

The Role of the WTO in the Transfer of Policy Knowledge on Trade and Competition

*Oliver Morrissey and Doug Nelson **

Abstract

This paper uses theories of policy learning and of policy-making to examine how global institutions such as the WTO can influence policy choices. In pure learning by doing, policy choices are based on information relating to the history of an active policy; there is no information on alternative policies. New information on priors provides an incentive to choose a different policy. In the case of social learning, policy-makers can observe the policies chosen by other actors, but the signals those other actors receive is unobserved. External agents (global institutions of knowledge transfer) can influence policy choice by altering priors, providing technical advice or providing information on the (unobserved) effects of the policy choices of others. If the aim is to promote competition, institutions should also show awareness of the global competition environment. Many policies are advocated on the implicit assumption of competitive markets. This gives a false impression of the state of nature, and increases the probability that the effect of reforms will be less beneficial than predicted. If so, the reputation of the WTO as a 'purveyor of global policy knowledge' may be undermined.

- Paper prepared for the conference on 'Innovation, Competition and Regulation in the Global Economy', Margaret River, Australia, 21-23 September 2001. Oliver Morrissey is Reader in Development Economics and Director of CREDIT, School of Economics, University of Nottingham. Doug Nelson is Professor, Murphy Institute of Political Economy, Tulane University, and an External CREDIT Fellow.

Contact: oliver.morrissey@nottingham.ac.uk.

1 Introduction

Globalisation and economic liberalisation over the past two decades have contributed to expanding flows of trade, technology and capital between countries in both the developed and developing world. Trade liberalisation at various levels has been a policy issue in almost all countries. Multilateral trade liberalisation has become an institutional part of the globalisation process, for good or bad. Important new issues were brought into the Uruguay Round and the World Trade Organisation (WTO) has a broader remit than the GATT it replaces. In particular, this remit extends to major ‘trade-related’ issues, impinging on investment, intellectual property rights, industrial policy, and the regulation of competition. Unilateral trade liberalisation has been implemented, to varying degrees, by almost all developing countries. The perceived benefits of liberalisation are to enhance growth prospects via integration into the global economy and increased efficiency in resource allocation. However, the evidence that trade liberalisation *per se* promotes growth is weak when exposed to careful scrutiny (e.g. Rodrik, 1999) so why, in the 1980s and 1990s, did so many countries ‘catch the reform bug’ (Rodrik, 1996: 11)?

This is the question we seek to answer in part. We limit attention to the role of external influences in shaping and supporting policy reform on trade and trade-related issues. Specifically, we are concerned with the role of the World Trade Organisation (WTO) in promoting global competition through multilateral liberalisation, and the influence this has had on the policies of individual countries towards the regulation of trade and competition. We will also refer to the World Bank, the other global institution that has played a major role in promoting liberalisation, especially trade policy reform. However, such references will be in passing as we here wish to focus on competition issues within the context of multilateral trade and investment liberalisation.

This paper represents a first attempt (by the authors) to provide a link between the largely theoretical economics literature on ‘policy learning’ and the somewhat more practical, albeit conceptual, political science literature on policy-making. While we believe such an exercise to be potentially fruitful, there are inherent difficulties. Given its theoretical foundation, the literature on policy learning tends to have strict assumptions and be formally restrictive (in terms of general applications). The policy-making literature, on the other hand, attempts to derive general concepts and principles from observed outcomes. To take a specific

example, a policy learning theory will start with a strict definition of and assumptions over the priors of policy makers. The policy-making literature will infer preferences (which are not the same as priors, about which there is usually no information) from observed behaviour. Consequently, there are few direct links between Sections 2 and 3, where we discuss the respective literatures. In the final section we attempt to illustrate how the two literatures can be linked with two examples, anti-dumping and competition policy within the WTO.

Essentially, our approach is to treat the WTO as a disseminator of ‘institutional policy knowledge’ that plays a direct role in encouraging, supporting and even coercing trade (and trade-related, including investment) policy reform. In what ways does this influence policy choice in developing countries, the recipients of advice for our purposes, and, in particular, does this increase the likelihood of countries adopting optimal policies? Strictly speaking, the WTO does not offer policy advice (although it does provide members with technical assistance in implementing commitments). However, it does represent, and through its functioning promulgates, a particular position on what constitutes optimal trade policy. In simple terms this could be described as non-discriminatory liberalisation, with the emphasis perhaps on non-discrimination. It is also a club, membership of which requires accepting this policy stance. It is the global institution in respect of trade policy.

The GATT 1994 is an agreement between governments in which they have made binding commitments. The WTO is the body that monitors and regulates compliance and implementation of this agreement. In this sense it clearly has an important influence on government policies towards trade and investment. However, most trade (at least two-thirds) is actually within or between multinational companies, and most foreign investment is by companies. The WTO does not have a remit to regulate the behaviour of companies; this is essentially a government function. Yet WTO commitments constrain what governments can do. Thus we will consider the role of the WTO in disseminating policy knowledge in two respects. First, the function through multilateral agreements in spreading policy knowledge on trade and investment. Second, the potential contribution it can play in spreading policy knowledge on regulation. Before doing this we need some framework in which to analyse how policy knowledge spreads and how this influences policy choices.

Section 2 reviews various theories on the policy learning (by governments as the agents making choices) and the spread of policy knowledge. Section 3 then relates this to the

political dimension of policy change and evolution within countries. Through what mechanisms do external actors such as the WTO influence policy choices, and does this lead to ‘better’ policies? Section 4 then illustrates our arguments in respect of two issues – anti-dumping, and competition policy. Section 5 concludes with an evaluation of the WTO against the criterion of spreading good policy advice.

2 Models of Policy Learning and Knowledge Transfer: Three Views of the Beast¹

Given the obvious importance of learning to virtually all forms of human endeavour, it is probably not surprising that the literature on learning is immense. In an effort to keep this discussion manageable, we will frame our discussion in terms of rational learning in a Bayesian environment.² Specifically, we will sketch three models of policy learning: essentially asocial, decision-theoretic learning; social learning; and hierarchical social learning. In each case, after sketching the basic model, we will suggest the implied role for policy research and the transfer of results from policy research by institutions like the WTO. We leave to the next section discussion of who is doing the learning, and the way that political considerations interact with institutional considerations to determine willingness and capacity to learn (in the specific meaning of learning being discussed here).

2.1 Decision-Theoretic Learning: Policy Experiments and Learning-by-Doing

Consider the case of a small, less-developed economy facing two policy options: import-substitution (IS) and export-orientation (XO). For now we will assume that these options are meaningful and exclusive. We begin with the case of pure learning by doing. That is, there is no possibility to learn from the experience of others. We suppose that the adoption of a policy results in an outcome, which we take to be either “good” or “bad”. The outcome provides some information about the effectiveness of the particular policy, but no information about the effectiveness of the other policy. However, the effectiveness of the policy is

¹ There are a number of good surveys of the economic literature on learning. For a convenient overview, see Sobel (2000).

² Thus we abstract from sizable bodies of research that seek to characterize and distinguish between knowledge and belief (e.g. Hintikka, 1962); those that deal with attempts to extend Bayesian logic to conditions of Knightian uncertainty (Bewley, 1998; Camerer and Weber, 1992); those that attempt to develop modern notions of rational behaviour under bounded rationality (Fagin, Halpern, Moses, Vardi, 1995; Rubinstein, 1998); and, finally, the enormous literature on philosophy/sociology of science that

determined by factors not under the control of the policymaker and this fact must be taken into consideration in evaluating the policy.

A bit of formalism may help here. Denote the state of the world (i.e. the wide range of things that are not under the policymaker's control but which affect the outcome of the policy experiments) by $\theta \in \Theta$. In each period, t , the policymaker chooses a policy $x_t \in X$ (in our case $X = \{\text{IS}, \text{XO}\}$). In a sense, this produces a state of the world in $X \times \Theta$, and results in a signal $y_t(x_t; \theta) \in Y$ (in our case $Y = \{\text{good}, \text{bad}\}$). We will suppose that policy $x(i)_t$ produces good states with unknown probability $p(i)$ and bad states with $[1 - p(i)]$ and that the policymaker begins with prior belief about the likelihood of a good outcome under policy i , $\rho_{i0} \in [0,1]$, which is commonly taken as deriving from a private signal of bounded accuracy that the policymaker receives in period $t = 0$. Knowing x_t and y_t , the policymaker can update his beliefs, r_{t-1} , using Bayes rule to get r_t . We assume that only the element of r_t referring to the active policy in period t changes in the updating, since there is no information about the effectiveness of a policy that is inactive.³ We suppose that the policymaker's objective is to maximize the expected number of good realizations.⁴ Specifically, if we let $y_t = \text{good} = 1$ and $y_t = \text{bad} = 0$, and assume that the policymaker applies geometric discounting with discount factor $\delta \in [0,1)$, we can write this objective as:

$$V(\mathbf{s}, \mathbf{r}) = E_{\mathbf{s}} \left[\sum_{t=0}^{\infty} \delta^t y_t(x_t; \mathbf{q}) \right]. \quad (2.1)$$

can be seen as treating precisely questions of how we do, or should, learn. We will also not pursue the rapidly growing literature on learning in the context of games (e.g. Fudenberg and Levine, 1998).

³ The assumption of independence is not entirely harmless, as it implies that import substitution could be less effective, more effective, or equally effective as export orientation. That is, policymakers may not assume that one policy is necessarily superior to the other. The case of dependent bandits is more complex, and we cannot convince ourselves that it is the more obviously applicable case. A useful overview of the dependent case can be found in Pressman and Sonin (1990).

⁴ The literature on the political economy of macroeconomic policy provides considerable warrant for this assumption. Alternatively we could follow much of the literature and, in addition to the signal $y(\bullet)$, we could introduce a reward function $r(x_t; \theta)$ incorporating any factors, e.g. income distribution, we might deem important to the policymaker's problem. However, since our interest in this paper is on learning *per se*, we focus on the information process and abstract from the reward process.

In the theoretical statistics literature, this is called a *Bernoulli two-armed bandit* problem, with the arms given by the policies.⁵ An intertemporally optimal policy takes into account both the one-period gain from a given policy and the gain from information that may be used in future plays of this game against nature.

In constructing an optimal strategy we need the notion of a *history at k* , a description of the policy used in each period up to $t = k$ and the signals observed: $h^k = \{x_t, y_t\}_{t=1}^{k-1}$. Let H^k be the set of all possible histories at k . A strategy, σ , for the policymaker specifies a policy choice to be made in any period as a function of initial beliefs and Bayesian updating on the history up to that point. Gittins and Jones (1974) proved a striking result for problems of this sort: to every policy (i.e. “arm” of a bandit) there is associated an index which depends only on the current prior on that arm, $\rho_{i|k}(\rho_{i0}, h^k)$, and the optimal strategy at time k is to adopt the strategy (“play the arm”) with the highest index. Furthermore, as Whittle (1982, chapter 14) makes clear, this index is essentially the value of a payment that would make the policymaker indifferent between stopping and continuing with strategy i . As a result, solving the policymaker’s problem involves solving an optimal stopping problem for each of the policies in X .

One of the fundamental questions that has been addressed in this framework is whether, with sufficient time, the policymaker would necessarily learn the best policy, i.e. the policy such that $p_i > p_j$ ($i \neq j \in X$), if such a policy exists.⁶ The usual answer to this class of question is that complete learning generically fails. Specifically, with strictly positive probability, the policymaker may eventually select and stay with the “wrong” policy—i.e. the policy j in the preceding inequality. Furthermore, in finite time, a policymaker might switch between policies many times.

⁵ A very clear introduction to bandit problems can be found in DeGroot (1970, chapter 14). For an excellent, considerably more extensive, overview see Berry and Fristedt (1985).

⁶ The actual characterization of this class of question is complex, involving at least two related questions: whether $\rho_i = p_i$ $\forall i \in X$; and whether $\rho_i > \rho_j$ if $p_i > p_j$ ($i \neq j \in X$). There is an extensive literature on this sort of question. Among the major papers in economics are Easley and Kiefer (1988); Feldman (1991); Aghion, *et al.* (1991); Banks and Sundaram (1992); Brezzi and Lai (2000). Kiefer (1988/9) provides a nice overview of the issues.

In this paper we are less interested in the implications of these results for observable policy histories, or the normative conclusions that can be drawn from any given policy history, than in the implications of learning theory for institutionalised policy advice. In this context we isolate two obvious, but important, roles highlighted by the simple model sketched above: technical support; and affecting the prior beliefs of the policymaker. While the model presented above is quite simple, it should be clear that a great deal of potential complexity is contained in and that the process of actually carrying out the analysis generating could be technically demanding. A substantial number of people trained (at many levels) as economists perform precisely this task. In this context, one of the important roles played by international agencies for least developed countries is the provision of precisely this sort of expertise. For example, in addition to direct provision of expertise, the World Bank has produced a number of briefing books to support developing country participation in trade rounds (e.g. Finger and Olechowski, 1987).

Somewhat more subtly, it should be clear that one fundamental role of policy advice is to affect the beliefs of policymakers. In her presidential address to the American Economics Association, Krueger (1997: 18) argues that “good policy-relevant theory provided blueprints for those windows of opportunity in which governments genuinely sought to improve economic performance ... [and] ... theory was invaluable when it showed why simple interpretations of received doctrine were in fact wrong”. In the context of the model above, a key role is played by θ_0 , the policymaker’s initial beliefs. There is a long tradition in Bayesian analysis of treating initial beliefs, like tastes, as primitive. However, there are a number of prominent examples of systematic argument affecting belief change. We mention two of the most prominent in the context of trade policy: the role of new developments in economic theory in Peel’s decision to repeal the Corn Laws (Irwin, 1989); and the role of leading intellectuals such as S.C. Tsiang and T.C. Liu in Taiwan’s transition to XO (Haggard, 1990). In both cases, change in beliefs of incumbent policymakers produced change in policy. These are striking examples, perhaps more commonly the beliefs of government change through a mix of changing incumbents and failure of old ideas in the face of policy

crisis (Harberger, 1993).⁷ International agencies are among the primary channels for transmitting current policy thinking to the policymaking community.⁸ This is especially the case for least developed countries with modest connection to the international academic sources of policy thinking. As an example, the World Bank Institute was developed to carry out training on a variety of policy relevant topics, including the analysis and implementation of trade policy. It may be that the Trade Policy Review Mechanism of the WTO also serves a useful information transfer function.

2.2 Social Learning: Learning from Others and Information Cascades

The discussion of the preceding section presumes that policymakers learn exclusively by doing. That is, a policy is adopted, an outcome occurs, and the policy is evaluated relative to the policymaker's beliefs about existing alternatives. It is surely the case that significant learning also occurs through observation of the policy experience of others. In this section we first offer a simple extension of the above framework and then consider the implications for policy transfer by international agencies.

We now suppose that there are a finite number of policymakers, in different countries, facing the problem sketched in the preceding section. In addition, we assume that these policymakers can observe the choices made by the other policymakers, but not the signals resulting from them. That is, denoting policymakers by superscript $a \in A$, everyone observes the vector \mathbf{x}_t of policy choices made at time t , but the $y_t^a(x_t^a; \cdot)$ are private information to each a . Furthermore, we assume that the $y_t^a(x_t^a; \cdot)$ depend only on the x_t^a and

⁷ The only place for this sort of belief change to occur in Bayesian analysis is with respect to the prior beliefs. More recent research on non-monotonic logic and belief change permits a more compelling analysis in which beliefs are defeasible at any point in learning process. For overviews see: Gärdenfors (1988), Gärdenfors and Rott (1995), and Schlechta (1997).

Although it is beyond the scope of this paper, it might be noted here that there is a rapidly growing literature in information economics that analyses the impact of, and optimal strategy toward, multiple and/or potentially biased experts. Among the many interesting papers here are: Milgrom and Roberts (1986); Dewatripont and Tirole (1999), Krishna and Morgan (2001); Morris (2001); and Sanchirico (2000).

⁸ In recent years, political scientists have become increasingly interested in the role of collective ideas, beliefs and knowledge in supporting and/or transforming policy. Most of this work has focussed on identifying these effects rather than the media by which they are transmitted, but international agencies clearly play an important role (Murphy, 1994). For useful reviews of the literature on the role of ideas, see: Jacobson (1995) and Berman (2001).

not on the full vector \mathbf{x}_t of policy choices made at time t .⁹ Now we must redefine our notion of *history at k* to be $h^{ka} = \{\mathbf{x}_t, y_t^a\}_{t=1}^{k-1}$, where the vector of policy choices at each t is public and the history of signals/realizations is private.¹⁰ Now each player updates not only with respect to the $y_t^a(\cdot)$ but also taking into account the information of others revealed in their policy choices.

Aoyagi (1998) presents an analysis of essentially this model, showing that all players eventually converge to the same policy. As in the private learning context, social learning will not generally be complete (i.e. while $p_i = p_j$ for some $i, j \in X$, this will only be true for the i finally selected, and $p_j > p_i$ for $j \neq i \in X$). It need not even be the case that $p_i > p_j$ if i is actually selected. Thus, herding occurs with probability 1 and what are essentially information cascades occur. That is, because policymaker's herd, potentially useful collective information is lost. It is important to note that the possibility of cascading or herding on an inefficient policy does not imply that social learning is in any sense worse than private learning. As we have already seen, both of these have equivalents in the private learning context.¹¹ The social learning case embodies two distinctive elements relative to the private learning case. First, every policymaker observes more information at each t , at least until a herd occurs. However, and this is the second point, where the private learner internalises the trade-off between expected current reward and accumulation of information (that is what the Gittins index does), in the social learning context only private learning is internalised in this fashion. That is, there is an information externality.¹²

⁹ This seems, in many ways, a doubtful assumption. However, it is the assumption underlying virtually all econometric research on the link between trade policy and economic performance.

¹⁰ An alternative, closely related, structure would follow the important paper on information cascades by Bikchandani, Hirshleifer, and Welch (1992), in which each country chooses its policy in a fixed order and all countries observe a private signal and the policy choice of all previous movers. The result is not essentially different and the structure above seems somewhat more natural.

¹¹ Aoyagi shows that, if each policymaker observes only a subset of A , then convergence need not occur. The important paper by Smith and Sørensen (2000), while dealing with the standard cascade model, provides useful ideas about directions of generalization for the model discussed above.

¹² Smith and Sørensen (2001), in their welfare analysis of informational herding in a cascade model, develop the notion of a team equilibrium in which agents collectively incorporate this externality. This paper also draws attention to the close relationship between social learning models and private learning models of the sort discussed in the previous sub-section.

The role of experts in general, and internationally organized experts in particular, is not qualitatively different between the private and social learning cases. With respect to initial beliefs, since these must be adopted before social learning occurs the role is identical. In a world with, say, 160 developing countries, the business of carrying out the updating implied by the above model is substantially more complex than that in the one country case. As a result, the need for expertise is that much greater. Krueger (1997) lays particular emphasis on the role of comparative research, especially large-scale projects such as those run by the OECD, NBER, and World Bank, in helping change prior beliefs on the relationship between trade policy and macroeconomic performance. In addition to assisting in the task of evaluating the evidence generated by the multi-country world, the international agencies play at least two additional roles: data collection; and evaluation of private research.

With respect to the first, the World Bank, the IMF, the WTO and UNCTAD, individually and in various joint projects, collect an enormous amount of information, in a relatively standard format, on the trade and industrial policies of the world's countries. These data are used by government researchers as well as private researchers to produce a truly massive quantity of output, much of which is at least potentially relevant to policymakers in industrial and developing countries. One of the tasks performed by the international agencies is the evaluation of this research. In publications like the World Bank's *World Development Report* and the WTO's annual reports, as well as occasional papers on specific topics, the results of this research are presented and evaluated. For industrial countries and even large developed countries, given the extensive economic bureaucracies with a particular focus on trade issues, the latter may not be particularly important. However, given the essentially public nature of data collection, the former is likely to be important even to the richest industrial countries.

2.3 Hierarchical Social Learning

In the previous two subsections international agencies have played a supportive, even subordinate, and essentially passive role in the determination of policy. With the exception of the possibility, noted in footnote 7, that experts might systematically mislead policymakers, their role has been completely positive to this point. In this section we consider the

possibility of a less obviously positive effect of such concentrated expertise. As a result, it will now be important that the expertise is associated with the potential for sanction in a way that we will make clear.

With reference to the literature on information cascades, Gul and Lundholm (1995) make a useful distinction between *statistical cascades* and *reputational cascades*. The framework of the previous subsection permits what are essentially statistical cascades—potentially useful information is lost as a result of herding which results strictly from the rational behaviour of individual agents. By contrast, a reputational cascade is driven by an agency relationship embedded in the sequential decision problem. The central reference here is Scharfstein and Stein (1990), who consider a model of investment by managers in an environment characterized by common prediction error in their decision to make one or another investment. This means that owners, in evaluating the performance of managers, consider both outcomes and whether a given manager did the same thing that other managers did. This creates an incentive for herding, even if there is no convergence in beliefs.

In the policy context, we now sketch a model (a hybrid of the previous two models) into which we introduce an international agency that can provide insurance against bad state realizations, as well as possessing information gathering and analysis capacity.¹³ That is, in t_0 , nature selects (θ) ; the policymakers have initial beliefs $\alpha_a (a \in A)$; and the international agency announces its initial beliefs and the terms of insurance against a bad realization. The model then proceeds as above: policymakers choose a policy $(x_{ta} \in X)$; receive a signal $(y_{ta}(x_{ta};\theta) \in Y)$; if the realization is bad, and they followed the preferred policy of the agency, they get a transfer; and update their prior to α_{ta} . It should be clear that this environment would induce herding, and an information cascade, without inducing convergence of beliefs. In fact, if the insurance were large enough it would induce a herd in t_1 , so there could be no social learning. Unless we are quite sure that the international agency's prior beliefs are accurate, then this sort of institutional environment is clearly harmful.

¹³ The World Bank and International Monetary Fund are more obvious referents here than the WTO.

What is left out of the above model is any role for policy research: the international agency is simply endowed with a fixed initial belief. Thus, we extend this model to incorporate policy research of the sort suggested in Krueger (1997). Suppose that, in addition to the international agency and the policymakers, there is now a finite set of economists. Now suppose that it is the economists, not the policymakers, who observe the vector of policies selected by the policymakers. Note that policymakers and economists observe different things: each of the former observes a country-specific signal, while each of the latter observes the full set of policies adopted in each period. In this extended version of the above model, the international agency forms its prior beliefs exclusively by aggregating the expressed conclusions of the economists.

If there were no international agency, neither the economists nor the policymakers would herd. If the international agency played a purely informational role, publicly reporting an aggregate prior based on the reports of the economists' work, both groups would herd in essentially the same fashion as the policymakers alone herd in the second subsection. However, now suppose that international agency offers to (partially) insure countries against bad realizations if an orthodox policy was pursued in the previous period. An orthodox policy will be a policy such that: 1) it is consistent with the international agency's current belief about the best policy; and 2) a majority of other policymakers are pursuing that policy. This again creates a strong reputational incentive to herd, and an incentive to herd on the agency's preferred policy, with a concomitant loss of socially valuable information.

If, as a result of elective affinity, common training, or some other factors, economists are more prone to herding than policymakers, the existence of an agency that enforces the beliefs of economists will have two effects. To the extent that, because they are aggregating information from a number of countries, their conclusions are more accurate, this should raise welfare by encouraging the adoption of better policies (think of this as the Krueger effect). Because this institutional arrangement encourages rapid herding, information will be lost, increasing the likelihood of a herd on an inferior policy (think of this as an anti-Gittins effect, reflecting that the institution tilts decision making toward current welfare and away from learning).

It seems worth noting that economists do appear quite prone to herding. The case discussed in detail in Krueger (1997) starts from a very tight collective prior on the benefits of first-stage IS. By some time in the 1980s there was an equally tight collective prior on the benefits of XO. What is striking is how little compelling empirical evidence was developed in the interim. As of the time that we are writing this paper, there seems to be a substantial reaction to precisely this fact (e.g. Rodríguez, and Rodrik, 2001). At this point, we do not have a particularly good story to explain how economists shift among quite tight collective priors on such apparently different policy conclusions, but the fact suggests the importance of taking into account the potential social costs implied by the above model.

3. The Political Context of Policy-making

The aim of this section is to add ‘political flesh’ to the concepts outlined in the previous section. Who are the policy-makers and what is the nature of the domestic policy environment that influences their priors and choices? For convenience we use the term ‘government’ to refer to the set of ruling politicians who are policy-makers (this may include senior Civil Servants and advisors), whereas ‘administration’ refers to the bureaucrats that implement policy. As the focus is on policy choice, we will say little about implementation (notwithstanding its evident importance). The policy choices actually made will depend on the way government functions, the power and influence of various interest groups, and the quality of technocrats involved in the process (as it is they who must identify the elements of a strategy to implement the policy chosen by leaders). While the previous section focussed on policy choice, we are particularly concerned with policy change (reform). The discussion will address factors influencing preferences for, and capacity to, change.

In the pure learning by doing model differences in policy choices are due to differences in information. If the signal from a policy choice is associated with a low Gittins index, priors on that policy will be revised downwards and in subsequent periods a different policy may be chosen. In practice, a new government may emerge with new priors or the domestic political environment may change, for example a new influential interest group emerges (e.g. civil

society). While we want to discuss these practical factors as being politically relevant to policy choice, note that the models of the previous section assumed that the objective function for policy-makers is fixed and the same for all, and the external environment (θ) is also fixed. In terms of the model of the previous section, information affects priors and thus affects policy choice, whereas in practice it may be agents with different priors who effect policy change. We will want to interpret the latter (for purposes of linking politics to the learning model) as information that alters priors. Thus, we can think of a set of political actors with policy preferences that provide information to influence the priors of policy-makers (conveniently, the WTO can be treated as one such actor, thereby introducing social learning). Formally, one could model this using bargaining games and negotiating strategies, but that is not necessary for our purposes. We first consider preference formation (of which priors are a component) and then political capacity (which can be interpreted as the influence of political actors on information and priors), finally summarising this in what we will term the policy environment.

To some extent the updating of priors will depend on the nature of the political regime; in practice, all policy-makers do not share the same objective function. At one extreme, ideological regimes will tend to have tight (nearly degenerate) priors, i.e. any updating will tend to occur very slowly if at all. These can change over time (e.g. China liberalised its trade and investment regime in the 1990s without altering the predominant ideological perspective; the same may be true of Vietnam). At the other extreme, liberal technocratic regimes will be inclined to search for the most appropriate policy; they are the most likely to be willing and able to update priors. Most governments are somewhere in between: they have priors, but these can be altered or refined in the face of a changing internal or external environment. Recognising this political reality, we will nevertheless assume that all policy-makers have the same fixed objective function (e.g. maximising the probability of being re-elected) and face the same fixed external environment (θ). Only information can elicit changes in policy choice.

In a technocratic regime the influence of vested interests (i.e. the emphasis attached to information they provide)) tends to be offset by a desire to maximise the performance of the economy (the objective function), and the latter is guided by technical arguments emphasising

management and economic efficiency. Technocratic regimes will embrace liberalisation (policy change) if they are convinced by the arguments (information) that liberal policies will improve economic performance (so priors are updated). Examples include countries as varied as South Korea, Thailand, Mauritius and Costa Rica. In these cases, preferences were conducive to reform and the capacity existed to ensure commitment and implementation. Even a government with a preference for reform will be slow to adopt politically risky policies (formally, this relates to the issue of rewards mentioned in fn4 and to the ‘insurance’ function of global institutions discussed in Section 2.3). The willingness to attempt reform will be constrained by political capacity, the ability to push through reforms in the face of opposition (from vested interests that may be within or associated with government rather than a political opposition).

Frey and Eichenberger (1994) emphasise the importance of attributability, the extent to which the government attracts credit for successful reform or blame for policy failures. Obviously, the government wants to maximise the former and minimise the latter. Some implications of this are discussed in Morrissey (1999); specifically, if the government has a preference for the reform but is unsure about how successful it will be, it can be advantageous to be able to attribute responsibility to external influences. Governments may be more willing to engage in the ‘trial’ of social learning if they will not be blamed for an ‘error’. Global institutions that offer insurance against a bad realization can thereby encourage trial by reducing the cost of error. While the WTO does not offer insurance against bad realizations, one could (for convenience) interpret the ‘rewards from being a member of the club’ (e.g. expectations of foreign investment) in this manner, and this is what we will do.

The ‘age’ of the regime can be quite important. Established regimes tend to have vested interests they will want to protect; this combined with hysteresis renders them less willing to update priors and adopt new policies (i.e. they are less receptive to new information). Many African countries, at least prior to the 1990s, fall into this category – the implementation of policy reforms was very gradual and frequently reversed (e.g. Kenya under Moi). This tendency would also apply to many Asian countries where (certain) policy preferences

change only slowly (e.g. India liberalised gradually in the 1990s; Korea only opened up to FDI in the mid-90s when preparing to join the OECD). One way of depicting this is that governments will stick with x_i as long as $y_i(x_i; \theta)$ is 'satisfactory'. While not modelled in our framework, it seems reasonable to suggest (in the context of, say, a model of bounded rationality) that if performance falls below some trigger level, the government will seek additional information. That is, governments may switch between learning by doing and social learning according to some rule.

New regimes may have weak priors, i.e. they have limited information, or history, on which to judge the value of $p(i)$. In Uganda, Museveni encouraged dialogue within the government and became convinced of the merits of liberalisation (Harvey and Robinson, 1995). This was under pressure from donors and in sight of a reward in the form of aid. The process of democratisation in Africa has given rise to intermediate cases. New governments emerge that, while they may not be very different from the previous regime, are more willing to experiment with policy reform (see Sandbrook, 1996). The transition of power in Tanzania after Nyerere retired was peaceful but only slowly did anything that could be termed a new regime emerge. Nevertheless, the Tanzanian government of the late 1990s was more reformist and market-oriented than that of the mid-80s. The shift to social learning and updating of priors is likely to be a gradual process: governments may be willing to engage in trial but will be reluctant to risk error. This highlights the importance of political capacity.

It is difficult to define political capacity, but the concept encompasses the presence of political actors with varying preferences and different degrees of influence on the choice made (hence an influence on x_i). In this sense political capacity represents the ability of the political system to institute policy evolution and policy change, or to incorporate new sources of information and update priors. This will depend on the nature of decision-making within the government itself and the relative strength of constituencies that support or oppose the direction of policy (the 'political economy' of policy). Preferences (here interpreted as priors that favour a specific policy) and capacity give rise to commitment to reform, but the ability to implement successfully will then depend on administrative capacity and institutional structures.

Much of the discussion of policy reform in developing countries has been concerned with the concepts of ‘ownership’ and/or ‘commitment’ (e.g. Sandbrook, 1996; Leandro *et al*, 1999), with ownership seen as necessary if policies are to be implemented successfully and sustained. Typically, the concept of ownership is used in a loose sense. Sandbrook, for example, argues that ‘ownership’, which is not defined, is necessary for commitment, apparently defined as requiring that ‘the executive authority must be [cohesive and] firmly convinced of the necessity of [reform]’ (Sandbrook, 1996: 5). Leandro *et al* (1999: 288) acknowledge that no clear and unambiguous definition of ownership appears in the literature, and consider it some combination of commitment and capacity to ‘conceive, negotiate and implement reforms’. Thus, an acceptable definition of ownership is that a government supports a policy that it has willingly chosen and designed itself, and openly expresses its commitment to the policy. In this sense, ownership is unlikely to be a necessary condition for effective reform (contrary to the claims made in much of the literature). The pure learning by doing model is one of ownership. This is less evidently the case for social learning *with* external influences. Without digressing into philosophical concerns, one can question how ‘willing’ the choice is of policies proposed by external actors, and thus one can question if commitment necessarily implies ownership.

Commitment can be seen as comprising two elements – preferences and political capacity. Preferences for reform are a sufficient condition to ensure an attempt at implementation (i.e. priors favour the reform), but are not sufficient to guarantee successful implementation, nor to guarantee that the government will make its intentions public. Preferences and capacity give rise to commitment to reform, but the ability to implement successfully will then depend on administrative capability and institutional structures. In this sense, we can define commitment as revealed preference. If a government favours a particular reform and believes it has the political capacity to advocate and try to implement the reform, it is willing to declare the commitment. If a government has a preference for reform but capacity is weak, it may choose not to declare its commitment. If there is no preference for the reform, there is no commitment by this definition (irrespective of what the government may declare).

Thus, we are concerned with commitment and its components – preference with adequate political capacity is the basic requirement for adopting policies. Commitment is especially important for policy change (reform) as it ensures that the ‘new’ policy will be advocated and attempted. It is now possible to consider the role of external influences and information. We introduce one further simplification: the set of policy options X includes the detail of policy design. For example, if $x_i = XO$ is chosen, there are many different ways of achieving this and one of these must also be chosen. We will, for convenience (as a more complex social learning model would be required to incorporate policy design as a sub-set of policies), treat policy choice as referring to the specific details of the chosen policy.

TABLE 1 ABOUT HERE

The discussion above is summarised in Table 1, which also indicates the various ‘dimensions’ that external actors can influence (the first three dimensions relate to preferences). If policy-making within government is relatively open and based on dialogue there is scope for developing new policies and the government may be receptive to external influences. In such cases, it is ‘easier’ to influence priors as policy-makers are more receptive to information. It is also easier to influence choices as policy-makers are more willing to accept technical assistance. At one extreme, external agencies can be ‘blamed’ for requiring governments to adopt unpopular policies (this is shown as helping strengthen capacity). More generally, the government may have priors in favour of the policy, but may have limited capacity to design an appropriate policy and mobilise support for it. External agencies can help with policy advice and technical assistance. General assistance ‘roles’ are listed in Table 1 (D-F), but we concentrate on influences on preferences (A-C).

External influences are often most important in shaping preferences; in our model, they do this by influencing priors. External actors can influence priors in a number of ways. Most obviously, they can provide information that can alter the belief set (τ_i) , including new information on θ that affects how governments interpret history. In other words, external agents may influence how the signal $y_i(x_i; \theta)$ is interpreted and hence the index value attached to $\rho_{ik}(\rho_{i0}, h^k)$, and can provide information on the strategies of others (ξ_j) to facilitate the

correct choice. They can also influence the importance attached to particular issues in the policy agenda. This is related to providing new information on policy options, expanding the policy set (X) that governments consider. In this sense, external agents encourage social learning by facilitating the transfer of policy knowledge.

In our discussion of social learning we noted the usual assumption that policy-makers can observe the actions of others but not the signals received. This is where external agents, especially if they have access to a research base and policy analysis, can play a very important role. They can provide information on the experiences of others and on what appears to have worked. In other words, they can provide an interpretation of the unobserved signals (y_t). This need not always be a ‘good thing’ as if global institutions exhibit herd behaviour they may simply compound information cascades and encourage governments to converge on sub-optimal policies. In this sense, global institutions that disseminate policy knowledge have a responsibility to ensure that they promote the ‘right’ policy option. Aggressive critics of institutions such as the WTO, such as the ‘anti-globalisation movement’, are effectively arguing that the policies are wrong and global institutions are engaged in herd behaviour. The discussion of social learning in sections 2.2 and 2.3 demonstrated that they at least have a point. Herding is the probable outcome and there cannot be a presumption of convergence on the optimal policy (although the likelihood of converging on a policy increases in its probability of yielding a good outcome).

3.1 Policy Environment for Reform

Equipped with the concepts above we can describe the ‘policy environment’ for reform on two dimensions (following Morrissey, 1999). Political commitment can be either low, where the desire and capacity to change policy is weak, or high, where preferences and capacity are strong. Similarly, administrative capability can be weak, such that only a few fairly simple reforms are feasible, or strong, such that the reform programme can be more ambitious. In our context, this capability can be interpreted in respect of the simplicity or complexity of the policy design. A merit of this approach is that the policies of concern, on our case trade and investment liberalisation, can be classified according to whether they are more demanding of political commitment or of administrative capability, or both. This is represented in the four

quadrants of Table 2. An appropriate objective in the design and sequencing of reforms is to keep the range of reforms narrow and increase complexity as commitment and capability are expanded. Successful implementation of even simple narrow reforms can promote commitment and enhance the capability for attempting more complex reforms.

TABLE 2 ABOUT HERE

If a country has neither the desire nor means to commence reform (cell I) then only minimal reforms are likely. In this situation dialogue is important as it can alter priors and increase the range of options considered. This tends to be the situation in the poorest countries, such as Malawi and Myanmar. Multilateral commitments play an important role by ‘locking’ a country into liberalisation (although support and encouragement for implementation will be required). High commitment, or at least preferences for the particular reform, but low capability (cell II) is common in many low-income countries such as Tanzania and Uganda. Governments may have adopted and embarked on the liberalisation process but administrative capability is the constraint that must be relaxed; the problems are design and implementation rather than policy choice. Technical assistance can be very important to enhance capability. Implementing unilateral reforms, especially in non-sensitive sectors, can be a useful way of initiating a reform programme. If benefits are realised, an impetus for reform is created.

The case of strong capabilities but low commitment (cell III) is most evident in established regimes such as Indonesia and Kenya. Even if there is little will to implement significant liberalisation, influence can be brought to bear to encourage initial import liberalisation and policies to promote FDI and exports. The requirement here is to shape preferences, and external influence can be effective. The case of strong capability and high commitment (cell IV), where extensive reforms can be implemented successfully, is quite rare in developing countries. Nevertheless, cell IV is the objective and the notion of the policy environment cautions for gradual sequenced policies that build administrative capacity and political commitment. The potential for extensive reform exists in middle - income countries, and tends to increase with incomes – Brazil, Korea and Mauritius are examples.

4 Two Illustrative Examples

The aim of this section is to provide two brief illustrations of how the proposed framework can be applied to aspects of trade policy choice relating to competition. The first considers information transfer on policy options, specifically anti-dumping as an example of ‘sub-optimal’ policy. The second addresses the design of competition policy within the framework of trade liberalisation. These correspond to the first three dimensions in Table 1. As our focus is on how external agents’ influence preferences, we are not specifically concerned with evaluating the empirical evidence on policy outcomes (y_t). Rather, we are concerned with the role of the WTO in determining the priors of policy-makers.

A specific issue is relevant to the WTO. Within a multilateral framework, it is the sum of individual country actions that determine the global trade environment. The model assumption of θ as fixed should be treated as predicated on the further assumption that all policy-makers implement the policy. This need not be the case in practice, and generally is not the case (suggesting a natural extension to the model). Non-compliance by a minor player has little effect, but non-compliance by major players has significant effects. For example, a widespread perception among developing countries with the implementation of the Uruguay Round agreement is that they opened up their economies but developed countries have not reciprocated (this is particularly true in respect of the *Agreement on Agriculture*). This is a non-trivial concern, at the least because it influences the signal received by a policy-maker. It should be kept in mind that when discussing the ‘real world’ of policies, many assumptions of the model are not maintained.

4.1 Anti-Dumping Actions

A core WTO trade policy is that members should open their markets to non-discriminatory access for other members. More specifically, discriminatory trade policy barriers should not be used to restrict access. Following decades of multilateral trade liberalisation under the GATT, the use of tariff and quota barriers has been greatly reduced, and almost eliminated for developed countries. As countries still want to exercise protection, this ‘has led to the emergence of the GATT-sanctioned instrument of anti-dumping as a key instrument of

protection' (Panagariya and Gupta, 1998: 1003). Economists have tended to be very critical of anti-dumping (AD) in practice; the genuine uses are few and it is used as an instrument of discriminatory protection. Rather than attempt to review the critical literature, a quote from two renowned experts will suffice.

The shortcomings of the AD system at the operational level have been well-documented... The developing and the newly industrialised (NICs), which were often at the receiving end ... are now filing about half of the total number of AD cases. The anti-dumping system is now poised to become the most important trade-restricting device in the post-Uruguay Round world. (Messerlin and Tharakan, 1999: 1251).

The argument is not that no case can be made for contingent protection. It may be true that the real reason contingent protection is permitted is political; in order to sign up to the GATT, members had to be allowed some discretionary protective measure to satisfy domestic political interests. Furthermore, strong economic arguments can be made for the principle of contingent protection, for example where genuine predatory dumping occurs. There are also justifiable health and safety criteria. The point is that, as practised, AD (if not other measures of contingent protection) has been misused and sub-optimal. One should expect that developing countries will want to use AD measures, but 'good' policy is to learn from the misuses so that the AD policies that spread adhere to the principle rather than current (mal)practice.

There can be cases where AD actions are justified, if the good is being sold to the importing country at a price below the 'normal' market price of the exporting country. In practice, the methods used to determine if dumping has occurred are biased towards a positive finding, thus AD is used where really it should not be. Consequently, 'the AD system is basically a flexible tool for preventing imports, whether dumped or not, from causing injury to domestic industry' (Messerlin and Tharakan, 1999: 1258). Quite simply, they are non-tariff measures to protect domestic firms against foreign competition, and they are discriminatory. What is more important from our perspective is, as noted above, they are being increasingly used by developing countries.

The WTO has not, to our knowledge, suggested that developing countries adopt AD legislation. It was not an element of the recommended set of policies (X). But its use has spread. The explanation is that developing countries have experienced AD actions against their exports by the developed countries (especially the EU and US). Prior to about the mid-1990s, NICs accounted for the majority of actions against developing countries. Many NICs then began to use AD themselves, often against other developing countries. Tharakan (2000) notes that over 1987-97, 14 poor and 18 lower middle-income countries were affected by AD actions, about half filed by NICs. Definitive AD measures have been applied against countries such as Bangladesh, Ecuador, Egypt, Kenya, the Philippines and Romania. As with the NICs earlier, countries on the receiving end of such protective barriers will begin to learn to use them.

To use our earlier notation, developing countries observe the policy $x(\text{AD})$ being used by others, and sometimes experience it being used against them. In this way, although they may not observe $y_j(\text{AD})$ they can infer that it yields a positive index value as they see their exports restricted, and they see persistent use. The WTO condones use, therefore countries can view AD as a legitimate element of X . Economic analysts would almost all agree that it is a sub-optimal policy as practiced. In this sense, the policy example of developed countries, especially as it is condoned by the WTO, encourages the spread of a sub-optimal policy. More generally, perceived failure of developed countries to liberalise access to their markets in line with WTO commitments undermines the legitimacy of the WTO as an agent of social learning.

It is worth remarking that Korea is one of the few countries making some attempt to lobby for more transparent criteria for AD within the WTO. Perhaps this is because it is so frequently a victim. Korean exports, as of end 1999, were subject to 87 anti-dumping actions, mostly applying to steel products and petrochemicals (Yang and Kim, 2000: 42-3). Of these, 23 were by the US (mostly concerning steel products), 13 with the EU (mostly electrical or electronic goods) and nine with each of India and South Africa. Korea has brought a dispute to the WTO alleging that the US used an inconsistent and unfair method to calculate the dumping margins for steel. More generally, Korea advocates restrictions on the

use of AD. Park (2000) argues that the US position on anti-dumping (i.e. its aggressive use of such contingent protection) is softening, partly because of so many disputes with major trading partners and partly because the US is increasingly on the receiving end. Korea may argue against AD, but will still find it optimal to adopt such policies itself; resisting the herd is not a dominant strategy.

In the context of our model, the point is that information has spread regarding a sub-optimal policy (in the sense of the AD policy design as practised). The reputation incentive is to herd on the WTO's preferred (contingent protection) policy, where it is preferred by implication (of being condoned). As AD policies are in practice misused, this cannot be assumed to raise welfare by encouraging the adoption of better policies (the Krueger effect is absent). The rapid herding implies that information will be lost and increases the likelihood of convergence on an inferior policy (the anti-Gittins effect is present). The result is to encourage the spread of AD policies based on maximising current welfare (of using countries, where a low weight is attached to consumers) rather than learning what a better AD policy would be. In this sense the WTO encourages, by condoning, discriminatory anti-competitive behaviour.

4.2 Investment Liberalisation and Competition Policy

The intention of the WTO trade policy agenda is to promote competition, yet it often effectively condones anti-competitive actions. 'The tension between the existing trade statutes and competition policy is probably most clearly evidenced by anti-dumping law' (Prusa, 1998: 1021). The reason is that competition law is concerned with protecting consumers, whereas AD is concerned with protecting producers. An important dimension of competition law is to regulate domestic competition, i.e. to regulate firm behaviour so as to prevent abuses of monopoly power. Conventional wisdom, supported by the WTO agenda, is that lowering protection and integrating with global markets is one means of reducing the monopoly power of domestic firms. This is true, but problems can arise if domestic monopolists are simply replaced by foreign firms with monopoly power. In the absence of domestic competition policy or regulatory powers, this can be to the disadvantage of consumers. We consider this issue in respect of foreign investors.

Morrissey (2001) considers the case of trade related investment measures (TRIMs) to examine the broader issue of the liberalisation of investment and the potential impact on developing countries. The TRIMs Agreement is usually interpreted as intended to prohibit the imposition of performance requirements on foreign firms present in the domestic market. However, the WTO Dispute Panel on Indonesia makes it clear that local content requirements are prohibited in themselves, whether applied to domestic or foreign firms. More generally, they argue that the 'term investment measures ... is not limited to measures taken specifically in regard to foreign investment' (Bora, 2001: 5). In this way the principle of National Treatment applied to the TRIMs Agreement implies the prohibition of performance requirements as an instrument of industrial policy. The economic impact of implementing the TRIMs Agreement, when taken in conjunction with broader liberalisation of investment, has significant implications that will constrain the policy options of governments and require issues of competition policy to be addressed.

To return to our notation, say governments face a policy choice $X = \{PR, IL\}$, where PR (performance requirements) refers to the use of industrial policy instruments to regulate the behaviour of firms, and IL (investment liberalisation) implies dismantling such instruments. If evidence on the benefits of trade liberalisation is contested, that on the benefits of investment liberalisation is almost non-existent. It is not a tried and tested policy (which rather begs the question of why it is so confidently promoted by global institutions). Thus, $x(IL)$ is being proposed when the prior on the effect, $y(IL)$, is essentially flat, i.e. there is widespread uncertainty (this helps to explain why institutions have been much less influential in altering preferences towards investment liberalisation as compared with trade liberalisation). The underlying argument is that IL promotes competition. This presupposes that global markets are competitive and/or countries have domestic competition policies to deal with situations where firms (domestic or foreign) have monopoly power.

Global markets are, at best, imperfectly competitive. There are many firms that have significant market power, if not at a global level then at least when they are the dominant entrants in relatively small markets (which applies to the majority of developing countries).

This implies a need for some global competition policy. ‘The main motivation for participating in a multilateral competition policy framework should be that each country recognises its own interest in a strong competition policy and the advantage that it can derive from co-operation in an increasingly globalised economy’ (Meiklejohn, 1999: 1248). However, the prospects for a global competition policy are slim. Multilateral competition policy would be difficult to agree and implement, and the WTO is not suitably placed to address competition policy. ‘The WTO is a system which deals almost exclusively with the actions of national governments whereas competition policy deals primarily with private actions’ (Lloyd, 1998: 1143). As firms, or investors more generally, are private agents the WTO does not have provisions to impose penalties or restrictions on their behaviour. It follows that a multilateral agreement on competition policy, a set of rules binding on all signatories (that would be governments), is not a feasible aim (Lloyd, 1998).

The principal problem facing developing countries is that their legal systems, in particular their capacity to implement competition policy and regulate firms with monopoly power, are limited. Measures to strengthen their capacity to implement effective domestic competition policy are essential (in the language of Section 3, both capacity and capabilities are severely limited in this area). The general principle is that government should promote domestic competition (Morrissey, 2001). The optimal strategy proposed is that governments should adopt measures to promote competition – between local firms and between multinationals investing in the economy. However, this is not the policy advice being proposed by the WTO. The present situation regarding IL is that the policy is being advocated although little is known about the optimal policy nor about the likely outcomes. In such a situation, the likelihood of converging on a sub-optimal policy is rather high.

The relevance of this discussion is two-fold. First, it points to what could be termed policy incoherence (not an issue addressed in the models). While the WTO promotes a pro-competition agenda as the broad policy thrust, individual policies can have (unintended) anti-competitive effects. This arises because while WTO policies impinge on government actions, there are no restrictions on the actions of private agents (such as multinationals), and these may be anti-competitive. Second, and related, the WTO is shifting priors towards specific

policies (IL) about which the effects are highly uncertain. The outcome for policy-makers could be a 'bad' realization, against which the WTO cannot insure. The combination of policy incoherence and bad outcomes can undermine the reputation of the WTO.

5 Conclusions

In pure learning by doing, policy choices are based on information relating to the history of an active policy; there is no information on alternative policies. Only if the policy fails or there is new information to alter priors will there be an incentive to choose a different policy. In the case of social learning, policy-makers can observe the policies chosen by other actors, but the signals those other actors receive is unobserved. External agents (global institutions of knowledge transfer) can influence policy choice by altering priors, providing technical advice or providing information on the (unobserved) effects of the policy choices of others. We have shown that this theoretical basis can be developed to illustrate how institutions such as the WTO and World Bank influence policies of developing countries. We suggest the way they have done this explains why so many countries adopted trade liberalisation policies since the mid 1980s, and why so many are now going further and adopting investment liberalisation policies.

However, social learning theory also predicts that there will tend to be policy herding. There is no presumption that the agents will not converge on the optimal policy, but they may not converge with a strictly positive probability. That is, there is no presumption that they will converge on the optimal policy. It is, however, the case that the likelihood that they will converge on a policy is increasing in its probability of yielding good outcomes. The issue then is how far from the optimum countries are likely to be, and what can be done to minimise the costs of such errors.

Consider two criteria that could be used to evaluate the 'optimality' of the WTO policies, $x(XO)$ and $X(IL)$ – outcomes in terms of economic growth, and outcomes in terms of promoting competition. On neither criterion is there convincing evidence that the policies, as implemented, are optimal. More importantly on the second criterion there are serious

weaknesses in the advice being given. Liberalisation policies are rarely if at all phrased in terms of the implications of imperfections in the competitiveness of global markets (information on which would be a valuable input in determining priors). Developed countries continue to restrict access to their markets, and anti-dumping as practised has been discussed as one of the most clearly anti-competitive practices. Liberalisation of trade, and more especially of investment, also gives rise to issues regarding the anti-competitive behaviour of multinational firms. These issues have not been adequately addressed by global institutions.

Policy advisors and international agencies, that tend to be the major proponents of liberalisation policies in developing countries, should show greater awareness of the prevailing policy environment. Persuasive economic arguments supported by relevant research can alter priors, shape preferences and build commitment to reform. If the aim is to promote competition, institutions should also show awareness of the global competition environment. Many policies are advocated on the implicit assumption of competitive markets. This gives a false impression of the state of nature, and increases the probability that the effect of reforms will be less beneficial than predicted. If so, the commitment of governments to sustaining reforms will be undermined. There is evidence that this is happening at present, and Doha may witness the backlash from developing countries. Institutions such as the WTO have been effective agents of social learning on trade and investment policy reform. But the ultimate success, the sustainability of reforms and of their influence, will be evaluated against outcomes.

REFERENCES

- Aghion, P., P. Bolton, C. Harris, and B. Jullien (1991). 'Optimal Learning by Experimentation' *Review of Economic Studies*; 58:4, 621-654.
- Aoyagi, M. (1998). 'Mutual Observability and the Convergence of Actions in a Multi-Person Two-Armed Bandit Model'. *Journal of Economic Theory*; V.82:2, pp. 405-424.
- Banks, J. and R. Sundaram (1992). 'Denumerable Armed Bandits'. *Econometrica*; 60:5, 1071-1096.
- Berman, S. (2001). 'Ideas, Norms and Culture in Political Analysis'. *Comparative Politics*; 33:2, 231-250.
- Berry, D. and B. Fristedt (1985). *Bandit Problems: Sequential Allocation of Experiments*. London: Chapman & Hall.
- Bewley, T. (1998). 'Knightian Uncertainty'. in D. Jacobs, E. Kalai, and M. Kamien, eds. *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*. New York: Cambridge University Press, 71-81.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). 'A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades'. *Journal of Political Economy*; 100: 5, 992-1026.
- Bora, B. (2001), 'Trade Related Investment Measures' forthcoming in B. Hoekman, P. English and A. Mattoo (eds), *Trade Policy, Economic Development and Multilateral Negotiations*, Washington DC: The World Bank.
- Brezzi, M. and T.L. Lai (2000). 'Incomplete Learning from Endogenous Data in Dynamic Allocation'. *Econometrica*; 68: 6, 1511-1516.
- Camerer, C. and M. Weber (1992). 'Recent Developments in Modeling Preference: Uncertainty and Ambiguity'. *Journal of Risk and Uncertainty*; 5: 4, 325-370.
- DeGroot, M. (1970). *Optimal Statistical Decisions*. New York: McGraw-Hill.
- Dewatripont, Mathias and Jean Tirole (1999). 'Advocates'. *Journal of Political Economy*; V.107-#1, pp. 1-39.
- Easley, D. and N. Kiefer (1988). 'Controlling a Stochastic Process with Unknown Parameters'. *Econometrica*, 56:5, 1045-1064.
- Fagin, R., J. Halpern, Y. Moses, and M. Vardi (1995). *Reasoning About Knowledge*. Cambridge: MIT Press.
- Feldman, M. (1991). 'On the Generic Nonconvergence of Bayesian Actions and Beliefs'. *Economic Theory*; 1:4, 301-322.
- Finger, J.M. and A. Olechowski, eds. (1987). *The Uruguay Round: A Handbook for the Multilateral Trade Negotiations*. Washington, DC: The World Bank.
- Frey, B. and R. Eichenberger (1994), 'The Political Economy of Stabilization Programmes in Developing Countries', *European Journal of Political Economy*, 10:1, 169-190.
- Fudenberg, D. and D. Levine (1998). *The Theory of Learning in Games*. Cambridge: MIT Press.
- Gärdenfors, P. (1988). *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Cambridge: MIT Press.

- Gärdenfors, P. and H. Rott (1995). 'Belief Revision'. in D. Gabbay, C.J. Hogger, and J.A. Robinson, eds. *Handbook of Logic in Artificial Intelligence and Logic Programming. Volume 4: Epistemic and Temporal Reasoning*. Oxford: Oxford Science Publications, pp. 35-132.
- Gittins, J. and D. Jones (1974). 'A Dynamic Allocation Index for the Sequential Design of Experiments'. In J. Gani, ed. *Progress in Statistics*. Amsterdam: North-Holland, 241-266.
- Greenaway, D. and O. Morrissey (1996), 'Multilateral Institutions and Unilateral Trade Liberalisation in Developing Countries', in V. N. Balasubramanyam and D. Greenaway (eds) *Trade and Development: Essays in Honour of Jagdish Bhagwati*, London: Macmillan, pp. 135-160.
- Gul, F. and R. Lundholm (1995). 'Endogenous Timing and the Clustering of Agents' Decisions'. *Journal of Political Economy*; V.103-#5, pp. 1039-1066.
- Haggard, S. (1990). *Pathways from the Periphery: The Politics of Growth in Newly Industrialized Countries*. Ithaca: Cornell University Press.
- Harberger, A. (1993). 'Secrets of Success: A Handful of Heroes'. *American Economic Review*, 83:2, 343-350.
- Harvey, C. and M. Robinson (1995), 'Economic Reform and Political Liberalisation in Uganda', *IDS Research Report 29*, IDS at University of Sussex.
- Hintikka, J. (1962). *Knowledge and Belief: An Introduction to the Logic of the Two Notions*. Ithaca: Cornell University Press.
- Irwin, D. (1989). 'Political Economy and Peel's Repeal of the Corn Laws'. *Economics & Politics*; 1:1, pp. 41-59.
- Jacobsen, J.K. (1995). 'Much Ado About Ideas: The Cognitive Factor in Economic Policy'. *World Politics*; 47:2, 283-310.
- Kiefer, N. (1988/9). 'Optimal Collection of Information by Partially Informed Agents'. *Econometric Reviews*; 7:2, 113-48.
- Krishna, Vijay and John Morgan (2001). 'A Model of Expertise'. *Quarterly Journal of Economics*; V.116-#2, pp. 747-775.
- Krueger, A.O. (1997). 'Trade Policy and Economic Development: How We Learn'. *American Economic Review*; 87:1, 1-22.
- Leandro, J., H. Schafer and G. Frontini (1999), 'Towards a More Effective Conditionality: An Operational Framework', *World Development*, 27:2, 285-300.
- Lloyd, P. (1998), 'Multilateral Rules for Competition Law?', *The World Economy*, 21:8, 1129-50.
- Meiklejohn, R. (1999), 'An International Competition Policy: Do We Need It? Is It Feasible?', *The World Economy*, 22:9, 1233-1251.
- Messlerin, P. and P. Tharakan (1999), 'The Question of Contingent Protection', *The World Economy*, 22:9, 1251-1270.
- Milgrom, Paul and John Roberts (1986). 'Relying on the Information of Interested Parties'. *Rand Journal of Economics*; V.17-#1, pp. 18-32.
- Morris, Stephen (2001). 'Political Correctness'. *Journal of Political Economy*; V.109-#2, pp. 231-265.

- Morrissey, O. (1995), 'Politics and Economic Policy Reform: Trade Liberalisation in Sub-Saharan Africa', *Journal of International Development*, 7:4, 599-618.
- Morrissey, O. (1999), 'Political Economy Dimensions of Economic Policy Reform', in M. McGillivray and O. Morrissey (eds), *Evaluating Economic Liberalisation*, London: Mcmillan, chapter 4.
- Morrissey, O. (2001), 'Investment and Competition Policy in the WTO: Issues for Developing Countries', *Development Policy Review* (forthcoming).
- Murphy, Craig (1994). *International Organization and International Change: Global Governance since 1850*. Cambridge: Polity Press.
- Panagariya, A. and P. Gupta (1998), 'Anti-dumping Duty v Price Negotiations', *The World Economy*, 21:8, 1003-1020.
- Park, S-H (2000), 'Issues of the WTO New Round and APEC's Role', *APEC Study Series 00-06*, Seoul: Korea Institute for International Economic Policy.
- Pressman, E. and I. Sonin (1990). *Sequential Control with Incomplete Information: The Bayesian Approach to Multi-Armed Bandit Problems*. (translated and edited by E. Medova-Dempster and M. Dempster). London: Academic Press.
- Prusa, T. (1998), 'Cumulation and Anti-dumping', *The World Economy*, 21:8, 1021-1034.
- Rodríguez, F. and D. Rodrik (2001). 'Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence'. *NBER Macroeconomics Annual 2000*. Cambridge: MIT Press, 261-325.
- Rodrik, D. (1996), 'Understanding Economic Policy Reform', *Journal of Economic Literature*, XXXIV (1), 9-41.
- Rodrik, D. (1999), *The New Global Economy and Developing Countries: Making Openness Work*, Washington, DC: Johns Hopkins University Press, ODC Policy Essay No. 24.
- Rubinstein, A. (1998). *Modeling Bounded Rationality*. Cambridge: MIT Press.
- Sanchirico, Chris (2000). 'Relying on the Information of Interested--and Potentially Dishonest--Parties'. *American Law and Economics Review*, forth.
- Sandbrook, R. (1996), 'Democratization and the Implementation of Economic Reforms in Africa', *Journal of International Development*, 8:1, 1-20.
- Scharfstein, D. and J. Stein (1990). 'Herd Behavior and Investment'. *American Economic Review*; V.80-#3, pp. 465-479.
- Schlechta, K. (1997). *Nonmonotonic Logics. Basic Concepts, Results, and Techniques*. Lecture Notes in Computer Science, 1187. Berlin: Springer-Verlag.
- Smith, L. and P. Sørensen (2000). 'Pathological Outcomes of Observational Learning'. *Econometrica*; 68: 2, 371-398.
- Smith, L. and P. Sørensen (2001). 'Informational Herding and Optimal Experimentation'. Ms: University of Michigan/University of Copenhagen.
- Sobel, J. (2000). 'Economists' Models of Learning'. *Journal of Economic Theory*; 94: 2, 241-261.
- Srinivasan, T. N. (1998), *Developing Countries and the Multilateral Trading System: From GATT to the Uruguay Round and the Future*, Boulder, Co: Westview Press.

- Tharakan. P. K. M. (1999), 'The Problem of Anti-Dumping Protection and Developing Countries', *WIDER Working Papers No. 198*, Helsinki: UNU/WIDER.
- Whittle, P. (1982). *Optimization Over Time: Dynamic Programming and Stochastic Control, Vol. 1*. Chichester: Wiley.
- Yang, J. and Kim, H-Y. (2000), 'Issues in Korean Trade 1999: Trends, Disputes and Trade Policy', *KIEP Working Paper 00-01*, Seoul: Korea Institute for International Economic Policy.

Table 1 External Influences on Policy Choice

POLICY DIMENSION	EXTERNAL INFLUENCES
A. Priors	Can influence r_t and provide evidence on θ to alter H^k Placing specific concerns high on the agenda
B. Options	Provide and interpret information on options in x_t and y_t Policy advice and knowledge transfer
C. Design	Technical assistance on elements of x_t Disseminate knowledge on policy design
D. Capacity	Support for policy choice strategies, s Taking responsibility for unpopular policies Providing evidence to build support or counter opposition
E. Commitment	Financial support for adopting policies Building policy-making capability
F. Administration	Technical support and assistance

Notes: Discussion in text. The aim is to identify the ‘entry routes’ of external influences on policy choice.

Table 2 **Policy Environment for Pro-Competition Policies**

POLITICAL COMMITMENT		
	Low (weak preference and capacity)	High (strong preference and capacity)
	I (minimal)	II (incremental)
Weak (narrow, simple)	dialogue multilateral 'lock-in' e.g. Malawi, Myanmar	unilateral reforms technical assistance 'non-sensitive' sectors e.g. Tanzania, Uganda
ADMINISTRATION		
	III (erratic)	IV (extensive)
Strong (wide, complex)	some import liberalisation export promotion promote FDI e.g. Indonesia, Kenya	regulatory reforms enacting new laws liberalise all sectors e.g. Brazil, Korea

Notes: Basic structure taken from Morrissey (1999: Table 4.3). Countries listed as examples refer to their policy environment in the 1990s. Obviously, countries can move between cells over time.