Modelling the role of domestic politics in international environmental negotiations

Joachim Fasting

Master thesis in political science

Department of Political Science, University of Oslo

Spring 2013

Contents

List of Figures

Preface

Acknowledgements

I am greatly indebted to my advisor, Jon Hovi, for his support and guidance, without which I am certain that I would not have been able to complete this work.

Meta information

Words in text: 22736 (64 pages at 354 words per page)

Chapter 1

Introduction

1.1 Problem statement

Modelling international negotiations as if states were unitary actors is substantively dissatisfying because it obscures important drivers of a state's foreign policy. Intuition demands attention be paid also to the role of domestic institutions and constituents. The approach recommended by Putnam (1988) is to augment models of international negotiations with a domestic ratification[1] stage, where essential domestic constituents determine whether an agreement reached at the international level is to be accepted as national policy[2].

The theory of two-level games suggests a generic mechanism relating domestic factors (broadly defined) to international outcomes. This thesis is concerned with how the abstractions introduced by formal developments of the theory of two-level games can be used to model the role of domestic politics in international environmental negotiations. I consider two related research questions:

1. what formal abstractions are used to capture the role of domestic in formal developments of the theory of two-level games; and

2. how well do these abstractions approximate the role of domestic politics in determining the outcome of international negotiations over environmental issues?

[1] Putnam uses the term "ratification" to "refer to any decision-process at Level 2 that is required to endorse or implement a Level 1 agreement, whether formally or informally" (Putnam 1988, 436) and notes further that "It is sometimes convenient to think of ratification as a parliamentary function, but that is not essential. The actors at Level 2 may represent bureaucratic agencies, interest groups, social classes, or even 'opinion"' (Putnam 1988, 436).

[2] See Iida (1993*b*); Mo (1994); Tarar (2005); Hovi, Sprinz and Bang (2012) for secondary (to Putnam (1988)) introductions to two-level games.

The specific abstractions considered here concern the role of information, political competition and the distribution of political power, and finally the influence of institutional features of the political system.

1.2 Method

To answer the research question I will proceed in two broadly defined steps. First I will present three formal developments of the theory of two-level games. Each presentation aims to identify key abstractions introduced by the model and to discuss its strengths and weaknesses as a general description of the domestic sources of foreign policy. Second, I will discuss applications to international negotiations over environmental issues by using U.S. climate policy as the empirical point of reference. Before continuing, some justifications of the choice of models, the general approach, and the choice of empirical reference point is in order.

The models covered here were all developed to answer specific questions about the impact of domestic constraints on the international bargaining position of the executive, which means that some of the trade-offs that have been made are not necessarily conducive to realistic accounts of how domestic politics operates. Each model illuminates an important aspect of the domestic sources of foreign policy, however: the role of information, political competition and the distribution of political power, and the political system. Together, these models provide examples of abstractions that can be used to formally analyse involuntary defection due to uncertainty about domestic constraints, how domestic groups leverage political power to shape foreign policy, and finally how the incentive to cooperate depends on the arrangement of domestic political institutions. Although far from covering every relevant aspect, the selection of models provides reasonable coverage for some of the most important aspects of the role of domestic politics.

Casual observation of real-world politics makes clear that there are fundamental differences in the way politics works across the policy spectrum, and it is not necessarily valid to apply a general formal statements about domestic politics to a specific policy area. The theory of two-level games as stated by Putnam (1988) does not differentiate between policy domains, and neither do the formal developments of the theory that are considered here. We may take this to mean that the postulates of the theory are intended to be generally valid. That is, it is assumed that the role of domestic politics is similar across the policy spectrum, so that the general formal statements about domestic politics can

be translated to the domain of environmental policy with a reasonable degree of confidence. It follows that the effectiveness of appealing to domestic constraints depends on information in much the same way regardless of the issue under negotiation, that the mechanisms that are used to leverage political power are fundamentally similar across the policy spectrum, and that the incentive structure created by political institutions is a stable background variable. Thus, to discuss these aspects of environmental policy it is reasonable first to assess carefully the general implications of the formal account of the domestic sources of foreign policy.

The purpose of the discussion is not to criticise the models as such, but rather to identify modelling decisions that are helpful for understanding the role of domestic politics in shaping the outcomes of international negotiations. A natural way of identifying helpful modelling decisions is to compare the accounts of domestic politics given by the models, highlighting expressive benefits or drawbacks that follow from the assumptions made in each model. The presentation of each model is intended to explain the details of how the model works, its general implications, and its overall description of the relationship between domestic politics and foreign policy. The formal presentation attempts to give an intuitive account of how the model works, and may be safely skipped by readers already familiar with the mechanics of the underlying game. Three points are emphasised for each model: the relationship between domestic politics and foreign policy outcomes as described by the model; what motivates cooperation; and under what conditions cooperation is likely to occur. Although drawn from a limited set of models, this general discussion of abstractions should provide a meaningful base for discussing the application of two-level models to the problem of capturing the role of domestic politics in international environmental policy.

The terms “abstraction”, “validity” and “expressive power” are defined as follows. Abstractions are constructs that capture the common structure of some set of concepts. I use a notion of abstraction inspired by computational systems. A system where several subsystems duplicate some feature can be improved by factoring that functionality into a common definition, which is referred to as an abstraction. For example, instead of reasoning about the voting procedures in each democratic state individually, we invent an abstract voting procedure that captures relevant (to our purposes) aspects of voting procedures in democratic states (or the particular set of states we have in mind) and use that to formulate general statements about voting in democratic states.

I consider the “validity” of an abstraction to be related to its functionality, not whether it “corre-

sponds" to some referent. That is, whether we use Newton's method to compute the square root is not important, so long as the output satisfies the specification (i.e., the stated functional requirements). I take a model to mean a system of abstractions designed to reason about a problem domain. A reasonable model of some class of phenomena has to be made up of abstractions that are functionally related in approximately the same way as the actual system, though the level of detail and generality will vary with the purpose of the model, and the validity of an abstraction depends on whether it helps the system fulfil its purpose.

When building systems, we have to begin with a set of basic abstractions, or "primitives", the fundamental nouns and verbs of our vocabulary for talking about the problem domain. I take the "expressive power" of some vocabulary to refer to the number of problems that can be reasoned about using only the definitions it contains weighted by the number of definitions. If adding a definition to the vocabulary enables the removal of two or more existing members, then the vocabulary has increased in expressive power. For the general problem of describing the role of domestic politics in international environmental, I will treat abstractions as if I were to build a system to simulate the relationship between domestic politics and international politics, and I will say that an abstraction is helpful to the extent that it furthers this goal.

The justification for choosing the U.S. position in the last rounds of the Kyoto proceedings as the empirical point of reference is straight-forward. For making theoretical points, any case exhibiting the relevant features would do. Substantively, however, cases involving the United States are important simply by virtue of the fact that U.S. participation is crucial to any international effort to mitigate global environmental problems — what most other states are doing is not significant in comparison. Choosing a less researched case would increase the likelihood of discovering new insights (pertaining to that case). By choosing a relatively old and well-researched case, new insights are traded for greater continuity with the literature.

1.3 Plan

The first three chapters are dedicated to presenting existing formal two-level models, discussing their strengths and weaknesses, and identifying key abstractions that can be used to describe the relationship between domestic and international levels of politics. Chapter 2 concerns the role of information, presenting the models of Iida (1993*b*). Chapter 3 concerns the formal account of political competition

and how the distribution of political power affects foreign policy, as modelled by Mo (1994). Chapter 4 concerns abstractions for capturing how certain aspects of the political system affect the incentives of executives and the workings of political power, as analysed by Tarar (2005). Chapter 5 considers the application of the abstractions discussed in previous chapters to international environmental negotiations, using the Kyoto proceedings as an empirical point of reference. The final chapter summarises the major points of interests and suggests specific avenues for future research into modelling the role of domestic politics in international environmental negotiations.

Chapter 2

The role of information

2.1 Introduction

Iida (1993*b*) is an early formalisation of the theory of two-level games, which addresses the role of information about domestic constraints in determining international negotiation outcomes. Treating states as unitary actors, however, precludes the possibility of uncertainty about domestic constraints, because unitary actors must know their own preferences, restricting state uncertainty to that of the "unobservable properties of the opponent" (Iida 1993*b*, 403). Including the legislature as an independent veto player allows for modeling the case where the negotiator[1] is uncertain about whether an agreement will be accepted and goes some way in relaxing the unitary actor view of states (Iida 1993*b*, 403). Iida (1993*b*) proceeds by specifying a basic model of negotiations among states as a two-stage process, which is then analysed under different informational structures to clarify the distinction between domestic and international asymmetric information. This endeavour amounts to investigating the formal relationship between informational structure and Putnam's (1988) hypotheses about win-sets[2] and negotiation outcomes, specifically concerning how domestic constraints enable the chief negotiator to extract greater concessions from the opponent.

This chapter is concerned with how the abstractions exemplified by the model proposed by Iida (1993*b*) can be used to model the role of domestic politics. I will proceed by giving a detailed presentation of the baseline model and its three variations, seeking to provide intuitive explanations of what drives the outcome in each case. Having established these preliminaries, I will then turn to a discus-

[1] I will interchange "negotiator" and "executive" throughout.

[2] Putnam defines the win-set as "… the set of all possible Level 1 agreements that would 'win' — that is, gain the necessary majority among the constituents — when simply voted up or down" (Putnam 1988, 437)

sion of what the model tells us about the role of domestic politics in general. Finally, I will discuss the assumptions that are made about the negotiators' access to and ability to process information. I will clarify what it means for the negotiator to be uncertain about the preferences of his constituency, what it means for two negotiators to agree about the likelihood of ratification, and discuss the plausibility of such agreement. There are two key points:

1. the model is designed in a way that severely limits expressive power; the range of institutional arrangements that it can be reasonably said to approximate is restricted and the description of how domestic politics operates shallow; and

2. while it is plausible that a negotiator can be uncertain about the preferences of his own constituents if the issue at hand is hotly contested, the model appears to imply an institutional setup that is difficult to reconcile with such uncertainty.

2.2 Formal description

2.2.1 Rubinstein bargaining with one-sided ratification

Iida (1993*b*, 406–409) describes negotiation over some issue between countries 1 and 2 as an instance of Rubinstein bargaining with one-sided ratification. An agreement is reached by alternating offers at the international level but must then pass ratification in country 1 before taking effect. Ratification is determined by voting among N (odd) legislators. Each legislator has his own evaluation of the status quo and never supports agreements he[3] considers to be worse than no agreement. The agreement is implemented if supported by a (simple) majority of the legislators in country 1, otherwise the negotiation ends in failure (the outcome is final).

Formally, negotiations at the international level (Level 1) are described as follows (Iida 1993*b*, 407–409). Let $N_P = \{1, 2\}$ be negotiators[4] representing country 1 and 2. The object of bargaining is the division of gains from cooperation, normalised to 1. A period of play consists of one negotiator making a proposal to which the other responds. Negotiators take turns proposing and responding to offers over an indefinite number of time periods ($T = \{1, 2, \ldots\}$) until an offer is accepted. An offer is a division $x = (x_1, x_2)$ of gains, where $x_{i \in N_P}$ refers to the share going to i under the proposed

[3] I will generally use "he" for all pronouns; I do this for convenience, not to further a sexist agenda.

[4] The negotiators representing country 1 and 2 are referred to as "1" and "2" or as "he" and "she"; "he" always refers to 1 and "she" always refers to 2.

agreement. Let X denote the set of all such divisions. Because agreements are timed, players' preferences are defined over $X \times T$. In this version of the bilateral bargaining game, negotiator 1 is subject to a ratification constraint[5]. Let z be the share going negotiator 1[6]. Let $\delta_{i \in N_P}$ denote the discount factors. An agreement z reached in period $t \in T$ gives the payoffs $\delta_1^t z$ and $\delta_2^t(1-z)$. That is, we have a Rubinstein bargaining game of alternating offers where the players have time preferences with a constant discount rate, with the caveat that payoffs obtain only after ratification.

Domestic ratification (Level 2) is formally described as follows (Iida 1993*b*, 408–409). There are N_L legislators who vote yes or no to the agreement z reached at Level 1. No abstention is allowed[7]. Each legislator's evaluation of the no agreement outcome is represented by the status quo payoff $s_{i \in N_L}$ in the range $\pm\infty$. Legislator i is indifferent[8] iff $s_i = z$ and prefers z iff $z \geq s_i$. Ratification occurs if at least $(N_L + 1)/2$ legislators support the agreement. If ratification fails, negotiators receive a payoff of 0 whereas legislators retain their status quo payoffs. Voting is sincere[9] and i will support z iff $z \geq s_{i \in N_L}$ in equilibrium. There exists a median legislator[10], with status quo payoff s_m, whose vote effectively decides the ratification outcome.

The main driver of the model is the cost of delay incurred by rejecting a proposal (captured by the discount factor), which motivates players to conclude agreements rather than continue to bargain forever (Rubinstein 1985, 1152). This mechanism relies (crucially) on several assumptions about the preferences[11] of the negotiator. The most basic assumptions are that the negotiators value what they bargain over, that more is better, that sooner is better than later (all else being equal), and that previous offers are irrelevant for the value of the current offer (there are no regrets) (Rubinstein 1982, 98–99).

A further assumption about the preferences of the negotiator is that he accepts an offer if indifferent between accepting it and continuing to bargain for another round. The negotiator accepts an offer if it is as least as good as the (discounted) expected value of continuing to bargain; he rejects it if he

[5]Because 1's share determines whether the agreement enters into force, it is the only relevant factor in this analysis and so is given special treatment in the presentation (e.g., by having all equilibrium proposals refer to the share going to 1 regardless of who is making the offer).

[6]Note that in the bilateral bargaining game, the agreement is completely defined by only specifying one of the halves. That is, an agreement z implies a division $(z, 1-z)$.

[7]Either politicisation is ignored or the issue achieves complete mobilisation, see Putnam (1988, 445).

[8]The assumption that i supports an agreement if he is indifferent between it and no-agreement is necessary to avoid indeterminacy.

[9]A legislator supports an agreement only if he prefers it to the status quo, see Iida (1993*b*, footnote 8, p. 409).

[10]Because N_L is odd.

[11]Formally these assumptions are statements about the properties of the negotiator's *preference relation*, commonly denoted by $\succeq_i$ for some player i. Here $\geq$ is used instead, because it carries the right connotation and is recognisable to most readers. The statements about the preference relation are true by definition and all arguments that follow from the model are true insofar as the preference relation of the actual people involved in bargaining have the same properties. See Rubinstein (1982) for a detailed discussion.

expects that the opponent is willing to accept a counter-offer that is good enough that it offsets the cost of delay (Rubinstein 1982, 102). Formally, i accepts some agreement x at time t rather than make the counter-offer x' iff $(x,t) \geq (x',t+1)$. By assuming stationary time preferences[12], this condition holds when $x_i \geq \delta_i x_i'$, where x_i is i's share.

Counter-intuitively, the standard Rubinstein model predicts that negotiations end immediately (Rubinstein 1982, 107). The reasoning behind this is that a negotiator who knows all strategically relevant information should be able to identify the minimum offer that the opponent is willing to accept[13] and make that offer immediately to avoid delay. The opponent, knowing what the proposer knows, can do nothing better than accept the minimum offer immediately.

The requirement that an agreement pass ratification in country 1 serves the function of an outside option (Iida 1993*b*, 410). An outside option means that one party may choose to end bargaining, ostensibly because a more profitable alternative is available to him (Osborne and Rubinstein 1990, 55). The canonical example is that of buyers and sellers in a market. A seller of a good may lose interest in continuing to bargain with a particular buyer once he realises that he will not get the price he wants. The existence of other potential buyers gives the seller the outside option of taking his business elsewhere. Making his partner aware of this fact might persuade her to be more forthcoming.

For the situation at hand, the outside option consists of continuing the current state of affairs[14]. This option is activated if 1 makes an agreement that the legislature refuses to ratify: in that case country 1 unilaterally terminates negotiations. The existence of an outside option impacts the outcome of bargaining if the value of opting out is sufficient to support a credible threat of withdrawal, otherwise it is of no consequence (i.e., the result is given by the regular Rubinstein model) (Osborne and Rubinstein 1990, 55). Here, the value of the option is the median legislator's status quo payoff. Obviously, it takes more to convince a satisfied legislator of the need for international cooperation than it takes to convince one eager for change. A content legislature — whether dominated by isolationists or simply

[12]The assumption of stationary time preferences, briefly, means that players prefer an agreement now rather than later but that they are indifferent to the number of rounds that have passed up to that point. That is, $(x,t) \geq (x',t+1)$ is the same as $(x,0) \geq (x',1)$.

[13]The negotiator is assumed to be aware of the structure and rules of the bargaining game and to formulate his strategy by applying the appropriate solution concept.

[14]The outside option exists only because negotiators get a single chance at formulating an agreement — legislators can not send the proposal back to the negotiating table (Iida 1993*b*, 417). This means that an unmodified version of the model is suitable for scenarios that are reasonably approximated by one-shot negotiations. For example, one may imagine that the current round of talks is the last in a protracted process that has left domestic decision-makers cynical and weary of continuing to pay the political cost of stale-mate, explaining their severe reaction to failure.

well-off under the current arrangement[15] — makes for a credible threat of negotiation breakdown if demands are not met.

2.2.2 Publicly known domestic constraints

Publicly known domestic constraints implies that legislators' preferences are publicly observable, that negotiator 1 and 2 are equally capable of assessing these preferences, and that their assessments are, in fact, true. When these assumptions obtain, both negotiators can know with certainty whether a particular agreement will pass ratification. Assuming, as the base model does, that negotiators do not derive any auxiliary benefits from bargaining itself and that they only gain by making agreements that end up being accepted by country 1's legislature, then neither negotiator can do any better than to leverage their respective bargaining power to the extent allowed by 1's constraints. Thus, the model states that the outcome of international negotiations is determined primarily by the demands of country 1's legislature and whether the legislature is willing to unilaterally terminate negotiations if their demands are not satisfied (i.e., constraints that are not sufficient to support a threat of unilateral termination are irrelevant).

Formally, the assumption of publicly known constraints entails that s_m is known by both 1 and 2 and that both know that 1's win-set is $[s_m, 1]$. Assume $s_m < 1$ and let the pair $\langle x^\star, y^\star \rangle$ denote proposals by 1 and 2 in equilibrium[16]. Then, the model predicts a division of gains as described by proposition 1 (see figure~2.1). Agreement is reached immediately, as in the standard Rubinstein model with complete information, and ratified with certainty. The implications of this result are that 1 derives an advantage from domestic constraints only if those constraints are severe, and only because the constraint is common knowledge (Iida 1993*b*, 409).

The benefit of cooperation (the "pie") is represented by 1, s_m represents the minimum benefit to country 1 necessary for an agreement to pass ratification, and δ_1 and δ_2 represent how much 1 and 2 value future benefits and thus their willingness to incur the cost associated with delay that follows from rejecting a proposal. That s_m is public knowledge means that the value of s_m is known by everyone with certainty, that this fact is known by everyone, and so on (Osborne and Rubinstein

[15]This is related to the discussion of the determinants of win-sets in Putnam (1988, 442–443), in particular the impact of a lower cost of no-agreement and the ratio of isolationists to internationalists in the legislature (i.e., legislators who, for some reason or other, devalue or value cooperation).

[16]Note that both $x^\star$ and $y^\star$ refer to the size of country 1's share.

$$\langle x^\star, y^\star \rangle = \begin{cases} \langle 1-\delta_2(1-s_m), s_m \rangle & \text{if } s_m > \delta_1(1-\delta_2)/(1-\delta_1\delta_2) \\ \langle (1-\delta_2)/(1-\delta_1\delta_2), \delta_1(1-\delta_2)/(1-\delta_1\delta_2) \rangle & \text{otherwise} \end{cases}$$

Figure 2.1: Proposition 1. The unique subgame perfect equilibrium with complete information (Iida 1993*b*, 409). $x^\star$ and $y^\star$ are the equilibrium proposals by 1 and 2, both refer to 1's share (2's share is $1-x$ for some x). The proposition states that the outcome is either $x^\star$ or $y^\star$ depending on whether 1 or 2 makes the first proposal; the conditional refers to the plausibility of the threat that 1's constituents will unilaterally withdraw from negotiations unless their demands are met.

1994, 95). Given this, we may reason about proposition 1 as follows[17].

The expression $s_m < 1$ states that the legislature in country 1 is willing to make at least some concessions to country 2 to secure agreement. This assumption is necessary for negotiations to occur in the first place. Clearly, there can be no mutually beneficial agreement between country 1 and country 2 if 1 is unable to accept anything less than all the benefits. To labour the point: setting $s_m = 1$ gives $x^\star = 1-\delta_2(1-1) = 1-\delta_2 0 = 1-0 = 1$ and $y^\star = 1$. That is, country 1 gets all the of the pie and country 2 gets none. Knowing this, country 2 has no reason to negotiate.

Provided that there is room for negotiation, proposition 1 states that there are three factors that determine the final outcome of bargaining: the severity of 1's constraints, who makes the first offer, and the parties' willingness to incur the cost of delay associated with rejecting an offer. The severity of 1's constraints chooses between two broad classes of outcomes. There are two possible "worlds" that the negotiators may find themselves in: one where 1 is severely constrained, one where he is not. Whichever world obtains, the order of events and the patience of negotiators determine the details of the agreement. The impact of these factors varies between worlds, however.

Because the model demands that negotiations end immediately, whoever makes the first offer effectively decides the size of country 1's share of the benefits. Being the first mover is always an advantage in the standard Rubinstein model with complete information, but here it is conditioned by the limit imposed by 1's constraints. The same is true for the impact of differences in patience. Greater patience should, according to intuition, translate to greater ability to extract concessions from the opponent. Here, patience may or may not be a benefit, again limited by the size of 1's constraints[18].

The expression $s_m > \delta_1(1-\delta_2)/(1-\delta_1\delta_2)$ is true whenever the demands of country 1's legislature

[17]This line of reasoning is directly analogous to the standard account Osborne and Rubinstein (1990); Rubinstein (1982).

[18]See Osborne and Rubinstein (1994, 125–) for descriptions of the first-mover advantage and the comparative statics of impatience under the standard model.

are sufficient to support a threat to unilaterally terminate negotiations. If true, country 1's legislature will be more than happy to continue the current state of affairs rather than pursue an agreement with country 2, in the case that 2 fails to meet their demands. Under this condition, 1 will be severely constrained in terms of what he will be able to concede to 2 without triggering the outside option.

If 1 is severely constrained, the outcome is given by reasoning as follows. When the negotiator representing country 1 makes a proposal, he knows that whatever agreement he proposes must at least give his own country s_m of the benefit from cooperation. He is unable to make an offer to country 2 that leaves less than s_m for country 1. The maximum concession negotiator 1 can make is $1 - s_m$, whatever remains of the whole after country 1 has taken the share necessary to achieve ratification. The concession country 1's negotiator *has* to make to secure 2's acceptance, however, depends on how patient 2 is. This notion is captured by the expression $\delta_2(1 - s_m)$. If 2 is impatient (δ_2 close to 0), 1 can get away with offering less than the maximum concession, because 2 values resolving the matter quickly over holding out for a better deal. Setting $\delta_2 = .2$, say, means that delay for a single round will reduce the value of an agreement to only 20% of its present value, and as the rounds accumulate so does the cost of delay.

For 2 to consider rejecting 1's offer, she must expect to be able to get 1 to accept a counter-offer that is large enough to offset the cost of delay. This is implausible, as it supposes that 1 is bluffing and that he is in fact willing to settle for less — so much less that it makes up for the cost incurred by 2. The possibility of bluffing is ruled out under complete information. Should 2 reject the proposal despite knowing that bluffing is impossible (say, by mistake), 1 will simply punish 2 by rejecting the counter-offer, and repeat his original offer. At that time, the offer made by 1 would be worth less to 2 than it was initially (assuming $\delta_2 < 1$), so that 2 would have been better served by accepting it immediately. Knowing this, 2 accepts 1's proposal immediately.

Proceeding in this manner it is easy to see that as δ_2 approaches 1, the more 1 must concede — but only up to the limit imposed by his constraints. The reason why 2 will not demand more than the maximum concession 1 can make is that they both know that the agreement will not pass ratification, so the expected value of demanding more than the maximum is 0. That is, even if 2 is extremely patient she can do no better than to accept the maximum concession.

We also see that the first-mover advantage applies mainly to 1 in this scenario. If 1 goes first, he can exploit an impatient 2 by reducing his concessions below the maximum as far as 2's impatience

allows, as already described. Negotiator 2, however, always makes the offer s_m and is unable to exploit a favourable difference in patience. It can be better for 2 to go first than to go second if she is very impatient, but this is more of a "defensive" benefit that comes from not being exploited. In summary, 1's constraints guarantee that he will get at least the minimum required by his constituency regardless of who goes first, although he surely prefers to go first if 2 is in fact impatient.

When 1 is not severly constrained, the outcome for 1 depends entirely on who goes first and the relative willingness of 1 and 2 to incur the cost of delay associated with rejecting a proposal, which is reflected in that s_m does not factor into determining the proposals that 1 and 2 will make. The outcome turns crucially on the constrained negotiator's outside option. When the home negotiator is heavily constrained, he can credibly claim that the best ratifiable agreement is that which gives him at least the minimum required to avoid unilateral termination of negotiations, otherwise the constraint is insignificant and the best ratifiable agreement is decided by their "raw" bargaining power deriving from the willingness to incur the cost of delay associated with rejecting an offer.

2.2.3 International asymmetric information

The second variant of the base model considers the more likely situation where one of the negotiators is unsure about the severity of the opponent's constraints. Relaxing the assumption of complete information is the same as ruling out direct observation of the facts, leaving open the possibility of bluffing and cheap talk, so that 2 will have good reason to doubt 1's claims of being constrained. Without direct observation, 2 must rely on rules of inference and learn the true price of 1's cooperation by observing his actions. This makes it attractive for 1 to not only speak as if constrained but to also act that way. The problem is that the person he is trying to convince is fully aware of this incentive and might not be so easily impressed. The bargaining power of 1 is then not so much dependent on the actual size of his constraints as it is on his ability to effectively communicate to 2 that he is constrained, whether that is the case.

The preferences of essential domestic actors are known only to the negotiator from the constrained country and this fact is common knowledge (Iida 1993*b*, 410). Formally, this is expressed as an extensive bargaining game between incompletely informed players (Rubinstein (1985) and also Osborne and Rubinstein (1990, chapter5)). Proposition 1 (see fig.~2.1) states that the outcome of bargaining is determined chiefly by the nature of negotiator 1's constraints. Refer to the nature of 1's constraints as

his "type". There are two types of negotiator 1: a weak type with nonbinding constraints (denoted 1_{w}) and a strong type with binding constraints (denoted 1_{s}). Let $b = \delta_1(1-\delta_2)/(1-\delta_1\delta_2)$. Then

$$s_m = \begin{cases} s_H > b & \text{if player} 1_s \\ s_L \leq b & \text{if player} 1_w \end{cases}.$$

Note that the condition is identical to that of proposition 1 (2.1). The only difference thus far is that the condition chooses between types of player 1, not between broad categories of outcomes of bargaining as in proposition 1. The core driver of the outcome here, as noted above, is 2's *beliefs* about 1's constraints, not what the constraints are.

The true type of 1 is chosen by Nature before the negotiations begin: 1_{s} is chosen with probability $0 < p^0 < 1$ and 1_{w} with probability $1 - p^0$[19]. These probabilities are known by both 1 and 2. Initially, 2 is uncertain about 1's true type (1 knows his own type with certainty, 2 only knows p^0). As 1 and 2 make proposals back and forth, 2 updates her beliefs about 1 according to Bayes' rule[20] whenever possible, treating 1's offers and counter-offers as evidence that she uses to fix her opinion about his true type.

There are three important assumptions about 2's beliefs. First, she is not allowed to change her beliefs freely. Updating occurs only if she observes behaviour that is inconsistent with her prior beliefs about 1's type (Osborne and Rubinstein 1990, 96). This means that as long as 1's behaviour conforms to 2's prior belief about his type, her belief remains unchanged. Second, she is not allowed to change her belief once she has become convinced of 1's type (i.e., if she sets $p = 1$ or 0). In technical terms, this means that 2's certainty triggers a transition from a bargaining game of incomplete information to one of complete information (see Osborne and Rubinstein (1990, 96) and the proof in Iida (1993*b*, Appendix)). Third, any observed deviation on 1's part will cause 2 to make an "optimistic conjecture" (see Iida (1993*b*, 420) and also Osborne and Rubinstein (1990, 96)) whereby she becomes convinced that 1 is of the weak type. For 1_{w}, this event amounts to 2 calling his bluff, whereas for 1_{s} it amounts to 2 misunderstanding the true nature of his constraint. Either way, 1 loses out when 2 drives a tough bargain, and both 1_{w} and 1_{s} have good reason to want to avoid it. The optimistic conjecture is commonly interpreted as a way for 2 to punish 1 for deviating from the equilibrium (and is crucial for justifying the outcome predicted by the model) (Osborne and Rubinstein 1990, 96).

[19]Note that, if $p^0 = 1$ or 0, then the outcome is as under complete information.

[20]The likelihood of observing some behaviour given the type weighted by the overall probability of facing that type.

Knowing that 1 is interested in making it look like he is in a better position than what he perhaps is, she should expect him to act as if he is the strong type regardless of his true type. If p^0 is high, then it is prudent for 2 to proceed as if 1 were strong. The constraints upon her beliefs imply that a weak 1 should be able to benefit from 2's caution simply by acting in a way that is consistent with having severe constraints. That is, there is a cut-off point for p^0 above which 2 will "give up", set aside her doubt, and treat 1 as if he was the strong type (Rubinstein 1985, 1153). As long as player 2 is uncertain, both types of player 1 will make identical equilibrium proposals (Iida 1993*b*, 314). This outcome confirms intuition: negotiator 1 will do his best to appear strong regardless of his actual position, and benefits if 2 is sufficiently close to being convinced that she is facing a severely constrained colleague.

There are two general outcomes: one where 2 believes 1 is strong and one where 2 believes 1 is weak. It is still assumed that $s_m < 1$, even in the case where $s_m = s_H$, i.e. even when 1 is strong there exists some possibility of a mutually beneficial outcome, with the same justification as under proposition 1.

The first general class of outcomes is known as the "pooling" equilibria, where 2 is sufficiently convinced that 1 is the strong type and induced to make and accept offers accordingly. In period 0, 1_{w} and 1_{s} proposes some $x^\star$ as per (2.1) (see figure~2.2) and player 2 accepts. If period 1 is reached without updating (i.e, 2 has no reason to change her prior belief that 1 is strong), then she will propose $y^\star = s_H$. This equilibrium occurs when p^0 is high, i.e., when 2 is nearly convinced that 1 is strong and observes no deviation that makes her change her mind. 1 is aware of 2's belief and can safely act accordingly, regardless of his true type, and 2, being nearly convinced, can do nothing better than to act as if 1 is strong. Note that this outcome is equivalent to the complete information outcome when 1 is heavily constrained, the only difference being that a weak 1 can achieve benefits beyond what is possible under complete information.

The second class of outcomes are the "separating" equilibria. Initially 1_w and 1_s propose some $x^\star$ per (2.2) (see figure~2.2) and 2 accepts. If period 1 is reached without updating, then 2 proposes $y^\star = \delta_1 x^\star$. If 1 turns out to be weak, then he accepts the offer. If 1 turns out to be strong, then he rejects the offer and makes the counter-offer $x_s = x^\star$ in period 2. This means that when 2 is more or less certain that 1 is weak, she will make an offer that is in accord with the discounting differential. If it turns out that her beliefs are wrong, 1 rejects and makes a new proposal, which signals to 2 that

$$\max\{s_H, (1-\delta_2)/(1-\delta_1\delta_2)\} \leq x^\star \leq 1-\delta_2(1-s_H) \tag{2.1}$$

$$\max\{s_H, (1-\delta_2)/(1-\delta_1\delta_2)\} \leq x^\star \leq [1+p^0\delta_2^2-d_2(1-p^0)]/[1+p^0\delta_2^2-\delta_1\delta_2(1-p^0)] \tag{2.2}$$

Figure 2.2: Proposition 2. 1's proposals in the pooling equilibria and the separating equilibria (Iida 1993*b*, 411).

he is in fact of the strong type so that she can do nothing better than to accept his counter-offer. The only way there can be delay in this model is if 2 misjudges 1's position.

When 2 is near convinced that 1 is heavily constrained, 1 may persuade her that the best ratifiable agreement is that which gives him at least the minimum demanded by his legislature, as under complete information. The difference is that 1 may persuade 2 of this even when not severely constrained, provided that 2 deems it highly likely that 1 is not bluffing. 2 may not be impressed by 1's posturing, however, and decide to drive a tough bargain. Then, if 1 is bluffing he is called on it, but if he is not bluffing he has to reject the agreement and incur the cost of delay. Iida (1993*b*, 416) notes that cooperation is not endangered even if 2 suspects 1 is bluffing; the only cost of uncertainty is the possibility of a slight delay.

This line of reasoning is identical to the standard version of alternating offers with incomplete information among individual market participants. Including the legislature serves mainly as a narrative device, an explanation of what gives rise to bargaining strength and weakness. The model is formally equivalent to treating 1's constraints as one of the "unobservable properties of the opponent" that 2 has to infer, as would be done under the unitary actor view (Iida 1993*b*, 413).

2.2.4 Domestic incomplete information

If the information about win-sets is incomplete there is a possibility of so-called "involuntary defection", whereby a negotiator signs an agreement in good faith that later turns out to fail ratification, forcing the signatory to renege on the agreement[21]. This implies that the home negotiator, to some extent, is uncertain about the likely outcome of ratification. Iida (1993*b*) assumes that the negotiators, although uncertain, agree about the probability that any given agreement passes ratification. Because the negotiators receive nothing if the agreement fails to achieve ratification, both have a strong in-

[21] "Involuntary defection [...] reflects the behaviour of an agent who is unable to deliver on a promise because of failed ratification" (Putnam 1988, 438)

$$\max \text{ s.t. } F(x^\star)(1-x^\star) = \delta_2 F(y^\star)(1-y^\star) \tag{2.3}$$

$$\max F(y)(1-y) \text{ s.t. } F(y^\star)y^\star \geq \delta_1 F(x^\star)x^\star \tag{2.4}$$

Figure 2.3: Equilibrium offers by 1 and 2 in the model with domestic incomplete information (Iida 1993*b*, 414) 1 and 2 attempt to identify the best ratifiable agreement, given a shared probability distribution over the set of possible agreements F. Bargaining ends immediately, so that the first offer effectively determines the outcome. Both the case where 1 goes first and that where 2 goes first is shown here.

centive to make proposals that maximise the likelihood of ratification. When in a position to make a proposal, the negotiator has to come up with an offer that meets the dual requirement that it maximises the probability of ratification and makes the opponent accept it immediately rather than continue to bargain (see proof Iida (1993*b*, 422)).

Formally, each legislator's status quo payoff is private information and the outcome of ratification is known to negotiators only in probabilistic terms. There is a commonly known function $F(z)$ that is the cumulative distribution function of the probability that $s_m \leq z$ (Iida 1993*b*, 414). Let $x^\star$ and $y^\star$ be proposals by 1 and 2 in equilibrium and assume that the (unconstrained) maximiser of $F(z)(1-z)$ is unique. Then, the stationary equilibrium proposals are as described by proposition 3, shown in figure~2.3. 1 accepts any offer y if $F(y)y \geq \delta_1 F(x^\star)x^\star$. 2 accepts any offer x if $F(x)(1-x) \geq \delta_2 F(y^\star)(1-y^\star)$. The equilibrium outcome is that 1's proposal is accepted by player 2 with certainty and ratified with probability $F(x^\star)$. See Iida (1993*b*, 422) for proof.

Proposition 3 tells us that negotiators are trying to maximise the expected value of an agreement $F(z)z$, i.e., the size of the share offered weighted by the probability of the agreement achieving ratification. Proposition 3 states that either negotiator will accept any agreement for which their payoff weighted by the probability of ratification is larger than the discounted value of the counter-offer weighted by the probability that *it* will be ratified.

The case where the home negotiator is uncertain whether a given agreement will pass ratification is similar to that where he has complete information. The only difference is that the value of an agreement must be weighted by the probability of ratification for that agreement. That is, the value of some agreement x is $F(x)x$. When $F(x)$ is close to 1 and x is high, then, clearly, x is an attractive offer. Conversely, if $F(x)$ is close to 0, then the value of x is small whatever x is.

The value of rejecting some offer x and making the counter-offer x' instead is given as $\delta F(x')x'$,

exactly as before, except that the probability of ratification is taken into account. That is, when ratification is uncertain, the value of continuing to bargain is the discounted value of the counter-offer weighted by the probability that it will be ratified. By the same logic as before, i must accept an offer x rather than make the counter-offer x' when $F(x_i)x_i \geq \delta_i F(x'_i)x'_i$.

Further, the decision-problem facing the negotiator is not much different from the complete information case. As before, he must find some agreement that leaves the opponent indifferent between accepting and continuing to bargain. The only difference is that he now must take into account the probability of ratification. So, he is tasked with finding some x for which $F(x)$ is at least as good for the opponent as any counter-offer x' she may credibly make.

Because F is known to the proposer, it should be easy enough to solve the problem of finding the x such that the opponent is indifferent between accepting it and continuing to bargain. That F is common knowledge implies further that the responder can easily verify that the offer in fact satisfies the constraint. It follows that whomever makes the first proposal should be able to figure out what the optimal offer is and make it immediately, to avoid delay. The responder, then, knowing what the proposer knows, can do no better than to accept the offer immediately, also to avoid delay.

When comparing two agreements, it is the weighted value of the agreements that matters. That is, a lower value agreement may win out over a higher value agreement if the difference in probability of ratification is high enough: a negotiator may accept a poor offer simply because it, after all, gives a higher weighted payoff when the risk of non-ratification incurred by insisting on a higher offer is large enough. That is, the specific outcome of bargaining depends on the definition of F.

Iida (1993*b*) illustrates this point with two examples. First is if $F(z) = z$, i.e., the probability of ratification is highly responsive to the bargaining outcome (Iida 1993*b*, 415). Intuitively, when 1 goes first, then he will be able to demand almost all of the pie, simply because 2 has no choice but to accept that conceding most of the benefits is the best way to ensure ratification. Second is the case where $F(z) = \alpha z$ for some $0 < \alpha < 1$ (Iida 1993*b*, 416). Then, the value of going first is reduced by α. When it is low enough, 1 will not be able to make a convincing argument that he needs a greater share to ensure ratification: they both know that ratification is more or less independent of the outcome, so that 2 will be safe in insisting on a share that is in accord with their respective cost of delay.

When ratification is uncertain, the best ratifiable agreement is that which gives the highest payoff and the greatest likelihood of ratification. The outcome is then heavily dependent on the exact

relationship between the bargaining outcome and the probability of ratification. Clearly, the less the outcome matters to the probability of ratification, the harder it will be for 1 to persuade 2 that he should get a larger piece to ensure ratification. Iida (1993*b*, 416) notes that this informational structure implies a risk of involuntary defection and thus is the most harmful to the prospect of cooperation (clearly, cooperation is ensured in both the previous cases so long as we assume that any non-zero benefit is sufficient to entice 2 to engage in cooperation).

2.3 Modelling the role of domestic politics

The description of the role of the executive and his relationship to the legislature within Iida's (1993*b*) model is rigid. Primacy is given to the decision-problem facing individual negotiators and the interaction between them, rather than the interaction between domestic and international levels of politics as such. Specifically, the model concerns itself with the problem of constructing ratifiable agreements under various levels of certainty about the preferences of individual legislators. The outcome of international negotiations is described as turning crucially on the negotiators' knowledge and understanding of the preferences of essential domestic actors, but does not specify a means for the home negotiator to communicate or interact with his constituency. Negotiators have no preferences of their own, seeking, under all circumstances, the best ratifiable agreement. Domestic actors express their desires only by voting on agreements after the fact. Taken literally, essential domestic actors are not actively involved in formulating foreign policy, and the substantive content of international agreements reflect the negotiators' subjective opinions about what their respective constituents want.

Two assumptions about the executive are of particular importance. First, the executive's ordering of available agreements is identical to that of the median legislator. Second, only the vote of the median legislator is allowed to influence whether cooperation occurs, the executive's explicit consent is irrelevant (the executive consents by default to any agreement that the legislature accepts). These assumptions amount to a description of an executive with limited agency. The executive is rewarded according to the size of his country's share of the agreement, but receives nothing if the proposed agreement is rejected by the legislature. Thus, the executive is motivated to seek the best ratifiable agreement, and he should want cooperation so long as there is non-zero benefit to be had. The model denies the possibility that the executive's incentives can have any bearing on whether cooperation is pursued, however. The only power that remains for the executive is the ability to initiate negotiations

if he believes agreement is possible (i.e., that there is a range of mutually acceptable agreements).

The model is restricted in its description of the way domestic actors shape foreign policy, in four important respects. First, the participation of domestic actors is strictly limited to voting on agreements after negotiations have ended; the model makes itself crucially reliant on the assumption that the executive's incentives force him to anticipate the outcome of voting, which makes for a thin description of how the domestic and international levels are related. Second, no effort is made to explicitly capture the influence of political competition or the distribution of political power, which by itself rules out most realistic negotiation scenarios. Third, domestic preferences are described abstractly in terms of the "status quo payoffs" of individual members of the legislature, to which proposed agreements are compared; essential stakeholders, according to the model, support international cooperation so long as they are better off under the proposed agreement than without, which is a rather shallow description of how politics operates. The model is sufficiently abstract to accommodate a wide variety of substantive narratives; by itself, however, the model gives an anaemic description of the relationship between domestic politics and foreign policy outcomes. Fourth, the way the model is setup greatly restricts the kinds of institutional arrangements that it can reasonably be said to approximate. The model does not afford the executive independent agency and there is no competition. A realistic interpretation of these assumptions is that the executive represents a majority government in a parliamentary system. Under a majority parliamentary government, it is reasonable to say that the prime minister's preferences are identical to the legislative block that he represents and that competition is made irrelevant by government holding a majority of the legislative seats (assuming sufficient discipline). It is difficult to see how the model could be said to approximate a presidential system, however, or a minority parliamentary government.

In summary, the model has two features of particular interest. First, that the substantive content of international agreements are shaped by domestic preferences to the extent that the foreign state's negotiator is sufficiently convinced that the home legislators are willing to unilaterally terminate negotiations unless their demands are satisfied. Second, that involuntary defection comes from the negotiators' inability to accurately assess the overall tendency among the legislators to support international cooperation on the issue under negotiation. This may occur if the issue is very hotly contested with no clear majority block on either side, as might be the case for issues that trigger complete political saturation (the model supposes full mobilisation of constituents). Then there is a non-zero probability

that the opposing side may win out, despite the negotiators' best efforts, forcing the home negotiator to renege on what he has promised. Overall, the model's description of how domestic politics operates and the incentives to cooperate lacks expressive power and is amenable to a restricted range of problems.

2.4 Discussion

2.4.1 What it means for the negotiator to be uncertain

The model with domestic uncertainty assumes that country 1's negotiator is uncertain about the probability that a given agreement will pass ratification, but it is not clear whether this is reasonable given the remaining assumptions that are made about the negotiator's incentives and the context of negotiations. The negotiator is said to be uncertain when negotiating over an issue that involves a great number of decision-makers and "[. . .] their preferences are unobservable" (Iida 1993*b*, 413). That the preferences of a legislator are not directly observable is true in the trivial sense that they exist within the mind of the legislator. The negotiator, strictly speaking, can not know the true preferences of individual legislators until after they have cast their vote (supposing the vote is public and that there is no strategic voting). What is at question is the *extent* to which the home negotiator may be uncertain about the probability of ratification.

Although the thoughts of a legislator are not directly observable, their actions are presumably guided by their thoughts, which means that their preferences can be (partially) inferred from public speeches and prior voting. Iida (1993*b*) explicitly disallows strategic voting on the part of legislators; we may interpret this to mean that legislators (as conceived by the model) also speak honestly and that their voting record is a reliable indicator of preferences over time. Iida (1993*b*) explicitly states that there is no abstention during voting, which must be taken to mean complete activation of constituents. Complete activation of constituents implies that the issue is heavily politicised and involving a great number of stakeholders (as Iida (1993*b*) states). With full mobilisation, it appears that there should be considerable amount of debate over what the country's position should be, during, and after international negotiations take place, giving the home negotiator ample opportunity to surmise which agreements are likely to garner sufficient support to achieve ratification, simply by inference from public speech (assuming some level of honesty).

The preferences of legislatures may further be inferred from the interests they are elected to represent. Inference can be based on, say, the economic activities that take place in the legislator's district, which is how Fisher (2006) interprets the voting patterns of legislators from coal producing states in the U.S. (Fisher 2006) (note that there are some problems with this particular example, which I shall discuss later). Iida (1993*b*) explicitly disallows the possibility of an interim election to the legislature between the time negotiations begin and voting takes place; what the home negotiator knows about the preferences of his constituency before negotiations begin will be true afterwards as well (in the sense that the legislators will represent the same interests and thus have the same underlying preferences; people can, of course, change their minds).

The ability of the home negotiator to assess the general tendency among domestic legislators is also dependent on the type of government that he represents. An executive that represents a majority government in a unicameral parliamentary system, for example, cannot reasonably be said to be "uncertain" about the probability of ratification, assuming party discipline.

For the executive to be uncertain about the preferences of his constituents other than in the most trivial sense, it appears that two conditions have to be simultaneously satisfied. First, the issue under negotiation has to be very hotly contested, so that no stable and sufficiently large majority exists. Second, the executive has to represent a government that is not a majority government in a unicameral parliamentary system or his party must have less than perfect party discipline (the degree of party discipline necessary to ensure "certainty" depends on the issue; the greater contention implies a greater dependence on discipline). The model appears to be inconsistent with this latter condition, however. The problem is that the model seems to imply that the executive in fact does represent a majority government, in that he has no independent agency and that there is no competition among legislators, which we may take to mean that there is an established majority block of legislators in place (the minority may disagree, but their preferences are strictly not relevant so that we can assume harmony without loss of validity).

2.4.2 What it means to agree about the likelihood of ratification

What does it mean that 1 and 2 "agree" about F? Technically, F covers all values in $[0, 1]$. Two F are strictly different if their point values deviate, but it makes little sense to differentiate $F_1(z) = .75$ from $F_2(z) = .80$, say, because both appear to indicate practical certainty of ratification for z. Practical

certainty implies either $F(z) \approx 1$ or $F(z) \approx 0$. Intuitively, we want to say that if $F_1(z) \approx 1$ and $F_2(z) \approx 1$ for some z, then 1 and 2 practically agree that z is likely to be ratified. Generally, that 1 and 2 "agree" about F could be stated as $F_1(z) \approx F_2(z)$ for all z, indicating that 1 and 2 have the same belief about practical likelihood of ratification for all agreements they consider.

While it is reasonable to say that .75 and .80 both indicate a belief that ratification of z is likely, a general statement about what it means for 1 and 2 to agree requires that we make explicit the point of practical certainty. The point of practical certainty is the threshold at which the negotiator will say that ratification is "likely", reflecting the negotiator's attitude to risk (e.g., a risk averse negotiator will require a higher probability before accepting practical certainty). If we are to allow negotiators to have their own F it also makes sense to allow them to have different attitudes to risk.

Formally, the general notion of practical agreement about the likelihood of ratification can be stated as follows. Let p_i be the point of practical certainty in the sense that $F_i(z) \geq p_i$ causes negotiator i to set $F_i(z) \approx 1$. Then, $F_1 \approx F_2$ if $F_1(z) \geq p_1 = F_2(z) \geq p_2$ for all z.

2.4.3 The plausibility of a commonly held F

The assumption of a commonly held F is not universally plausible because it makes demands on the negotiators that are not reasonable. The assumption that F is common knowledge means that negotiators 1 and 2 know that they agree about the probability of ratification for agreements they consider. Now, executives that only gain by achieving ratification have an incentive to harmonise their beliefs about F, but to do so requires equal access to and ability to process information about the preferences of the median legislator. That is, to accept a common F implies that proximity to the legislature makes no difference and that there is no variation in skill among the negotiators.

Assuming symmetry here is counter to intuition, which is that the home negotiator should have a clear advantage both in terms of access to and in terms of assessing the general tendency among his own constituents. If we accept that the model approximates a majority government in a parliamentary system, then the home negotiator is certainly in a privileged position, it may in fact be hard to accept that he is uncertain at all, as argued previously. Eliding this potential problem, the home negotiator is in a much better position to gauge the range of agreements that his constituents will accept, by the methods of inference discussed previously, simply by virtue of proximity. Legislators have reason to make their preferences known to the chief negotiator/executive before negotiations and are in a

much better position to communicate with their own negotiator than with the foreign negotiator. The executive, in turn, has reason to be well informed about what the legislature will accept (his payoff, after all, depends on it), and is in a much better position to do so than his foreign counter-part is. Even if we grant that it is reasonable that the home negotiator is meaningfully uncertain about the outcome of ratification, it is not reasonable to assume that he is equally uncertain as his foreign counter part. The negotiator may see the preferences of legislators through a glass, darkly, but the home negotiator is closer and should see more clearly than his foreign colleague who has to observe from afar.

Although intuition suggests that F should differ for the negotiators involved, it is important to not overstate the case. Given the formal definition of a common F the difference in proximity and skill has to be great enough that $F_1 \approx F_2$ no longer holds, which may be harder to justify than the previous paragraph made it out to be. Most importantly, a difference in proximity alone is not necessarily enough to yield diverging F. Accepting that there is no auxiliary benefit to negotiation implies that states have a natural incentive to compensate for differences in skill and access to information. The greater the skill of the negotiator the less proximity matters, supposing diplomatic skill includes the ability to acquire and process information about another state's domestic affairs (the model is supposed to account for cases where deliberations among legislators are public, and so available also to a foreign observer). A state is likely to send its most skilled negotiators and spend the resources necessary to acquire relevant information if it considers the issue to be important. The negotiator, in turn, is likely to use whatever resources he is given to the best of his abilities Still, assuming a shared F requires some assessment of whether it is reasonable to say that proximity and skill makes no significant difference in the particular case under analysis.

Chapter 3

The distribution of domestic political power

3.1 Introduction

At the most abstract level, policy may be conceived as an allocation of some fixed pool of resources. Those who stand to gain or lose by the choice of policy on some issue are the stakeholders, each having their own preferred policy in mind. An allocation may be taken literally as the division of wealth or, more generally, of some underlying benefit that each stakeholder derives from the implementation of a particular policy option — say, under a particular version of a law for regulating the use and ownership of maritime resources. The purpose of politics is to determine policy and so impose one's preferred allocation on others. A conflict of interest between two stakeholders occurs when there is no single policy that gives both what they would get under their own most preferred policy option. Naturally, the more one stands to lose under the policy most preferred by the other, compared to his own, the greater the intensity of conflict between them. As the number of stakeholders increases, so does the number of possible allocations, which implies a greater likelihood of incompatible policy preferences and conflicting interests. The costs and benefits of participating in international cooperation are typically unevenly dispersed among domestic stakeholders, producing relative winners and losers regardless of the absolute level of gains. Even supposing that cooperation imposes no direct costs, there is still the question of how the gains should be distributed. That is to say that most issues will involve some level of conflict of interest: conflict is the natural order of things. A general understanding of foreign policy outcomes, then, must include some way of accounting for the way domestic stakeholders resolve their differences.

Mo (1994) analyses whether the executive benefits from greater domestic constraints depending on the distribution of political power in negotiations over issues characterised by domestic conflict of interest. This chapter is concerned with the modelling account of how political power works, as well as the general relationship between domestic politics and foreign policy outcomes implied by the model. The chapter proceeds by presenting the technical details of the assumptions, general structure, solution concept, and predicted outcomes of the model. With these preliminaries in place, the general implications for the relationship between domestic and international politics are discussed. The final section concerns specifically the notion of political power defined by Mo (1994), discussing the origins of power, the general notion of veto, and comparing the definition of political power which includes the possibility of veto with the approach that equates power with satisfaction. Two points are made:

1. the description of the role of domestic politics in Mo (1994) is so high-level that it is difficult to relate to actual political systems and so lacks the ability to express institutional variation;

2. Mo (1994) shows that a formal conception of power has to include a general-purpose veto mechanism (for both legal and de facto veto); without veto, power is impotent.

3.2 Formal description

3.2.1 Modelling domestic political competition

Domestic political competition is typically modelled as an instance of legislative bargaining (Baron and Ferejohn 1989) with majority rule and no amendments, as is done by Iida (1993*b*); Mo (1994); Tarar (2005). Briefly, n legislators bargain to distribute a unit of benefit among their respective districts. At the beginning of each session, a legislator is picked at random and asked to make a proposal, an n-way allocation of benefits, while the remaining legislators are asked to vote on that proposal. Bargaining continues either for a fixed number of sessions, or until a majority is formed. No benefits are allocated if the legislature adjourns without having formed a majority (Baron and Ferejohn 1989, 1183). The closed rule version of legislative bargaining disallows amendments to proposals during a session, so that legislators make their decision to support or not support a proposal based only on their expected value of continuing to bargain. The justification for disallowing amendments is that it removes the possibility of delay and that it is the rule legislators *should* choose, although in practice,

of course, open rules are common (Baron and Ferejohn 1989, 1198–1199). Under majority rule, a proposal to distribute benefits needs only include a majority of the legislators (i.e., the minimally winning coalition necessary to pass the vote), leaving the minority with nothing. The motivation to support a proposal rather than continue to bargain derives, then, from the fear of being excluded from proposals made in future sessions (Baron and Ferejohn 1989, 1185). Mo (1994) reflects the general logic of the model, but the structure includes both the domestic and the international level, as described in the following.

3.2.2 Players and preferences

There are two countries: foreign (F) and domestic (D). F represents a single group[1] and is a unitary actor. D consists of three competing groups: D_1, D_2, and D_3. By convention, D_1 is D's negotiator, D_2 and D_3 are her constituents. An agreement enters into force if supported by F and a majority of groups in D (Mo 1994, 406).

The object of bargaining is the division of benefits from cooperation. An agreement takes the form of an allocation $x = (x_1, x_2, x_3, x_f)$, subject to $x_i \geq 0, i = 1, 2, 3, f$ and $\sum x_i \leq 1$, where x_i is group i's share of benefits and x_f is F's share. Each group is concerned only with the size of its own share of the benefits and prefers more to less. This is formalised by the assumption that group i prefers x to x' if $x_i > x'_i$ (Mo 1994, 406–407).

Groups care about when agreement is reached. Let $t \in T = \{0, 1, 2, \ldots\}$ represent time periods and let $u_i = \delta_i^t x_i$ be the utility function of group i. A greater δ_i implies that future benefits are valued higher, and reflects i's ability to hold out for a better offer. This is the basis of group i's preference based power (Mo 1994, 407).

Political power is represented by a triple (δ_i, p_i, v_i), where v_i indicates whether i has veto power and p_i is the probability that i will be assigned agenda setting power, reflecting the idea that no group has monopoly over agenda setting power (though D_1 is designated as the executive by convention). For F this is $(\delta_f, 1, 1)$. Mo (1994, 409) makes the simplifying assumption that D_1 and D_2 have identical political power[2] and that D_3 is the only group that can have veto power[3].

[1] The generic term "group" is used for any relevant domestic actor, which may refer to interest groups, politicians, political institutions, and so on. See Mo (1994, 407) for uses and interpretations.

[2] $(\delta_1, p_1, v_1) = (\delta_2, p_2, v_2) = (\delta, \frac{1-p}{2}, 0)$ (Mo 1994, 409).

[3] In that case, $(\delta_3, p, v_3) = (\delta_3, p, 1)$ (Mo 1994, 409).

3.2.3 Extensive form

International negotiations are described as an instance of bilateral bargaining by alternating offers, with the caveats that each domestic group in D is a participant in bargaining and that outcomes are contingent on ratification by majority rule with or without veto (Mo 1994, 408). To summarise, bilateral negotiation over some issue with domestic conflict of interest in one of the countries is described in extensive form as follows.

F makes an initial proposal at $t = 0$ which is then voted over by D (who goes first is not important, see Mo (1994, 414)). There are two ratification processes depending on the presence of veto power. If no group has veto power, then F is free to specify any two groups as his coalition. If there is a group with veto power, then F has to include that group in his coalition. If the proposal is rejected, then D makes a counter-proposal at $t = 1$ (Mo 1994, 408).

D formulates a national proposal by domestic political competition over agenda-setting power. Agenda-setting power means that a group can make its own proposal the national proposal. No group has monopoly over agenda-setting power. This is modelled by randomly assigning agenda-setting power to one of the groups, with p_i denoting the probability that i receives agenda-setting power. Once agenda-setting power has been assigned to D_i, it makes a proposal α_i which becomes D's national proposal (Mo 1994, 408).

D_i seeks the support of F and at least one other domestic group. If the national proposal is rejected, then play proceeds to the next round where F now makes a proposal (Mo 1994, 409). This repeats until agreement is reached. Utilities are discounted each time a proposal is defeated.

3.2.4 Solution concept

All subgames after rejection by D are structurally identical (have identical strategy sets and possible sequence of moves). This allows the solution to be stated naturally as a stationary Nash equilibrium: at every point where a group chooses an action, the strategy from that point is the best response to the other groups (NE) and groups use the same strategies for structurally identical subgames ("stationarity") (Mo 1994, 410). For simplicity, attention is restricted to symmetric equilibria where D_1 and D_2 have the same discount rate. The implication is that purchasing the support of D_1 or D_2 is equally expensive (Mo 1994, 410).

A strategy for a group has two components: the proposal it makes when it has the chance to

propose and whether to support proposals made by any of the other groups or of F. F always makes the same proposal because every subgame after the initial proposal are structurally identical (Mo 1994, 410).

Groups decide to support or oppose proposals made by another group based on their "continuation value". This is the (discounted) value the group expects to get from rejecting the proposal and continuing the bargaining process. During ratification, a group will compare the benefits it gets from the agreement proposed by F to what it could hope to extract from the others in exchange for its support in the domestic bargaining that would follow from rejecting F's offer. During domestic bargaining, a group will compare a proposal by another group to what it would get from the agreement proposed by F (Mo 1994, 410).

Let F's proposal be $\alpha_f \in X$. Mo (1994, 410) states that, under complete information, α_f is accepted by D and becomes the equilibrium policy outcome (Mo 1994, 410). The reason is intuitive: F uses his knowledge of the structure and rules of the game to figure out the allocation of benefits he has to propose to make it unprofitable for at least two domestic groups to continue bargaining. More precisely, F has to make an offer an that makes the members of his chosen coalition indifferent between accepting the offer and continuing to bargain. Knowing that F's offer never changes in equilibrium, the domestic groups that F has included in his policy compromise can do nothing better than to support the initial proposal. The group not included in the coalition has no influence on the outcome.

3.2.5 Outcomes

There are two general classes of outcomes predicted by the model: outcomes without veto power and outcomes with veto power. The outcomes differ mainly in the coalition patterns that occur and in the implications of increased constraints on the negotiator for his bargaining position. There are two key facts: that a proposer gains by choosing the cheapest available coalition and that non-coalition members receive nothing (Mo 1994, 410).

Outcomes without veto power are characterised by the fact that the proposer, F or D_1 depending on the phase of negotiation, may choose coalition members freely. A proposer maximises his own reward by specifying the "cheapest" coalition available. The less a group stands to gain by continuing to bargain, the less it may credibly demand in exchange for its support, which makes it more attractive

for the proposer. F seeks the support of the two weakest groups, whereas a domestic proposer seeks the support of F and the weakest other group (Mo 1994, 410).

Outcomes in the presence of veto power are characterised by the fact that all coalitions must include the domestic group that has veto power. Any proposed coalition has to include the veto group, regardless of its demands. If the veto group makes demands that are too high, cooperation is impossible. When the price of the veto group increases, so does the constraints upon the negotiator, but he does not benefit from the increase itself. In all symmetric equilibria, $D_{3's}$ price has no bearing on F's choice of coalition and an increase will not improve $D_{1's}$ chances of being selected by F. Because D_1 has no choice but to purchase $D_{3's}$ support, an increase in $D_{3's}$ price weakens $D_{1's}$ position and reduces the value of being a coalition member for D_1 (see Mo (1994, 413) for a more detailed discussion). When D_1 is the second most expensive group it is always excluded from F's coalition and thus gains nothing from cooperation.

3.3 Modelling the role of domestic politics

The model describes a pluralistic system where no single entity can dictate policy, requiring instead that interested parties compete to form minimally viable legislative majorities. All policy outcomes are conceived as coalition formation. A coalition is specified by allocating one unit of benefits, representing the gross gains from cooperation. Under majority rule, a proposer of a particular policy must attract a certain number of votes but only has to allocate benefits to would-be supporters, giving the minority nothing. The coalition patterns that emerge, and thus the substantive content of policy, depends on the structure and distribution of political power. With increasing preference-based power comes the ability to make greater demands to join a coalition, but when there is no veto the proposer will simply opt to exclude the most expensive domestic groups and form a coalition of the weak, to retain as much of the gains as possible for himself. Groups that are powerful by virtue of being satisfied with the status quo have little interest in international cooperation, but will simply be excluded from any coalition forming process. With veto, however, the proposer is forced to include the veto holders regardless of their price. Powerful actors with veto will be included in any coalition and will gain from it unless their demands are so great that cooperation is made impossible. The executive, when making a proposal, will be rewarded according to what remains of the gross share of the gains from cooperation after he has made the necessary concessions to achieve ratification. Having agenda-

setting power, of course, means that the executive is guaranteed to be part of the coalition that results from the domestic effort to define a national proposal. If there is veto and the veto holder makes large demands, however, the executive's share of the gains are reduced and so diminishing the value of international cooperation to the executive.

The constitutional setup implied by the model is unclear. Procedurally, there is no clear separation between international and domestic stages of negotiation, because any allocation proposed at either level implicitly refers to players at both levels simultaneously. When making a national proposal, the executive has to specify the share going to the foreign nation as well as his domestic constituents, so that he is making a coalition with himself, one other domestic group and the foreign country. Likewise, when the foreign state is making a proposal, it specifies a coalition including itself and two domestic groups in the home country. The model appears to imply that domestic groups are directly involved in international negotiations, in the sense that offers from the foreign country include a specification of how the gains are to be distributed among the domestic groups (Tarar 2005, 392), thus forming a "coalition" with the domestic groups that are afforded a share in the benefits. This may be taken to capture the notion that most agreements imply costs and benefits that are unevenly distributed among domestic stakeholders, and that proposals made by the foreign state are no different. It is not that the domestic groups are actually involved in negotiations with the foreign state, but because an agreement requires majority support among the domestic groups to enter into force it is as if they were involved — F uses his knowledge of the game to identify what he has to give to each to secure their support without having to negotiate with them directly (complete information is assumed, after all).

Although it is not hard to justify what appears at first to be diffusion of constitutional boundaries, the role of the executive vis-a-vis the other groups is still quite vague. Because the game specifies that agenda setting power (and thus the role of executive) is assigned only during domestic bargaining and because ratification of proposals by F can be achieved by any two groups in D, Tarar (2005, 391) argues that the game has to be interpreted as saying that the home country *has no executive* at the time when the foreign state makes its offer. Another interpretation is that the executive does not have any other choice than to accept whatever the majority decides, but this interpretation only works if we suppose that D goes first and that the group that is designated as the executive continues to be the executive in the following round (where F makes its proposal). That the executive in the model has to induce ratification by formulating a policy compromise that a majority of relevant political

actors support is an intuitive mechanism for describing how domestic groups leverage their power to influence foreign policy, but the way the model is setup makes it difficult to see how it could account for institutional variation or indeed what it tells us about the influence of institutional variables, which is addressed by Tarar (2005).

3.4 Discussion

3.4.1 The origin of power

Preference based power is relevant in all phases, as it represents the ability to reject an offer to join a coalition and continue to bargain, but the model does not explain where it comes from. Groups set their price by how well they expect to do if play passes to the next phase. How well a group expects to do in the next phase derives from their power level in the current phase. Because the game is infinite, there is no "no-agreement" outcome[4] Clearly, it is necessary to set some initial value for each group's power for this to work, which means that a group's initial power level is arbitrary from the point of view of the model. The model explains how preference based power works, but not its origin. A finite horizon version of the game makes this more intuitive. When the game is finite, there exists a round of play (i.e., the final round) where failure to form a coalition results in a no-agreement outcome. The motivation to join a coalition comes from the possibility of being rejected from coalitions specified in future rounds. The role of proposer is assigned at random at the beginning of each session, which means that a group that is included in the current proposer's coalition has no guarantee that it will be included in the coalitions proposed by future proposers. Because minority groups get nothing, exclusion is identical to the no-agreement outcome. Thus, a group's willingness to reject an offer to join a coalition is fundamentally tied to its valuation of the no-agreement outcome. This is yet another way of saying that domestic actors who are satisfied with the status quo are harder to entice into pursuing international cooperation. The finite horizon version of this argument is easier to understand in practical terms, but does no better in terms of explaining what makes a group better or worse off under no-agreement. This is addressed by Tarar (2005), described in the next chapter.

[4]Which is why stationarity is necessary to achieve stability, as explained by Baron and Ferejohn (1989, 1190).

3.4.2 The meaning of veto

The formal notion of veto power (Mo 1994) is generic and can accommodate legal veto power as well as de facto veto power. Generally, a "veto group" has to be included in any policy compromise on some issue and the policy proposer has to make whatever concession is necessary to gain the veto group's consent (if the veto group makes demands that are impossible to satisfy, then international agreement is impossible). There is no technical reason that the groups in Mo's (Mo) model of political competition have to be of the same type, though the group with agenda setting power serves the function of the government in that it is the country's international representative. The remaining two groups can be thought of, say, as a political party and a powerful non-political actor such as a large national labour union. This generality could be indirectly achieved for a model that only included legislative blocks, of course, by supposing that each block is a shallow proxy for some group interest (e.g., labour). To illustrate, suppose that we want to model negotiations over an issue that is relevant to vital labour interests. Suppose the existence of a coordinated national labour union with the ability and incentive to threaten and carry out a retaliatory strike which would impose significant economic and political costs on the government. Although not legally included in decision-making, the labour union in this example can impose costs that are so large that the government cannot afford to ignore their demands when considering policy options. To approximate this situation, let D_1 be the government, D_2 be a political party, and D_3 be the national labour union. The political power of a group is represented by a 3-tuple indicating agenda setting power (binary), preference-based power (discount factor), and veto power (binary). For D_1 the political power is $(\delta_1, 1, 0)$, for D_2 it is $(\delta_2, 0, 0)$, and for D_3 it is $(\delta_3, 0, 1)$. Thus, groups with de facto veto power can be included in the formal narrative of political competition. It is not clear that Mo (1994) can support this notion of veto without modification, however. The problem is the lack of executive agency during the ratification phase, which suggests that there is no-one to reject proposals by F that would trigger retaliatory measures by the de facto veto group. That is, it is not clear, substantively, why F would be compelled to include a de facto veto group in its coalition given the way the model is setup. Still, the concept of de facto veto power remains useful, and there is no technical impediment to using it. For practical application to de facto veto groups, it is necessary, of course, to justify why the group has to be included in any policy compromise and on which issues, whereas for a legally instituted veto holder the source and extent of the veto is self-evident.

3.4.3 The meaning power in the absence of veto

A useful conception of political power has to imply the ability to influence policy outcomes, but this conception is not well approximated by modelling power as satisfaction with the status quo. The discount factor captures the value to a legislator of continuing the current status affair, and so represents his willingness to hold out for a better deal. This is the definition of political power implied by Iida (1993*b*) and the basis for what Mo (1994) calls preference-based power. A dissatisfied legislator is eager to accept agreements, a satisfied legislator is harder to convince. There is a clear difference between being merely satisfied with the way things are, and thus demanding a greater price for supporting change, and the ability to actively influence policy. The latter is what is most commonly associated with political power.

Power that derives solely from satisfaction is substantively uninteresting. The preferences of groups that are "content" but have no means of resisting change are irrelevant for understanding policy outcomes: there is no way that their desires could affect the domestic political process. The model confirms this intuition by predicting that (the most) content groups are excluded from decision-making; in the absence of veto, all policy is championed by coalitions of the weak. The preferences of a legislator that is excluded from decision-making are of no consequence. A conception of power that is derived from satisfaction alone is of little substantive interest. Defining political power to include the possibility of veto approximates the notion of power as the ability to produce favourable policy. A group that is assigned veto power must be included in any policy compromise, regardless of the price it demands. That is, a veto group's preferences are always of consequence.

Although the abstraction of veto power provided by Mo (1994) is useful, the way the model is constructed restricts real-world applications. The model predicts that merely satisfied groups are excluded from policy compromises if their demands are too high. Groups that are powerful in the sense that they place a high premium on the current state of affairs but lack the power to resist change, can have no influence on whether international cooperation occurs. As argued, the common notion of power has to include the ability to influence policy, and most real-world applications of the model will end up having to include a veto player. According to the model, however, executives never benefit from increasing constraints deriving from demands made by a veto group because they have no choice but to pay the price demanded, thus executives become less and less interested in pursuing international cooperation as their constraints increase. As discussed by Tarar (2005), the effect of constraints on

the executive can vary depending on institutional factors, which means that the model of Mo (1994), as stated, is restricted in the range of possibilities it can express.

Chapter 4

The political system

4.1 Introduction

Politics is about the imposition of one's preferred policy in a peaceful manner; power is mediated by institutions: voting rules, constitutional provisions for the separation of power, and so on. To make useful statements about the relationship between domestic politics and foreign policy it is desirable to go beyond the abstract workings of power and attempt to capture the way it is mediated by the institutional context. Tarar (2005) examines how the executive's position in the two-level game varies with "breadth of the executive's constituency and the source of his ratification constraint" (Tarar 2005, 393). Technically, this amounts to varying the extent to which the executive's payoff from international cooperation depends on the specific domestic allocation of gains that he has to make to induce ratification. The model has two attractive features: it allows the executive and members of the legislature to be beholden to particular domestic groups and it describes the no-agreement outcome as the result of domestic bargaining rather than as a given. The chapter proceeds by presenting the details of the modelling account of how institutional aspects of the political system affect the incentives of the executive, before discussing the role of domestic politics implied by the model. The final section suggests a possible informal extension to include veto groups, akin to Mo (1994), and explores the expressive power of allowing the no-agreement outcome to vary. Two points are made:

1. the modelling account of domestic politics is the most realistic of the three models considered here in that it can account for conflict and harmony in addition to differentiating between broad classes of political systems; and

2. the most powerful abstraction provided is the variable no-agreement outcome, which enables modelling variations in the incentive to cooperate across broad classes of political systems.

4.2 Formal description

Formally, the general structure of Tarar's (2005) model is as follows (Tarar 2005, 387–388). There are three possible phases: international bargaining, ratification, and no-agreement policy formulation. First, two executives bargain over the division of gains from cooperation by alternating offers, as per Rubinstein (1982), resulting in a partition $x = (x_1, x_2)$, summing to 1. Once an agreement is accepted, it is passed on for ratification by the two states' respective legislatures. Executives negotiate under complete information and face procedurally identical ratification constraints.

Then follows the ratification phase (Tarar 2005, 390). Executive i specifies an n_i-way allocation z_i of x_i. Legislator j_n receives $\delta_{iL}^t z_i$ where δ_{iL}^t is the common discount factor for legislators in country i. The value of x_i to legislator j depends only on his share under the proposed agreement, reflecting the assumption that legislators are motivated primarily by political gain. In model 1, the executive is rewarded independently of the specific allocation of gains from cooperation he makes to induce ratification. This is formalised by setting his payoff to $\delta_i^t x_i$, the discounted value of the share of the gains that executive i is able to secure for his country.

In the case of ratification failure, the executive receives a 0 payoff, while legislators bargain to determine the no-agreement outcome, conceived as the division of the aggregate value of not pursuing international cooperation. The no-agreement policy bargaining is setup similarly to the no amendments version of legislative bargaining (Baron and Ferejohn 1989). Each session begins by randomly selecting a legislator who is tasked with proposing a division of the no-agreement pie. A proposal to divide a pie is an n-way allocation $z = (z_1, z_2, \ldots, z_n)$, where z_i is legislator i's share of the pie. There is no abstention and proposals are accepted if supported by a simple majority. If no majority is achieved, then the game moves on to the next period. All legislators are assumed to have a common discount factor, giving the payoff $\delta^t z_i$ for some proposal z. In the finite horizon version of the game, there is a final round where the alternative to accepting a proposal is that bargaining ends without policy having been determined, giving 0 to all legislators. Then, the players can look ahead to the final round and deduce what the optimal strategy is. In the infinite version of the game, legislators use the notion of a stationary equilibrium to figure out the best course of action. The predicted out-

come is that the first chosen legislator specifies a minimally winning coalition where he offers each member their discounted continuation value (expected value of continuing to bargain). Each member of the coalition gets δ/n, whereas the proposer keeps the remainder $1 - ((n-1)/2)(\delta/n)$, and non-members get 0. Legislators bargain over a pie sized $\Delta_i \geq 0$ where Δ_i is the over-all value of the no-agreement outcome. The pie is variable to reflect the idea that i may be worse off on its own than with cooperation. Without taking differences in political power into account, the expected value of the no-agreement outcome for each legislator is Δ_i/n_i Tarar (2005, 390).

Legislators set their price for supporting the executive's proposal based on the expected outcome of the no-agreement policy bargaining game. To induce ratification, executive i offers the members of his coalition a share corresponding to their discounted payoff in the domestic bargaining game, $(\delta_{iL}\Delta_i)/n_i$. This is only possible if x_i is large enough (so that there is enough to go around). When n is odd, it is necessary that $x_i \geq ((\delta_{iL}\Delta_i)/n_i)((n_i+1)/2)$. Then $C_i = ((\delta_{iL}\Delta_i)/n_i)((n_i+1)/2)$ is the effective domestic ratification constraint facing i. The better legislators expect to do in the no-agreement outcome, the higher their price and the greater the constraints on the executive. An international agreement (x_1, x_2) is implemented if $x_1 \geq C_1$ and $x_2 \geq C_2$ (Tarar 2005, 390); trivially, when true both 1 and 2 can and will induce ratification.

Model 2 is as model 1 except that the executive, like the legislators, has a local constituency and his reward depends on how much of the gains that he can allocate to his core constituency. Now, the payoff for executive i is $\delta_i^t(x_i - C_i)$ where C_i is the share demanded by legislators from constituents outside the executive's core (Tarar 2005, 391). Assuming that the executive only gains by satisfying the preferences of his own group, then any concession made to outside legislators is a loss to himself. That is, when the executive needs support from outside legislators, he may no longer induce ratification "for free".

After paying the cost of ratification, executives are left with $1 - C_1 - C_2$ to divide among themselves. The predicted outcome is that i gets C_i plus the share obtained by bargaining over $1 - C_1 - C_2$ by alternating offers. For this to occur it is necessary that $C_1 + C_2 < 1$, otherwise executives have nothing to gain from cooperation and no reason to engage in negotiation (Tarar 2005, 391).

The substantive interpretation provided by Tarar (2005) is that the situation described here mirrors that of a minority or coalition parliamentary government with members beholden to different/opposing constituencies. Then C_i is the share demanded by coalition partners not from the prime

minister's party (Tarar 2005, 393). It also describes the situation where a single-party government faces a ratification requirement that is higher than the requirement for forming government (Tarar 2005, 393) (e.g., a simple majority government facing a requirement of 2/3 majority).

4.3 Modelling the role of domestic politics

The model implies a political system comprising self-interested legislators and an independent executive whose payoff from international cooperation varies depending on the institutional position of the executive vis-a-vis the legislature. International cooperation occurs only when both the legislators and the executive gain from it. The preferences of domestic actors influence policy via the domestic allocation of the gains from cooperation that the executive makes to induce ratification. If a proposal by the executive fails ratification, then the legislature determine the no-agreement policy by legislative bargaining. The value a legislator expects to get from legislative bargaining determines how much he will demand in return for supporting the executive's proposal. By making the no-agreement outcome endogenous, Tarar's (2005) model provides a more realistic description of what individual legislators are comparing the executive's proposal to when deciding whether to support it (as opposed to supposing that it is compared to some future offer made by the foreign state, as is done in Mo's (1994)). The likelihood of cooperation also depends on the executive's payoff. An executive with a local constituency facing a ratification constraint that requires him to seek support from constituents outside his core group gains less from cooperation with increasing demands from non-core constituents, thus weakening his incentive to pursue international cooperation. An executive who is rewarded according to the gross share of gains from cooperation and not the specific domestic allocation of those gains, however, benefits from increasing constraints, up to the point where no agreement is possible, and should be more eager to pursue cooperation (see Tarar (2005, 391)).

The implications of Tarar's (2005) model pertaining to coalition patterns and the distribution of power are superficially similar to the non-veto version of Mo's (1994) model. An executive that has to pay a price for inducing ratification naturally seeks to maximise his own payoff by choosing the cheapest available coalition. Legislators set their price for support based on their expected payoff from domestic bargaining to divide the "value" of no-agreement. Because there is no veto, however, there is nothing that prevents the executive from excluding the most demanding legislators (unless the executive is bound to a local constituency that happens to be made up of the most expensive

legislators). The most powerful parties outside the executive's core group are punished for making high demands by not being including in the ratifying coalition. Further, as the strongest party becomes stronger, all else equal, the remaining parties become weaker in the sense that they can expect to do worse in the no-agreement policy game and so will reduce their price for supporting the executive's proposal (Tarar 2005, 406). The expectation of doing well in no-agreement bargaining is equivalent to being satisfied with the way things are (or will be in the absence of cooperation), and power derived from this is impotent, as discussed in chapter 3.

Structurally, the model differs from that of Mo (1994) in that it is symmetric and that it allows the executive independent agency. Domestic groups do not make agreements directly with the foreign state (as under the literal interpretation of Mo (1994)) and cooperation occurs only with the executive's consent. Substantively, the model is sensitive to differences in institutional make-up, differentiating between presidential systems and majority and minority parliamentary governments. By making the no-agreement outcome endogenous and letting the overall value of no-agreement vary, the model is also able to differentiate between issues that are more or less reliant on international cooperation (such as transboundary pollution). The model is limited in its description of political competition and the distribution of political power compared to Mo (1994), however. There is no way to differentiate between the merely satisfied and the truly powerful; the willingness to hold out for a better deal is grounded in satisfaction with the status quo and the model predicts that domestic actors that are the most willing to hold out may simply price themselves out of decision-making.

This concludes the presentation of the formal two level models. Although each model is designed with different goals in mind, there is a clear line of progression from Iida (1993*b*) to Tarar (2005) in terms of capturing relevant aspects of the relationship between domestic politics and foreign policy. In Iida (1993*b*), domestic actors are represented only in the most abstract sense, their impact being reduced to the status quo payoff of the "median legislator". Mo (1994) extends this notion by incorporating legislative bargaining into the two-level model, although eschewing the possibility of incomplete information. Finally, Tarar (2005) presents a model that adds further expressive power by allowing the executive's position to vary depending on the nature of his ratification constraint. The presentation so far has concerned only the general relationship between domestic and international politics. The next and final chapter shall discuss specific applications to problems of international environmental governance.

4.4 Discussion

4.4.1 The informal logic of veto groups

There is no technical reason that veto groups could not be included in the model, if necessary. By informal extension of the logic of the model, the presence of a veto group would impact the executive depending on the nature of his constituency. A popularly elected president should have an unaltered payoff, so that the coalition pattern is the only variable that changes with the addition of a veto group (which is included regardless of its price). That is, supposing that the veto group's demands are not so great that cooperation is impossible, in which case the executive gets 0. The prime minister, however, is impacted by the presence of a veto group according to whether the veto group is a member of his core constituency. Under traditional left-right allegiances, a right-wing government facing a strong labour union will lose, whereas a left-wing government may gain as a result of increasing demands from the union (again, so long as the demands are not too great).

The informal logic of adding veto groups to the model implies that the alignment of government along the relevant axis also impacts the likelihood of cooperation. Using labour-capital as an example, the model implies that a labour government improves its position in the international bargaining *and* its reward with increasing demands from a strong labour union, up to the point where cooperation is made impossible. A right-wing government improves its position, but is forced to concede policy provisions, such as protections for domestic labour, that are not in line with their core constituency's interests. In the presence of a strong labour union, then, a labour government may in fact be more interested in international cooperation than a right-wing counter-part (assuming traditional left-right allegiances). The effect, of course, will vary from issue to issue with the extent that the interests of domestic labour are affected.

4.4.2 A variable no-agreement pie and the threshold for cooperation

Making the no-agreement outcome endogenous and allowing the aggregate value of no-agreement to vary, as done by Tarar (2005), are among the most potent abstractions introduced by the models considered here, and are of particular interest for modelling the role of domestic politics in international environmental negotiations. All legislators are presumed to attach some quantifiable value of no-cooperation. The aggregate value of continuing the current state of affairs, denoted Δ, represents

what legislators will have to bargain over to formulate policy instead of an international agreement (hereafter "no-agreement bargaining"). That Δ may be allowed to vary (between 0 and 1) implies that the state may be worse off than with cooperation or that it may be indifferent to it. Clearly, when Δ is 1, then there is nothing to be gained from cooperation, and little reason to expect that negotiations will take place (i.e., legislators enter into domestic bargaining to determine the no-agreement allocation immediately). The less that a single state can hope to achieve on its own, however, the smaller Δ will be and thus the incentive to cooperate will be higher. The reason is that a small Δ means that legislators will have less to gain from the no-agreement bargaining because the pie that is to be divided is smaller, reducing their price for supporting the executive's proposal to cooperate. A variable Δ makes it possible to formally differentiate between issues based on how "international" they are, i.e., how much each state can hope to achieve alone, with smaller values indicating a weaker ability to address the problem alone.

A relevant example is transboundary pollution, assuming emissions are more or less evenly distributed among the affected states so that no single state can effectively solve its pollution problem by domestic reductions alone. Obviously, the larger a country's share of the overall level of pollution, the greater the impact of emission reduction, but the costs of those reductions will likely be greater as well (provided the correlation between economic activity and emissions). Although there is a cost to be paid to enter into an agreement of joint emission reductions, legislators may be willing to pay that cost provided that Δ is low enough (i.e., that the pollution problem is sufficiently severe). This is, of course, assuming that each legislator's local constituency is equally affected by the pollution — in practice it usually occurs that a minority of districts suffer greatly from pollution whereas the others either gain or are indifferent to it. That is, the size of Δ is the aggregate value of the no-agreement outcome and does not capture the fact that costs/benefits of no-agreement may be severely skewed.

We may conjecture that there exists some critical value α for which the legislature in country i is willing to carry the cost entailed by entering into a legally binding commitment with the other affected states to abate emissions rather than attempt to address the problem by domestic policy alone. A national α is not immediately plausible, however. The distribution of costs and benefits is likely skewed, and it is more reasonable to say that each legislator i has an individual α_i. The size of α_i depends on the local impact of the environmental degradation that results from pollution weighted by the local benefits of the economic activity that contributes to the pollution.

The notion of a variable Δ implies the existence of a critical value α at which the legislature is motivated to accept cooperation rather than attempt to address the issue alone. From this follows a coarse formal statement of the boundary condition on the legislature's willingness to engage in cooperation over environmental issues. The calculus of costs and benefits has to favour cooperation on aggregate, formalised by stating that Δ has to be lower than some α. Unfortunately, the cost and benefits are typically severely skewed, which suggests a modification that introduces a local α_i for legislators $i \in N$ such that the national α is the sum of $\alpha_{i \in N}$. It is unclear, however, what has been gained by this over simply stating that legislators support abatement if the costs of the economic activity that causes it exceeds the benefits generated by that activity. Further, the fact that the economic activity is ongoing makes it likely that any negative impacts are mostly externalised to other districts (otherwise, it would have been shut down already, following Coase (1960)), so that one may expect that a critical cost-benefit ratio never obtains in those districts where harmful economic activity occurs. If this is true for all districts where harmful activity is ongoing, the national α will never be in favour of cooperating to abate, if that implies taxing or otherwise stifling harmful economic activity. This line of reasoning relies on a crucial assumption, however, that the number of districts that capture value from harmful activity exceeds those who suffer negative externalities from that activity. The best possible scenario is one where a minority of districts have harmful economic activity, from which they capture a great deal of benefit (or at least enough to offset any negative effects), whereas the majority of districts suffer the externalities while reaping none of the benefits (more precisely, not reaping enough to offset the cost). Then, it is possible that a majority of legislators at the national level could agree to impose restrictions on the minority of producing districts. In conclusion, this is only yet another way of stating the fundamental problem of negative externalities.

Chapter 5

Discussion

5.1 Introduction

The presentation so far has been largely theoretical and concerned with the general implications of the formal developments of the theory of two-level games. This final chapter concerns how the abstractions that have been discussed previously can be applied to problems of international cooperation over environmental issues. The discussion comprises a specific part, where empirical examples are used to illustrate the use of specific abstractions to environmental policy, and a general part concerning the fundamental question of what motivates the chief negotiator in negotiations over environmental issues.

The discussion of specific abstractions follows two lines of evidence related to U.S. climate policy. Hovi, Sprinz and Bang (2012) presents analysis of interviews with delegates to the Kyoto negotiations about the reasons why the U.S. first signed and then repudiated from the Kyoto protocol. I will use this analysis as a reference-point for discussing the formalisation of involuntary defection presented in chapter 2. Fisher (2006) presents evidence of the relationship between coal dependence and climate policy in the United States. I will use the claims made by Fisher (2006) as a reference-point for discussing how the domestic sources of national climate policy can be expressed in terms of the formalisation of power and the notion of local constituency, discussed in chapters 3 and 4. Three points are made:

1. the formal notion of involuntary defection is not well-suited to explain the U.S. rejection of Kyoto; although the case may at first appear to be a clear-cut case of involuntary defection, it is difficult to reconcile this interpretation with the formal description; and

2. the hypothetical link between coal extraction by state and voting on national climate legislation is neatly expressed in terms of local constituencies, but the evidence that is taken to support this hypothesis may be too weak to justify doing so; and

3. the influence of the coal industry on national policy can be captured by stipulating a coordinated national Coal Union with de facto veto power on CO2 legislation, but I will argue that the credibility of the threat that supports the de facto veto may be weaker than the overall consumption of coal energy suggests.

Having covered the application of specific abstractions to environmental issues, I will turn to addressing a more general assumption, shared by all the formal models considered here, concerning the motivations of the chief negotiator. Hovi, Sprinz and Bang (2012, 130) point out that the assumption that negotiators only consider ratifiable agreements is not borne out by Kyoto. Formally, the mechanism that ensures that the executive only considers ratifiable agreements is the assumption of a zero payoff for failure to ratify. I will argue that disallowing the possibility that the executive may derive some auxiliary benefit from engaging in negotiations excludes many substantively interesting interpretations of real-world negotiation outcomes, e.g., the Kyoto case. From this assumption follows also that a shared belief in a range of mutually beneficial agreements is the only reason why international negotiations occur, precluding other interesting interpretations of the occurrence of negotiations. At worst, the occurrence of negotiations may be taken as evidence that profitable cooperation is possible.

5.2 U.S. repudiation of Kyoto as involuntary defection

To illustrate the application of the formal abstractions for describing the role of information, I will consider the case of U.S. repudiation of the Kyoto protocol as an instance of involuntary defection. I will proceed by briefly describing the Byrd-Hagel resolution and the explanations for the Kyoto outcome presented by Hovi, Sprinz and Bang (2012), which I then discuss first in terms of domestic uncertainty about the preferences of legislators and then using the model of international asymmetry (chapter 2). The following may be considered partly as an explicitly formal variant of the discussion of information in Hovi, Sprinz and Bang (2012, 135).

The so-called Byrd-Hagel resolution specifies conditions for the United States to sign on to “any international agreement on greenhouse gas emissions under the United Nations Framework Conven-

tion on Climate Change" (*Senate Resolution 98. Report No. 105-54* 1997). The Resolution states that the United States "should not be a signatory to any protocol" that exempts developing countries from reductions or that causes "serious harm" to the U.S. economy (the increasing industrial capacity of developing nations is mentioned as the cause for concern over the inclusion of exemption clauses) (*Senate Resolution 98. Report No. 105-54* 1997). The Kyoto protocol violated both of these conditions (see Hovi, Sprinz and Bang (2012, 130)). The Resolution was unanimously passed by the Senate prior to the Kyoto round of talks, sending a clear message about the kind of agreement that domestic legislators in the United States were willing to accept (Hovi, Sprinz and Bang 2012, 133).

Briefly, the explanations considered by Hovi, Sprinz and Bang (2012) are:

1. Kyoto delegates (the Europeans in particular) underestimated or did not understand the implication of Byrd-Hagel;

2. European delegates were not concerned with the effectiveness of Kyoto nor U.S. participation; they may have understood but did not care about the implications of Byrd Hagel.

3. Clinton/Gore signed the Kyoto agreement despite Byrd Hagel for domestic political reasons.

Explanation 1 is naturally expressed as an instance of ineffective communication of constraints (Hovi, Sprinz and Bang 2012, 134); formally we may say that European delegates made an "optimistic conjecture" and chose to punish what they saw as a bluff by driving a hard bargain (see chapter 2). Unlike the outcome in the model of asymmetric information about constraints, however, the Kyoto negotiations ended with failure. If explanation 1 is true, then, this constitutes an example of how lack of information *can* jeopardise cooperation. Many delegates present at the Kyoto negotiations have stated in interviews that Byrd-Hagel was misunderstood or underestimated. Some non-U.S. delegates reportedly believed the President could ensure ratification despite Senate opposition (possibly reflecting a lack of understanding of how the U.S. political system works), whereas other delegates simply saw it as an attempt at feigning constraints (Hovi, Sprinz and Bang 2012, 137–139). Explanations 2 and 3 are not easily reconciled with the two-level model, because they violate the essential assumption that executives only consider ratifiable agreements (Hovi, Sprinz and Bang 2012). Explanation 2 suggests that Europeans considered the Kyoto a stepping stone towards a comprehensive climate regime (perhaps akin to the incremental development of the Montreal protocol): they "looked even *further* ahead than to the ratification phase" (Hovi, Sprinz and Bang 2012, 135). Explanation 3 suggests that the

U.S. had no intention or hope of ratifying Kyoto, but that they nevertheless hoped to gain something from signing the agreement, i.e., that "[t]he US Kyoto delegation acted on the basis of instructions motivated by other considerations than the agreement's attractiveness to the Senate." (Hovi, Sprinz and Bang 2012, 136). Clearly, only explanation 1 is suitable for direct assessment of the formal definition of involuntary defection, and it is this explanation I shall focus on in what follows, though I shall return to explanation 3 later in this chapter.

Generally, the executive of the United States can plausibly be said to be uncertain about the likely outcome of ratification, provided no stable and sufficiently large majority exists (see discussion in chapter 2, see Bang, Hovi and Sprinz (2012, 3) for the United States specifically). The U.S. system is characterised by low party cohesion and party discipline and opportunities for obstructive tactics, which can make it difficult to predict whether the conditions for ratification will obtain (i.e., advice and consent by the Senate and the passage of enabling legislation in the House) (Bang, Hovi and Sprinz 2012, 3–4). To illustrate how to apply the formal notion of domestic uncertainty about the preferences of legislators, let F_{EU} refer to the European delegates' beliefs about the probability of ratification and F_{US} refer to the U.S. delegation. To reiterate, the formal model assumes that both delegations receive 0 for failure to ratify and that F is common knowledge, which is true when $F_{\text{US}} \approx F_{\text{EU}}$ (see chapter 2). Suppose explanation 1 is true and that the European delegates consistently underestimated the constraints implied by Byrd Hagel. The assumption of a shared F requires us to believe that $F_{\text{EU}}(\text{Kyoto}) \approx F_{\text{US}}(\text{Kyoto}) \approx 1$ despite the Byrd-Hagel resolution clearly stating that $F(\text{Kyoto}) = 0$. We must assume that the all delegates were familiar with Byrd-Hagel and at least understood the conditions set forth by the Resolution and that the Kyoto protocol violated both. Unless the U.S. delegation believed that a majority of the Senate could be convinced to change their mind after the Kyoto agreement had been signed, the assumption of practical agreement that Kyoto could be ratified is simply not plausible. The Byrd-Hagel was clear evidence of a stable and sufficiently large majority on the issue of U.S. accent to the Kyoto, so that it makes little sense to say that the President was meaningfully uncertain about the preferences of the median legislator: he should have set $F_{\text{US}}(\text{Kyoto}) \approx 0$. Several U.S. delegates interviewed by Hovi, Sprinz and Bang (2012) claim that Clinton understood very well the implications of Byrd-Hagel, which is reflected in the fact that he failed to present the agreement to the Senate for advice and consent. Though it may be possible that $F_{\text{EU}} \approx 1$, that implies a significant disparity in skill in assessing the implications of Byrd-Hagel; that

is, the difference in skill was such that it violated the condition for a shared F. But allowing for F to diverge brings us no closer to explaining why the U.S. signed on to the agreement. Clearly, attempting to analyse the U.S. refusal to join Kyoto as an instance of involuntary defection caused by domestic uncertainty about constraints is not reasonable — unless both U.S. and European delegations were severely lacking in skill.

If we relax the requirement of a shared F (which is the same as transitioning to a game of international asymmetry), however, there is an interpretation of the outcome that appears to be consistent with explanation 1 and the assumption of a zero payoff for failure to ratify, which is that the European delegates simply bargained too hard. The outcome was a form of involuntary defection in the sense that the European delegates' strategy restricted the win-set of the U.S. delegation to 0, forcing the U.S. to sign on to an agreement it had no hope of ratifying. A constrained negotiator with an empty win-set probably cannot do anything better than sign an agreement even if he knows it will not be ratified. This also explains why Clinton did not seek the Senate's advice and consent after signing Kyoto: he knew it would be pointless (Bang, Hovi and Sprinz 2012, 3). The model does not say anything about the executive withdrawing on his own accord, but we may imagine that doing so could imply a negative payoff and so be worse than signing an agreement that fails to ratify, if for no other reason than that it is more politically convenient to sign an agreement than to outright refuse it. Had European delegates believed in the claims of U.S. domestic constraints, they might have had motivation to offer the U.S. improved terms. Unless the Europeans had been willing to concede the exemption of developing nations from commitments to reduce greenhouse gases, however, it is unlikely that the Senate would have accepted an improved agreement in any case — $F_{\text{US}}(\text{Kyoto}) = 0$ for all variations of Kyoto with those provisions included. The puzzle, then, is why the European delegates did not make an offer they knew the U.S. would ratify, which they should according to the model of international asymmetric information — the Europeans should update their beliefs about the U.S. negotiator's true type following the rejection of the initial proposal. Again, the attempt to formally explain the outcome as involuntary defection appears to break down.

The result of attempting to model the outcome of Kyoto as involuntary defection does not recommend using this abstraction to understand the U.S. position. A simpler explanation may be that there simply was no overlapping win-set. According to the theory of two-level games, however, that belief could not have been in place initially, because then there would be no reason to negotiate. For

the empty win-set explanation to work, we must suppose that the U.S. delegation believed that the Europeans could be persuaded to remove the exemptions for developing nations. This explanation is supported, somewhat, by senator Byrd's statement that he intended for the Resolution to strengthen the U.S. delegation's position (Hovi, Sprinz and Bang 2012, 134), presumably with regard to changing the European position on exemption. In addition, the Europeans must initially have believed that an agreement that ensured U.S. participation was possible despite Byrd-Hagel. Then, there could have existed a shared belief in a mutually beneficial outcome, which simply turned out to be incorrect — as revealed when the Europeans did not adjust their offer in response to U.S. rejection. At that point, the U.S. delegation could do no better than to sign the agreement, despite knowing that it could not be ratified. However, all these assumptions are only made necessary by the single assumption that there are no auxiliary benefits to negotiation, and it is much more attractive to simply relax that assumption than to introduce additional unknowns. The benefit of relaxing the assumption of zero payoff is that it enables addressing explanation 2 and 3 directly, which are intuitively more plausible than having to insist on involuntary defection as the only possible cause for the U.S. negotiation behaviour, a topic I will consider in the final section of this chapter.

5.3 Extending the notion of involuntary defection

Iida (1993*a*) distinguishes between strategic and analytic uncertainty. The notion of involuntary defection captured by F is restricted to strategic uncertainty, whereas environmental problems are characterised by both strategic and analytic uncertainty. Strategic uncertainty refers to variables that influence the bargaining position of negotiators, such as the true nature of their domestic constraints and the willingness to incur the cost of delay. Analytic uncertainty refers to disagreements about the way the world is, and about future states of the world. Putnam (1988) states that involuntary defection occurs when a state reneges on an agreement because of failure to ratify and that the fear of involuntary defection, if present, can jeopardise international cooperation. The size of the win-set, or the true preferences of the median legislator, indicates the bargaining strength of the constrained negotiator and so concerns the strategic aspects of negotiations. Another important source of contention in negotiations over environmental issues, in particular long-term environmental problems, is uncertainty about the future state of the world. This is clearly different from uncertainty about the preferences of the median legislator or disagreement of the true definition of F.

There are at least two major types of analytic uncertainty at work in negotiations that seek to address long-term environmental problems, such as global climate change. First, and most obvious, is the uncertainty in predictions about how the current course will influence the future. That is, will *catastrophic* climate change occur and if so, when? If it occurs, how will the gains and losses be distributed globally? As concerns involuntary defection, these questions are possibly not that important, although it may be that a state fears involuntary defection because climate change turns out to have positive effects for the opponent, so that he will renege on the agreement when the time comes to implement its provisions (see also Hovi, Sprinz and Underdal (2009) for time lag as a general problem for cooperation over long-term environmental problems). More relevant to involuntary defection is uncertainty related to future economic activity. It is well known that emission levels are correlated with economic activity (e.g., Congleton (1992)). Economic activity, in turn, is determined in large part by global (or regional) economic conjunctures that are outside any national government's control.

Hovi, Sprinz and Underdal (2009) argue that states are reluctant to make binding agreements if they are uncertain about their ability to deliver on the terms, with greater time lag between agreement and implementation increasing the uncertainty about the factors that determine the ability to deliver on the terms. This illustrates how international cooperation to address long-term environmental problems can be jeopardised by fear of involuntary defection due to analytic uncertainty about the future state of the world, which is not well approximated by F. A formal description of involuntary defection that also covers analytic uncertainty, then, would be a useful complement to the definition provided by Iida (1993*b*).

5.4 The domestic sources of U.S. climate policy

The overall claim of Fisher (2006) is that the U.S. position on climate change is related to domestic energy needs and production patterns, specifically related to extraction and burning of domestic coal. The argument, briefly, is that the local dependence on coal for energy and the auxiliary economic and social benefits of coal extraction influence the position of state representatives who vote on climate legislation at the national level. Because coal releases relatively more CO2 than oil and natural gas, regulation of CO2 emissions would be particularly harmful to the coal industry and, by extension, its employees, their social network, and all others who rely on coal for energy or work (Fisher 2006, 477–480). The wide dispersal of coal deposits and the proportion of domestic energy produced by

burning coal ensures that a majority of states (Fisher 2006, 480) are affected by the well-being of the coal industry, which in turn implies that a majority of representatives at the national level are friendly to coal interests and reluctant to support legislation to stifle domestic CO2 emissions. The cost of transitioning away from coal is simply too high, both economically and politically (Fisher 2006, 480); the interests of Big Coal run too deep into the U.S. political and social structure to be uprooted without causing severe damage. I will consider two aspects of the argument for a relationship between coal dependence and national legislation, electoral pressure and union influence, and whether the formal abstractions of de facto veto power and local constituencies can be reasonably applied to them.

Generally, the U.S. system is very well approximated by the model of self-interested legislators bargaining over the division of benefits from cooperation over environmental issues. U.S. legislators tend to favour constituency interests over party politics, and the allocation of federal subsidies is an important part of national debates over U.S. participation in international environmental agreements (Bang, Hovi and Sprinz 2012, 3). Both characteristics suggest that a state representative will be primarily concerned with his own state's share of any policy compromise proposed by the executive. The thesis of Fisher (2006) is that coal production affects national climate policy partly because it is labour intensive and dispersed, which means that the interests of the Coal Industry affect a great number of voters and thus binding a majority of state representatives to vote in favour of coal interests nationally. To assess the findings of Fisher (2006), it is useful to consider how states are categorised according to their level of coal production and what this implies for voting on CO2 legislation at the national level. Fisher (2006, 484) differentiates "coal producing states" according to the number of million tons of coal extracted in the year 2000. The cut-off between major and minor coal producers is set at 25 million tons of coal extracted, with the latter category including any state that extracted more than 0 tons of coal in 2000. Categorising coal producers in this way yields 13 major and 13 minor coal producing states (Fisher 2006, 484). We may loosely define "politically relevant" levels of coal extraction activity as obtaining when the number of people employed is great enough that it affects the general voting pattern in the state in favour of the coal industry[1] (i.e., voters will reward candidates who cater to the interests of coal). Then, politically relevant levels of extraction is trivially expressed formally using the state representative's local constituency. It is unlikely, however, that coal extrac-

[1]That the general voting pattern can be affected even if the number of people directly employed by the coal industry is relatively low (Fisher 2006, 479) can be supported by appeal to the social multiplier effect (Fisher 2006, 486), which allows counting friends and family of employees as well as the employees themselves among voters likely to favour the coal industry.

tion itself could be a good explanation for aggregate voting on climate policy at the national level, and so imposing restrictions on what the executive can agree to internationally. The problem is that the separation between major and minor coal extracting states is too coarse: the lower bound is too low to reasonably assume that political relevancy obtains in all 26 states. The aggregate political effect of coal extraction as such may not be what Fisher (2006) is claiming, though the state-level effect of having significant levels of coal extraction activity may be plausible.

The immediate economic benefits to coal extraction activity is only part of coal dependence, as states without significant activity are independently reliant on coal as a cheap source of energy (Fisher 2006, 480). The dependency of a energy user is weaker than a state where extraction is a major employer, however. A user state has no reason to consider anything but the price of energy from some source. If and when a cheaper alternative to coal becomes available, states that depend on coal for energy should switch to that alternative — assuming the infrastructure allows it. A user of coal energy can switch to another source of energy easily, but a coal worker may not be able to find alternative employment as easily. Coal dependence should influence voting in user states only to the extent that no significantly cheaper alternative to coal exists, whereas voting in states with significant coal extraction activity should be less sensitive to the development of alternatives. When deciding whether to include the coal industry as part of a particular state representative's local constituency, then, it is necessary to show either that politically relevant levels of coal extraction occurs in that state or that switching to an alternative energy source would increase the average price of electricity so much that it would reduce the representative's chances of reelection.

There is some evidence to suggest a correlation between coal dependence and voting on climate legislation. The vote on Byrd-Hagel itself offers little evidence (Fisher 2006, 483), both because it was unanimous and second because its main concern is exemption for developing countries not CO2 reduction as such, which was made clear by its main sponsors (Hovi, Sprinz and Bang 2012, 135). The votes on the Climate Stewardship Act — which would legislate reductions to 2000 levels by 2010 using a market for tradable allowances — offers some supporting evidence, with the correlation between coal dependence and voting claimed to "... present a clear relationship between coal dependence and opposition to domestic climate change regulation" (Fisher 2006, 485). A similar pattern is reported for CO2 reducing legislation at the state level (Fisher 2006, 487). The number of cases and the strength of the observed correlations are such that the evidence appears fairly weak, how-

ever, which casts doubt on whether this evidence can be taken as supporting the local constituency explanation of U.S. climate policy.

Although the number of coal workers in the U.S. may be too small to influence national U.S. climate policy via voting, there is still an argument to be made that a coordinated Coal Union has the potential to influence policy by the threat of retaliatory strike. Policy targeted at reducing domestic CO2 emissions, say by taxation, would have a negative impact on coal interests, because coal emits large amounts of CO2 when burnt for energy (Fisher 2006). Proposals to tax CO2 emissions should prompt the Coal Union to threaten retaliatory strike, and supposing the threat to carry out that strike was credible and the cost high enough, then policy makers would have to comply. This logic is approximated by the formal notion of veto power (see chapter 3); if we say that the U.S. Coal Union is a cohesive unit, a "group" in the generic terminology of Mo (1994), then the Union influence on climate policy is expressed by assigning veto power to it within a model of international negotiations over CO2 emission reductions. If the coal union has de facto veto power on issues that involve policy to reduce domestic CO2 emissions, then that could help explain the U.S. position on global climate change policy in general and in the Kyoto proceedings specifically. Furthermore, the Coal Union should want to leverage their de facto veto power to resist efforts to transition from coal to alternative energy sources, such as natural gas or renewable energy sources. There are reasons to believe, however, that the Coal Union should not be treated as a de facto veto power on these issues.

To reasonably say that the U.S. Coal Union acts as a de facto veto player, it is necessary to justify the extra-legal source of the Union's influence. For the Coal Union, influence has its material basis in the demand for electricity. Burning coal contributed around 40% of all electricity consumed in the United States in 2012 (Broderick and Anderson 2012). That means that a coordinated shut-down of coal plants nation-wide would cause a large drop in the domestic supply of electricity in the United States. The primary cost of such a strike is the temporary switch to more expensive alternatives to make up for the deficiency in supply, which would be incurred by private industry and business, state owned operations, and private citizens, ensuring not only economic costs but also the political costs borne by legislators for allowing such a thing to occur instead of yielding to the demands of the Coal Union. The changes that have occurred in energy production in the United States over the last decade suggests that the Union lacks this ability, however.

The so-called "shale gas revolution" in the United States, associated with the development of

improved methods of extraction and the discovery of large domestic deposits of shale gas, has contributed to a beginning shift away from coal to gas as a source of electricity in the United States (Broderick and Anderson 2012) Broderick and Anderson (2012), citing reports by the US Energy Information Administration (EIA), show that coal has become less important as a source of electricity in the United States, with gas becoming more important. This development is correlated with an increase in the price of coal relative to gas. The material basis of the Coal Union's supposed power is directly tied to the degree of coal dependency in the United States — the greater the dependence, the greater the cost of a strike. The "shale gas revolution", then, undercuts the power of the Coal Union by making the U.S. less dependent on coal. Furthermore, the fact that the transition from coal to gas has been "allowed" to occur is in itself evidence that the Coal Union does not have de facto veto power on this issue.

The argument against assigning de facto veto power to the U.S. Coal Union can also be stated as a question of credible threats. A threat is not credible unless it is profitable to carry it out, otherwise the recipient of the threat will deduce that the threat is without consequence and ignore it. A retaliatory strike among coal workers would impose some costs in the short term, but the long-term effect of a strike would be to reduce the supply of coal and thus increase its price relative to alternative sources of energy. Assuming that the demand for coal is elastic (which it will be in the presence of viable alternatives), then consumers will respond to an increase in the price of coal by switching to alternatives. A coordinated strike among coal workers is not a credible threat because it reduces the price of alternatives and accelerates the transition away from coal in the long term.

5.5 The assumption of no auxiliary benefits

A fundamental assumption of the theory of two-level games is that the executive derives no auxiliary benefit from negotiation itself. This is the mechanism that ensures the link between level 1 and 2, it is the motivational justification of the assumption that the executive anticipates the outcome of ratification when making agreements at the international level. Without this assumption, the causal link between the domestic and international levels suggested by the theory of two-level games dissipates. Furthermore, because it is (typically) the executive that holds the formal power to initiate and conduct negotiations, the assumption of no auxiliary benefits also supports the contention that states negotiate because there is a belief in a range of mutually profitable agreements. Neither of these assumptions

are unreasonable, though they exclude several interesting interpretations of real-world negotiation outcomes. By stating that the executive receives a zero payoff in the case of ratification failure, the analyst is forced to reach certain conclusions about the beliefs of the negotiator and his motivations for acting in a certain way. A zero payoff from ratification failure implies that the negotiator always believes that the agreements he signs on to will enter into force, which leads to the somewhat disappointing conclusion that negotiations break down only because the executive has misjudged the possibility of overlapping win-sets or the likelihood of legislative support.

A major premise in two-level modeling is that negotiators only consider ratifiable agreements, which is not borne out by the Kyoto case (Hovi, Sprinz and Bang 2012, 131). Despite Byrd-Hagel, the U.S. signed on to the Kyoto agreement, most likely knowing that it would not pass ratification if presented to the Senate (which it was not). All the two-level models presented here assume that the executive's primary motivation is to achieve the best ratifiable agreement. This is reflected in the assumption that the executive gains, if he gains, only in the case that an agreement is ratified. Formally, all models indicate a zero payoff to the executive in the no-agreement outcome; the way his reward is calculated varies but is always derived in some way from the gross share of the gains from cooperation. The formal account suggests that the U.S. decision to sign on to Kyoto was informed by the mistaken belief that ratification could be achieved despite the Byrd-Hagel act. Then, Kyoto is simply an instance of involuntary defection on the part of the U.S., as described by Iida (1993*b*). As I have already argued, this interpretation is not very convincing, and it excludes the equally plausible possibility that the U.S. delegation knew that ratification was likely to fail. What motivation could the U.S. delegation have for signing on to Kyoto with the belief that it would fail ratification? Hovi, Sprinz and Bang (2012) suggests that the U.S. decision to sign on to Kyoto was motivated by domestic political concerns. That is, there may have been some auxiliary benefit for Clinton and Gore to signing on to Kyoto, regardless of the outcome of ratification. That the president is popularly elected (has a national constituency) implies greater freedom to pursue auxiliary goals such as improving his chances of reelection and securing a positive political legacy (Bang, Hovi and Sprinz 2012, 4). To capture this notion, however, one must set the no-agreement payoff for the executive to some non-zero value, reflecting the political or personal gain that they get for participating in negotiations and for signing on to a non-ratifiable agreement.

The second major assumption that rests crucially on there being no auxiliary benefit from nego-

tiations is that states only negotiate when there is a shared belief in a range of mutually profitable agreements. Assuming a zero payoff for failure to ratify implies that negotiations require at least a "… misperception that there is a mutually acceptable range of agreement" (Iida 1993*b*, 405). A corollary to this assumption is that negotiations are evidence that there is a belief in mutually profitable agreements. It is difficult to reconcile this interpretation with the fact that international negotiations on global climate change mitigation are held at regular intervals despite never producing tangible results (nor can they be expected to do so in the current format, see Victor (2011)). Insisting that negotiations only occur provided a shared belief in mutually profitable agreements suggests that the participants are incompetent and unable to learn from the experience of previous rounds. An executive who derives some auxiliary benefit from negotiation, however, does not necessarily have to believe that a substantial agreement is possible to pursue negotiation. That is, the fact that negotiations over environmental issues occur is not necessarily evidence that the participants believe that substantial agreement is realistic, if we allow for executives to derive auxiliary benefits. This opens the way for interpretations that are more interesting than an appeal to incompetence.

From a rational choice perspective it is tempting to construct some argument by which it is rational to enter into negotiations absent the belief that a substantial and mutually acceptable agreement exists. To do so, it is sufficient to provide some set of conditions under which the executive can reasonably be said to gain from negotiations in themselves independently of the ratification outcome. An obvious way of doing so is to imagine that the executive is interested in the domestic political benefits from negotiations, specifically as concerns his own reelection probabilities. Suppose that domestic constituents create "demand" for policy and that the primary goal of the executive is to "supply" policy in such a way that it maximises the executive's political survival (de Mesquita et al. 2002). For the Kyoto case, explanation 3 illustrates this idea: "Signing an agreement with relatively ambitious emissions reduction targets, but with little or no chance of Senate ratification, allowed the administration, according to explanation 3, to look climate-friendly without committing the United States to costly emissions reductions" (Hovi, Sprinz and Bang 2012, 136). That is, the U.S. executive viewed the Kyoto agreement, and participating in the negotiations leading up to its signing, as a cheap source of policy. If the leader uses international negotiation as a cheap way of satisfying domestic demand for policy, or derives some other auxiliary benefit from negotiating (e.g., diversion), then it is reasonable to say that pursuing negotiations for their own sake can be rational. If entering into international

negotiations in itself is sufficient to satisfy domestic demand for policy (in particular if the demand comes from the broader constituency), then pursuing negotiations may be the most economical means of supplying policy and so the best choice. Suppose there is support for climate policy in the broader electorate, but that there is strong opposition among core constituents or veto groups, such that implementing actual policy would be politically costly. Suppose that the broader constituency can be satisfied by (appearing to) taking part in international efforts to address climate (say, by negotiating a comprehensive and binding agreement); and finally, suppose that opponents of international cooperation will not retaliate unless substantive, costly policy is implemented. Under these conditions, the leader may reasonably be said to derive some auxiliary benefits from participating in negotiations (satisfying policy demand, cheaply) while having no belief in a feasible agreement or a real will to reach substantive agreement. It is then "rational" to enter negotiations, even if there is no prospect of reaching substantive agreement.

Chapter 6

Conclusion

6.1 What was done

The theory of two-level games is presented in the literature as a general description of the relationship between domestic politics and foreign policy. This thesis has considered three formal developments of the theory of two-level games and discussed applications to environmental policy. Specifically, this thesis has attempted to identify formal abstractions that are useful in modelling the role of domestic politics in international negotiations over environmental issues. Each model was discussed in-depth to point out some of the model's strengths and weaknesses as a general description of the domestic sources of international negotiation behaviour and to identify particularly useful abstractions exemplified by the model. Having identified key formal abstractions, I discussed applications of the theory of two-level games to environmental issues by using the U.S. position in the Kyoto negotiations as an empirical reference point. The remainder of this chapter summarises what I found and suggests avenues for further research.

6.2 What was found

The models provide examples of abstractions for describing the general relationship between domestic politics and foreign policy outcomes. Three abstractions are of particular interest for modelling the role of domestic politics in international environmental negotiations: a probabilistic ratification event (Iida 1993*b*), veto power (Mo 1994), and a variable no-agreement outcome (Tarar 2005). Briefly, the benefit of each abstraction is summarised as follows. A probabilistic ratification event captures in-

voluntary defection that is due to uncertainty about the true preferences (and thus voting behaviour) of legislators. Although the specific assumptions made about the access to and ability to process information about individual legislators' preferences may or may not be plausible (depending on the case at hand), the approach itself is clearly reasonable — it covers the essential idea without introducing too much complexity. Veto power is a useful abstraction that makes it possible to differentiate between domestic political actors that are merely satisfied from those who are powerful in the traditional sense of having influence over policy. As argued, the approach of deriving political power from satisfaction with the status quo is impotent and substantively uninteresting. A truly powerful domestic actor (e.g., a large national labour union) is more aptly described as a player that has to be included in any policy compromise on some issue, i.e., as a player that has (de facto) veto power. Finally, the addition of a variable no-agreement outcome enriches a model of political competition with the notion that there are varying degrees to which a state will be dependent on cooperation with others. For problems that are truly "international", such as transboundary pollution, individual states may not achieve much alone, which is naturally expressed by a low no-agreement value. From this follows the possibility of modelling states' willingness to cooperate as a function of the no-agreement value, by postulating some critical value at which legislators will prefer cooperation over addressing the issue by domestic policy alone. This critical value may be seen as convenient notation for describing the micro-foundations for environmental legislation. Substantively, however, this approach results in yet another way of articulating the fundamental problem of externalities.

Applying these abstractions to the empirical case of the U.S. position in the Kyoto proceedings is not very satisfying. It appears that none of the explanations considered here, derived from the formal two-level models, can give a reasonable narrative of what caused the United States to first sign on to and then repudiate from the Kyoto protocol. The primary reason for this deficiency appears to be that all the two level models assume a zero payoff for the executive in the case of ratification failure. That assumption restricts the range of possible interpretations of real-world negotiation outcomes, including the Kyoto case. According to the zero payoff assumption, executives only consider ratifiable agreements and only negotiate if there is a belief in mutually acceptable agreements, which means that the Kyoto outcome has to be explained as a result of involuntary defection. The facts of the case, however, suggest that the U.S. delegation knew full well that U.S. accession to the Kyoto, or agreements like it, would be impossible. Most telling, perhaps, is that Clinton did not send the Kyoto

agreement to the Senate for consideration. Explaining the U.S. position in terms of the influence of coal extraction activity on voting at the national level can easily be accommodated by the models considered, but does not appear to be well supported by the evidence. Another approach to capturing the link between the interests of the coal industry and national climate policy is to treat coal workers as a unitary group with (de facto) veto power on the issue of climate legislation, supported by the threat of a coordinated strike. Although coal remains a major contributor to the domestic energy supply, the potential power of the Coal Union may not be realisable, however, because a strike could accelerate the transition to alternative energy sources such as gas, which has been ongoing for the last decade following the discovery of large shale gas deposits in the U.S. and improvements to the technology used to process it.

6.3 What could be done

Much remains to be done on the topic of modelling the role of domestic politics. The discussion suggests four possible avenues of future research. First, given that uncertainty about future states of the world is such an important element of long-term environmental policy, models of international negotiations on the issue should better accommodate the concept of analytic uncertainty as well as strategic uncertainty, as described by Iida (1993*a*).

Second, the range of interpretations that can be supported by formal models would be increased by relaxing the assumption of a zero payoff to the executive in the case of ratification failure. As argued for the Kyoto case, the assumption of a zero payoff commits the analyst to a certain mode of thinking about what could explain the negotiation behaviour of states which excludes explanations that are theoretically interesting and supported empirically. Careful assessment of the motivation of the negotiator is necessary to ensure that the formal assumptions about what drives bargaining can be reasonably applied to the case at hand.

Third, formal models of the role of domestic politics in international environmental negotiations should take into account that negotiators differ in their access to and ability to assess information about relevant domestic political factors in the different countries involved in the negotiations. The assumption of a zero payoff can support the notion that negotiators have an innate interest in harmonising their beliefs about the general probability of ratification, but much expressive power can be gained by allowing each negotiator a personal assessment reflecting his access to information and

ability to process it. This difference is one of proximity but also of skill. Allowing the latter to vary suggests exploring models based on bounded rationality, perhaps using techniques inspired by Rubinstein's (1993) model of consumers who differ in their ability to process price information. Fourth, an issue that has not been explicitly addressed here, is that of procedural realism. A bilateral bargaining model is far removed from the reality of international negotiations, and is certainly quite unlike the format of climate negotiations (see e.g., Victor (2011)). Substantively, multilateral extensions are important for the basic reason that a great number of real-world negotiations are multilateral. Further, the number of participants in negotiations may be a relevant variable in explaining the outcome of some negotiations, as suggested by Victor's (2011) argument about the detrimental effect of the inclusive nature of climate negotiations. Theoretically, multilateral extension is necessary to verify that the results from the bilateral model can in fact be generalised to the multilateral case, which is implicitly assumed by the models considered here and which is required for the two-level model to be a general description of international negotiations. Krishna and Serrano (1996) obtain a unique and stable solution to multilateral bargaining by allowing satisfied parties to "exit" from bargaining, under the requirement that the overall outcome predicted by the model be the same for reduced versions of the game (i.e., following exit by a satisfied participant). Introducing the possibility of "exit" and justifying its substantive meaning in the context of bargaining over environmental issues, say the division of CO2 quotas, might be a suitable starting point for a multilateral theory of two-level games and the role of domestic politics in international environmental negotiations.

Bibliography

Bang, Guri, Jon Hovi and Detlef F. Sprinz. 2012. "US presidents and the failure to ratify multilateral environmental agreements." *Climate Policy* 12(6):755–763. DOI: 10.1080/14693062.2012.699788.

Baron, David P. and John A. Ferejohn. 1989. "Bargaining in Legislatures." *American Political Science Review* 83(4):1181–1206.
URL: *http://www.jstor.org/stable/1961664*

Broderick, J. and K. Anderson. 2012. Has US Shale Gas Reduced CO2 Emissions? Technical report Tyndal Consortium Institution.
URL: *http://tyndall.ac.uk/publications/technical-report/2012/has-us-shale-gas-reduced-co2-emissions*

Coase, Ronald. 1960. "The Problem of Social Cost." *Journal of Law and Economics* 3:1–44.

Congleton, Roger D. 1992. "Political Institutions and Pollution Control." *The Review of Economics and Statistics* 74(3):412–421.

de Mesquita, Bruce Bueno, James D. Morrow, Randolph M. Siverson and Alastair Smith. 2002. "Political Institutions, Policy Choice and the Survival of Leaders." *British Journal of Political Science* 32(4):559–590.

Fisher, Dana R. 2006. "Bringing the Material Back In: Understanding the U.S. Position on Climate Change." *Sociological Forum* 21(3). DOI: 10.1007/s11206-006-9026-2.

Hovi, Jon, Detlef F. Sprinz and Arild Underdal. 2009. "Implementing Long-Term Climate Policy: Time Inconsistency, Domestic Politics, International Anarchy." *Global Environmental Politics* 9(3):20–39. DOI: 10.1162/glep.2009.9.3.20.

Hovi, Jon, Detlef F. Sprinz and Guri Bang. 2012. "Why the United States did not become a party to the Kyoto Protocol: German, Norwegian, and US perspectives." *European Journal of International Relations* 18(1):129–150. DOI: 10.1177/1354066110380964.

Iida, Keisuke. 1993*a*. "Analytic Uncertainty and International Cooperation: Theory and Application to International Economic Policy Coordination." *International Studies Quarterly* 37(4):431–457.
URL: *http://www.jstor.org/stable/2600840*

Iida, Keisuke. 1993*b*. "When and How Do Domestic Constraints Matter? Two-Level Games with Uncertainty." *Journal of Conflict Resolution* 37(403). DOI: 10.1177/0022002793037003001.
URL: *http://jcr.sagepub.com/content/37/3/403*

Krishna, Vijay and Roberto Serrano. 1996. "Multilateral Bargaining." *Review of Economic Studies* 63:61–80.

Mo, Jongryn. 1994. "The Logic of Two-Level Games with Endogenous Domestic Coalitions." *Journal of Conflict Resolution* 38(3):402–422. DOI: 10.1177/0022002794038003003.
URL: *http://jcr.sagepub.com/content/38/3/402*

Osborne, Martin J. and Ariel Rubinstein. 1990. *Bargaining and Markets*. San Diego, CA: Academic Press, Inc.

Osborne, Martin J. and Ariel Rubinstein. 1994. *A Course in Game Theory*. Cambridge, MA MIT Press. ISBN: 0-262-65040-1.

Putnam, Robert D. 1988. "Diplomacy and Domestic Politics: The Logic of Two-Level games." *International Organization* 42(3):427–460.

Rubinstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50:97–109.

Rubinstein, Ariel. 1985. "A Bargaining Model with Incomplete Information About Time Preferences." *Econometrica* 53(5):1151–1172.
URL: *http://www.jstor.org/stable/1911016*

Rubinstein, Ariel. 1993. "On Price Recognition and Computational Complexity in a Monopolistic Model." *Journal of Political Economy* 101(3):473–484.

Senate Resolution 98. Report No. 105-54. 1997. http://www.nationalcenter.org/KyotoSenate.html.

Tarar, Ahmer. 2005. "Constituencies and Preferences in International Bargaining." *Journal of Conflict Resolution* 49(383). DOI: 10.1177/0022002705276567.
URL: *http://jcr.sagepub.com/content/49/3/383*

Victor, David G. 2011. *Global Warming Gridlock. Creating More Effective Strategies for Protecting the Planet*. New York: Cambridge University Press. ISBN: 978-0-521-86501-2 Hardback.