February 7, 2007

Drift Control of International Reserves

by

Avner Bar-Ilan

Nancy P. Marion

David Perry

Abstract: We develop a model of optimal reserve holdings where the reserve authority controls the upward and downward drift of international reserves and chooses the trigger points that induce changes in drift. We argue that this drift control model better describes the dynamic behavior of reserves than does the popular buffer stock model. We present an innovative mathematical tool for analyzing the drift control based on martingale stopping theory. Our model shows that a reserve authority facing shocks prefers to let the drifts in the reserve path take on most of the burden of adjustment. Further, since the reserve authority has more instruments with drift control than with a buffer-stock strategy, it can manage reserves at significantly lower cost. We observe that a country may wish to accumulate reserves over a long period of time if the cost of changing the drift is high relative to the cost of holding reserves.

Avner Bar-Ilan
Dept. of Economics
University of Haifa
Haifa, Israel

Nancy P. Marion
Dept. of Economics
Dartmouth College
Hanover, NH, USA

David Perry
Dept. of Statistics
University of Haifa
Haifa, Israel

JEL classification: F3, F4 Key words: international reserves, drift control

We would like to thank Joshua Aizenman, Avinash Dixit, Robert Flood and Andrew Rose for helpful comments. We thank Mahmod Kaa'dan for research assistance.

1. Introduction

In many economic applications, the optimal control model adopted is impulse control, also called the (S,s) model or buffer stock model. Impulse control has been applied to problems of capital accumulation, labor demand, price setting and inflation, the demand for money and international reserves, inventories, and many others. The essence of this model is lumpy control of the state variable, usually at its boundaries. For instance, when the state variable is "too high" or "too low" given some payoff function, it is controlled to some intermediate level.

Impulse control assumes the drift of the state variable between controls is exogenous. However, it is often the case that the drift is in fact an endogenous variable that the decision maker can control. For example, by adjusting its output price, a firm can control the rate of expansion of output and thus capital and labor growth. A household can control its rate of expenditures and not just the level where an asset reallocation is required. A central bank can use monetary or exchange-rate policy to influence the rate of accumulation of international reserves.

In this paper we introduce another model of optimal control -- a drift control model -- and present an economic application to the demand for foreign reserves. The new model extends and nests the buffer stock model. The decision maker, whether it is a firm, household, or a central bank, can control both the boundaries of the state variable and its drift. Given this policy, we present an innovative mathematical tool for analyzing the drift control. This analytical tool, based on martingale stopping theory, is used to derive analytic expressions of the payoff functions. Control variables such as drifts and boundaries that minimize the payoff can be straightforwardly found using these payoff functions.[1]

We present the model of drift control in the context of foreign reserve holdings. To the best of our knowledge, this is the first application of a drift control methodology to an economic problem. To date, the most popular model for explaining reserve holdings has been the buffer stock model, illustrated in Figure 1.[2] In the buffer stock model, reserves are initially set at their target level by the reserve authority. They decline smoothly until they reach an exogenous trigger (usually zero), at which time the authority immediately restocks reserves. The model postulates that the authority chooses a target level of reserves to minimize its total expected costs. Total costs consist of the opportunity cost of holding reserves and the adjustment cost incurred at the time of restocking.

While the buffer stock model is simple and straightforward, it has some shortcomings. First, the shark-tooth pattern implied by the buffer stock model-- of a gradual stochastic decline in reserves followed by an abrupt increase at the time of restocking—is not evident in the data. Figure 2 illustrates the reserve pattern for twelve countries over the 1985-2001 period. We find a reserve path of gradual declines and increases that are bounded from below and above. Observe that reserves are not restocked

---

[1] Recent examples of impulse control are Altman (1999) and Bar-Ilan et al. (2004), while examples of drift control are Perry (1997) and Ata et al. (2005).

[2] For examples of buffer stock models, see Heller (1966), Heller and Khan (1978), Frenkel and Jovanovic (1981), Edwards (1983, 1985), Lizondo and Mathieson (1987), and Flood and Marion (2002). For examples of models that stress a precautionary demand for reserves, see Van Wijnbergen (1990), Ben-Bassat and Gottlieb (1992) and Aizenman and Marion (2004).

immediately after they hit a lower bound. In fact, we find that it takes about twice as long for reserves to accumulate from trough to peak as to decline from peak to trough.

A second shortcoming of the buffer stock (BS) model is that the policy required of the reserve authority to generate the shark-tooth pattern is not the policy that countries generally adopt. The BS model literally describes a situation in which an injection of reserves, perhaps from some external source like the IMF, immediately restores reserves to their target level when reserves hit a lower bound. This description may be appropriate for a household that transfers funds from a savings account to a checking account when its checking account approaches a lower bound. It is not the usual response of a reserve authority to a low level of reserves. A more important and common reaction is to change monetary or exchange-rate policy. The policy change alters the reserve drift. It alters financial and current account flows that gradually increase the level of reserves. The reserve authority does not restock reserves to their target level in the abrupt manner characterized by the BS model.

The BS model also implies an unchanged policy environment before and after restocking. A policy setting may lead to a gradual decline in reserves. (Speculation and other shocks may cause sudden changes in reserves around this negative trend.) When reserves reach their lower bound and are restocked at their optimal level, they again start to decline because the same policy environment persists. In practice, it is much more common for policy to change when reserves hit their lower bound or when they reach an upper bound. Indeed, it is this policy change that reverses the direction of the drift in reserves.

In this paper, we develop an alternative model of international reserve holdings, the drift control model. We argue that is more consistent with the dynamic behavior of reserves and the actual tools used by reserve authorities.[3] Its key feature is that the reserve authority controls the average rates of reserve accumulation and depletion (the upward and downward drifts). It does so by its choice of exchange-rate policy or monetary/fiscal policies. Reserves can still change stochastically as upward and downward Brownian motions, but their mean rates of change are under the control of the authority.

The reserve authority also decides *when* to apply a policy adjustment to change the drift in reserves. It sets the lower bound on reserves that will trigger an increase in the drift and the upper bound that will trigger a reduction in the drift. Using our model, we show how to obtain an explicit solution for the expected total discounted cost of managing reserves. We then minimize this cost and derive the steady-state distribution of reserves and their mean level.[4]

The drift control model substantially extends the BS model by allowing the reserve authority to choose four policy variables instead of only one. In addition to controlling the upper bound for reserves as in the BS model, the reserve authority optimally chooses the value of the lower drift rather than having it fixed at some negative value; it optimally sets the value of the upward drift rather than having it fixed infinitely

---

[3] The drift control model also uses a more fully-specified demand for reserves and a richer stochastic structure than simple models of precautionary demand for reserves. See Aizenman and Marion (2004) for a precautionary demand model.

[4] For additional technical references on drift control, see the appendix.

high; and it optimally selects the lower bound on reserves rather than having it exogenously given. The BS model is just a constrained version of drift control.

The drift control model predicts that some countries might allow their reserves to build up over a long period of time. This strategy is optimal when the cost of holding reserves is low but the cost of adjusting policy is high. In this case there may be little or no intervention to adjust downward the drift in reserves. Drift control also allows the reserve authority to put its international reserves policy on “automatic pilot”. By choosing a small, positive drift rate, the authority can keep reserves relatively constant, or rising slowly on average, and thus reduce the need for policy adjustments. As such, the reserves policy will not interfere with other goals the authority might pursue.

The structure of the paper is as follows. Section 2 motivates the drift control model by presenting some empirical evidence on country reserve holdings. Section 3 describes the drift control model. Section 4 provides explicit solutions for the expected total cost of managing reserves, the stationary distribution of reserves, and their mean level. Section 5 finds the parameters that minimize the cost of managing reserves. It shows how the drift levels, the triggers, and the average level of reserve holdings respond to changes in cost parameters and the volatility of the reserves processes. Section 6 looks at the cost savings of managing reserves by drift control rather than a buffer-stock strategy. It also discusses other advantages of drift control and draws some conclusions. Most technical details are relegated to the appendix.

2. Empirical Evidence

To motivate the drift control model, we present two pieces of empirical evidence. The first concerns the dynamic behavior of reserve holdings, specifically the duration of reserve accumulation and depletion. The second illustrates the association between changes in the direction of reserve drift and explicit decisions by the reserve authority to incur the fixed cost of switching the drift by changing policy.

To document the dynamics of international reserves, we examine the empirical properties of monthly international reserves for 145 countries over the sixteen-year period 1985:1 - 2001:6.[5] Figure 2 previously illustrated the reserve dynamics for a few of these countries.

A cursory examination of Figure 2 suggests that the dynamics are somewhat symmetrical with respect to the direction of reserve movements. The duration of the reserve build up and the rate of reserve accumulation are similar to the duration and rate of reserve depletion. This pattern contrasts sharply with the asymmetric dynamics predicted by the BS model. The BS model says the upward drift should be much larger than the downward drift, infinitely larger in fact, and consequently it should take zero time to accumulate reserves relative to the finite time for reserve depletion.

In order to document the degree of symmetry in the dynamics of international reserves, we compute for each reserve cycle of a country the number of months of upward drift (Nu), the months of downward drift (Nd), and the ratio (Nu/Nd).[6] Table 1

---

[5] International reserves are defined as total reserves minus gold. The data are from the IMF's *International Financial Statistics*.

[6] For each country, we first identify the months in which reserves are at a local minimum or maximum. In order to avoid short-term fluctuations and concentrate on long cycles,

displays averages and standard deviations of these variables for various country groupings of our 145-country sample.

Table 1 shows that Nu/Nd is about two. That is, on average it takes about twice as long to accumulate reserves as to deplete them over a cycle. This result holds regardless of how countries are grouped. When we look at industrial countries alone, developing countries alone, developing countries by geographic region, emerging markets, or country groups based on exchange-rate regime choice, Nu/Nd is about two.[7] Recall that the BS model implies that the accumulation time is much shorter (technically zero) than the depletion time. If anything, we find that the reserve cycle is asymmetric in the opposite direction—the accumulation time is longer than the depletion time.

To document evidence of policy changes at local peaks and troughs in international reserve holdings, we use a case study approach. We find the association between policy changes and turning points is particularly pronounced for developed countries. Moreover, these policy changes are motivated primarily or in part by concerns about reserve levels. For developing countries, over the time period we study, the switch from an upward to a downward drift in reserves may also be triggered by an external shock rather than an explicit policy change. In the aftermath of recent financial crises, some developing countries appear to have raised the upper bound for reserve holdings in

---

we smooth the country's reserve data by taking the moving average over a twelve-month period. We define a local minimum (maximum) as a low (high) turning point. Nu is the number of months between a minimum level of reserves and the consecutive maximum. Nd is the number of months between a maximum and the next minimum. The results are similar when we use real reserves.

[7] Income and geographic classifications follow the IMF's *International Financial Statistics*. The emerging market classification comes from the International Finance Corporation. Exchange-rate classifications are based on Reinhart and Rogoff (2004).

their desire to accumulate reserves. To appreciate the nature of this evidence, we present some case studies drawn from countries highlighted in Figure 2.

Italy. In the years 1985-2001, Italy's reserve holdings reached local peaks in 1990:8 and 1997:12, and local troughs in 1992:8 and 1999:3. A few months before the first peak, the Italian government cited Italy's strong balance of payments position and sizeable reserve holdings to justify its decision to adjust the lira's value and adopt the narrower 2.25% band of fluctuation used by fellow members in the European Exchange-Rate Mechanism (ERM). Just after the trough in 1992:8, Italy expressed concern about declining reserves in the face of strong speculation against the lira. It dropped out of the ERM and let the lira float. In the months prior to the 1997:12 reserve peak, the Italian authorities noted that a stronger lira and stronger reserve position had loosened the constraints on monetary and exchange-rate policies. Italy rejoined the ERM and started a period of interest-rate reductions, including one in the month when reserves reached their local peak. Reserves then started to fall, reaching a new low at the start of 1999 when Italy joined the European Monetary Union (EMU).

Netherlands. The Netherlands pursued a de facto monetary union with Germany from 1983 until it joined the EMU in 1999. Reserve holdings in the Netherlands began to increase dramatically in 1992:8, when the Dutch authorities joined Germany in defending weak ERM currencies against a speculative attack. Reserves continued to climb in the face of sustained tight monetary policy through 1994. When reserves peaked in early 1995, the guilder was the strongest currency within the ERM. The strong currency and reserve position relaxed the constraint on monetary policy. The Dutch central bank followed Germany and cut interest rates, and later in 1995, the Dutch central bank went solo and cut interest rates below their German equivalent for the first time since the guilder-DM peg was established as the key instrument of monetary policy in 1983.

United Kingdom. Reserve holdings grew over the 1985-1988 period, reaching a local peak in 1988:12. After that, they declined somewhat but remained relatively stable into 2001. In the year prior to the peak, the U.K. informally shadowed the ERM currencies. When the U.S. stock market crashed in 1987:12, the U.K. initially intervened in foreign currency markets to keep the pound-DM rate stable. In 1988, however, the authorities decided to allow the pound to appreciate against the DM. The change in exchange-rate policy ended the period of reserve accumulation. The policy change was motivated primarily by the desire to control inflation, but the growing stock of reserves contributed to unwelcome pressure on the money supply at a time when inflation control was needed.

Significant changes in U.K. policy around local reserve peaks and troughs are also evident in longer time series data going back to 1960. British international reserve holdings (scaled by GDP or months of imports to eliminate the effects of inflation) reached local reserve peaks in 1971 and 1977 and reserve troughs in 1976 and 1984. The peak in 1971 corresponded to the breakdown of the Bretton Woods arrangement. The massive accumulation of reserves by the British authorities as investors dumped their dollar holdings forced the U.K. to end its defense of the fixed exchange rate and float the pound in August, 1971. In the months preceding the second local peak in 1977, the British again altered policy. Worried that the increase in reserves was putting expansionary pressures on domestic money, the Bank of England abandoned its policy of pegging the pound/dollar rate in July and switched to a basket peg. But by the fall, the authorities felt they needed to take additional action to stem the capital inflows and rapid expansion of reserves that was pushing up the money supply. At the end of October, 1977, the Bank of England announced the pound would float for the time being. Reserve accumulation ended. The government repaid external debts in 1978, in some cases making early repayment, and those actions pushed down reserve holdings.

Reserve lows in the 1970s also triggered changes in British policy. In 1976, pound depreciation and declining reserves forced the UK to request a record $3.9 billion standby credit from the IMF on top of a standby credit made available by the Bank of International Settlements a few months earlier. In January, 1977, an international agreement was reached on a special facility to shield British reserves from a further rundown of sterling holdings by foreign central banks and to help in the gradual reduction in the use of sterling as a reserve currency. The next reserve low was the end of 1984. In January, 1985, the British authorities abandoned their exchange-rate policy of benign neglect. They sought to strengthen the pound by tightening both monetary and fiscal policies. The policies had the effect of reversing the reserve decline.

<u>Korea</u>. Reserve holdings increased in Korea from 1985-1996, reaching a local peak in 1996:6, just days before Thailand announced a devaluation and triggered the Asian financial crisis. The switch to a downward trend in Korean reserves was triggered by an external shock rather than a domestic policy change. There was a significant change in policy around the date of the local trough in Korean reserves, however. In the fall of 1997, turbulence in the Hong Kong market spilled over to Korea and led to Korean reserve losses as the authorities tried to defend the won against panic selling. Finally, in

November, 1997, with reserves at a new low, the Korean authorities decided to float the won and seek a $57 billion standby agreement with the IMF. Since the end of 1997, there has been a rapid and massive build up of Korean reserve holdings, consistent with Korea's stated goals of restoring investor confidence and protecting itself against future crises.

Singapore. Over the years 1985-2001, Singapore reserve holdings have displayed a continual positive drift. This upward trend is particularly pronounced in the data when reserves are scaled by GDP or by months of imports. There was a pause in the upward drift around the time of the Asian financial crisis, when the Singapore Monetary Authority responded to weakening aggregate demand conditions by adopting a wider band within which the Singapore dollar could float. In the fall of 1999, the monetary authority narrowed the exchange-rate band to its pre-crisis width and reserves again began their upward drift.

Brazil. After a period of low and stable reserves, Brazil returned to a fixed exchange rate in the fall of 1991 and reserve holdings began to rise dramatically. Reserves continued to drift upwards until the end of 1998, interrupted by three sharp drops in reserves associated with contagion from external financial crises-- the Mexican crisis (December, 1994), the Asian crisis (July, 1997), and the Russian crisis (August, 1998). After all three financial crises, Brazil reacted by increasing interest rates sharply to reduce capital flight and improve its reserve position. After the Russian crisis, however, capital flight from Brazil continued and reserves fell further, causing the authorities to seek a large support package from the IMF in the fall of 1998. The support package did not stem the pressure and Brazilian reserves continued to fall. Worried about its reserve position, Brazil increased the exchange-rate band on January 13, 1999. When that did not stem the capital flight, it announced it would float the currency. Since 2001, Brazilian reserves have trended upward, aided by periodic increases in domestic interest rates and external financing extended by the IMF to protect Brazil from adverse spillovers from Argentina's crisis ($15 billion in 2001:9 and $30 billion in 2002:8).

To summarize, a drift control model that associates changes in the direction of reserve drift with explicit decisions by the reserve authority to incur the cost of switching the drift by changing policy is strongly supported by the data. A model that allows for finite periods of reserve accumulation and depletion and finite drift rates in both directions also finds support in the data.

3. Describing the Drift Control Model

Consider a reserve authority that holds foreign-exchange reserves R(t), where R(t) denotes the level of reserves at time t. Reserves follow upward and downward Brownian motions. The two reserve drifts are controlled by the reserve authority at a fixed cost per control as described below. One drift is set at γ_0; the other at γ_1, with $\gamma_1 \leq \gamma_0$. Without loss of generality, we assume that the drift control is of the (0, a, b) form, where a and b are trigger points set by the authority that initiate a change in the drift and $0 \leq a < b < \infty$.

At time 0 the reserve level is R(0)=a and the drift is γ_0. The drift is switched to γ_1 the first time reserves hit level b. The drift is controlled back to γ_0 as soon as reserves hit level a again, and so forth. Figure 3 illustrates the dynamics of R(t). When the reserves level is $R \leq a$, the drift is the high drift γ_0; when $R \geq b$, the drift is the low drift γ_1. When $b > R > a$, the drift is γ_0 (γ_1) if R last hit the level a (b).

Even when the drift in reserves follows the higher rate γ_0, a series of bad shocks can push reserves below a. The reserve authority will not intervene, however, unless reserves are pushed below zero. By choosing a high drift value γ_0 and a high target value a, the reserve authority can reduce the probability that reserves will fall to zero, but the authority cannot eliminate the possibility altogether. Consequently, we must consider how the reserve authority responds to this contingency.

When a bad shock pushes reserves below zero, the reserve authority can either do nothing, in which case it incurs costs from having a negative reserve position for a period of time, or the reserve authority can intervene immediately to ensure non-negative reserve holdings, in which case it incurs a cost related to intervention. Modeling the costs associated with a negative reserve position involves unnecessary complications, so we instead assume the reserve authority intervenes to prevent this outcome. Such intervention may take the form of obtaining additional reserves from the IMF or another country, for instance. In Section 5, we interpret this intervention cost as part of the cost of a financial crisis.

To account for the cost of intervention, we define the barrier 0 to be a reflecting barrier. That is, the reserves process R(t) is reflected (only from below) at level 0.

The reflected process is defined as follows. Let $X^0 = \{X^0(t) : t > 0\}$ be a Brownian motion (BM) with drift γ_0, variance σ_0^2, and initial value $X^0(0) = a$. Define L(t) as

$$L(t) = -\min[0, \min_{s \le t} X^0(s)] \tag{1}$$

L(t) is the minimal amount of regulation (foreign reserves injection) necessary to keep the reserves level R(t) from falling below the boundary 0 up to time t.[8]

Define the stopping time T_0 as the first time when the drift is controlled to γ_1. It is defined as

$$T_0 = \min\{t>0: X^0(t)+L(t)\geq b\}. \tag{2}$$

Given T_0, let $X^1 = \{X^1(t): t > T_0\}$ be a BM with drift γ_1, variance σ_1^2, and $X^1(T_0) = b$. The stopping time T_1 is then defined as the first time when the drift is controlled back to γ_0 and is given by

$$T_0 + T_1 = \min\{t > T_0 : X^1(t) \leq a\}. \tag{3}$$

The reserve level $\{R(t): t \geq 0\}$ is a regenerative process with cycle $T_0 + T_1$ such that for $t \leq T_0 + T_1$,

$$R(t) = \begin{cases} X^0(t) + L(t) & t \leq T_0 \\ X^1(t) & T_0 < t \leq T_0 + T_1 \end{cases} \tag{4}$$

Note that $\{R(t): 0 < t \leq T_0\}$ is a one-sided regulated BM with parameters (γ_0, σ_0^2) and that $\{R(t): T_0 < t \leq T_0 + T_1\}$ is a BM with parameters (γ_1, σ_1^2); also, $R(0) = a$ and $R(T_0) = b$.

[8] Technically, the diffusion process R is reflected from below by the local time L(t) which is a non-decreasing, adapted and non-anticipating process with respect to R. The general form of control is (A,a,b,B), where $A \leq a < b \leq B$ and A and B are reflecting barriers. Without loss of generality we assume A=0 and $B \rightarrow \infty$. The latter assumption implies that $\gamma_1 < 0$, otherwise R(t) does not have a well-defined stationary distribution. Given the cost structure assumed here, the optimality of the (A,a,b,B) is intuitively appealing. Formal proof is beyond the scope of this paper. See also Ata et al. (2005) on this point.

Having described the dynamics of drift control, we now model the costs associated with managing reserves. Policy is optimal when these costs are minimized. We identify three types of costs -- the cost of holding reserves, the cost of regulation, and the cost of controlling the drift.

Let the cost of holding reserves be hR(t), where h is the cost of holding \$1 of reserves per unit of time. The expected discounted cost of holding reserves is

$$A_1 = hE_a \int_0^{\infty} e^{-\beta t} R(t) dt, \tag{5}$$

where β denotes the discount rate and $E_z(*) = E(* \mid R(0) = z)$.

Since R(t) is a regenerative process, we can express A_1 in terms of a cycle. Let

$$\theta_0(\beta) = E_a(e^{-\beta T_0})$$

and

$$\theta_1(\beta) = E_b(e^{-\beta T_1}).$$

We show in the appendix that

$$A_1 = h \frac{E_a \int_0^{T_0} e^{-\beta t} R(t) dt + \theta_0(\beta) E_b \int_0^{T_1} e^{-\beta t} R(t) dt}{1 - \theta_0(\beta)\theta_1(\beta)}. \tag{6}$$

Next assume that there is cost k per \$1 of regulation at the boundary 0. There are an infinite and uncountable number of times that reserves hit the boundary level 0. To evaluate the regulation cost, we make use of L(t) defined above. The expected discounted cost of regulation is

$$A_2 = kE_a \int_0^{\infty} e^{-\beta t} dL(t). \tag{7}$$

This cost can be expressed in terms of a cycle as

$$A_2 = \frac{kE_a \int_0^{T_0} e^{-\beta t} dL(t)}{1-\theta_0(\beta)\theta_1(\beta)}. \quad (8)$$

Finally, assume a cost π_1 is incurred every time the drift is switched from γ_0 to γ_1, and a cost π_0 is incurred when the drift is switched from γ_1 to γ_0. The expected discounted cost of controlling the drift is

$$A_3 = \frac{\pi_1\theta_0(\beta) + \pi_0\theta_0(\beta)\theta_1(\beta)}{1-\theta_0(\beta)\theta_1(\beta)}. \quad (9)$$

Adding together the three costs, the total expected discounted cost of managing reserves is therefore

$$C(\beta) = A_1 + A_2 + A_3. \quad (10)$$

This completes the description of the drift control model of international reserves.[9]

4. Solving the Drift Control Model

We now compute the total expected discounted cost, $C(\beta)$ in (10). To do so, we first need to derive explicit solutions for the functions $\theta_i(\beta)$, $E_z \int_0^{T_i} e^{-\beta t} R(t)dt$, i=0,1, and $E_z \int_0^{T_0} e^{-\beta t} dL(t)$ that determine the A's in (10). We present these derivations in the appendix and provide here only a brief description of our methods.

Let x_0 and y_0 be the two roots of the quadratic equation $(\sigma_0^2/2)\alpha^2 - \gamma_0\alpha - \beta = 0$, so that

---

[9] The long-run average cost per unit of time C can be derived from $C(\beta)$ as $C = \lim_{\beta \to 0^+} \beta C(\beta)$.

$$(x_0(\beta), y_0(\beta)) = \frac{\gamma_0 \pm \sqrt{\gamma_0^2 + 2\beta\sigma_0^2}}{\sigma_0^2}. \tag{11}$$

For interest rate $\beta > 0$, both x_0 and y_0 are real numbers and $x_0 y_0 < 0$. Using the definition of (x_0, y_0) in (11), the Appendix proves that

$$\theta_0(\beta) = \frac{y_0 e^{-ax_0} - x_0 e^{-ay_0}}{y_0 e^{-bx_0} - x_0 e^{-by_0}}, \tag{12a}$$

$$\eta_0(\beta) \equiv E_a \int_0^{T_0} e^{-\beta t} dL(t) = \frac{e^{-ax_0 - by_0} - e^{-ay_0 - bx_0}}{x_0 e^{-by_0} - y_0 e^{-bx_0}}, \tag{12b}$$

$$E_a \int_0^{T_0} e^{-\beta t} R(t)\, dt = \frac{[a - b\theta_0(\beta) + \eta_0(\beta)]\beta + \gamma_0(1 - \theta_0(\beta))}{\beta^2}. \tag{12c}$$

The analogous equations for $T_0 < t \le T_1$ are, [10]

$$(x_1(\beta), y_1(\beta)) = \frac{\gamma_1 \pm \sqrt{\gamma_1^2 + 2\beta\sigma_1^2}}{\sigma_1^2}, \tag{13a}$$

$$\theta_1(\beta) = e^{-x_1(b-a)}, \tag{13b}$$

$$E_b \int_0^{T_1} e^{-\beta t} R(t)\, dt = \frac{[b - a\theta_1(\beta)]\beta + \gamma_1(1 - \theta_1(\beta))}{\beta^2}. \tag{13c}$$

The appendix also presents the derivations of the steady-state density of international reserves R and their mean level. Here we provide a brief description of the required calculations.

We first compute the expected amount of regulation (foreign exchange injection) within one cycle, between time 0 and time $T_0 + T_1$. We find that

---

[10] The expressions derived for $T_0 < t \le T_1$ are simpler than those for $0 \le t \le T_0$ because they do not have to account for an upper reflecting barrier. They can also be found in standard references to stochastic processes, such as Harrison (1985), section 3.2.

$$\eta_0(0) = EL(T_0) = \frac{\exp(-ax_0(0)) - \exp(-bx_0(0))}{x_0(0)}, \tag{14}$$

where

$$x_0(0) = \frac{2\gamma_0}{\sigma_0^2} \;; \qquad x_1(0) = \frac{2|\gamma_1|}{\sigma_1^2} \;.$$

The expected levels of T_0 and T_1 are

$$ET_0 = \frac{b - a - \eta_0(0)}{\gamma_0}, \qquad \gamma_0 \neq 0, \tag{15}$$

and

$$ET_1 = \frac{b-a}{|\gamma_1|}. \tag{16}$$

If the reserve authority optimally chooses a positive or negative value of γ_0, it is straightforward to see that (15) and trivially also (16) are positive. In fact, (16) is the well-known formula for the first-passage time. The expression in (15) takes into account the regulation at level 0 that shortens the expected time for reserves to move from the low trigger a to the higher level b.

If the reserve authority optimally chooses $\gamma_0 = 0$ (recall that the possibility of choosing $\gamma_1 = 0$ is excluded; see footnote 7), we use l'Hospital rule to obtain the expected amount of regulation within one cycle and the expected level of T_0:

$$\lim_{\gamma_0 \to 0} \eta_0(0) = b - a, \tag{17}$$

and

$$\lim_{\gamma_0 \to 0} ET_0 = \frac{b^2 - a^2}{\sigma_0^2} \;. \tag{18}$$

Expressions (17) and (18) show that when the authority chooses a zero upward drift, the expected level of intervention when reserves fall below zero is (b-a), and the first-passage time is positive and finite.

The steady-state probability density function of reserves, *f*(R), is given by the time-weighted average of the steady-state densities over the time periods $(0,T_1)$ and (T_1,T_2):

$$f(R)=\begin{cases}\frac{ET_0}{ET_0+ET_1}\frac{\exp(-x_0(0)(a-R))-\exp(-x_0(0)(b-R))}{b-a}, 0\le R\le a\\ \frac{ET_0}{ET_0+ET_1}\frac{1-\exp(-x_0(0)(b-R))}{b-a}+\frac{ET_1}{ET_0+ET_1}\frac{1-\exp(-x_1(0)(R-a))}{b-a}, a\le R\le b\\ \frac{ET_1}{ET_0+ET_1}\frac{\exp(-x_1(0)(R-b))-\exp(-x_1(0)(R-a))}{b-a}, b\le R\end{cases} \quad (19)$$

The mean of the steady-state density of reserves, denoted ER, is

$$ER=\frac{1}{2}(a+b)-\frac{ET_0}{ET_0+ET_1}[\frac{1}{x_0(0)}-\frac{\eta_0(0)}{b-a-\eta_0(0)}\frac{a+b}{2}]+\frac{ET_1}{ET_0+ET_1}\frac{1}{x_1(0)}. \quad (20)$$

When the reserve authority optimally chooses a positive or negative value of γ_0, ER is positive. When $\gamma_0\to 0$, it is straightforward to see that

$\lim_{\gamma_0\to 0}\eta_0/(b-a-\eta_0)=2/(a+b)(1/x_0(0)$, so the bracketed expression in (20) vanishes and ER is again positive and equal to:

$$\lim_{\gamma_0\to 0}ER=\frac{1}{2}(a+b)+\frac{\sigma_0^2}{\sigma_0^2+(a+b)|\gamma_1|}(\frac{1}{x_1(0)}) \quad (21)$$

The mean level of reserves *ER* is therefore always positive, regardless of the chosen value for the upward drift.

5. <u>The Reserve Authority's Response to Shocks</u>

Having obtained analytic solutions for the cost functions and the steady-state mean level of international reserves, we can study how changes in model parameters affect variables of interest. The important model parameters are the various costs of managing reserves and the variances of the two stochastic processes for reserves. The endogenous variables of interest are the minimum total cost of holding reserves, the four variables that the central bank controls (the two drifts and the two triggers), and the mean level of international reserves.

For each set of parameters, we find numerically the vector of drifts and triggers that minimizes the total discounted cost, and we compute the mean of the steady-state density of reserves that is implied by minimizing cost.

We choose the following set of parameters as our baseline. Cost parameters are (h, k, π_0, π_1) = (0.01, 0.4, 0.1, 0.1), variances are (σ_0^2, σ_1^2) = (1, 1), and the interest rate is β=0.04. We first examine the effects of changing cost parameters.

Table 2 shows the change in the minimum total discounted cost, $C(\beta)$, the two drifts and triggers, $(\gamma_0, \gamma_1, a, b)$, and mean reserves, ER, when the cost of holding reserves (h) varies between 0.01 and 0.35. As expected, when holding costs increase, the total cost of managing reserves increases and the reserve authority ends up holding fewer reserves on average.

We can also get some insight into *how* the reserve authority reduces its reserve holdings. Table 2 shows that the reserve authority moderates the upward drift and makes the downward drift more negative. Both adjustments prolong the time that reserves are close to the low trigger and shorten the time that reserves are at the high end. At the same time, the reserve authority reduces the upper trigger monotonically, as expected. It allows the lower trigger “a” to stay relatively constant or even increase, a counterintuitive response. However, the dramatic reductions in the two drifts result in smaller average reserve holdings even when the lower trigger does not decrease.

Interestingly, the decline in average reserve holdings is achieved largely by adjusting the drifts rather than the triggers. To see this in a simple way, consider the separate contributions of trigger changes and drift changes to the reduction in mean reserves from 3.60 to 0.73 in Table 2. If both drifts remain at the values corresponding to h=0.01 (namely 1.333 and -0.312), while the triggers change to the values that correspond to h=0.35 (a=0.548 and b=2.064), average reserves fall from ER=3.60 to ER=2.565. Similarly, if both triggers stay at their initial levels (0.555 and 4.154) while the drifts adjust to the values that correspond to h=0.35 ($\gamma_0 = .012$ and $\gamma_1 = -116.3$), average reserves fall much more, from ER=3.60 to ER=1.378. When the reserve authority sets both drifts and triggers, the drifts take on most of the burden of adjustment to higher holding costs.

Table 3 illustrates what happens when the cost of regulation at the zero boundary (k) is allowed to vary. Even more than in the previous case, we see that drifts, not triggers, take on most of the adjustment burden.

With a higher cost of regulation, the reserve authority tries to reduce the likelihood that reserves will fall below zero and generate this cost. Intuitively, we might expect the central bank to react to higher regulation costs by raising the low trigger and perhaps raising the upper trigger as well. Instead, Table 3 shows that these two trigger levels actually tend to fall. The reason is that an increase in the upward drift and a more gradual decline in the downward drift can push up reserves and prevent them from getting close to zero even when trigger level a gets closer to zero.

When regulation costs rise, average reserve holdings increase. On net, the combination of lower triggers that decrease average reserves and higher drifts that increase reserves serve to increase mean reserves monotonically.

The total cost of managing reserves $C(\beta)$ increases with higher regulation costs, but at a declining rate. When regulation costs are low, increasing them at first raises total cost but after awhile has little additional effect. That is because the probability of reserves hitting the zero boundary and generating a regulation cost becomes very small when the upward drift is increased.

The regulation cost (k) can proxy for some of the cost of a financial crisis. To see this, consider ET_0 and ET_1 for various values of k. When k=0.13, ET_0 = 38.988 and ET_1=0.809. When k=0.45, ET_0=0.682 and ET_1=13.33. When the expected cost of a crisis is small, the monetary authority allows reserves to wander around the lower trigger for a long time and permits reserves to return quickly from the upper trigger to the lower one. Such a strategy minimizes the cost of managing reserves. The opposite holds true when the cost of a crisis is high. In this case, the monetary authority wants to build up reserves quickly and to keep reserves around the upper trigger for a long period of time. Such behavior accords with observed reserve dynamics after crises. Moreover, it is impossible to capture this behavioral response using the more restrictive BS model.

We can also get a sense of how large the expected cost of a crisis might be in our model. The ratio $k\eta_0(\beta)/ER$ measures the expected cost of a crisis as a percentage of the mean level of reserves. This ratio will be a function of all the parameters of the model, not just k. For instance, suppose that k=0.4, the baseline level. For the parameter values that correspond to Table 2, the ratio will vary from 1% when h=0.01 to 80% when h=0.35. The higher ratio describes a case where holding reserves is costly, fewer reserves are held, and the expected cost of a crisis relative to reserves is large. For South Korea, an expected crisis costing between 1%-80% of its reserves in 2002 would amount to 0.25% -18.5% of its GDP in 2002. For China, a crisis costing up to 80% of its 2002 reserves is equivalent to almost 19% of its 2002 GDP. For Mexico, it is equivalent to 6.8% of its 2002 GDP.

The observation that it takes on average twice as long to accumulate reserves as to deplete them yields empirically realistic values for key parameters. Equations (15) and (16) give $ET_0 / ET_1 = |\gamma_1|(b-a-\eta_0(0))/(\gamma_0(b-a))$. For instance, if the accumulation-depletion time ratio is in the range $2 \geq ET_0 / ET_1 \geq 1.5$, then from Table 2 the cost of holding reserves is $5\% \geq h \geq 4.5\%$ per dollar per unit of time and the expected cost of a crisis, $k\eta_0(\beta)$, is about 10% of the average reserve level ER.

Tables 4 and 5 show what happens when there is an increase in the cost of changing the drift, perhaps because it becomes more difficult to modify the exchange-rate or interest-rate policies that influence the drift. When changing the drift is more costly, the reserve authority tries to maintain the same drift for a longer period of time. The reserve authority does so by increasing the gap between the lower and upper triggers. The reserve authority also dampens the upward drift to extend the time until reserves hit their upper bound. Because the upward drift is smaller, the authority also increases the lower trigger to reduce the chance that reserves fall below zero and generate a regulation cost.

The mean level of reserve holdings and the total cost both increase monotonically with higher costs of changing drifts.

Next consider the impact of increased uncertainty on optimal policy. Uncertainty is measured by the variances of the two Brownian motion processes for reserves. Recall that in a BS model, more uncertainty increases "optimal" reserve holdings. The reserve authority raises the upper target, the model's measure of optimal reserves and the one variable under its control.[11] The reason is straightforward. Higher variance in the downward drift increases the probability that reserves will hit their lower trigger level and force the authority to pay the cost of restocking reserves. To avoid paying this restocking cost too often, the policy maker raises the target level of reserves.

With drift control, the reserve authority must deal with two BM processes and hence two variances. Tables 6 and 7 report results when the economy faces increased uncertainty due to increases in σ_0^2 and σ_1^2, respectively, while Table 8 reports the case where the two variances change by the same amount ($\sigma_0^2 = \sigma_1^2$).

When drift control is followed, uncertainty is also costly. An increase in either or both of the variances leads the authority to increase its average reserve holdings and face a monotonically increasing cost of managing them.

While the increase in reserve holdings in the face of greater uncertainty mimics the prediction of the buffer stock model, the mechanism for acquiring reserves differs in the two models. In the buffer stock model, reserves increase because the policy authority raises the target level of reserves, the upper trigger. In the drift control model, the reserve

[11] Frenkel and Jovanovic (1981) and others provide empirical support for this result. Flood and Marion (2002) show that this result may be a statistical artifact if reserves are not normally distributed.

authority has more instruments at its disposal and focuses more attention on changing the drifts in reserve holdings.

Table 6 demonstrates exactly what happens under drift control. As the variance of the upward drift increases, there is a higher probability that reserves will fall below zero and generate regulation costs. To offset this risk, the policy authority finds it optimal to increase the upward drift γ_0 and moderate the downward drift. Unlike the buffer stock model and intuition, the reserve authority does not always raise the triggers in response to more uncertainty.

In contrast to the results of Table 6, the results of Table 7 show that when the variance of the downward drift σ_1^2 increases, it is optimal to accelerate the downward drift while moderating the upward drift. Accelerating the downward drift lessens the time that reserves will follow a BM process with high variance. Moreover, adopting this policy will not increase the chance of incurring regulation costs because the reserve authority always switches to an upward drift when reserves reach level a. The reserve authority also increases both triggers and the gap between them, as in the BS model. The increase in the lower trigger, like the acceleration of the downward drift, allows the authority to reduce the time it must follow a policy of negative drift when that drift has a lot of variance.

Table 8 illustrates what happens when the variances of the upward and downward BM processes increase by the same amount. Again, more uncertainty increases average reserve holdings, and the reserve authority increases its holdings largely by adjusting the drifts rather than the triggers.

6. Comparisons and Conclusion

The buffer stock model of international reserves is a constrained version of drift control. To see this, constrain the lower bound in the DC model to be zero. Set the upward drift in reserves to approach infinity and set the downward drift at some exogenous negative value. Ignore the cost of regulation because reserves never fall below zero; should they hit zero, they are restocked. Set the cost of changing the drift at the upper bound at zero. These constraints transform the drift control model into the buffer stock model. Now reserves drift down at an exogenous rate. When they reach the lower barrier of zero, they are adjusted immediately back to the upper trigger level b. The cost of adjustment when reserves hit the lower barrier is the cost of adjusting the drift upwards when reserves hit the low trigger. The other cost of managing reserves, common to both the drift control and buffer stock models, is the opportunity cost of holding reserves.

Because drift control gives the authority more policy instruments than a buffer stock strategy, it allows the authorities to manage reserves at less cost. Table 9 shows the total cost of managing reserves when the reserve authority can optimize over four controls (Column 2). The total cost is computed for various values of the holding cost and the baseline values for the rest of the exogenous parameters. The table also shows the total cost when the reserve authority can only optimize by choosing the upper trigger b, which is the restocked value of reserves under the buffer stock model (Column 3). The fourth column of Table 9 shows the additional cost (in percentage terms) of managing reserves when the reserve authority must rely on one tool instead of four.

We find that constraining the reserve authority increases the total cost of managing reserves. Moreover, the extra cost is substantial. For the range of holding costs shown in Table 9, the cost of managing reserves is 20-40% higher when the reserve authority has only one tool instead of four. Not only does drift control better capture the dynamic behavior of reserve holdings, it allows a reserve authority to manage reserves at less total cost than more constrained strategies.

Column 5 of Table 9 shows the cost of managing reserves when the authority can optimize over two triggers, a and b. Column 6 shows the extra cost involved in limiting the authority to two instruments instead of the four allowed with drift control. We observe here that even if the reserve authority is allowed to control both boundaries rather than just the upper one as in the BS model, the cost saving from drift control is still substantial. Control over the drifts rather than control over the triggers is key to reducing the cost of managing reserves.

Drift control has some other advantages as well. It makes explicit how the reserve authority responds to crises. The DC model associates a crisis with a low reserve level and identifies an explicit cost that is incurred once reserves fall to this level. The reserves authority can reduce the probability of a crisis, but it cannot eliminate the chance of crisis altogether given the stochastic nature of the environment. When the cost of a crisis increases, the reserve authority responds by trying to reduce the probability of crisis. It resets its instruments, both drifts and triggers. As a result, the country may go for a very long period of time and many reserves cycles without a crisis. Note that the story is quite different in the BS model. There, a crisis occurs on a regular basis, indeed every time reserves hit the lower trigger (corresponding to our level a). Raising the upper trigger can reduce the frequency with which reserves hit the lower trigger, but they will still hit that lower trigger in every reserve cycle.

The drift control model is compatible with long periods of reserve accumulation without intervention by the reserve authority. The average rate of reserve change is under the control of the reserve authority. Depending on cost parameters, the authority might choose a relatively small positive drift and high upper trigger. This choice would allow a long period of reserve accumulation. Such an outcome is optimal if holding reserves is not very costly but reversing course is relatively expensive.

We close the paper with an overview of several issues left for further research. First, our drift control strategy for managing international reserves could be embedded in a macroeconomic model where the government uses a set of tools to achieve multiple objectives. Such a macroeconomic model would make explicit the linkages between exchange-rate policy and the management of international reserves. For example, it might specify how exchange-rate policy endogenously affects the drifts. A macro model could also help identify any range of inaction over which the policy maker puts reserves policy on automatic pilot while it focuses on other priorities.

Second, even without a full-fledged macroeconomic model, it would be interesting to investigate alternative types of drift control policies. For example, international reserves might follow a drift that has both a fixed component and a policy-influenced component, where the cost of changing the drift depends on the current values of the drift and the reserve level.[12] When reserves are relatively low, it may be more costly to turn them around. Consequently, the reserve authority may wish to make more

[12] We thank Avinash Dixit for this suggestion.

frequent and increasingly severe adjustments in the drift as reserves move towards the lower barrier.

Third, the drift control methodology can be applied to other economic problems. For example, it would be useful to explore the implications of having a firm adopt a drift control strategy for capital accumulation when the rate of capital depreciation can be regulated. Another straightforward application is to forestry and natural resources in general where the rate of depletion is under control.

References

Aizenman, Joshua and Nancy Marion (2004). "International Reserve Holdings with Sovereign Risk and Costly Tax Collection," *Economic Journal* 114 (July), 569-591.

Ata, Baris, Michael J. Harrison and Lawrence A. Shepp (2005). "Drift Rate Control of a Brownian Processing System," *Annals of Applied Probability* 15. 1145-1160.

Altman, Eitan (1999). *Constrained Markov Decision Processes,* Chapman and Hall/CRC.

Bar Ilan, Avner, David Perry and Wolfgang Stadge (2004). "A Generalized Impulse Control Model of Cash Management," *Journal of Economic Dynamics and Control* 28 (March), 1013-1033.

Ben-Bassat, Avraham and Daniel Gottlieb (1992). "Optimal International Reserves and Sovereign Risk," *Journal of International Economics*, 345-362.

Chung K.L. and R.J. Williams (1990). *Introduction to Stochastic Integrals,* 2nd edition, Birkhauser: Boston.

Edwards, Sebastian (1983). "The Demand for International Reserves and Exchange Rate Adjustments: The Case of LDCs, 1964-72," *Economica* 50 (August), 269-80.

________________ (1985). "On the Interest-Rate Elasticity of the Demand for International Reserves: Some Evidence from Developing Countries," *Journal of International Money and Finance* 4 (June), 287-95.

Flood, Robert and Nancy Marion (2002). "Holding International Reserves in an Era of High Capital Mobility," *Brookings Trade Forum 2001*, Brookings Institution: Washington, DC, 1-68.

Frenkel, Jacob A. and Boyan Jovanovic (1981). "Optimal International Reserves: A Stochastic Framework," *Economic Journal* 91 (June), 507-14.

Harrison, Michael J. (1985). *Brownian Motion and Stochastic Flow Systems*, Wiley: NY.

Heller, Robert (1966). "Optimal International Reserves," *Economic Journal* 76 (June) 296-311.

Karlin, Samuel and Howard M. Taylor (1974). *A First Course in Stochastic Processes*, 2nd edition, Academic Press: NY.

Lizondo, Jose Saul and Donald J. Mathieson (1987). "The Stability of the Demand for International Reserves," *Journal of International Money and Finance* 6 (September), 251-82.

Perry, David, (1997). "A Double Band Control Policy of a Brownian Perishable Inventory System," *Probability in the Engineering and Informational Sciences* 11, 361-373.

Perry, David and Wolfgang Stadje (1999). "Heavy Traffic Analysis of a Queuing System with Bounded Capacity for Two Types of Customers," *Journal of Applied Probability* 36, 1155-1166.

Reinhart, Carmen and Ken Rogoff (2004). "The Modern History of Exchange-Rate Arrangements: A Reinterpretation," *Quarterly Journal of Economics* 119 (February), 1-48.

Van Wijnbergen, Sweder (1990). "Cash/Debt Buy-Back and the Insurance Value of Reserves," *Journal of International Economics* 29 (August), 123-131.

Appendix

I. Computing the total expected cost of managing reserves

Given that R(t) is a regenerative process with a cycle $T_0 + T_1$, we can write the total expected discounted cost of managing reserves $C(\beta)$ as,

(A.1)

$$C(\beta) = hE_a\int_0^{T_0} e^{-\beta t}R(t)dt + h\theta_0(\beta)E_b\int_0^{T_1} e^{-\beta t}R(t)dt + kE_a\int_0^{T_0} e^{-\beta t}dL(t) + \theta_0(\beta)\pi_1 + \theta_0(\beta)\theta_1(\beta)(\pi_0 + C(\beta)).$$

Grouping the $C(\beta)$'s on the left-hand side of (A.1) gives the expression for the total cost, equation (10) in the text. What remains on the right-hand-side of (A.1) is then the sum of the three costs associated with reserve management—the holding cost, the regulation cost, and the cost of changing the drifts. These three costs are called $A_1, A_2,$ and A_3, respectively, and are equations (6), (8), and (9) in the text.

To compute the functional forms of $\theta_i(\beta)$, $E_z\int_0^{T_i} e^{-\beta t}R(t)dt$, $i = 0,1$, and $\eta_0(\beta) = E_a\int_0^{T_0} e^{-\beta t}dL(t)$, we generalize the technique used in Bar-Ilan et al. (2004), Perry and Stadje (1999), and Perry (1997). The main tool of our analysis is a martingale M(t). It follows from Ito's Lemma (see chapter 5 of Chung and Williams (1990)) that if U is a BM with exponent $\varphi(\alpha) = (1/2)\sigma^2\alpha^2 - \gamma\alpha$, $V = \{V(t): t \geq 0\}$ is an adapted process of bounded variation on finite intervals, and $W = \{W(t): t \geq 0\}$ satisfies $W(t) = U(t) + V(t)$, then

(A.2) $$M(t) = \varphi(\alpha)\int_0^t e^{-\alpha W(s)}ds + e^{-\alpha W(0)} - e^{-\alpha W(t)} - \alpha\int_0^t e^{-\alpha W(s)}dV(s)$$

is a martingale. We use this martingale as follows. Since $\{R(t): t \geq 0\}$ is a regenerative process with cycle $T_0 + T_1$, we divide the cycle into two parts and analyze each of them separately. The first part is $\{R(t): t \leq T_0\}$, which is one sided reflected BM (RBM) with $R(0) = a,\ R(T_0) = b$, drift $\gamma_0 \in (-\infty, \infty)$ and variance $\sigma_0^2 > 0$. The second part is $\{R(t): T_0 < t \leq T_1\}$ which is regular BM with $R(T_0) = b,\ R(T_0 + T_1) = a$, drift $\gamma_1 < 0$ and variance $\sigma_1^2 > 0$.

To use the martingale (A.2) on the first part of the cycle, set $\varphi(\alpha) = \varphi_0(\alpha) = (1/2)\sigma_0^2\alpha^2 - \gamma_0\alpha$, $U(t) = X^0(t)$, $V(t) = L(t) + (\beta/\alpha)t$, and $V(t) = L(t) + (\beta/\alpha)t$, and $W(t) = R(t) + (\beta/\alpha)t$. Then

$$\text{(A.3)} \quad M_0(t) = \varphi_0(\alpha)\int_0^t e^{-\alpha R(s)-\beta s} ds + e^{-\varepsilon a} - e^{-\alpha R(t)-\beta t} - \alpha\int_0^t e^{-\alpha R(s)-\beta s} d(L(s) + \frac{\beta}{\alpha}s)$$

is a martingale. Since $\{R(t): t \leq T_0\}$ is bounded in $[a,b]$, it is straightforward to see that the conditions hold for Doob's optional sampling theorem (Karlin and Taylor (1974)). By setting $E_a M_0(0) = E_a M_0(T_0)$, we obtain

$$\text{(A.4)} \quad \varphi_0(\alpha)E_a\int_0^{T_0} e^{-\alpha R(s)-\beta s} ds = -e^{-\varepsilon a} + E_a e^{-\alpha R(T_0)-\beta T_0} + \alpha E_a\int_0^{T_0} e^{-\alpha R(s)-\beta s} d(L(s) + \frac{\beta}{\alpha}s).$$

Since L(t) increases only when R(t)=0, we have

$$\int_0^{T_0} e^{-\alpha R(s)-\beta s} dL(s) = \int_0^{T_0} e^{-\beta s} dL(s).$$

Rearranging terms in (A.4) and using $d(L(s) + (\beta/\alpha)s) = dL(s) + (\beta/\alpha)ds$ and $R(T_0) = b$ yields

(A.5)

$$(\varphi_0(\alpha)-\beta)E_a\int_0^{T_0} e^{-\alpha R(s)-\beta s}ds = -e^{-\varepsilon a} + E_a e^{-\alpha R(T_0)-\beta T_0} + \alpha E_a\int_0^{T_0} e^{-\beta s}dL(s) = -e^{-\varepsilon a} + e^{-\alpha b}\theta_0(\beta) + \alpha\eta_0(\beta)$$

with $\theta_0(\beta)$ and $\eta_0(\beta)$ defined earlier as $\theta_0(\beta) = E_a(e^{-\beta T_0})$ and $\eta_0(\beta) = E_a\int_0^{T_0} e^{-\beta s}dL(s)$.

Let x_0 and y_0 be the positive and negative roots, respectively, of the quadratic equation $\varphi_0(\alpha)-\beta = (\sigma_0^2/2)\alpha^2 - \gamma_0\alpha - \beta = 0$, so that

(A.6) $$(x_0(\beta), y_0(\beta)) = \frac{\gamma_0 \pm \sqrt{\gamma_0^2 + 2\beta\sigma_0^2}}{\sigma_0^2}.$$

Equation (A.6) is equation (11), section 4. Substituting $\alpha = x_0(\beta)$ and $\alpha = y_0(\beta)$ into equation (A.5) makes the left-hand-side equal to zero. Equation (A.5) therefore yields two equations in the two unknowns, $\theta_0(\beta)$ and $\eta_0(\beta)$:

(A.7) $$\theta_0(\beta) = \frac{y_0e^{-ax_0} - x_0e^{-ay_0}}{y_0e^{-bx_0} - x_0e^{-by_0}},$$

(A.8) $$\eta_0(\beta) = \frac{e^{-ax_0-by_0} - e^{-ay_0-bx_0}}{x_0e^{-by_0} - y_0e^{-bx_0}}.$$

Equations (A.7) and (A.8) are equations (12a) and (12b) in section 4.

Now substitute equations (A.7) and (A.8) into (A.5), divide both sides by $\beta - \varphi_0(\alpha)$, take the derivative with respect to α and set $\alpha = 0$. This yields

(A.9) $$E_a\int_0^{T_0} e^{-\beta t}R(t)\,dt = \frac{[a - b\theta_0(\beta) + \eta_0(\beta)]\beta + \gamma_0(1-\theta_0(\beta))}{\beta^2},$$

which is equation (12c).

The solution technique for the second part of the cycle is similar and yields equations (13a)-(13c).

II. Computing the steady-state density of reserves

The steady-state density of reserves $f(R)$ is obtained as follows. Let $f^0(R)$ [$f^1(R)$] be the conditional steady-state densities of R given that the second [first] part of each cycle is deleted. Then, $f(R)$ is a weighted average, such that

$$\text{(A.10)} \quad f(R) = \frac{ET_0}{ET_0 + ET_1} f^0(R) + \frac{ET_1}{ET_0 + ET_1} f^1(R).$$

In terms of Laplace Transforms, denoted LT, equation (A.10) is

$$\text{(A.11)} \quad \tilde{f}(R) = \frac{ET_0}{ET_0 + ET_1} \tilde{f}^0(R) + \frac{ET_1}{ET_0 + ET_1} \tilde{f}^1(R),$$

where $\tilde{f}(R)$, $\tilde{f}^0(R)$, and $\tilde{f}^1(R)$ are the LTs of $f(R)$, $f^0(R)$, and $f^1(R)$, respectively.

To compute the LT $\tilde{f}^0(R)$, let $\beta = 0$ in equation (A.5) and divide both sides by $ET_0\varphi_0(\alpha)$. Since $\theta_0(0) = 1$, this yields

$$\text{(A.12)} \quad \frac{E_a \int_0^{T_0} e^{-\alpha R(s)} ds}{ET_0} = \frac{-e^{-\alpha a} + e^{-\alpha b} + \alpha\eta_0(0)}{ET_0\varphi_0(\alpha)}.$$

By renewal theory the left-hand-side of (A.12) is $\tilde{f}^0(R)$. The solution for $\eta_0(0) = EL(T_0)$ can be computed from equation (12) and is given by equation (14). Finally, by letting $\alpha \to 0$ in equation (A.12) and using l'Hospital rule we get

$$\text{(A.13)} \quad ET_0 = \frac{b - a - \eta_0(0)}{\gamma_0},$$

which is equation (15) in the text. Thus,

$$\text{(A.14)} \quad \tilde{f}^0(R) = \frac{-e^{-\alpha a} + e^{-\alpha b} + \alpha\eta_0(0)}{(1/2)\alpha^2\sigma_0^2 - \alpha\gamma_0} \frac{\gamma_0}{b - a - \eta_0(0)},$$

where $\eta_0(0)$ is in equation (14).

In a similar manner we get,

$$\text{(A.15)} \qquad ET_1 = \frac{b-a}{|\gamma_1|},$$

and

$$\text{(A.16)} \qquad \tilde{f}^1(R) = \frac{e^{-\alpha a} - e^{-\alpha b}}{(1/2)\alpha^2\sigma_1^2 + \alpha|\gamma_1|}\frac{|\gamma_1|}{b-a}.$$

Substituting (A.14) and (A.16) into (A.11) gives $\tilde{f}(R)$, the Laplace transform of the steady-state density of reserve holdings.

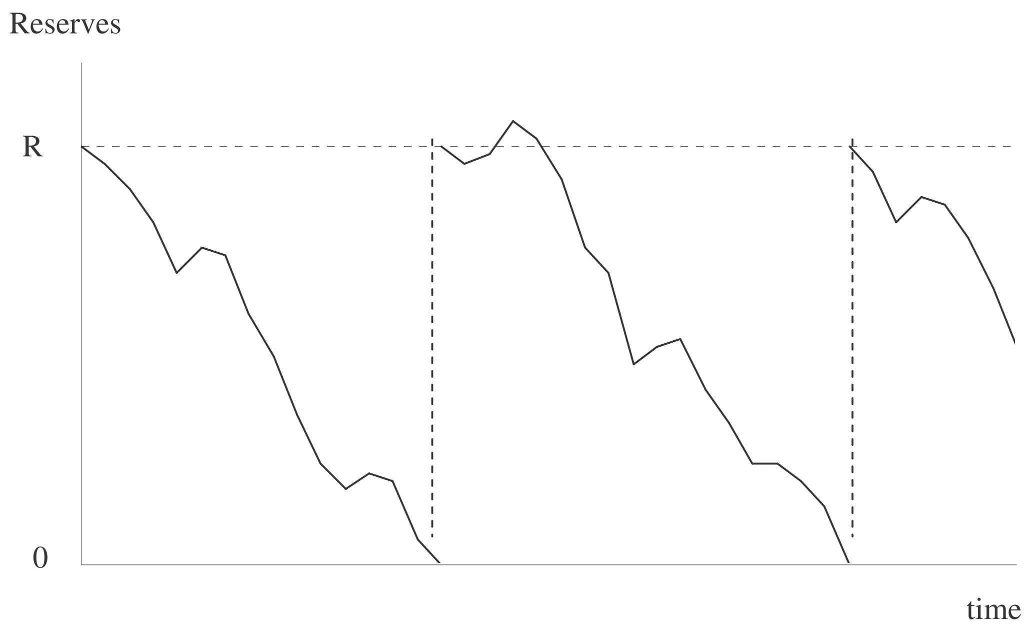

Figure 1: Buffer stock model of international reserves

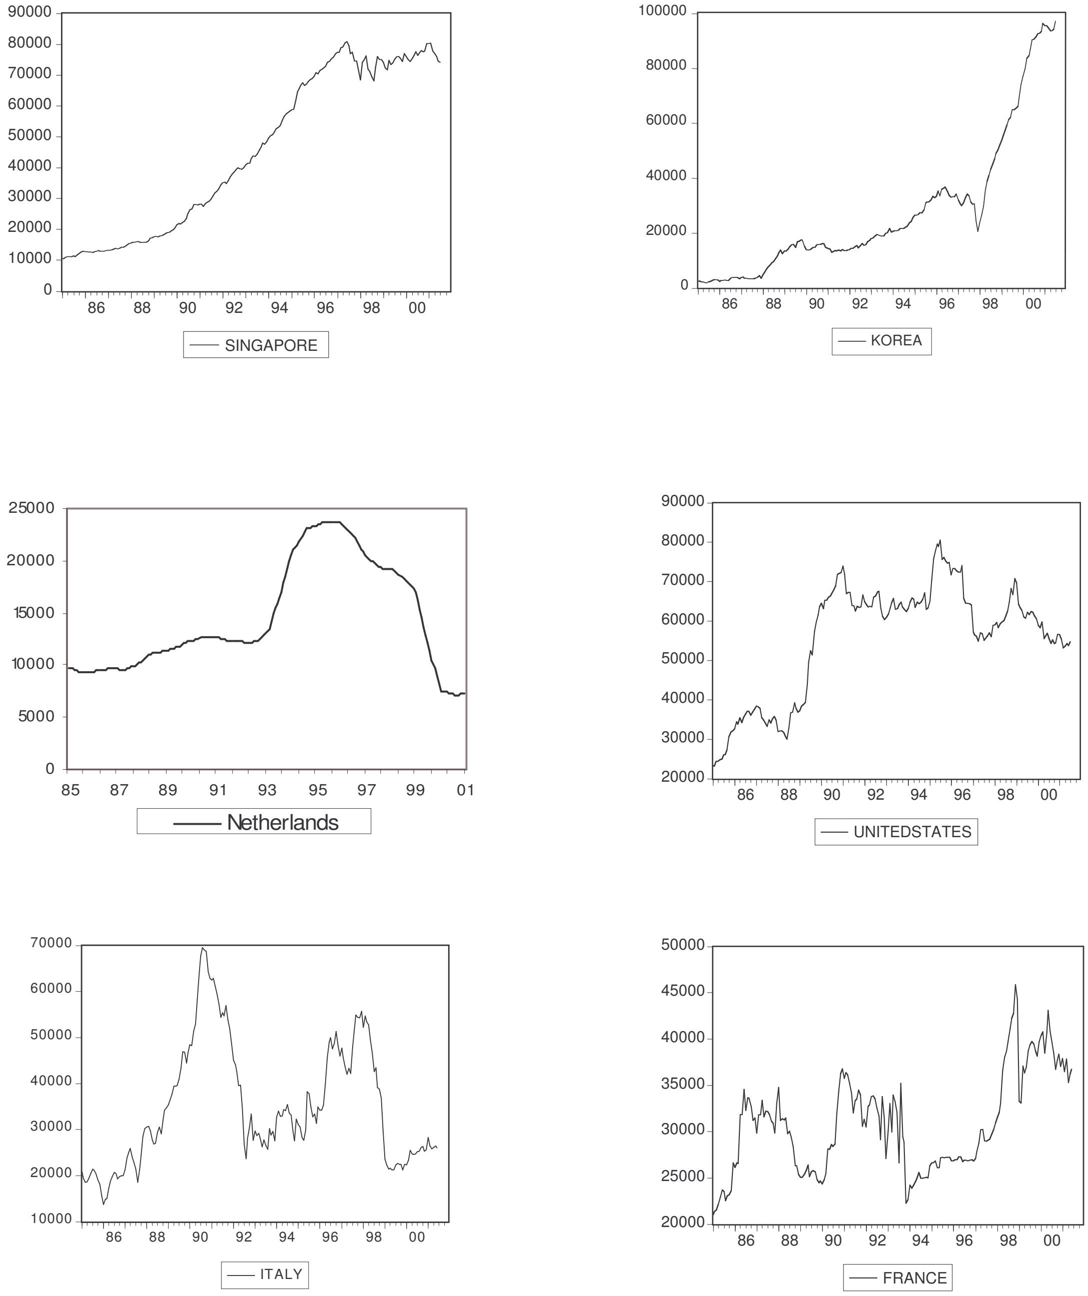

Figure 2: Reserves in Individual Countries

Note: The vertical axis measures reserves in millions of dollars; the horizontal axis is the year.

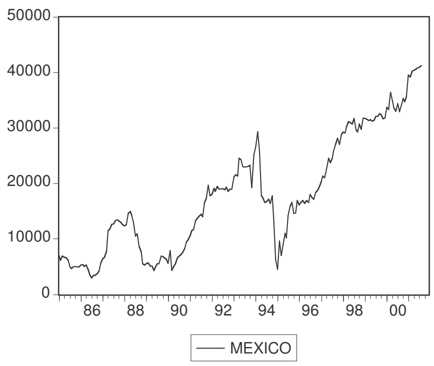
50000
40000
30000
20000
10000
0
86
88
90
92
94
96
98
00
MEXICO

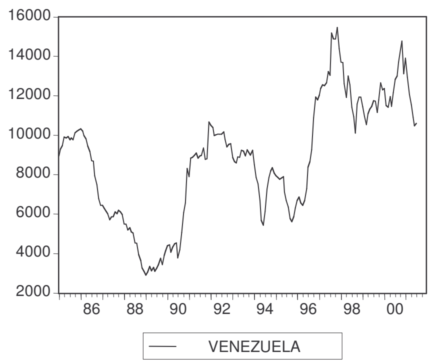
16000
14000
12000
10000
8000
6000
4000
2000
86
88
90
92
94
96
98
00
VENEZUELA

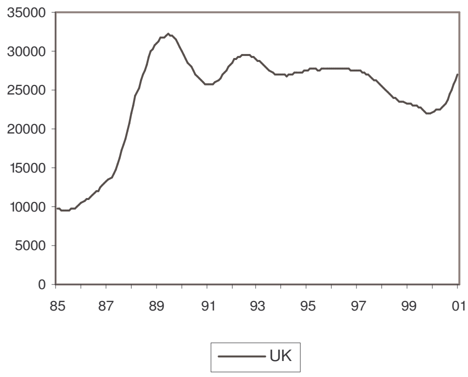
35000
30000
25000
20000
15000
10000
5000
0
85
87
89
91
93
95
97
99
01
UK

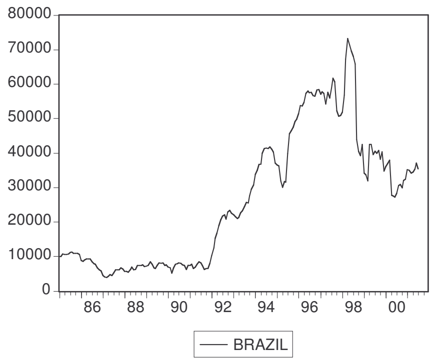
80000
70000
60000
50000
40000
30000
20000
10000
0
86
88
90
92
94
96
98
00
BRAZIL


1000
800
600
400
200
0
86
88
90
92
94
96
98
00
KENYA

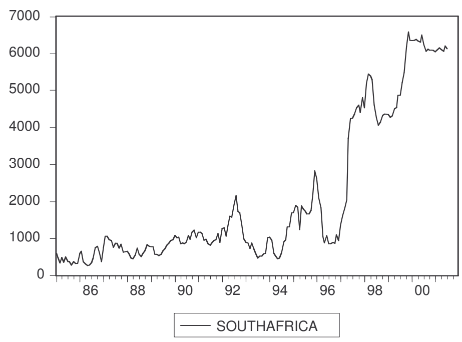
7000
6000
5000
4000
3000
2000
1000
0
86
88
90
92
94
96
98
00
SOUTHAFRICA

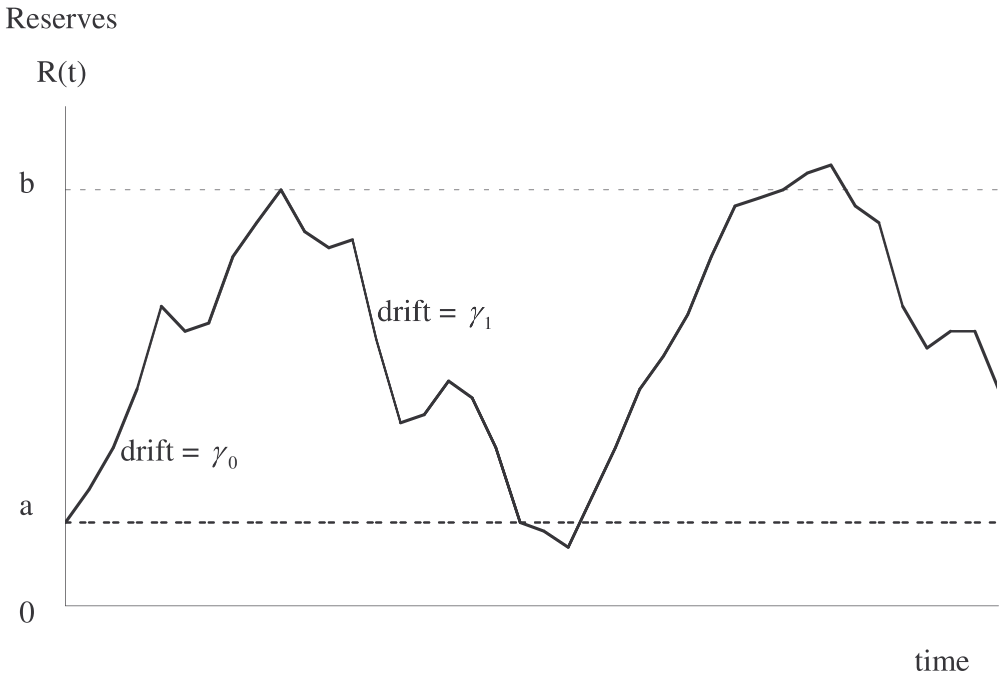

Figure 3: Drift control model of international reserves

Variable	All Countries (145 countries)	Industrial (22)	Developing					
			All (123)	Africa (42)	Asia (24)	Middle East (12)	Eastern Europe (16)	South/Central America (29)
Nu/Nd	2.02 (1.92)	1.88 (1.62)	2.05 (1.98)	1.64 (1.51)	2.44 (2.31)	1.84 (1.45)	2.03 (1.44)	2.46 (2.55)
Nu	31.30 (23.04)	28.49 (17.28)	31.93 (24.14)	27.06 (19.14)	34.82 (29.35)	29.52 (19.83)	33.91 (22.46)	36.81 (27.50)
Nd	20.68 (13.69)	20.53 (13.76)	20.71 (13.69)	20.76 (12.20)	20.25 (15.48)	20.73 (12.32)	19.37 (18.25)	21.44 (13.11)

Variable	Fixed/Managed Exchange Rate (126 countries)	Floating Exchange Rate (19 countries)	Emerging Markets (32 countries)
Nu/Nd	1.97 (1.83)	2.34 (2.47)	2.67 (2.32)
Nu	31.57 (23.67)	28.77 (16.89)	40.69 (30.92)
Nd	20.81 (13.14)	19.96 (17.20)	19.66 (15.22)

Table 1: Reserve Accumulation Times, Depletion Times and Drifts

Nu (Nd) is the number of months of reserve accumulation (depletion).
Numbers are averages. Standard deviations in parentheses.

h	$C(\beta)$	a	b	γ_0	γ_1	ER
0.01	1.313506	0.555081	4.154258	1.333219	-0.31155	3.602403
0.03	2.648767	0.551258	3.327752	0.883108	-0.60606	2.287094
0.05	3.64313	0.586882	3.10086	0.649506	-0.90957	1.832378
0.07	4.475442	0.607702	2.958336	0.515916	-1.25482	1.574156
0.09	5.206216	0.617067	2.842381	0.42715	-1.65371	1.401343
0.11	5.865611	0.619228	2.741223	0.361914	-2.11799	1.275296
0.13	6.471318	0.617241	2.651373	0.310345	-2.6638	1.178197
0.15	7.034759	0.612869	2.570907	0.267492	-3.31425	1.10046
0.17	7.563808	0.607176	2.498414	0.230548	-4.10198	1.036425
0.19	8.064162	0.600802	2.432766	0.197827	-5.076	0.982496
0.21	8.5401	0.594108	2.37298	0.168272	-6.31045	0.936272
0.23	8.994931	0.587284	2.318242	0.14118	-7.92363	0.896081
0.25	9.431274	0.58062	2.267981	0.116012	-10.1373	0.860712
0.27	9.85125	0.573816	2.221394	0.092492	-13.297	0.829271
0.29	10.2566	0.567563	2.178658	0.069898	-18.3144	0.8011
0.31	10.64878	0.560327	2.137212	0.049302	-26.4484	0.775529
0.33	11.02904	0.551063	2.095604	0.031295	-35.9628	0.752648
0.35	11.39833	0.547617	2.0644	0.011598	-116.316	0.731755

Table 2: Changing the Holding Cost

k	$C(\beta)$	a	b	γ_0	γ_1	ER
0.05	0.736138	1.377287	6.769530	-0.203998	-452.420	1.696641
0.09	0.922713	1.658907	6.939021	-0.05759	-141.785	2.229923
0.13	1.043024	1.783878	6.990555	0.033466	-6.43556	2.599382
0.17	1.127463	1.746435	6.764256	0.115699	-1.82395	2.883464
0.21	1.188722	1.637038	6.444516	0.19951	-1.06604	3.095734
0.25	1.233721	1.470708	6.044453	0.296957	-0.74281	3.25319
0.29	1.266542	1.257575	5.574653	0.424729	-0.5586	3.369837
0.33	1.289852	1.01199	5.058713	0.611726	-0.43917	3.460453
0.37	1.305642	0.753819	4.537286	0.917873	-0.35752	3.540625
0.41	1.315465	0.485672	4.02507	1.554897	-0.29794	3.624714
0.45	1.320197	0.169084	3.463841	4.803835	-0.24713	3.736782
0.49	1.320740	0.006374	3.189463	134.7241	-0.22599	3.806718
0.53	1.320760	0.003621	3.185305	245.9576	-0.22575	3.807264
0.57	1.320778	0.002664	3.183578	343.6737	-0.22544	3.809589
0.61	1.320795	0.002893	3.187182	334.8247	-0.22581	3.807766

Table 3: Changing the Regulation Cost

π_0	$C(\beta)$	a	b	γ_0	γ_1	*ER*
0.01	0.599299	0.003891	1.564982	365.4228	-0.59974	1.61677
0.02	1.143899	0.003284	2.74154	258.0055	-0.28971	3.09634
0.03	1.170082	0.003261	2.807142	266.5706	-0.2796	3.191611
0.04	1.194831	0.003941	2.870825	215.1837	-0.27005	3.286541
0.05	1.218303	0.006254	2.937629	133.304	-0.26134	3.381418
0.06	1.240602	0.041523	3.0502	19.60156	-0.25833	3.455964
0.07	1.261217	0.19555	3.363594	4.055119	-0.27344	3.486666
0.08	1.280057	0.333051	3.654933	2.324766	-0.2877	3.522475
0.09	1.297416	0.450731	3.915213	1.678788	-0.30019	3.561851
0.1	1.313506	0.555081	4.154258	1.333219	-0.31155	3.602403
0.11	1.328485	0.650401	4.378567	1.113395	-0.32226	3.642936
0.12	1.342476	0.739244	4.591927	0.958889	-0.33261	3.682894
0.13	1.355579	0.822969	4.796218	0.843336	-0.34271	3.722026
0.14	1.36788	0.902324	4.992441	0.753218	-0.35261	3.760232
0.15	1.379451	0.977681	5.181007	0.680836	-0.36229	3.797466
0.16	1.390356	1.049251	5.362166	0.6214	-0.37171	3.833707
0.17	1.400654	1.117196	5.536107	0.571726	-0.38084	3.868947
0.18	1.410394	1.181668	5.70309	0.529605	-0.38965	3.90319
0.19	1.419623	1.242814	5.863346	0.493449	-0.3981	3.93645
0.2	1.428381	1.300804	6.01721	0.462077	-0.40619	3.968745
0.25	1.466267	1.54953	6.702453	0.351918	-0.44099	4.116761
0.3	1.496573	1.744439	7.277024	0.285142	-0.46713	4.24557
0.4	1.54217	2.030873	8.209021	0.206881	-0.49948	4.461602

Table 4: Changing the Cost of Switching to the Upward Drift

π_1	$C(\beta)$	a	b	γ_0	γ_1	*ER*
0.01	1.026131	0.002595	2.678959	355.1233	-0.30212	2.994342
0.02	1.064004	0.003269	2.733289	319.9431	-0.28823	3.101423
0.03	1.100158	0.002747	2.805398	321.5173	-0.27947	3.191639
0.04	1.134912	0.003224	2.868998	263.1136	-0.26986	3.287013
0.05	1.168385	0.00409	2.929771	202.9626	-0.26116	3.378968
0.06	1.200689	0.004394	2.987028	187.0929	-0.25296	3.469631
0.07	1.231985	0.007342	3.046311	108.7138	-0.24602	3.554626
0.08	1.262116	0.107192	3.263098	7.287499	-0.25242	3.597954
0.09	1.289257	0.350309	3.733932	2.170768	-0.28227	3.589696
0.1	1.313506	0.555081	4.154258	1.333219	-0.31155	3.602403
0.11	1.335341	0.741699	4.553198	0.970035	-0.34275	3.625442
0.12	1.355053	0.918645	4.941292	0.760085	-0.37738	3.654847
0.13	1.372872	1.085779	5.315226	0.622976	-0.41562	3.688628
0.14	1.389019	1.240569	5.668642	0.527486	-0.45704	3.724989
0.15	1.403701	1.381515	5.997981	0.457921	-0.50116	3.762353
0.16	1.417111	1.508669	6.302984	0.405315	-0.54778	3.799597
0.17	1.429416	1.623056	6.585374	0.364236	-0.59692	3.836057
0.18	1.440757	1.726055	6.847576	0.331265	-0.64881	3.871395
0.19	1.451253	1.819081	7.092093	0.304181	-0.70382	3.905478
0.2	1.461003	1.903429	7.321193	0.281494	-0.76241	3.938292
0.25	1.501101	2.229149	8.294987	0.206595	-1.1341	4.085566
0.3	1.531087	2.451353	9.079478	0.163818	-1.75904	4.211884
0.4	1.573301	2.735714	10.33562	0.115385	-8.86599	4.427501

Table 5: Changing the Cost of Switching to the Downward Drift

σ_0^2	$C(\beta)$	a	b	γ_0	γ_1	ER
0.1	0.508982	0.907681	3.568472	0.02739	-5.31079	1.472981
0.15	0.611103	1.065724	4.047864	0.043249	-3.70752	1.727699
0.2	0.694437	1.187849	4.420716	0.059336	-2.88971	1.936353
0.25	0.765907	1.286012	4.724436	0.07575	-2.37958	2.115957
0.3	0.828998	1.365943	4.97717	0.092637	-2.02291	2.275096
0.35	0.885745	1.430712	5.18894	0.110177	-1.75459	2.418781
0.4	0.937446	1.481989	5.365576	0.128591	-1.54215	2.550204
0.45	0.984975	1.520569	5.510389	0.148157	-1.36743	2.671522
0.5	1.028948	1.54661	5.624948	0.169241	-1.21944	2.784236
0.55	1.0698	1.559713	5.709424	0.192339	-1.09105	2.889411
0.6	1.107839	1.558906	5.762672	0.218153	-0.97739	2.987793
0.65	1.143273	1.542552	5.7821	0.247719	-0.87496	3.079883
0.7	1.176219	1.508142	5.763314	0.282644	-0.78115	3.165978
0.75	1.206705	1.451926	5.699434	0.325589	-0.69389	3.246227
0.8	1.234659	1.368334	5.579938	0.381305	-0.61146	3.320713
0.85	1.259871	1.249119	5.388967	0.459159	-0.53242	3.389746
0.9	1.281935	1.082556	5.103663	0.58017	-0.45574	3.454737
0.95	1.300161	0.854886	4.697526	0.800561	-0.38151	3.520778
1	1.313506	0.555081	4.154258	1.333219	-0.31155	3.602403
1.05	1.320368	0.138074	3.41563	5.736926	-0.24327	3.745834
1.1	1.32077	0.00516	3.187924	161.9714	-0.22607	3.805137
1.15	1.320795	0.003162	3.183969	279.3701	-0.22571	3.80697
1.2	1.320873	0.00371	3.185023	245.0373	-0.22568	3.807815
1.25	1.320826	0.002065	3.182677	455.1487	-0.22549	3.808625
1.3	1.320793	0.001275	3.181149	770.8404	-0.22542	3.808632

Table 6: Changing Uncertainty About the Upward Drift

σ_1^2	$C(\beta)$	a	b	γ_0	γ_1	*ER*
0.2	0.791572	0.000802	1.773077	728.1559	-0.06131	2.517728
0.4	0.984934	0.000805	2.286547	765.5151	-0.10959	2.968429
0.6	1.120976	0.001914	2.648898	376.6372	-0.15127	3.307838
0.8	1.229183	0.001417	2.937933	515.4386	-0.18993	3.574981
1	1.313506	0.555081	4.154258	1.333219	-0.31155	3.602403
1.2	1.343857	1.255971	5.575442	0.560803	-0.56559	3.609333
1.4	1.35673	1.550765	6.183123	0.439198	-0.80073	3.631987
1.6	1.363989	1.711605	6.518393	0.389637	-1.02588	3.644924
1.8	1.368694	1.813843	6.733263	0.362468	-1.24638	3.652801
2	1.372007	1.884924	6.883542	0.345232	-1.46438	3.657992
2.2	1.37447	1.937352	6.994885	0.333294	-1.68089	3.661635
2.4	1.376375	1.97768	7.080823	0.324524	-1.89646	3.664314
2.6	1.377894	2.009694	7.149229	0.317804	-2.11138	3.666359
2.8	1.379134	2.035738	7.205	0.312488	-2.32585	3.667966
3	1.380166	2.057353	7.251368	0.308176	-2.53998	3.669261
3.2	1.381038	2.075579	7.290528	0.304607	-2.75384	3.670323
3.4	1.381785	2.091162	7.324054	0.301603	-2.96751	3.67121
3.6	1.382431	2.104639	7.353081	0.299041	-3.18103	3.671963
3.8	1.382997	2.116411	7.378459	0.296829	-3.39443	3.672606
4	1.383496	2.126783	7.400837	0.2949	-3.6077	3.673163

Table 7: Changing Uncertainty About the Downward Drift

$\sigma_0^2=\sigma_1^2$	$C(\beta)$	a	b	γ_0	γ_1	ER
0.1	0.507869	0.857556	3.358752	0.031603	-0.4592	1.468928
0.2	0.690184	1.039406	3.96328	0.076131	-0.40984	1.928419
0.3	0.820606	1.104439	4.279347	0.130227	-0.3881	2.259183
0.4	0.924515	1.104479	4.442876	0.196679	-0.37175	2.523933
0.5	1.01157	1.060338	4.505302	0.28008	-0.35722	2.747854
0.6	1.086676	0.984965	4.496483	0.3871	-0.34407	2.94488
0.7	1.152751	0.88925	4.440182	0.526974	-0.33273	3.124008
0.8	1.211729	0.782568	4.356761	0.712831	-0.32366	3.291114
0.9	1.26498	0.670826	4.25996	0.966505	-0.31683	3.449784
1.0	1.313506	0.555081	4.154258	1.333219	-0.31155	3.602403
1.1	1.358044	0.432346	4.036342	1.92675	-0.30679	3.751324
1.2	1.39913	0.298025	3.89961	3.111835	-0.30161	3.899437
1.3	1.437157	0.148416	3.73875	6.893655	-0.29547	4.05008
1.4	1.472437	0.015501	3.607502	72.099	-0.29292	4.191532
1.5	1.505745	0.006929	3.682071	175.6072	-0.30706	4.282736
1.6	1.537589	0.004407	3.761888	296.4685	-0.32181	4.366385
1.7	1.568148	0.00441	3.843831	317.7739	-0.33675	4.44554

Table 8: Changing Uncertainty About Both Drifts

(1)	(2)	(3)	(4)	(5)	(6)
h	***C(a,b,γ_0 , γ_1)***	***C(b)***	**[C(b)-C(a,b,γ_0 , γ_1)/C(b)]*100**	***C(a,b)***	**[C(a,b)-C(*a,b,*γ_0 , γ_1 *)* / C(a,b)]*100**
0.001	0.327528556	0.556186288	41.11171688	0.5542431	40.90524956
0.002	0.484400911	0.77680926	37.6422327	0.7728065	37.31924774
0.003	0.617985611	0.950129125	34.95772372	0.944028	34.53736124
0.004	0.73795446	1.098764601	32.83780172	1.0905406	32.33131822
0.005	0.848575885	1.231566445	31.09783982	1.2212018	30.51305035
0.006	0.952214013	1.353096379	29.62703709	1.340577	28.9698406
0.007	1.050350938	1.466072481	28.35613848	1.4513869	27.63122308
0.008	1.143992771	1.572269385	27.23939153	1.5554077	26.4506166
0.009	1.233862851	1.672921616	26.24502908	1.6538753	25.39565366
0.01	1.320501711	1.768929059	25.35021662	1.7476905	24.44304782
0.011	1.404327763	1.860971675	24.53792919	1.8375342	23.57542204
0.012	1.485673514	1.949578114	23.79512759	1.9239356	22.77945576
0.013	1.56480701	2.035169042	23.11169354	2.0073158	22.04480242
0.014	1.64195239	2.118085711	22.47941707	2.0880167	21.36306072
0.015	1.717291453	2.198609485	21.8919292	2.1663198	20.72770477
0.016	1.790982836	2.276975595	21.34378426	2.2424609	20.13315223
0.017	1.863164601	2.353383086	20.83037342	2.3166392	19.57467528
0.018	1.933950416	2.42800216	20.3480768	2.3890252	19.04855462
0.019	2.003440828	2.500979707	19.89375912	2.459766	18.55156698
0.02	2.07172732	2.572443547	19.46461479	2.5289895	18.0808271
0.021	2.138887778	2.642505726	19.05834839	2.5968082	17.63397089

Table 9: The Expected Discounted Cost of Managing Reserves