

For each player in each period, what is the set of best-responses to all combinations of other player actions? Conceptually, this is straightforward: because we have a finite number of periods, we solve by backward induction. (For simplicity, we will assume throughout that, if a player is indifferent between accepting and refusing, (s)he will accept. This changes nothing of the underlying structure of the game.)

Start by considering the final period (note that this is equivalent to a game with a single period: $T = 1$). In this case, the unique SPE is:

- Player 1 proposes $x = 1$; and
- Player 2 accepts any proposal $x \leq 1$.

Note that this description includes a description of play that is off the equilibrium path; it says much more, for example, than 'Player 1 proposes $x = 1$ and Player 2 accepts.' Unless we describe Player 2's behaviour if $x < 1$, we haven't actually described an equilibrium. Nonetheless, 'Player 1 proposes $x = 1$ and Player 2 accepts' is an accurate description of play in the final period; therefore, should the game reach period $t = 1$, Player 1 receives utility of 1, and Player 2 receives utility of 0.

Now suppose that $T = 2$. Discounting, the present value of Player 1's payoff, should the game proceed to period 1, is δ_1 . Therefore, in the first period of a two-period game, the SPE is:

- If the game reaches period 1, the players behave as above.
- In period 2,
 - Player 2 proposes $x = \delta_1$; and
 - Player 1 accepts if $x \geq \delta_1$ and rejects otherwise.

Therefore, Player 2 proposes $x = \delta_1$ in period 2, and this is accepted; Player 2 therefore receives a payoff of $1 - x = 1 - \delta_1$.

Now what if $T = 3$? Well, this is now identical to the subgame where $T = 1$, except that the payoff to Player 2 from refusing is now $1 - \delta_1$, rather than zero. Therefore, for $T = 3$,

- If the game reaches period 2, the players behave as above.
- In period 3,
 - Player 1 proposes $x = 1 - \delta_2 \cdot (1 - \delta_1)$; and
 - Player 2 accepts any proposal $x \leq 1 - \delta_2 \cdot (1 - \delta_1)$.

Therefore, if $T = 3$, Player 1 proposes $x = 1 - \delta_2 \cdot (1 - \delta_1)$, and Player 2 accepts. Note that, for each period, there is a unique SPE — and therefore, there will be a unique SPE for the game as a whole, on any finite horizon T .

What about $T = 4$? From here, we should be able to see a pattern emerging. Rather than step tediously through the same reasoning over and over again, let's now denote $x(t)$ as the proposal made, in equilibrium, in period t . Therefore, we can say:

$$\begin{aligned}
 & x(1) = 1; \\
 & \text{in even-numbered periods... } x(t) = \delta_1 \cdot x(t-1); \\
 & \text{in odd-numbered periods... } x(t) = 1 - \delta_2 \cdot [1 - x(t-1)] = 1 - \delta_2 + \delta_2 \cdot x(t-1); \\
 & \text{therefore... } x(2) = \delta_1; \\
 & x(3) = 1 - \delta_2 + \delta_1 \delta_2; \\
 & x(4) = \delta_1 - \delta_1 \delta_2 + \delta_1^2 \delta_2; \\
 & x(5) = 1 - \delta_2 + \delta_1 \delta_2 - \delta_1 \delta_2^2 + \delta_1^2 \delta_2^2; \\
 & x(6) = \delta_1 - \delta_1 \delta_2 + \delta_1^2 \delta_2 - \delta_1^2 \delta_2^2 + \delta_1^3 \delta_2^2; \\
 & x(7) = 1 - \delta_2 + \delta_1 \delta_2 - \delta_1 \delta_2^2 + \delta_1^2 \delta_2^2 - \delta_1^2 \delta_2^3 + \delta_1^3 \delta_2^3; \\
 & \vdots
 \end{aligned}$$

This iterative process defines the unique SPE to this game.¹⁶

I think this is a very nice solution. Unlike the simple game at the beginning of the lecture, the alternating-offers game provides a unique prediction as to the outcome of the bargaining process: namely, the game ends after a single offer, where the allocation to Player 1 is given by $x(t)$. Note that the solution is Pareto efficient — the players are able to exploit the gains from bargaining, and do so without any costly delay.

What's more, that solution provides useful insights as to the determinants of the agreed division. To see the intuition of this, ask yourself three questions:

1. What is the effect of δ_1 and δ_2 on the offers made by Player 1? Why?
2. Does the game have a first-mover advantage, or a first-mover disadvantage? (We know that Player 1 makes the final offer, in period 1; would Player 1 prefer there to be an odd or even number of periods? What about Player 2?)
3. Define \tilde{x}_1 as $\lim_{t \rightarrow \infty} x(t)$, for odd-numbered periods (that is, the limit for proposals made by Player 1). In this limit, we have:

$$\tilde{x}_1 = 1 - \delta_2 + \delta_1 \delta_2 \cdot \tilde{x}_1; \quad (7.2)$$

$$\therefore \tilde{x}_1 = \frac{1 - \delta_2}{1 - \delta_1 \delta_2}. \quad (7.3)$$

Define \tilde{x}_2 as $\lim_{t \rightarrow \infty} x(t)$, for even-numbered periods (that is, the limit for proposals made by Player 2); then it follows from the previous reasoning that:

$$\tilde{x}_2 = \delta_1 \cdot \tilde{x}_1; \quad (7.4)$$

$$= \frac{\delta_1 \cdot (1 - \delta_2)}{1 - \delta_1 \delta_2}. \quad (7.5)$$

¹⁶ Of course, it is also possible to solve generally the form of $x_1(t)$ and $x_2(t)$, rather than, as here, having to rely on iterative expressions — though I am not sure that doing this really adds much to our intuition. Look at the proposals made by Player 1; we can say generally that, for odd-numbered periods:

$$x(t) = 1 - \delta_2 + \delta_1 \delta_2 \cdot x(t-2).$$

This is a ‘linear first-order difference equation with constant coefficient’ — and, given the starting point of $x(1) = 1$, is solved like this:

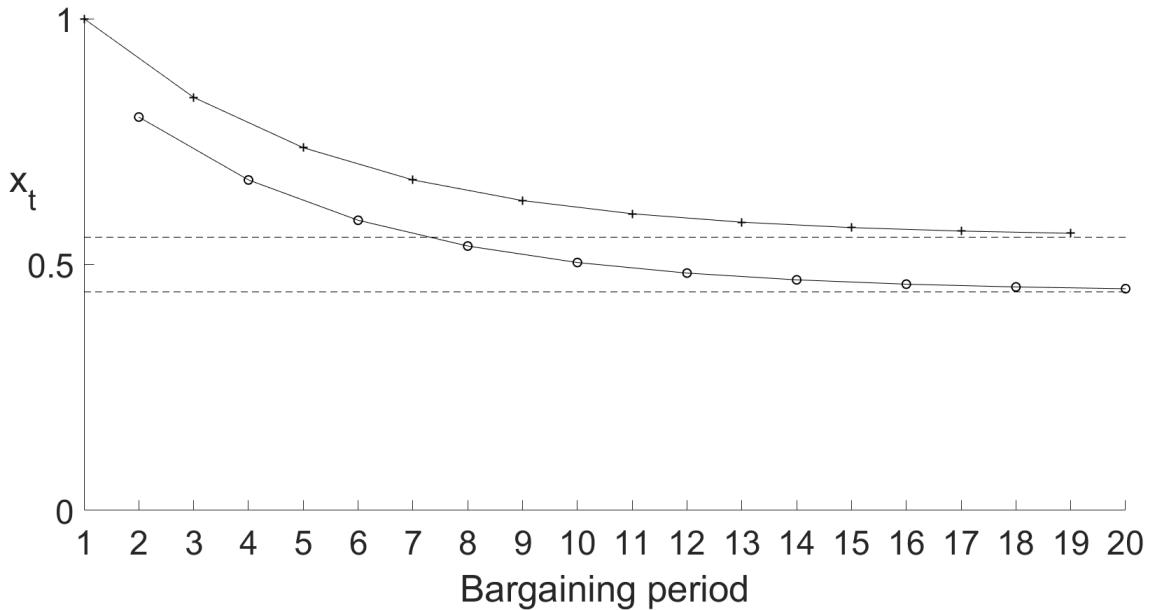
$$x(t) = (\delta_1 \delta_2)^{0.5t-0.5} + (1 - \delta_2) \cdot \sum_{j=0}^{0.5t-1.5} (\delta_1 \delta_2)^j.$$

This equation works for odd-numbered periods; for even-numbered periods, we can simply say:

$$x(t) = \delta_1 \cdot x(t-1) = \delta_1 (\delta_1 \delta_2)^{0.5t-1} + \delta_1 (1 - \delta_2) \cdot \sum_{j=0}^{0.5t-2} (\delta_1 \delta_2)^j.$$

The following graphs illustrate these principles for several specific cases. (In each case, ‘+’ refers to the offers made by Player 1; ‘o’ refers to offers made by Player 2. Values for $\lim_{k \rightarrow \infty} \tilde{x}_1(k)$ and $\lim_{k \rightarrow \infty} \tilde{x}_2(k)$ are superimposed as dotted lines.) Figure 7.4 shows the offers made for a game with $T = 20$ periods, with $\delta_1 = \delta_2 = 0.8$.

Figure 7.4: **Solution illustration:** $T = 20$; $\delta_1 = 0.8$; $\delta_2 = 0.8$



Now, let’s see what happens if Player 1 is more impatient. Figure 7.5 shows the path of offers if we shift δ_1 from 0.8 to 0.7.

Finally, let’s see what happens if both players are more impatient: Figure 7.6 keeps $\delta_1 = 0.7$, and reduces δ_2 from 0.8 to 0.7. You should check that you understand the intuition for these comparative statics. Why, for example, does Player 1 make more generous offers — and receive less generous offers — in Figure 7.5 than in Figure 7.6?

Figure 7.5: **Solution illustration:** $T = 20$; $\delta_1 = 0.7$; $\delta_2 = 0.8$

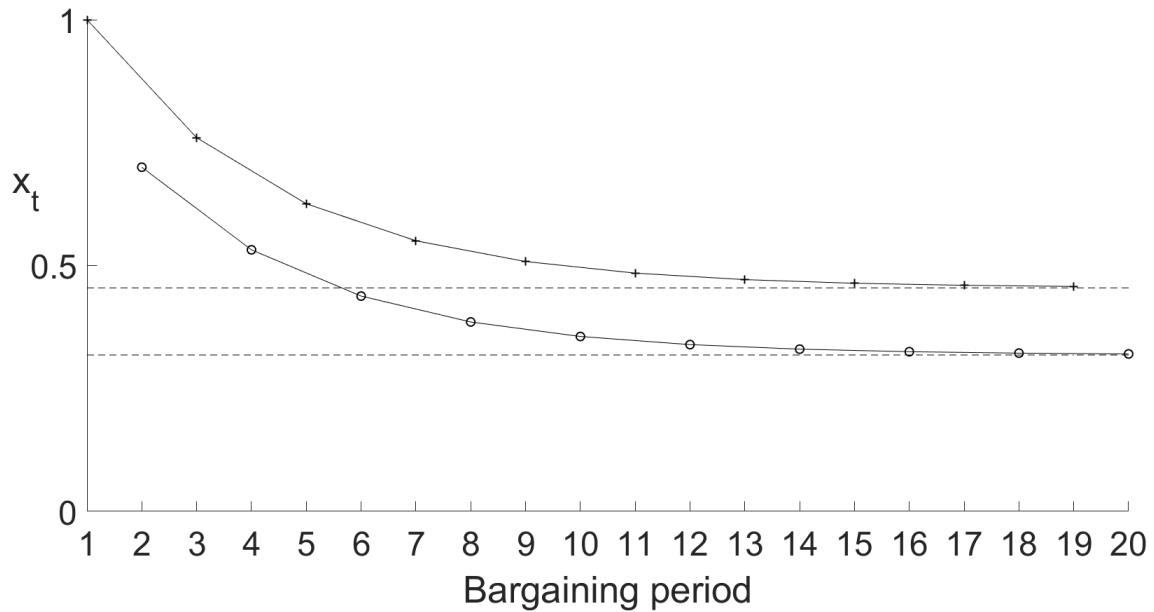
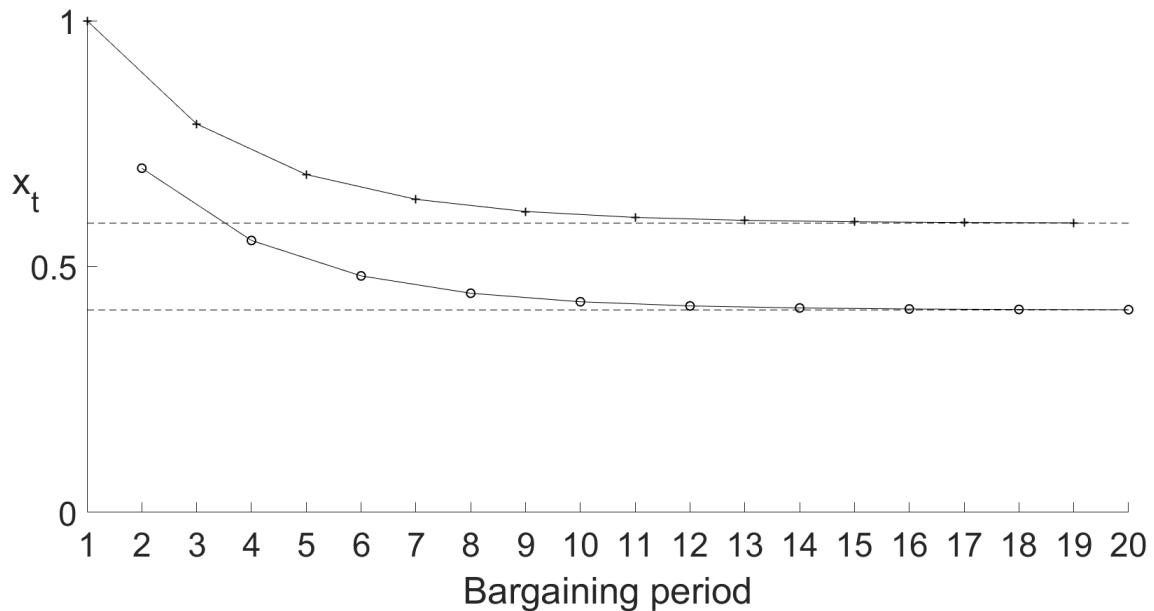


Figure 7.6: **Solution illustration:** $T = 20$; $\delta_1 = 0.7$; $\delta_2 = 0.7$



7.2.2 An infinite horizon

In this lecture, we have studied a *finite*-horizon alternating-offers game. I think this is a very nice way of capturing the intuition of alternating-offer games, and for making useful predictions about how a surplus can be divided through bargaining. However, it is certainly not the only kind of alternating-offer game; in particular, a popular alternative model considers players interacting on an *infinite* horizon.

It turns out that the predictions of such a model are very close to those of the finite-horizon game considered here. In particular, the limiting solution described by equations 7.3 and 7.5 holds in every period of the infinite-horizon game (depending, of course, on whether it is Player 1's or Player 2's turn to offer). However, it's important to note that this conclusion is *not* implied merely by the fact that equations 7.3 and 7.5 are the limiting cases of the finite-horizon game. The infinite-horizon game requires a completely different method of solution, because we cannot use backward induction (instead, we must rely on the 'one-shot deviation principle' — something that is beyond the scope of our lectures).

If you are interested in reading more on the infinite-horizon alternating-offers game, I strongly recommend Meg Meyer's excellent lecture notes — and, if you like, the seminal original paper by Rubinstein (1982).

8 Bargaining: The axiomatic approach

Reading:

- Lecture 1 of Dr Meyer's course on 'Bargaining, Contracts and Theories of the Firm'.

8.1 The Nash Bargaining Solution

In the previous lecture, we considered two different models for thinking about how two players might bargain over the division of a £1 coin. Both models were extremely specific about the mechanism by which this would occur — either through submitting sealed proposals (in the first model), or through a very rigid pattern of offer and counter-offer (in the second model).

But what if you and your classmate *actually* had to divide a £1 coin? Or, for that matter, decide how to split a cab fare, or develop a business opportunity, or share a textbook? Only the most dedicated economics students would choose to rely upon a very rigid mechanism, specified exactly in advance.¹⁷ In the real world, you would *talk*. . . you would try to thrash out an agreement, hoping that any deal would make you both better off, and would not waste resources in doing so.

That is what the axiomatic approach to bargaining theory is about: in effect, the axiomatic approach says, 'Never mind the details of *how* a deal is done; let's see what predictions we can make by focusing purely on the substance of the deal itself.' The Nash Bargaining Solution is the most famous form of this axiomatic approach to bargaining.

So what might we say about the division of a £1 coin, if we focus purely on players' preferences and the properties of the division itself? Let's start by reframing our earlier problem, so that it is amenable to an axiomatic approach.

A 'bargaining problem': The Nash Bargaining Solution only applies to 'bargaining problems'. In axiomatic bargaining theory, this has a very particular meaning. Specifically, we use the term 'bargaining problem' to describe a pair, $\langle U, d \rangle$ — where U is a *set of pairs of payoffs to agreements*, and d is a single *pair of payoffs from disagreement*.

It is worth pausing at this point to emphasise clearly what is going on... *Every other graph that we have seen in my part of the course has been in 'action space'* — that is, the two dimensions each represent specific actions that each player could take. In axiomatic bargaining theory, we make a subtle but fundamental shift: namely, we represent outcomes in a space of possible *utilities*.

¹⁷ This is not intended as an endorsement.

Figure 8.1: Simple bargaining game: Utility space

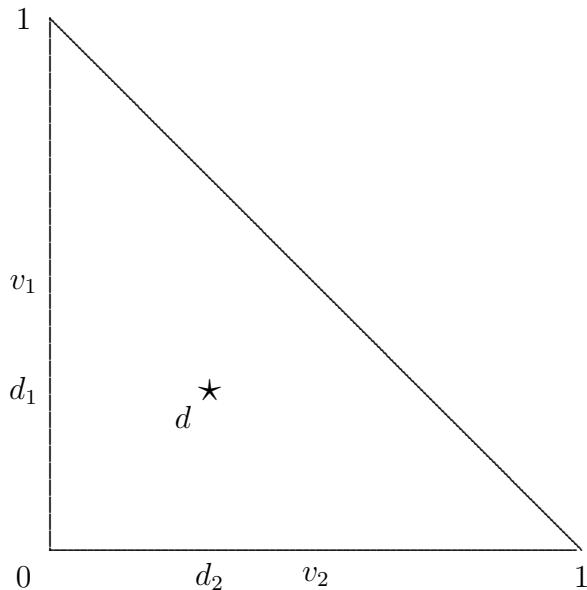


Figure 8.1 shows our earlier bargaining game, graphed against the utility obtained from possible outcomes (v_1 and v_2). (Remember that we assumed earlier that each player has a linear utility function: this is why, in this particular example, the ‘utility space’ and the ‘action space’ look so similar.) The diagram shows a pair $\langle U, d \rangle$ — that is, the triangle shows a set of pairs of possible payoffs to agreements, and the point d shows the payoffs from disagreement.

In order to be a ‘bargaining problem’, our pair $\langle U, d \rangle$ needs to meet four specific requirements; fortunately, each is met in our case. They are:

1. *d must be a member of U* : This is clearly fulfilled, as the diagram shows.
2. *There must be some point in U preferred to d by both players*: That is, d cannot lie on the Pareto frontier. This is also clearly fulfilled in our case.
3. *U must be a convex set*: This means that, for any two points in U , a straight line between them will never pass outside of U . This is also clearly fulfilled in our case.
4. *U is bounded and closed*: This is also fulfilled in our case; we will not spend much time discussing this requirement.

Having reassured ourselves that we can think of our problem as a ‘bargaining problem’, we can now discuss the four axioms for a Nash Bargaining Solution — that is, the four properties that we will demand of our solution.

Pareto efficiency: First, we impose that the solution must be Pareto efficient; that is, we assume that the players do not agree on an outcome that leaves potential gains unexploited. In our case, this means simply that any solution must lie on the ‘-45 degree’ line.

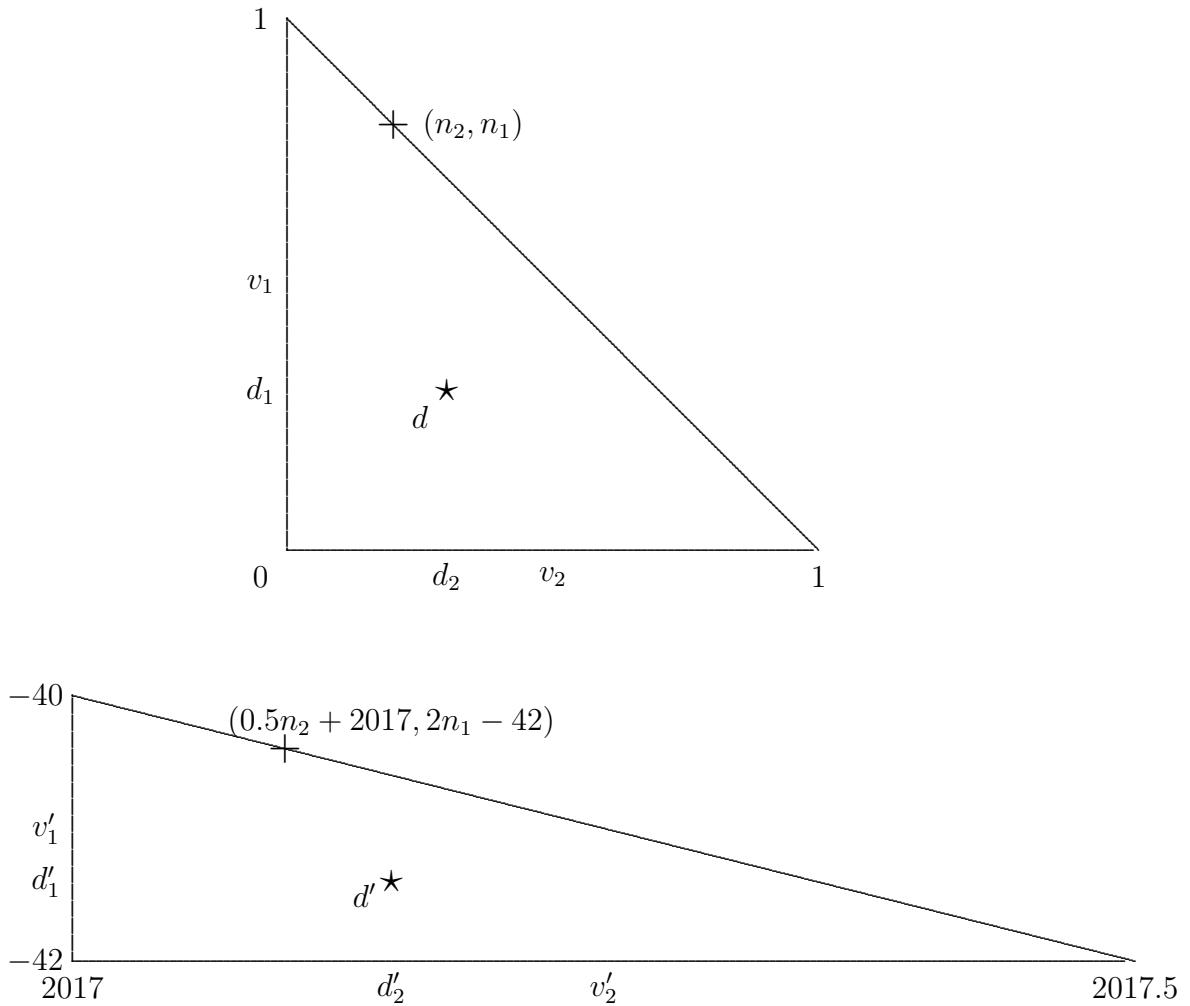
Symmetry: Second, we impose that, *if* the players are identical in their preferences and their possible outcomes, the solution must provide the same utility to each of them. The axiom of symmetry applies if *and only if* the problem is symmetrical — that is, if the set U is symmetric around a 45-degree line, *and* if the disagreement point d lies on that line. In that case, the axiom requires that the solution should also lie on the 45-degree line; note that, coupled with Pareto efficiency, this uniquely solves the problem — if the problem is actually symmetrical.

Invariance to equivalent payoff representations: Bargaining problems are expressed using von Neumann-Morgenstern (‘vNM’) utility. Remember that preferences under vNM utility are invariant to positive linear transformations of the utility.¹⁸ Therefore, any solution to a bargaining problem should also be invariant to positive linear transformations of the players’ utility. That is, consider a bargaining problem $\langle U, d \rangle$, and let the solution be (n_1, n_2) . Now consider a different bargaining problem, $\langle U', d' \rangle$, obtained by a positive linear transformation of $\langle U, d \rangle$. That is, if U is given by the set of pairs (v_1, v_2) , then U' is given by the set of pairs $(\alpha_1 + \beta_1 v_1, \alpha_2 + \beta_2 v_2)$, where $\beta_1 > 0$ and $\beta_2 > 0$; similarly, if d is the point (d_1, d_2) , then d' is the point $(\alpha_1 + \beta_1 d_1, \alpha_2 + \beta_2 d_2)$. In that case, the solution to the bargaining problem $\langle U', d' \rangle$ must be the point $(\alpha_1 + \beta_1 n_1, \alpha_2 + \beta_2 n_2)$.

Figure 8.3 shows a simple example. In the top graph, we have a candidate bargaining solution, (n_2, n_1) . In the bottom graph, we have the same bargaining problem, but distorted through a positive linear transformation in each dimension — such that $v'_1 = 2v_1 - 42$, $d'_1 = 2d_1 - 42$, $v'_2 = 0.5v_2 + 2017$, and $d'_2 = 0.5d_2 + 2017$. ‘Invariance to equivalent payoff representations’ then requires that the solution should follow the same transformation: thus, the candidate solution must now lie at $(0.5n_2 + 2017, 2n_1 - 42)$.

¹⁸ Thus, for example, if you take a positive linear transformation of a utility function, you do not change the CRRA or the CARA.

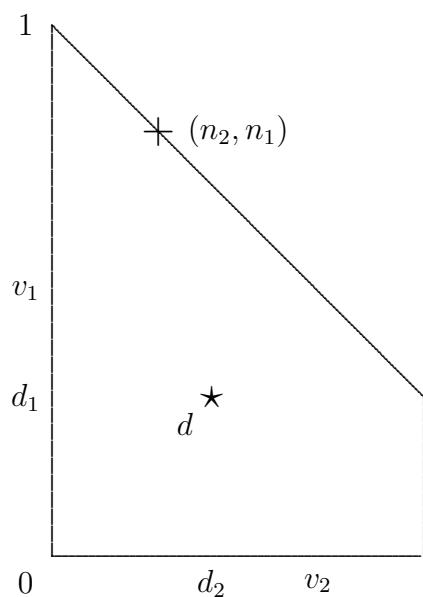
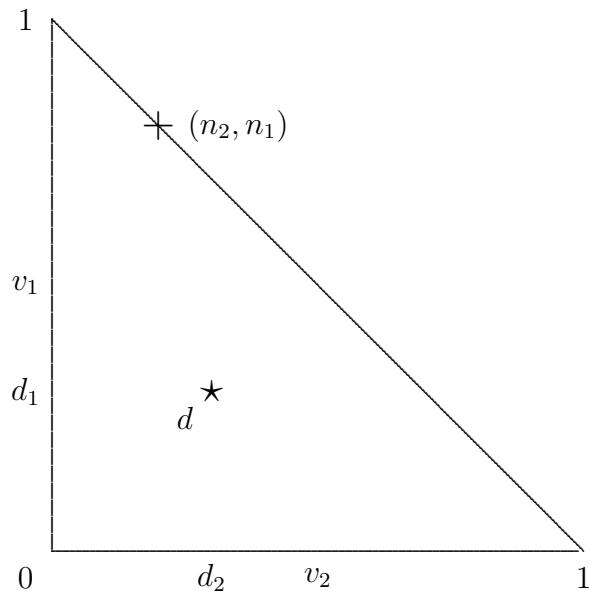
Figure 8.2: Invariance to equivalent payoff representations



Independence of irrelevant alternatives: The final axiom is ‘independence of irrelevant alternatives’. Put informally, this requires that the solution to a bargaining problem is not affected by deleting other possible outcomes.¹⁹ Figure 8.3 illustrates; ‘independence of irrelevant alternatives’ requires that, if point (n_2, n_1) is the solution to the bargaining problem in the top graph, it must also be the solution to the bargaining problem in the bottom graph — because (n_2, n_1) does not lie in the area that was deleted.

¹⁹ We will discuss ‘Independence of Irrelevant Alternatives’ again — when we discuss multinomial choice, in our final lecture on ‘Limited Dependent Variables’. It is worth noting that this is a *different* concept. In econometrics, the Independence of Irrelevant Alternatives is a statement about the ratio of choice probabilities; this is very different to the IIA concept considered by Nash. In econometrics, we should really use the term ‘Luce’s Independence of Irrelevant Alternatives’, or ‘Luce IIA’.

Figure 8.3: Independence of Irrelevant Alternatives



The Nash Bargaining Solution: Nash's famous bargaining result states that, for any given bargaining problem, there is a unique bargaining solution that satisfies these four axioms — and it is the solution to the following problem:

$$\max_{(v_1, v_2)} (v_1 - d_1) \cdot (v_2 - d_2), \quad (8.1)$$

subject to

$$(v_1, v_2) \in U; \quad (8.2)$$

$$v_1 \geq d_1; \quad (8.3)$$

$$v_2 \geq d_2. \quad (8.4)$$

The proof of this axiom is well beyond the scope of our course. Nonetheless, this is a beautiful result: it shows that, if we have a bargaining situation that can be expressed formally as a ‘bargaining problem’ — and if we are willing to accept the four axioms just discussed — then we can predict the bargaining outcome by solving a very simple maximisation problem.

For example, in our single-good bargaining case, we can work out how the coin will be split by solving:

$$(s_1, s_2) = \arg \max_{(s_1, s_2)} (s_1 - d_1) \cdot (s_2 - d_2), \quad (8.5)$$

subject to

$$s_1 \in [0, 1]; \quad (8.6)$$

$$s_2 \in [0, 1]; \quad (8.7)$$

$$s_1 + s_2 \leq 1. \quad (8.8)$$

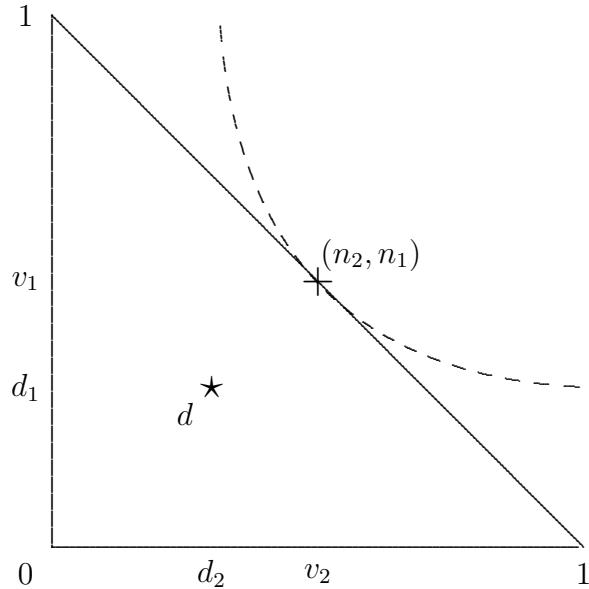
Figure 8.4 illustrates: it adds an ‘indifference curve’ for the ‘utility function’ $(s_1 - d_1) \cdot (s_2 - d_2)$; the Nash Bargaining Solution is the point at which that curve tangents the Pareto frontier.

Note that, if we set $d_1 = d_2 = 0$, the Nash Bargaining Solution predicts $s_1 = 0.5; s_2 = 0.5$. Go back to equations 7.3 and 7.5. Assume that $\delta_1 = \delta_2 = \delta \dots$ what happens in the limit as $\delta \rightarrow 1$?²⁰ Should this surprise us? Why or why not?²¹

²⁰ Remember L'Hôpital's Rule...!

²¹ This is a special case of a much more general result, due to Rubinstein (1982). The general result is beyond the scope of our course — but is something that you should explore further if you are interested in bargaining theory.

Figure 8.4: The Nash Bargaining Solution



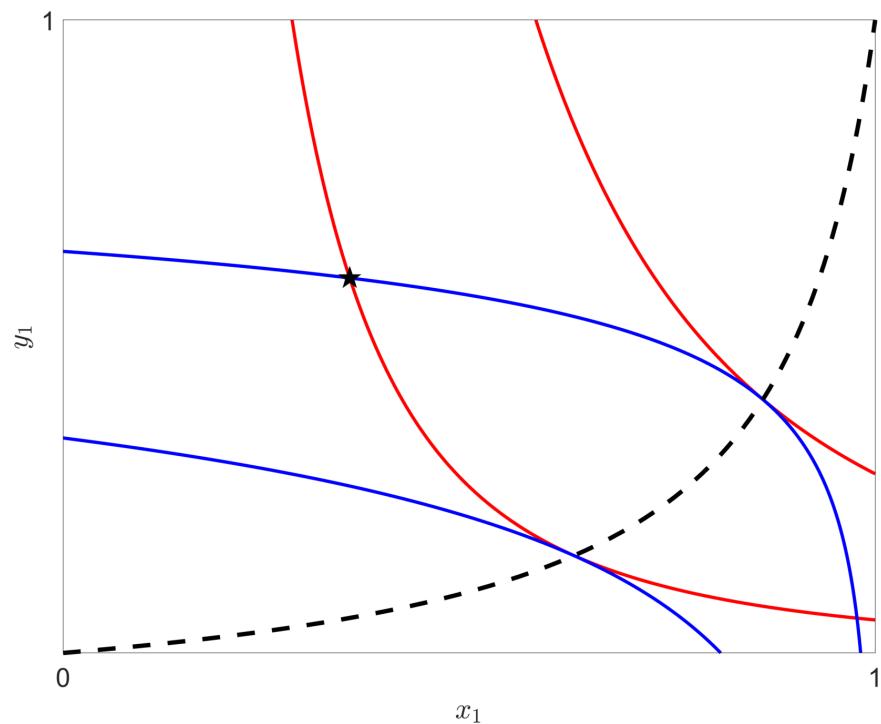
8.2 Illustration: The Edgeworth Box and axiomatic bargaining

So what? Of course, it's very *nice* that the axiomatic approach and the Nash Bargaining Solution can make a specific prediction about a simple one-good bargaining problem. However, the real value of the Nash Bargaining Solution is in providing predictions for bargaining problems in higher dimensions.

To illustrate, we will consider a two-good bargaining problem, as illustrated by the Edgeworth Box in Figure 8.5. The Edgeworth Box is often used to illustrate principles of general equilibrium, as the number of players increases towards infinity; in our case, we will stick with two players, and will think instead about the application of the Nash Bargaining Solution. We will denote the two goods as x and y , with a maximum quantity for each good of 1. We denote x_1 and y_1 as the respective amounts assigned to Player 1; therefore, $x_2 = 1 - x_1$ and $y_2 = 1 - y_1$ are the amounts assigned to Player 2.

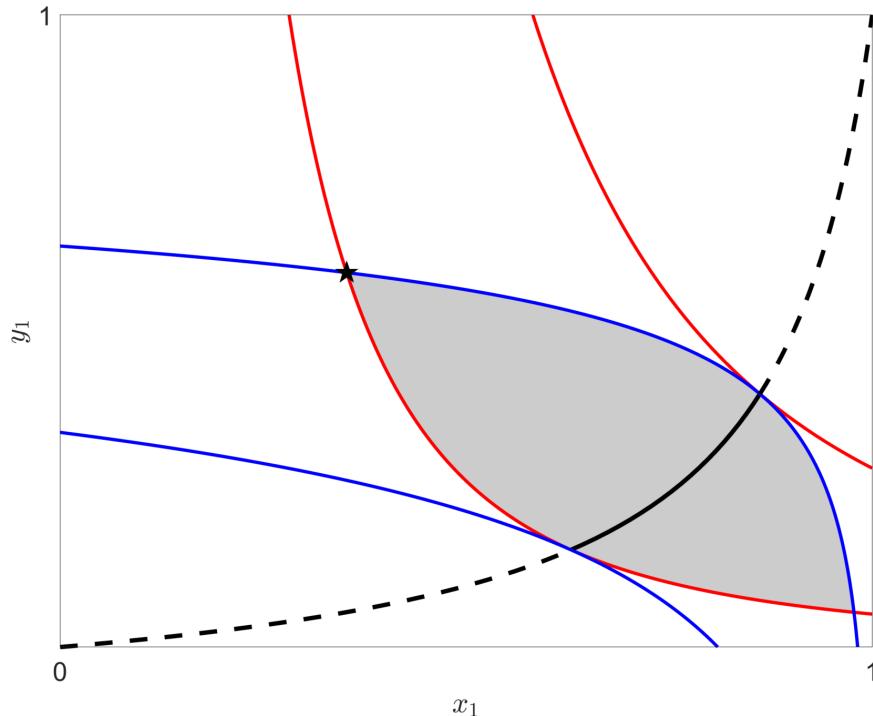
You should be familiar with all of the components of Figure 8.5: Player 1's indifference curves are in red, Player 2's indifference curves are in blue, and the dotted line shows the 'contract curve' — that is, the set of points that are Pareto efficient. The black star ('*) shows the 'initial endowment' — that is, the amounts that each player can take away if the bargaining process fails.

Figure 8.5: A two-good bargaining problem: Edgeworth Box



In Figure 8.6, we highlight the set of allocations that are ‘individually rational’ for both players: that is, the set of allocations that *both* players would prefer instead of the initial allocation. Within the set of individually rational allocations, the contract curve has a special name: ‘the core’. In a two-player game, we can think of the core as *the set of allocations that is individually rational and Pareto optimal*.²²

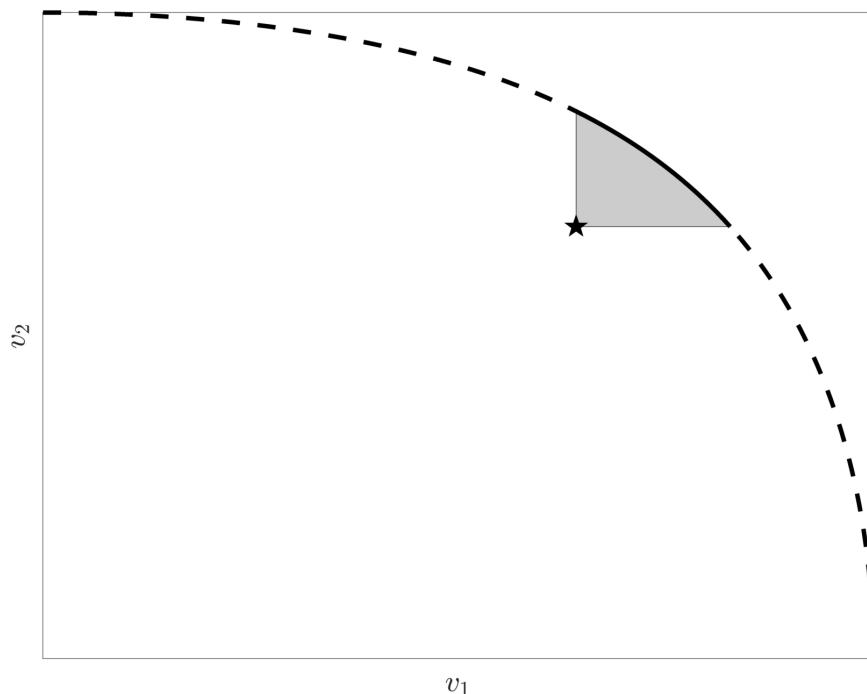
Figure 8.6: A two-good bargaining problem: individual rationality and the core



²² More generally, we can think of the core as *the set of allocations that cannot be blocked by any coalition of players*. In this two-player game, there are three possible ‘coalitions’: Player 1 alone, Player 2 alone, and both players acting together.

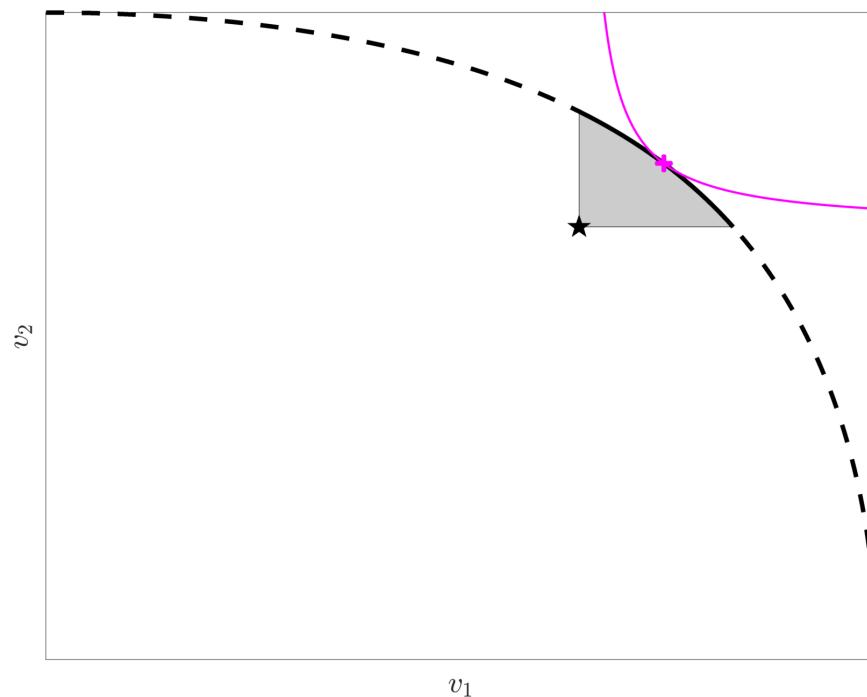
What does this look like in the axiomatic bargaining framework? Figure 8.7 represents the same problem in utility space. Note that the ‘initial endowment’ of the Edgeworth Box is conceptually identical to the ‘disagreement point’ in bargaining theory. Figure 8.7 uses the same notation as Figures 8.5 and 8.6 — the dotted line shows the Pareto frontier, ‘ $*$ ’ shows the disagreement point, the shaded area is individually rational for both players, and the solid line is the core.

Figure 8.7: Individual rationality and the core in utility space



In Figure 8.8, we add the ‘indifference curve’ for the objective function required by the Nash Bargaining Solution; ‘+’ shows the solution.

Figure 8.8: The Nash Bargaining Solution



Finally, in Figure 8.9, we translate the same point back to the Edgeworth Box.

Figure 8.9: The Nash Bargaining Solution in the Edgeworth Box

