International financial crises and public policy: some welfare analysis

By Michael Chui, Prasanna Gai and Andy Haldane of the Bank's International Finance Division.

This article describes a model of financial crisis and explores its implications for public policy. The framework nests the key features of earlier models but is better able to address international architecture questions in a welfare setting. In particular, this framework is used to assess the welfare costs of creditor coordination failure and several recent public policy proposals on reforming the international financial architecture. The costs of creditor coordination failures are found to be high. But policies that improve sovereign liquidity management or that stall creditor runs—such as payments standstills—can mitigate these costs.

Introduction

During the 1990s, a number of emerging market economies experienced well-publicised financial crises: Mexico in 1994/95; South East Asia during 1997; Russia in 1998; and, most recently, Brazil in 1999. On some estimates, the frequency of financial crisis has increased since the 1980s. For example, the World Bank documents 69 instances of 'systemic' crisis since the late 1970s.(1) These crises have afflicted developed and developing countries alike.

There have been a number of recent attempts to measure the output costs of these crises—either the direct fiscal costs (such as the cost of recapitalising banks), or the indirect opportunity costs (of below-trend growth) associated with crisis. These cost estimates are large, often lying between 10% and 20% of annual pre-crisis GDP. The GDP contractions are also often protracted, averaging—on some estimates—more than four years for industrial countries and around three years for emerging economies.(2)

The cost and frequency of financial crises suggests that crisis prevention and crisis resolution are major international public policy concerns. In recent years, this has been reflected in a debate on what has become known as the reform of the 'international financial architecture'.(3) There are many facets of this debate. What are the causes of financial crisis? What public policy measures best address these frictions? And what are the welfare implications of crisis and of different approaches to dealing with them?

Rigorous answers to such questions require an analytical evaluation of the determinants of crises and a quantitative assessment of the welfare implications of policy measures to resolve them. In the next section, some existing analytical models of financial crisis are outlined. The subsequent section sketches an alternative model, which builds on earlier models but which is better able to assess the welfare implications of crisis and public policy intervention.(4) We then assess, from a welfare perspective, various recently proposed public policy measures for averting or resolving crises, including improved sovereign liquidity management and better data disclosure.(5) A final section suggests some research avenues for the future.

Models of financial crisis

Broadly speaking, there have until recently been two strands of the literature on financial crises.(6) Both have tended to focus on models of currency crisis, though the same framework can often be applied generically to liquidity crises in any financial market.

'First-generation' models were motivated by the financial crises of the late 1970s and 1980s, in particular in Latin America. These crises were often preceded by over-expansive macroeconomic (in particular fiscal) policies, which eventually served to prompt the collapse of an exchange rate peg. First-generation models provided an analytical foundation for this phenomenon.(7) In these models, the actual and expected deterioration of fundamentals—say, domestic credit expansion—*pushes* an economy into crisis. Macroeconomic policy in the medium term is inconsistent with maintaining the peg. And with rational expectations about these fundamentals among atomistic investors, the currency collapse is anticipated and so brought forward to today.

(1) See Caprio and Klingebiel (1996, 1999).
(2) For example, Hoggarth, Reis and Saporta (2001), and IMF (1998).
(3) For a summary evaluation, see, for example, Eichengreen (1999).
(4) This draws on Chui, Gai and Haldane (2000).
(5) Drage and Mann (1999) provide a summary of the many recent public policy initiatives aimed at reforming the international financial architecture.
(6) See Flood and Marion (1998).
(7) For example, Krugman (1979), and Flood and Garber (1984).

Such models provide a set of fairly conventional policy recommendations. The best way of lowering the probability of crisis is to pursue prudent monetary and fiscal policies. This policy prescription has clearly been taken on board by many national authorities and the international financial institutions over the past two decades. It is questionable, however, whether monetary and fiscal prudence is a *sufficient* condition to avert a currency collapse, even if it is a *necessary* one. For example, the Asian crisis countries had, in the main, pursued a course of monetary and fiscal prudence ahead of their recent problems. Broader sets of 'fundamentals'—embracing micro-prudential as well as macroeconomic policies—might also need to be included to make sense of these crises.

A second strand of the crisis literature—'second-generation' models—suggests that fundamentals, on any definition, may be neither sufficient nor indeed necessary conditions to determine the likelihood of a crisis.(1) According to second-generation models, crises can occur even with robust fundamentals. The crisis mechanism is instead a coordination failure among creditors, whose expectations and actions are affected importantly by the actions of other creditors. In other words, creditors behave strategically rather than atomistically. If some random event is sufficient to alter adversely these collective expectations, then they can become self-fulfilling. In this way, an economy can be *pulled* into crisis by the actions of fleeing creditors, independently of fundamentals. An economy can be subject to a 'run' in much the same way as a bank. Because countries can be driven into crisis independently of fundamentals, these types of models admit multiple equilibria. There is a range of fundamentals over which an economy is susceptible to liquidity crisis.

There are at least two problems with models of this second-generation variety. First, they are silent on precisely why and when a crisis might strike. The trigger for crisis is a random, unpredictable event—a 'sunspot'. This hinders public policy analysis somewhat because it is difficult for these models to determine what policy measures might best be put in place to avert crisis. Second, with multiple equilibria, it is difficult to conduct meaningful welfare analysis of crisis or of public policy measures to resolve crisis, because equilibrium is not precisely identified.

The two generations of crisis model can be illustrated schematically in a diagram, as in Chart 1. The parameter θ is a summary measure of fundamentals, which are assumed to be random and normally distributed. Below θ_1, the economy is assumed to be 'fundamentally insolvent'. So the zone to the left of θ_1 defines the range of fundamentals over which the economy might be subject to a first-generation crisis, with θ_1 the trigger value for such a crisis. It is the zone of *solvency* or *fundamentals-based* crisis.

The area to the right of θ_2 defines the range of fundamentals over which the economy is solvent irrespective of investors' expectations—the economy is 'strongly solvent' in that it can withstand a run. Between θ_1 and θ_2 lies the range of fundamentals within which self-fulfilling expectations might result in crisis, even though the economy's fundamentals by themselves suggest solvency. With fundamentals in the range $\{\theta_1, \theta_2\}$, an economy is susceptible to *liquidity* or *beliefs-based* (rather than solvency or fundamentals-based) crisis. The fact that this is a range reflects the possibility of multiple equilibria.

Chart 1
A classification of fundamentals

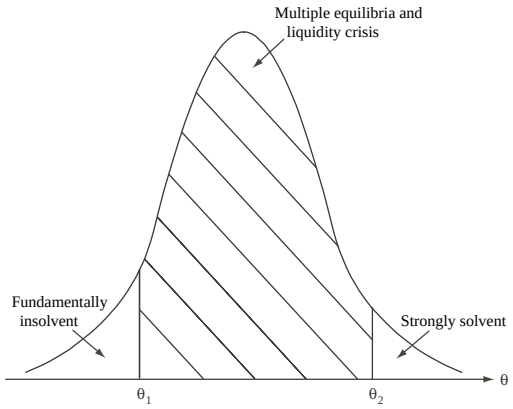

Most recently, a 'third generation' of crisis models has emerged.(2) These aim to mitigate some of the problems of the first two generations of crisis model outlined above. For example, some third-generation models define fundamentals more broadly, to include micro-prudential policies. Accordingly, they allow explicitly for a banking and/or corporate sector, which is subject to frictions such as moral hazard induced by government guarantees. Other models allow for an explicit interaction between fundamentals and beliefs, so that crises can be partly fundamentals-based and partly beliefs-based, rather than one or other in isolation. Most of these models still result in multiple equilibria, which limits their usefulness for policy analysis. By making different informational assumptions, however, some recent crisis models are able to resolve this problem.(3) Using the same informational assumptions, the model developed below has a unique equilibrium and so is more amenable to policy analysis, while at the same time embracing some of the key features of third-generation models.

An alternative model of crisis

We sketch a model of sovereign liquidity crisis that builds on the insights of earlier models, but which addresses specific questions in the architecture debate.(4) There are assumed to be two sets of agent: a single debtor, and a set of creditors that is large in number. The debtor can be thought of as a sovereign borrower (in an emerging market economy) and the creditors a set of international lenders.

(1) For example, Obstfeld (1996).
(2) For example, Krugman (1999), and Chang and Velasco (1999).
(3) See Morris and Shin (1998).
(4) Technical details of the model are given in Chui *et al* (*op cit*).

The debtor invests in an investment project that takes two periods to complete. The project is financed from the debtor's own resource endowment (illiquid assets, E) and from foreign borrowing (L). Both of these inputs are fixed prior to the investment project commencing. The returns to the investment project depend on the factor inputs (E and L) and on the outcome of some random productivity shock. Since productivity is the only random fundamental factor in the model, we denote it θ, as in Chart 1. So gross income from the project (y) is given by:

$$y = \theta (E + L) \qquad \textbf{(1)}$$

Creditors in the model lend to the debtor at an interest rate of r_L. The debt contract between the debtor and creditors is assumed to take a particular form. Specifically, it gives creditors the option to withdraw their funds after one period—that is, before the investment project is completed. In other words, the project is financed with short-term loans that need to be rolled over. If creditors choose to exercise their option and refuse to rollover their loan ('flee'), they face an exit cost, c. If creditors choose to stay for the full two periods ('stay'), then they receive repayment with interest if the debtor is solvent ('repay'), but nothing if the debtor is insolvent and forced to 'default'. The payoff matrix for each representative creditor under the four possible scenarios is shown in Table A.

Table A
Payoff matrix for creditors

			Debtor action	
		Time of payoff	'repay'	'default'
Creditor action	'flee'	Stage 1	$L(1-c)$	$L(1-c)$
	'stay'	Stage 2	$L(1+r_L)$	0

Some of the assumptions underlying this model are worth emphasising because they are important to the outcome of the debtor/creditor game. First, the quantum of foreign lending is fixed up front, together with the other endowments. Second, the model assumes that the monies leaving the project when creditors flee cannot be replaced; there is no secondary market in the debt contracts. Third, the model assumes that the debtor does not default strategically, so will repay if able to do so.(1)

In the model, the debtor's ability to pay depends on the returns to the investment project. This, in turn, depends crucially on two factors: the outcome for the productivity shock, θ; and the proportion of creditors that flee at the intermediate stage, denoted λ. In the event of creditors fleeing, the debtor meets these payments by drawing down its liquid reserve assets, A.(2) But fleeing also causes disruption to the investment project. This can be thought to be the cost of prematurely liquidating the investment project—a half-built bridge or abandoned factory. The marginal cost of this disruption is denoted k. So the solvency constraint facing the debtor at the end of the game, which determines the ability to repay, is:

$$\theta (E + L) - k \lambda L + (1 + r_A)(A - \lambda L) \geq (1 - \lambda) L (1 + r_L) \qquad \textbf{(2)}$$

The left-hand side of equation **(2)** defines the debtor's return on the project at the end of period two, while the right-hand side defines the debtor's debt repayments. Default will only occur when the inequality in equation **(2)** is violated, namely when gross repayments exceed gross income.

We can also use the solvency constraint in equation **(2)** to determine the regions of 'fundamental insolvency' (below θ_1) and 'strong solvency' (above θ_2, where the debtor is solvent irrespective of creditors' expectations and actions), as defined in Chart 1. For example, the trigger value for 'fundamental insolvency', θ_1, is given by:

$$\theta_1 = [(1 + r_L) L - (1 + r_A) A] (E + L)^{-1} \qquad \textbf{(3)}$$

In essence, this insolvency trigger is determined by the debtor's gross gearing and gross reserve asset ratios—or, more generally, by the debtor's net liquidity position. This underlines the importance of adequate liquidity management by borrowers, which is discussed as a public policy measure below.

If we assume that the debtor and creditors all have the same information on the random fundamental, θ, then this model is simply a hybrid first/second-generation model. Below θ_1 the economy behaves as in first-generation models. Between θ_1 and θ_2 the economy behaves as in a second-generation model: there are multiple equilibria and even fundamentally solvent borrowers can be driven to default by a beliefs-based crisis resulting from a creditor coordination failure. Because of this multiplicity of equilibria, the model with perfect information about fundamentals across creditors cannot reach very precise welfare conclusions.

But a slight modification of the basic model helps to sharpen these conclusions. Specifically, assume instead that there is *imperfect* information across creditors about the state of fundamentals. This seems to be a reasonable assumption because in practice common knowledge across creditors is unlikely. With imperfect information across creditors, the model has a unique equilibrium within the fundamentals range $\{\theta_1, \theta_2\}$.(3) Creditors' views converge on a particular equilibrium θ^*. The result is illustrated in Chart 2. Here θ^* denotes the unique value for fundamentals at which crisis is triggered. This lies above θ_1, the value at which a fundamentals-based solvency problem would occur. So the shaded area between θ^* and θ_1 defines the zone where beliefs-based liquidity crises strike.

(1) A different strand of the literature considers the effects of strategic sovereign default (see Eaton, Gersovitz and Stiglitz (1986))—willingness rather than ability to pay.
(2) Which pay a rate of interest r_A.
(3) See Morris and Shin (1998) for a general derivation, and Chui *et al* in the context of the model presented here.

Chart 2
Unique equilibrium

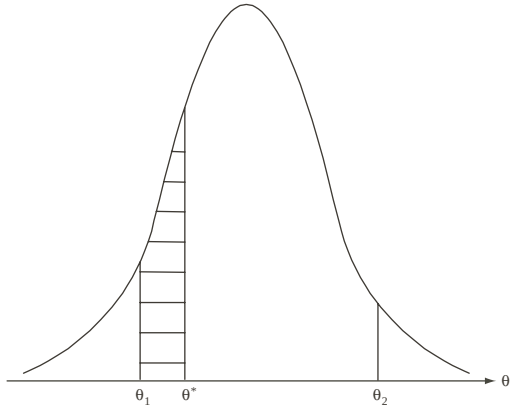

The framework combines elements of beliefs and fundamentals-based crises. Indeed, expectations and fundamentals are not independent, but instead now interact in important ways. For example, the probability of pessimistic investor expectations becoming self-fulfilling is greater, the weaker is the state of fundamentals. So an economy is more susceptible to a creditor run—is more financially fragile—the weaker the underlying macroeconomic outlook. Crises are not the product of poor fundamentals or pessimistic expectations, but a subtle interaction of the two. This better squares with the evidence from recent crises, where both fundamentals and expectations seem to have played a role.(1) It also means that we are better able to define the types of shock that might trigger creditor runs in the first place; they are no longer unpredictable 'sunspots'.

This type of framework also allows us to address public policy questions. For example, it allows us to assess the welfare costs to the debtor of creditor coordination problems. This welfare loss is based on expected income and is related directly to the shaded zone in Chart 2, where beliefs-based crises operate. The model also allows us to assess the welfare implications of different public policy measures, and to decompose these welfare effects into their impact on the probability of a fundamentals-driven (first-generation) crisis and of a beliefs-driven (second-generation) crisis. Specifically, the effect on welfare (W) of a policy change can be decomposed thus:

$$W - W' = \alpha\,[(\theta^{*\prime} - \theta^{*}) - (\theta_1' - \theta_1)] \qquad \textbf{(4)}$$

where ' denotes values of parameters after the policy change.(2) The first term on the right-hand side of equation **(4)** quantifies the impact of the policy change on the probability of a beliefs-based liquidity crisis, and the second term the impact on the probability of a fundamentals-based solvency crisis. The model thus nests both types of welfare friction and allows a decomposition of their effects.

Public policy proposals

In this section we attempt a quantification of the welfare effects of various policy measures, using illustrative values of the model's parameters. Clearly any precise quantification of costs is difficult, as welfare effects are sensitive to the parameterisation of the model. Nevertheless, some broad conclusions can be reached.

The welfare costs of creditor coordination failures

The welfare costs depends importantly on the parameter k, which measures the marginal disruption cost of creditor runs. This parameter is difficult to gauge, so we consider a range of values. When $k = 0.06$—that is, every dollar withdrawn by creditors reduces the return on investment by 6 cents—the welfare costs of creditor coordination failure are around 10% of *ex ante* income, taking illustrative values of the other parameters. If $k = 0.4$, the welfare cost rises to 66% of *ex ante* output. These costs are non-trivial. Although difficult to pin down precisely, they suggest that the welfare effects of creditor coordination failures are significant. Policy measures that reduce creditor panics are potentially valuable from a welfare standpoint.

One possible proposal in this regard is for countries to establish 'country clubs'. These are standing committees of creditors that might serve as a coordination device for creditors' actions. They can also be used to share information between the debtor and creditors and among creditors themselves. The official sector has recently supported the introduction of country clubs by emerging market borrowers.(3) If these helped creditor coordination problems, they could deliver a potentially significant welfare benefit according to the model.

Sovereign liquidity management

A number of theoretical models are based on the belief that (lack of) foreign currency liquidity played a key role in the genesis and propagation of recent financial crises.(4) In parallel work, a number of recent empirical studies have shown that various measures of foreign currency liquidity serve as a good in-sample predictor of crisis—as good, in fact, as most other macroeconomic variables.(5)
Policy-makers have also recently emphasised the importance of prudent liquidity management in averting crisis. The G22 working group on strengthening financial systems, which reported in October 1998, and the recent Financial Stability Forum working group on capital flows, which reported in March 2000, both proposed a risk-management framework for national balance sheet monitoring and management. More specifically, Greenspan (1999) has proposed that, as a rule of thumb, countries should hold enough foreign exchange reserves to cover a year's maturing foreign currency obligations.

(1) For example, Fischer (1999).
(2) See Chui *et al* for a derivation.
(3) For example, the communique by G20 Finance Ministers and Central Bank Governors, October 2000.
(4) For example, Chang and Velasco (1999).
(5) For example, Berg and Pattillo (1999), and Bussiere and Mulder (1999).

To illustrate the point, Table B considers the ratio of short-term debt (with a residual maturity of one year or less) to foreign exchange reserves for a selection of countries that have recently experienced crisis. The ratio is shown on two dates, immediately prior to crisis and at the end of 1999. It is striking that, for each of these countries, the short-term debt/reserves ratio stood at or above one—sometimes considerably so—immediately prior to crisis. Inadequate foreign currency liquidity was a harbinger of currency and in some cases banking crisis. Looking at the ratios more recently, a number of countries, most notably Korea, have clearly made considerable efforts to improve their net liquid foreign currency position, most often by stockpiling reserves. There has been active management of foreign currency liquidity.

Table B
Short-term debt/reserves ratio for crisis countries

	Before crisis	End-1999
Mexico	5.00	0.74
Korea	1.96	0.25
Thailand	1.18	0.41
Indonesia	1.75	0.70
Russia	1.47	0.88
Brazil	0.91	0.95

Sources: BIS and national sources.

The model underlines the importance of these policies in mitigating the costs of crisis.(1) In the model, a lower ratio of short-term debt to reserves has a dual effect. It improves fundamentals, because the trigger for solvency crisis depends importantly on net liquidity. It also reduces the probability of beliefs-based crises, however, by positively shaping expectations of eventual repayment. Chart 3 decomposes the welfare benefits of an improvement in the debt/reserves ratio into these two components. It suggests two conclusions.

Chart 3
Welfare effects of changes in the debt/reserves ratio

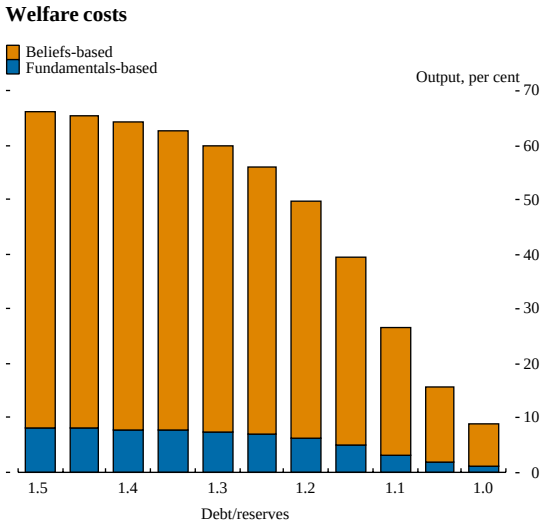

First, the welfare benefits of even relatively modest improvements in the debt/reserves ratio can be sizable. For example, lowering the ratio from 1.5 (around its level for some of the countries in Table B before their crisis) to around 1.0 (as suggested by Greenspan) lowers welfare costs significantly, by a factor of around seven.(2) Second, most of this welfare gain derives from a fall in the probability of beliefs-based crisis. So, naturally enough, improvements in liquidity management serve to reduce significantly the risk of a liquidity run. This would seem to help explain the importance attached to sound country liquidity management by policy-makers in recent years; and why central banks and supervisory agencies more generally have for many years emphasised prudent liquidity management by banks.

Data disclosure and transparency

Improved information provision and transparency have been at the heart of recent attempts to improve the international financial architecture. The G22 working group on transparency and accountability published its report in October 1998. And since then there have been significant strides forward: through the IMF's Special Data Dissemination Standard (SDDS); through codes of transparency for monetary, fiscal and financial policies; through pilot publication of IMF Article IV country reports; and, most recently, through pilot Reports on the Observance of Standards and Codes (ROSCs).(3)

But how do improvements in data availability and transparency affect the welfare costs of crisis? Chart 4 plots these welfare costs against the degree of informational imperfection across creditors—one obvious measure of transparency.(4) The effect of reducing informational imperfections across creditors is to raise welfare. In the stylised example, doubling the precision of creditor information (relative to fundamentals) succeeds in reducing the expected output loss. But the effects are small. Transparency helps, but is no panacea for financial crisis in the model. The point here is a general one. If crisis is rooted in a coordination failure, greater information provision, by itself, need not increase the probability of coordination and hence reduce the probability of crisis.(5) Turning on the lights will not necessarily stop creditors running for the door. Indeed, in theory, transparency could even hasten their exit. When there is perfect information across creditors, we are back to the multiple equilibria, second-generation world described earlier. This set-up can deliver outcomes that may be worse, in a welfare sense, than

(1) The model does not distinguish between domestic and foreign currency liquidity. But if we interpret the debtor as a sovereign and the creditors as foreign lenders, then liquidity is most naturally thought to be foreign currency denominated.
(2) The precise size of the welfare cost depends on the parameterisation of the model. In Chart 3, we set $k = 0.4$.
(3) See King (1999).
(4) More precisely, it takes as the transparency measure the ratio of the variance of information across creditors to the variance of fundamentals.
(5) See Morris and Shin (1999).

Chart 4
Welfare effects of transparency

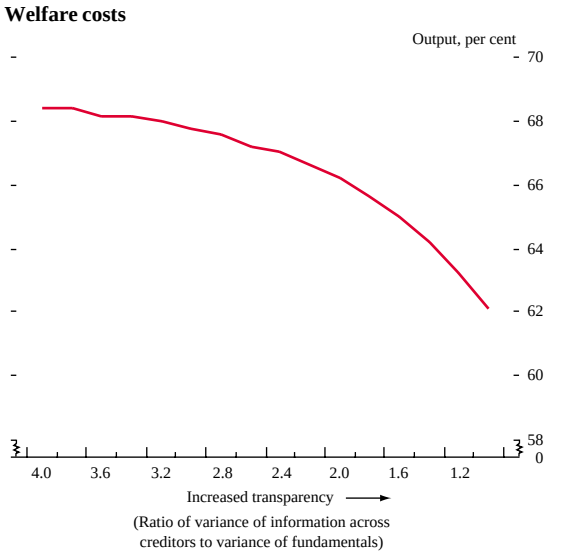

the model with informational imperfections. So, in general, policies seeking greater information disclosure are unlikely, by themselves, to be decisive in averting sovereign liquidity crises induced by creditor coordination failure. Different models and/or different definitions of transparency might, however, deliver a different answer.

Capital controls and payments suspensions

There is an active public policy debate about the efficacy of capital controls.(1) Most of this debate has focused on the effects of imposing restrictions or taxes on capital *in*flows as a means of pre-empting potential liquidity crises, or imposing an orderly queue of otherwise indigestibly large inflows. Chile has operated controls of this type.(2) Among academics, the jury is still out on the usefulness of these types of control. There has been rather less academic and official sector support for controls on capital *out*flows in the face of a liquidity crisis. Payments suspensions, or standstills, can be thought to be the limiting case of controls on capital outflow, where the effective tax rate is unity ($c = 1$). There has been some recent discussion of the case for international payments suspensions, in both academic and official circles.(3) Some have argued that standstills can play a useful role in mitigating the effects of creditor panics.

So can controls on capital outflow—and, in the limiting case, payments standstills—be potentially welfare-enhancing? The model can provide only partial answers to this question, because it does not consider the effects of controls on the initial lending decision nor other potential spillover effects of controls.(4) The model does, however, capture the potential merits of controls in stemming a creditor panic once it has taken hold. In these circumstances, controls or standstills can enforce creditor coordination through quantitative restrictions on portfolio behaviour.

Chart 5 shows the effect on the welfare costs of crisis of changes in the exit tax, c, for a given parameterisation of the model. Small values of the exit tax deliver only small welfare benefits. At high values of the exit tax, however, the welfare gains become substantial. A payments standstill ($c = 1$) completely offsets the *ex post* welfare costs of coordination failure (in this example equal to around two thirds of *ex ante* output). While these quantitative estimates need to be interpreted cautiously, the qualitative implications of the model—that taxes on outflows or payments suspensions can be useful in mitigating the coordination costs of creditor panics—is clear-cut. There is more work to be done on whether the potential (*ex ante* and *ex post*) spillover costs of standstills could offset these benefits.(5)

Chart 5
Welfare effects of exit taxes

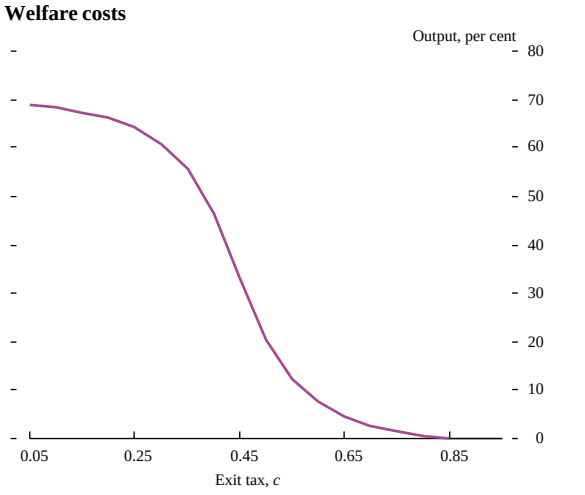

Conclusions

Analytical models can be useful in assessing public policy means of preventing and resolving crisis. They allow quantified, welfare-based policy analysis. We have outlined one particular model of crisis and used it to explore the welfare costs of crisis and the implications of certain policy measures to resolve crisis. The results of this exercise are only as robust as the model from which they are drawn. But that is of course true of all public policy analysis. The merits of the model outlined are that it is spelt out explicitly, builds on existing models of crisis and, as a result, nests their most important features.

(1) See, for example, Cooper (1999).
(2) See, for example, Edwards (1998).
(3) See Eichengreen (2000) and Gai, Hayes and Shin (2000) on the former, and IMF (2000) and Clementi (2000) on the latter.
(4) Gai *et al* (*op cit*) consider this issue in the context of a model of standstills. They find that, although standstills may result in lower *ex ante* lending, they can lead to higher *ex ante* welfare.
(5) See IMF (2000) on this point.

Future research might usefully consider relaxing some of the more restrictive assumptions in the model. First, we assume that the quantum of debt and the form of the debt contract is fixed in advance. Debt size and debt structure might be affected importantly by some of the public policy measures considered here. Second, the model uses a simple measure of welfare and side-steps difficult issues about the distribution of gains and losses between different parties. Third, only a sub-set of the myriad policy proposals currently on the table are considered here. It would be useful to explore these and other extensions in a quantitative, welfare-theoretic, setting.

References

Berg, A and Pattillo, C (1999), 'What caused the Asian crises: an early warning system approach', *mimeo*, IMF.

Bussiere, M and Mulder, C (1999), 'External vulnerability in emerging market economies—how high liquidity can offset weak fundamentals and the effects of contagion', IMF WP/99/98.

Caprio, G and Klingebiel, D (1996), 'Bank insolvencies: cross-country experience', *World Bank Policy and Research WP 1574*.

Caprio, G and Klingebiel, D (1999), 'Episodes of systemic and borderline financial crisis', *mimeo*, World Bank.

Chui, M K, Gai, P S and Haldane, A G (2000), 'Sovereign liquidity crises: analytics and implications for public policy', *Bank of England Working Paper*, No 121.

Chang, R and Velasco, A (1999), 'Financial crises in emerging markets: a canonical model', in Bernanke, B and Rotemberg, J (eds), *NBER Macroeconomics Annual 1999*, Vol 14, MIT Press, Cambridge, Mass.

Clementi, D C (2000), 'Crisis prevention and resolution—two aspects of financial stability', *Bank of England Financial Stability Review*, forthcoming.

Cooper, R N (1999), 'Should capital controls be abolished?', *Brookings Papers on Economic Activity*, pages 89–125.

Drage, J and Mann, F (1999), 'Improving the stability of the international financial system', *Bank of England Financial Stability Review*, Issue 6, pages 40–77.

Eaton, J, Gersovitz, M and Stiglitz, J (1986), 'The pure theory of country risk', *European Economic Review*, Vol 3(30), pages 481–513.

Edwards, S (1998), 'Capital flows, real exchange rates, and capital controls: some Latin American experiences', *NBER Working Paper*, No 6800.

Eichengreen, B (1999), 'Toward a new international financial architecture: a practical post-Asia agenda', Institute for International Economics.

Eichengreen, B (2000), 'Can the moral hazard caused by IMF bailouts be reduced?', Centre for Economic Policy Research.

Fischer, S (1999), 'Reforming the international financial system', *Economic Journal*, Vol 109(459), F557–76.

Flood, R P and Garber, P M (1984), 'Collapsing exchange rate regimes: some linear examples', *Journal of International Economics*, Vol 17, August, pages 1–13.

Flood, R P and Marion, N P (1998), 'Perspectives on the recent currency crisis literature', *IMF Working Paper 98/130*.

Gai, P S, Hayes, S and Shin, H (2000), 'The role of sovereign debt standstills in crisis management', *Bank of England Working Paper*, forthcoming.

Greenspan, A (1999), 'Testimony before the Committee on Banking and Financial Services, May', US House of Representatives.

Hoggarth, G, Reis, R and Saporta, V (2001), 'Costs of banking system instability: some empirical evidence', *Bank of England Working Paper*, forthcoming.

International Monetary Fund (1998), *World Economic Outlook*.

International Monetary Fund (2000), 'Involving the private sector in the resolution of financial crises, attachment 2—standstills', www.imf.org/external/np/pdr/sstill

King, M (1999), 'Reforming the international financial system: the middle way', *Bank of England Financial Stability Review*, Issue 7, pages 203–11.

Krugman, P (1979), 'A model of balance of payments crises', *Journal of Money, Credit and Banking*, Vol 11, pages 311–25.

Krugman, P (1998), 'Saving Asia: it's time to get radical', *Fortune Magazine*, (7).

Krugman, P (1999), 'Balance sheets, the transfer problem and financial crises', *mimeo*, MIT.

Morris, S and Shin, H (1998), 'Unique equilibrium in a model of self-fulfilling currency attacks', *American Economic Review*, Vol 88(3), pages 587–97.

Morris, S and Shin, H (1999), 'Coordination risk and the price of debt', *mimeo*, Oxford University.

Obstfeld, M (1996), 'Models of currency crises with self-fulfilling features', *European Economic Review*, Vol 40, pages 1,037–47.