Characterizing International Policy Learning

Oliver Morrissey

School of Economics
&
CREDIT
University of Nottingham

Doug Nelson

Murphy Institute
Tulane University
&
Leverhulme Centre for Research on Globalization and Economic Policy
University of Nottingham

This paper was prepared for presentation at the joint Leverhulme Centre for Research on Globalization and Economic Policy/Murphy Institute Conference on the Political Economy of Policy Transfer, Learning and Convergence; at Tulane University 11 and 12 April 2003.

Characterizing International Policy Learning:
A Statistical Decision Theoretic Approach

Most analyses of policy choice and change presume that the policy environment is fixed and known to the policy-maker, so that the optimal policy can be identified.[1] By contrast, many memoirs and journalistic accounts of policy making emphasize uncertainty and learning. Furthermore, most policy-oriented texts represent the lessons they contain as the result of some unspecified learning process. In this paper, we draw on statistical learning theory to characterize several types of learning that appear relevant to issues of policy transfer and convergence.[2] In an effort to keep this discussion manageable, we will frame our discussion in terms of rational learning by Bayesian rational agents.[3]

2 Models of Policy Learning and Knowledge Transfer: Three Views of the Beast[4]

In this section, we sketch three models of policy learning: essentially asocial, decision-theoretic learning; social learning; and hierarchical social learning. Throughout this section we will presume that policy leaning is carried out by a single, rational actor we will refer to as the state. Furthermore, we will also assume throughout this section that all such agents face the same policy problem. We leave to the next section issues of heterogeneity, and domestic and international political economy, as they relate to the characterization of learning.

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

Suppose the state faces a choice between two policy options. We are thinking here of fairly broadly defined options: Classical v. Keynesian macroeconomic policy; fixed v. flexible

[1] This description includes stochastic and dynamic environments as long as the decision-maker knows the laws governing the stochastic or dynamic components of the environment.

[2] It strikes us that Knightian uncertainty plays a particularly significant role here as well, but this is beyond the scope of the current paper, not least because the methods are beyond the abilities of the authors. Bewley (1998) is a useful non-technical introduction. Also see surveys by: Camerer and Weber (1992); Ghirardato (1993); Dekel, Lipman, and Rustichini (1998); and Epstein (2001).

[3] Thus, in addition to Knightian uncertainty, we abstract from sizable bodies of research that seek to characterize and distinguish between knowledge and belief (e.g. Hintikka, 1962); 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 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).

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

exchange rates; import substitution v. export orientation. More broadly we could consider any number of policy states characterized by mixes of policies, e.g.: {Keynesian macro policy, fixed exchange rates, import substitution}, {Keynesian macro policy, fixed exchange rates, export orientation}, et cetera. The focus on a pair of policies is for concreteness only. Suppose for concreteness that the state is considering a tax policy involving either progressive taxation (*P*) or a flat tax rate (*F*). 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 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 $X = \{P, F\}$). 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 $Y = \{good, bad\}$). 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, ρ_{t-1}, using Bayes rule to get ρ_t. We assume that only the element of ρ_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.[5] We suppose that the policymaker's objective is to maximize the expected number of good realizations.[6] Specifically, if we let y_t = good = 1 and y_t = bad = 0, and assume

[5] 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).

[6] 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

that the policymaker applies geometric discounting with discount factor $\delta \in [0,1)$, we can write this objective as:

$$V(\sigma, \rho) = E_{\sigma}\left[\sum_{t=0}^{\infty} \delta^{t} y_{t}(x_{t};\theta)\right]. \qquad (0.1)$$

In the theoretical statistics literature, this is called a *Bernoulli two-armed bandit* problem, with the arms given by the policies.[7] 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, $\rho_{ik}(\rho_{i0}, h^{k})$, 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

---

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.

[7] 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).

$p_i > p_j$ ($i \neq j \in X$), if such a policy exists.[8] 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.

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 θ and that the process of actually carrying out the analysis generating σ 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 is the provision of precisely this sort of expertise. The World Bank, the International Monetary Fund, the World Trade Organization, and many others support substantial research missions, and provide policy support ranging from provision of information, through direct advice, to financial support of policy.

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 $\boldsymbol{\rho}_0$, 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

[8] 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.

systematic argument affecting belief change. For example, Paul Krugman (1999) argues that the idea of globalization has been more effective in inducing certain sorts of market conforming policy choices than globalization itself. That is, change in beliefs of incumbent policymakers produced change in policy. 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).[9] Certainly the usual story about global change in macroeconomic policy commitments involves shocks such a mix: the role of the depression in the adoption of Keynesian policies (Hall, 1989); and the role of stagflation in the turn to more conservative macroeconomic policies. International agencies are among the primary channels for transmitting current policy thinking to the policymaking community.[10] 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.

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

[9] 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).

[10] 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).

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 $y_t^a(x_t^a;\theta)$ are private information to each a. Furthermore, we assume that the $y_t^a(x_t^a;\theta)$ depend only on the x_t^a and not on the full vector $\mathbf{x}_t$ of policy choices made at time t.[11] Now we must redefine our notion of *history at k* to be $h^{ka} = \left\{\mathbf{x}_t, y_t^a\right\}_{t=1}^{k-1}$, where the vector of policy choices at each t is public and the history of signals/realizations is private.[12] Now each player updates not only with respect to the $y_t^a(\bullet)$ 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 $\rho_i = p_i$ for some $i \in X$, this will only be true for the i finally selected, and $\rho_j \neq p_j \; \forall \, j \neq i \in X$). 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.[13] 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.[14]

---

[11] 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.

[12] 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.

[13] 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.

[14] 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

A more natural learning environment might seem to be one in which states gather to discuss the results of their policy experiments. With respect to various aspects of macroeconomic policy, we might see G-7 meetings, the annual Bank-Fund meetings, and WTO ministerial meetings in this light. Specifically, we might imagine meetings of this sort as occasions on which states exchange information about realizations (i.e. the y_t^a) as well as policies (i.e. the x_t). Beginning with papers by Shiller (1995), Ellison and Fudenberg (1995) and Banerjee and Fudenberg (1995), a literature has developed discussing just this sort of environment under the label of "word-of-mouth learning" or "conversation". Banerjee and Fudenberg, in particular, develop a model in which successive generations of agents are able to observe subsets of the choices and payoffs of members of previous generations.[15] In this model, under the common prior assumption, and assuming (as we do throughout this paper) that agents do not misreport if each agent samples at least two

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

---

attention to the close relationship between social learning models and private learning models of the sort discussed in the previous sub-section.

[15] Banerjee and Fudenberg's model relies on each generation being a continuum of agents; however, Smith and Sorenson (1996) develop a finite generalization of the Banerjee/Fudenberg model, with many of the same results.

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.[16] That is, in t_0: nature selects θ ($\in \Theta$); the policymakers have initial beliefs ρ_{0a} ($\forall\ a \in A$); 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 ρ_{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.

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)

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

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.

2.4 Extending Social Learning Models, 1: State Heterogeneity

To this point we have assumed that state preferences are identical, but over quite a wide range of policies this seems like a highly dubious assumption. For example, there is a large body of research on macroeconomic policy suggesting that optimal policy varies with such things as: labor and product market structure; political center of gravity (e.g. social democrat; Christian democrat, liberal); political institutions (e.g. veto points); and degree of openness. Furthermore, while these dimensions plausibly vary, not only is the way they affect policy controversial, but even the mapping of countries into categories is controversial. Thus, there may be substantial

noise in the informational environment. We can draw on an exceptional paper by Smith and Sørenson (2000) to extend the social learning model to this sort of environment. Suppose we go back to the social learning model, but now we assume: *Preferences over policy are heterogeneous*, so we need to include *type* in the characterization of agents; and we include some *crazy types* (i.e. states that adopt a given policy independent of information). In a model of this sort, Smith and Sørenson (2000) show that cascades remain a robust property, but they identify another important phenomenon: *confounded learning*—an interior rational expectations dynamic steady state in which: It may be impossible to draw any clear inference from history even while it continues to accumulate privately-informed decisions; and This incomplete learning informational pooling outcome exists even with unbounded beliefs, when an incorrect herd is impossible

In the social learning context this does not seem to change any of the things that policy advisors do. However, in the hierarchical social learning case, this sort of heterogeneity could raise the costs of "instant herding" by the research authority. This would seem to be one way of representing one line of criticism of the IMFs actions in the Asian crisis.

2.5 Extending Social Learning Models, 2: Political Economy Problems

To this point we have assumed that, while states are purely self-interested, their behavior is non-competitive. Furthermore, we have also assumed that the research authority's only goal is to maximize the number of "good" realizations. Neither of these assumptions command universal assent. Dealing with competitiveness would seem to take us into the enormous realm of game learning, which is simply beyond the scope of this paper. Dealing with the hierarchical learning model in the state heterogeneity context seems most obviously modeled by repeated principal-agent models, which are also beyond our scope. We hope to address these issues in later work.

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.

Banerjee, Abhijit and Drew Fudenberg (1995). "Word-of-Mouth Learning". Harvard University Economics Discussion Paper #22.

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.

Cao, H. Henry and David Hirschleifer (2000). "Conversation, Observational Learning, and Informational Cascades". Ms: Berkeley/Ohio State.

DeGroot, M. (1970). *Optimal Statistical Decisions*. New York: McGraw-Hill.

Dekel, Eddie, Barton Lipman, and Aldo Rustichini (1998). "Recent Developments in Modeling Unforseen Contingencies". *European Economic Review*; V.42-#?, pp. 523-542.

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.

Ellison, Glenn and Drew Fudenberg (1995). "Word-Of-Mouth Communication and Social Learning". *Quarterly Journal of Economics*; V.111-#1, pp. 93-125.

Epstein, Larry (2001). "Subjective Probabilities on Subjectively Unambiguous Events". *Econometrica*; V.69#2, pp. 265-306.

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.

Ghirardato, Paolo (1993). "Non-additive Measures of Uncertainty: A Survey of Some Recent Developments in Decision Theory". *Rivista Internazionale di Scienze Economiche e Commerciali*; V.40-#3, pp. 253-276.

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.

Hall, Peter, ed. (1989). *The Political Power of Economic Ideas*. Princeton: Princeton 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.

Krugman, Paul (1999). "". In Susan Collins and Robert Lawrence, eds., *Brookings Trade Forum–1999*. Washington, DC: Brookings Institution.

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.

Messerlin, 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 (2001). 'Relying on the Information of Interested--and Potentially Dishonest--Parties'. *American Law and Economics Review*, V.3-#2, pp. 320-357.

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.

Shiller, Robert J. (1995). "Conversation, Information, and Herd Behavior". *American Economic Review*; V.85-#2, pp. 181-185.

Smith, Lones and Peter Sørensen (2000). 'Pathological Outcomes of Observational Learning'. *Econometrica*; 68: 2, 371-398.

Smith, Lones and Peter Sørensen (2001). 'Informational Herding and Optimal Experimentation'. Ms: University of Michigan/University of Copenhagen.

Smith, Lones and Peter Srrensen (1996). "Rational Social Learning with Random Sampling". ms: MIT.

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.