History Dependent Financial Claims: A Monte Carlo Approach

Robert A. Jones Rodney L. Jacobs*

August 1986

The volatility of interest rates over the last decade has spawned a broad array of interest rate contingent claims — securities whose payoffs depend on future interest rates. An important subclass of securities are those whose payouts depend on the history of interest rates prior to a given payment date. Examples include floating rate notes with coupons specified as some average of lagged Treasury yields,[1] adjustable rate mortgages and average rate cap contracts. Techniques are needed to assess the value and riskiness of such claims.

This paper develops a Monte Carlo method for valuing default free interest rate claims. It is an alternative to direct numerical procedures for valuing any claim; but its main advantage lies in its ease of incorporating history dependent payoffs. We illustrate its application to the pricing of average rate cap contracts. The evolution of interest rates over time is assumed to follow the two factor equilibrium process of Jacobs and Jones [7].

Section I of the paper describes rate cap contracts of various types. Section II lays out the stochastic process followed by interest rates and associated valuation equation for interest rate instruments. Section III discusses direct numerical methods for solving the equation and difficulties associated with history dependent payouts. A Monte Carlo solution is proposed. Various details of its implementation are laid out in IV. Section V tests the procedure against direct solutions for a variety of claims with history independent payouts: discount bond, European interest rate options, and instantanous rate caps. Section VI explores the properties of average rate caps and contrasts them with instantaneous rate caps.

*Simon Fraser University and Wells Fargo Bank respectively.

[1] A number of such instruments are examined in [12]. Additionally, to the extent that the administered commercial bank prime rate can be modelled as a lagged average of open market rates, most prime rate linked contracts can be viewed as history dependent claims.

1 Interest Rate Cap Contracts

The major portion of commercial bank lending is done at rates that float with some measure of market rates, such as CD or LIBOR rates, or with the administered prime rate. Depending on the rate structure of its remaining assets and liabilities, a floating rate loan may expose the borrower to unacceptable risk due to interest rate fluctuations. This interest rate risk could be eliminated with a fixed rate loan or, equivalently, with an interest rate swap. However it is difficult to accommodate the flexible or uncertain takedown schedules of most commercial loans with a fixed rate loan. Also, many borrowers are reluctant to 'pay up' the yield curve to obtain fixed rate financing, and are reluctant to give up the potential to gain from declines in market rates.[2] What these borrowers need (and is also desirable from the standpoint of default risk borne by the bank) is insurance against significant increases in interest rates.

An interest rate cap is a contract that provides insurance against rate increases. What we call an *instantaneous rate cap* is a contract stipulating that the periodic rate on outstanding loan balances will be the lesser of the floating rate prevailing at that time and the cap rate. [3] For example, assume that the loan rate is reset every month based on one-month LIBOR at that time, with a cap level of 10%. This rate cap is equivalent to a series of European put options on one month discount bonds, with expiry dates one month apart and exercise prices to provide a 10% yield. If LIBOR is below 10% on the reset date, that particular option is not exercised and the loan rate is based on the then current LIBOR. If LIBOR is above 10% the expiring put option is exercised and the loan rate is based on the LIBOR cap of 10%. The borrower's value of LIBOR for a hypothetical time path of market rates is illustrated in Figure 1.

An instantaneous rate cap pays out whenever the floating rate exceed the cap rate. The borrower's concern, however, may not be with occasional spikes in borrowing costs, but rather with high rates sustained over long periods. Protection against this type of occurrence is more specifically provided by an *average rate cap*. An average rate cap guarantees that the average rate paid over the

[2]Borrowers would generally prefer to have fixed rate loans as an option. If rates go up they would stay on the fixed-rate line, but if rates decline they would prepay the fixed-rate line and borrow at the then lower current rate. This difficulty (from the bank's point of view) with fixed-rate lending can be eliminated by imposing mark-to-market prepayment fees, rather than just not offering fixed rate loans. The existence of such fees limits fixed rate lending to borrowers with a real need for such financing.

[3]Such contracts were introduced in the early 1980's by Citibank, Wells Fargo and other financial intermediaries. Although they have received less publicity than rate swaps, they have played an important role in hedging real estate construction financing and leveraged buyouts. The premium for the contract is typically paid as an up front fee (which may simply be added to the loan balance). The rate cap may be provided in conjunction with a floating rate loan, or sold on its own to be conbined with floating rate funds from other sources. In this latter form, it is a claim to a stream of payments equal to the greater of zero and the difference between the floating rate and the cap rate at each point in time, applied to a fixed hypothetical balance schedule.

Figure 1: Borrower's loan rate with instantaneous rate cap

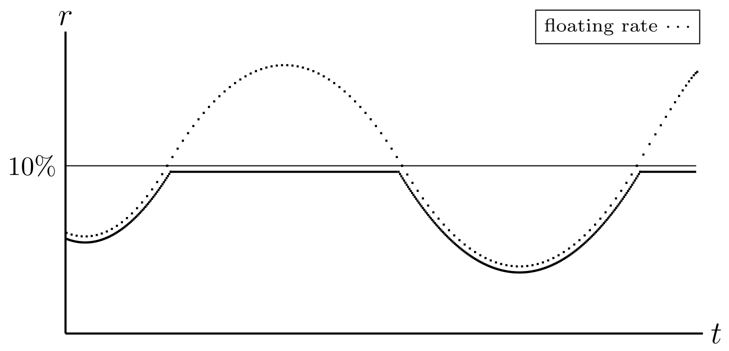

Figure 2: Borrower's loan rate with average rate cap

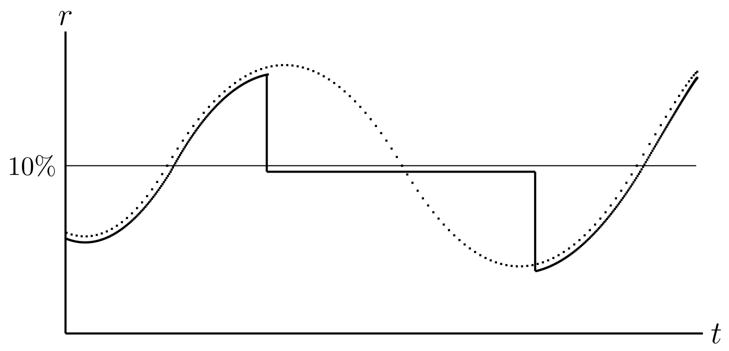

term of the loan (weighted by balances outstanding) will be no greater than the cap rate. Periods of above cap rates are offset by periods of below cap rates before benefits accrue under the contract.

Figure 2 depicts the borrower loan rate for an average cap. As market rates rise above the cap level, the borrower's LIBOR rate increases until the average LIBOR rate to date equals the cap level. At that time, the borrower's LIBOR rate drops to the cap level. As rates decline, the borrower's LIBOR rate remains at the cap level until the average rate to date falls below the cap. At that point the borrower's rate drops to the current market rate. Since payouts under the average cap are always less than or equal to those under the instantaneous cap, the premium or fee required for such protection is lower.

Figure 3: Borrower's loan rate with hybrid rate cap

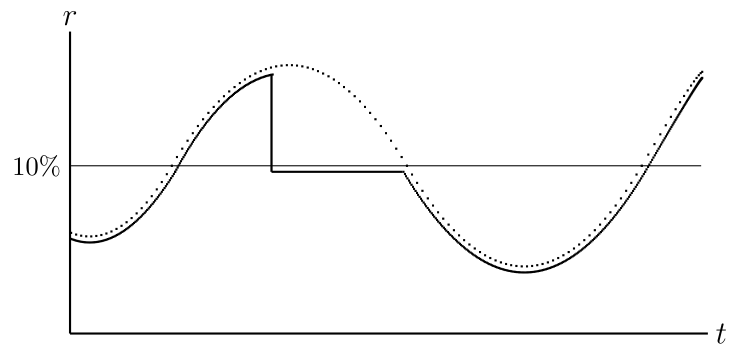

With the average cap, the borrower's effective LIBOR can exceed market LIBOR when rates fall below the cap level. The borrower may prefer a cap contract in which any decline in rates below the cap level is reflected immediately in the rate he is charged. This type of rate protection is provided by a *hybrid rate cap*. With a hybrid cap, the borrower's LIBOR is the lesser of market LIBOR and a rate such that his average LIBOR to date does not exceed the cap rate. Although the cost for such protection is higher than for the average cap, the contract has the advantages of no default risk, since the borrower is never obliged to pay above market rates, and elimination of incentives to prepay the loan when rates fall. Figure 3 illustrates the borrower's loan rate over time.

Although the instantaneous rate cap and the average and hybrid rate caps appear to be similar products, the options associated with the average and hybrid caps are significantly more complex because they are history dependent. That is, the payoff on these instruments depends on the time path of interest rates since the cap was initiated, whereas the payoff with the instantaneous cap depends only on the current state. In this paper we investigate the valuation of average and hybrid caps and the sensitivity of that value to changes in current interest rates. Solution to the first problem provides a basis for pricing such contracts offered to customers. Solution to the second provides a basis for hedging the risk taken on through positions in traded securities whose sensitivities to the same underlying factors are known. The success of such instruments hinges on the premium being sufficiently reasonable to attract borrowers, which depends in turn on the ability of the originating institution to accurately value and hedge the contracts.

2 Valuing Interest Rate Contingent Claims

The value of rate caps and other interest rate related claims depend on how interest rates move over time. We assume that the instantaneous risk free rate $r(t)$ follows the two factor continuous time stochastic process derived and estimated in [7]:

$$\begin{aligned} dr &= \kappa_1 r(\ln\mu - \ln r)\, dt + \sigma_1 r\, dz_1 \\ d\mu &= \kappa_2 \mu(\ln\gamma - \ln\mu)\, dt + \sigma_2 \mu\, dz_2 \end{aligned} \tag{1}$$

In the above, dz_1 and dz_2 are the increments of a standard Weiner process with zero mean, variance one per unit time and instantaneous correlation coefficient ρ. The parameters κ_1, κ_2, γ, σ_1, σ_2 and ρ are constant. Short term interest rates are treated as moving toward some current target $\mu(t)$, while μ in turn regresses toward some long run fixed level γ. Both are subject to stochastic shocks whose magnitudes are proportional to their current levels. A two factor representation of interest rate movements is the minimum needed to allow both shifts and twists in the term structure.

The equilibrium prices of all default-free interest rate related claims are assumed to be deterministic functions of the state variables r and μ. In the general equilibrium context of Jacobs and Jones, the value at time t of a security that matures with value $P(r,\mu,T)$ at time T, and promises a flow of payments at rate $q(r,\mu,t)$ in the meantime, is the function $P(r,\mu,t)$ satisfying the partial differential equation (pde)

$$\begin{aligned} 0 &= \frac{1}{2}r^2\sigma_1^2 P_{rr} + \rho r\mu\sigma_1\sigma_2 P_{r\mu} + \frac{1}{2}\mu^2\sigma_2^2 P_{\mu\mu} + \big(\kappa_1 r\ln(\mu/r) - \lambda_1 r^{3/2}\big)P_r \\ &\quad + \big(\kappa_2\mu\ln(\gamma/\mu) - \lambda_2\mu r^{1/2}\big)P_\mu + q(r,\mu,t) + P_t - rP \end{aligned} \tag{2}$$

with the required value at maturity. The two additional parameters λ_1 and λ_2 reflect the market risk premiums required for exposure to unexpected fluctuations in the underlying factors r and μ respectively. This valuation equation applies equally to discount bonds, coupon bonds, interest rate futures and interest rate options.

Valuation equation (2) applies to any rate cap contract whose payout can be expressed solely in terms of the current level and structure of interest rates. This is the case with instantaneous rate caps. An instantaneous rate cap has maturity value 0 and promises a flow of payments

$$q(r,\mu,t) = \max\{\, 0\,,\ \phi(r,\mu) - \bar{\phi}\}\, B(t) \tag{3}$$

where $\bar{\phi}$ is the agreed upon cap rate, $\phi(r,\mu)$ is the state dependent floating rate that would otherwise be paid, and $B(t)$ is the loan balance at time t.

The situation is more complex for an average rate cap. The payout of the cap at time t depends not only on the above factors, but also on the history of interest rates since the contract's inception. The contract only pays out if the average rate paid to date equals the cap rate $\bar{\phi}$. Clearly, one or more additional state variables must be introduced to keep track of the relevant history.

The appropriate history variable differs slightly for the average cap and the hybrid cap. We treat each in turn, beginning with the average cap. Let $s(t)$ be the difference between the interest payments that would have been made to date if $\bar{\phi}$ had been charged and the payments that would have been made if the uncapped floating rate ϕ had been charged. One can view this as the balance in a hypothetical reserve account of funds available for periods of higher rates. If $s > 0$ then the average floating rate to date is less than the cap rate, and there is no payout under the contract. But if $s < 0$ then the floating rate has averaged more than the cap rate, and a payment under the cap contract is due. The amount paid to the borrower (which may be negative) is $\big(\phi(r,\mu) - \bar{\phi}\big)B(t)$ That is, the borrower pays the cap rate $\bar{\phi}$ as long as floating rates have averaged more than that rate, even though they may currently be below it. An average cap is thus a claim to the flow of payments

$$q(r,\mu,s,t) = \begin{cases} 0 & \text{for } s > 0 \\ \big(\phi(r,\mu,s) - \bar{\phi}\big)B(t) & \text{for } s \leq 0 \end{cases} \tag{4}$$

The dynamics of $s(t)$ are given by

$$ds(t) = \big(\bar{\phi} - \phi(r,\mu)\big)B(t)\,dt \tag{5}$$

The effect of this third state variable must be incorporated into the valuation equation. Since s is locally non-stochastic, it does not give rise to any further second order terms beyond those already in equation (2). However a first order term equal to P_s times $\mathrm{E}[ds]/dt$ must be added. The valuation pde becomes

$$\begin{aligned} 0 \quad = \quad & \frac{1}{2}r^2\sigma_1^2 P_{rr} + \rho r\mu\sigma_1\sigma_2 P_{r\mu} + \frac{1}{2}\mu^2\sigma_2^2 P_{\mu\mu} + \big(\kappa_1 r\ln(\mu/r) - \lambda_1 r^{3/2}\big)P_r \\ & + \big(\kappa_2\mu\ln(\gamma/\mu) - \lambda_2\mu r^{1/2}\big)P_\mu + \big(\bar{\phi} - \phi(r,\mu)\big)B(t)P_s \\ & + q(r,\mu,s) + P_t - rP \end{aligned} \tag{6}$$

The particular solution is determined by the terminal condition $P(r,\mu,s,T) = 0$ and the boundary condition $P(r,\mu,0,t) = -1$. This last condition embodies the notion that if the cap contract is just on the borderline of paying out (i.e., $s = 0$), then a \$1 increase in interest paid to date gives rise to an immediate additional \$1 payment from the contract.

For the hybrid cap, let $s(t)$ denote the accumulated difference between the total interest payments that would have been made to date if the rate $\bar{\phi}$ had been charged and the actual total interest paid. Again, if $s > 0$ then the average rate paid to date is less than the cap rate, and no payout is due under the contract. But if $s = 0$ and $\phi(r,\mu) > \bar{\phi}$, then a payout is required to keep s just equal to

0. The hybrid cap contract is thus a claim to the flow of payments

$$q(r,\mu,s,t) = \begin{cases} 0 & \text{for } s(t) > 0 \\ \max\{\,0\,,\ \phi(r,\mu) - \bar{\phi}\,\}B(t) & \text{for } s(t) \le 0 \end{cases} \tag{7}$$

The dynamics of $s(t)$ are given by

$$ds(t) = \begin{cases} \big(\phi(r,\mu) - \bar{\phi}\big)B(t)\,dt & \text{for } s(t) > 0 \\ \max\{\,0\,,\ \bar{\phi} - \phi(r,\mu)\,\}B(t) & \text{for } s(t) = 0 \end{cases} \tag{8}$$

Upon combining the P_s times $\mathrm{E}[ds]/dt$ implied by (8) with q as given by (7), the valuation pde becomes

$$\begin{aligned} 0 \;=\; & \frac{1}{2}r^2\sigma_1^2 P_{rr} + \rho r\mu\sigma_1\sigma_2 P_{r\mu} + \frac{1}{2}\mu^2\sigma_2^2 P_{\mu\mu} + \big(\kappa_1 r\ln(\mu/r) - \lambda_1 r^{3/2}\big)P_r \\ & + \big(\kappa_2\mu\ln(\gamma/\mu) - \lambda_2\mu r^{1/2}\big)P_\mu + \big(\bar{\phi} - \phi(r,\mu)\big)B(t)P_s + P_t - rP \end{aligned} \tag{9}$$

The boundary conditions are the same as for the average cap.

Comparison of (6) or (9) with (2) suggests the practical problem in implementing average caps. For a two factor interest rate model, valuing the instantaneous cap entails solving a pde with two state variables plus a time dimension. Valuing the average cap entails solving a pde with three state variables plus a time dimension. Since the cost of solving such equations by direct methods rises geometrically with the number of state variables, valuing an average cap of moderate duration by these means becomes prohibitive.

Equations (4)-(9) describe caps on loans with continuous interest payments. In practice, loan payments are made at predetermined intervals. The analogues of the above relations when payments are due at discrete intervals are given below. Let the payment dates be $t_1 < t_2 < \ldots < T$. We make the following assumptions:

1. the loan balance $B(t_i)$ is constant in the interval $[t_i, t_{i+1})$
2. the payment intervals are of equal length
3. the rates $\bar{\phi}$ and $\phi(r,\mu)$ are simple interest rates per period
4. rates payable are determined at the beginning of each interest period
5. interest for each period is payable at the end of the period

Let $\phi_i \equiv \phi(r(t_i), \mu(t_i))$ denote the simple interest rate, determined at time t_i, that is to be charged over the interval $[t_i, t_{i+1})$. Let $B_i \equiv B(t_i)$ be the loan balance to which it applies. Similarly let $q_i \equiv q(t_i)$ and $s_i \equiv s(t_i)$. With interest paid not in advance, an instantaneous cap is a claim to the discrete stream of payments

$$q_i = \max\{\,0\,,\ \phi_{i-1} - \bar{\phi}\,\}B_{i-1} \tag{10}$$

The average cap is a claim to the stream

$$q_i = \max\big\{\,0\,,\,(\phi_{i-1}-\bar{\phi})B_{i-1} - \max\{0, s_{i-1}\}\\ + \min\big\{\,0\,,\,(\phi_{i-1}-\bar{\phi})B_{i-1} + \max\{0, s_{i-1} - (\phi_{i-1}-\bar{\phi})B_{i-1}\}\big\} \tag{11}$$

with the dynamics of the state variable s given by a difference equation

$$s_i = s_{i-1} + (\bar{\phi} - \phi_{i-1})B_{i-1} \tag{12}$$

The hybrid rate cap is a claim to the stream

$$q_i = \max\big\{\,0\,,\,(\phi_{i-1}-\bar{\phi})B_{i-1} - s_{i-1}\,\big\} \tag{13}$$

with the dynamics of the state variable s given by the difference equation

$$s_i = \max\big\{\,0\,,\,s_{i-1} + (\bar{\phi} - \phi_{i-1})B_{i-1}\,\big\} \tag{14}$$

Equations (10)-(14) are the discrete time analogues of (3)-(5) and (7)-(8) respectively. Notice that (11) and (13) call for q to be the smallest payment necessary to keep the average rate paid from exceeding $\bar{\phi}$.

Between payment dates q and ds are both 0. The value of the rate caps satisfy equation (2), with q omitted, within each interval (t_i, t_{i+1}). On payment dates the value drops discontinuously to reflect receipt of the previously determined $q(t_i)$. We thus have a sequence of boundary conditions augmenting (2):

$$\lim_{t\to t_i^-} P(r,\mu,s,t) = \lim_{t\to t_i^+} P(r,\mu,s,t) + q(t_i) \tag{15}$$

Again, the instantaneous cap's value is given by a two state variable equation, while the average and hybrid caps' are given by a three state variable equation.

3 Numerical Valuation

The valuation pde has no known closed form solution and must be solved numerically. [13] provide an introduction to some of the techniques available. Finite difference methods, which discretize the state space into a regular grid and replace differentials with differences, are well suited to our problem. We use an alternating direction implicit (ADI) scheme for the two state variable problem (see [11]). With a grid of n points in each of the r and μ directions, one solves $2n$ systems of n simultaneous linear equations at each time step. Each system is tridiagonal and can be solved with just $6n$ multiplications/divisions plus $3n$ additions. For the parameter estimates given in the previous section acceptable accuracy is achieved with a grid of 41 points in each state direction and a time step of 1 week. Valuing a five year security thus involves solving 21320 tridiagonal systems. A third state variable, however, would increase computation requirements by a factor of at least the number of grid points in the third direction. We turn to Monte Carlo techniques as a potentially economical means for valuing average rate caps.

Monte Carlo techniques have been used to solve differential equations since very early in digital computing (see [4]). The method is described in [8] and [6]. Lattes [10] applies it to boundary value problems arising in mathematical physics. Boyle [1] demonstrates how it can be used to value options on dividend paying stocks. Its application in the context of stochastic processes seems natural. Yet the method applies equally to problems arising from deterministic systems. The technique involves interpreting the solution, or 'potential' at a given point in the region where the equation applies, as the average 'score' obtained by taking a random walk to the region's boundary. In finance contexts the method is closely linked with the notion of risk neutral valuation, although its derivation in no way presupposes such ideas.

The situations we examine all involve payments at discrete points in time. Let us focus on the equation describing a security's value between payment dates. Consider the pde

$$aU_{\mu\mu} + bU_{rr} + cU_{\mu r} + dU_{\mu} + eU_r + fU + U_t = 0 \tag{16}$$

in which the coefficients a–f may be functions of μ, r, t. We assume the pde is parabolic — i.e., $a, b \geq 0$ and $4ab \geq c^2$.

Discretize the state space by placing on it a net of mesh size g in the μ direction, h in the r direction and k in the t direction. Let V_{ij}^t denote the solution value at the ij^{th} gridpoint — i.e., $U(ig, jh, t)$. We focus on the ij^{th} gridpoint, and let $V_{+-} \equiv V_{i+1,j-1}^{t+k}$, $V_{0+} \equiv V_{i,j+1}^{t+k}$, *etc.*. Write centered difference approximations for the derivatives appearing in (16):

$$\begin{array}{lclclcl} U_{\mu} & = & (V_{+0} - V_{-0})/2g & & U_r & = & (V_{0+} - V_{0-})/2h \\ U_{\mu\mu} & = & (V_{+0} - 2V_{00} + V_{-0})/g^2 & & U_{rr} & = & (V_{0+} - 2V_{00} + V_{0-})/h^2 \\ U & = & (V_{00}^{t+k} + V_{00}^t)/2 & & U_t & = & (V_{00}^{t+k} - V_{00}^t)/k \end{array} \tag{17}$$

We give three discretizations of $U_{\mu r}$ applicable when $c \leq 0$. The analysis for $c > 0$ would proceed in similar fashion.

$$U_{\mu r} = \begin{cases} (V_{+0} + V_{0+} + V_{0-} + V_{-0} - V_{+-} - V_{-+} - 2V_{00})/2gh & \text{case 1} \\ (V_{+0} - V_{00} - V_{+-} + V_{0-})/gh & \text{case 2} \\ (V_{0+} - V_{00} - V_{-+} + V_{-0})/gh & \text{case 3} \end{cases}$$

Case 1 is the average of 2 and 3. For a given gridpoint, one is selected depending on the size and relative magnitude of the coefficient functions a–c in order to maintain non-negative probabilities of moving to adjacent gridpoints (derived

Figure 4: Gridpoints used when $c < 0$

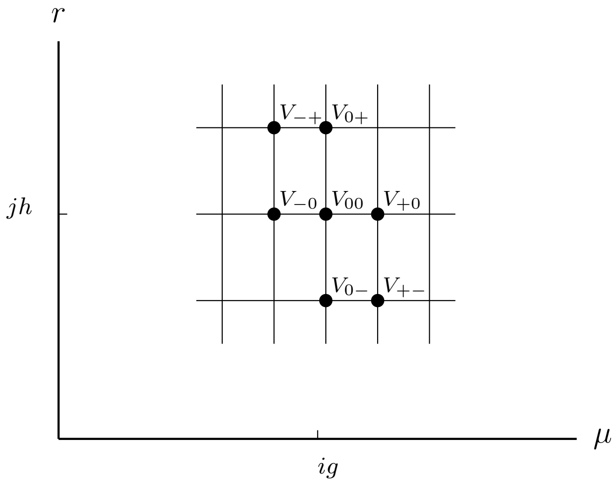

below). Figure 4 depicts the seven gridpoints in the (μ, r) plane utilized at points when $c \leq 0$.

Substituting the difference expressions (17) into (16), using case 1 for $U_{\mu r}$, results in

$$
\begin{aligned}
0 \quad = \quad & \frac{a}{g^2}(V_{+0} - 2V_{00} + V_{-0}) + \frac{b}{h^2}(V_{0+} - 2V_{00} + V_{0-}) \\
& + \frac{c}{2gh}(V_{+0} + V_{0+} + V_{0-} + V_{-0} - V_{+-} - V_{-+} - 2V_{00}) + \frac{d}{2g}(V_{+0} - V_{-0}) \\
& + \frac{e}{2h}(V_{0+} - V_{0-}) + \frac{f}{2}(V_{00} + V_{00}^t) + \frac{1}{k}(V_{00} - V_{00}^t) \qquad (18)
\end{aligned}
$$

Unless otherwise indicated, V values are at time $t + k$. Now solve (18) for V_{00}^t, removing as a factor the sum of the coefficients on the values of V at time $t + k$. This gives

$$
\begin{aligned}
V_{00}^t \quad = \quad & \frac{1 + fk/2}{1 - fk/2}\Big(\pi_{00} V_{00}^{t+k} + \pi_{+0} V_{+0}^{t+k} + \pi_{-0} V_{-0}^{t+k} + \pi_{0+} V_{0+}^{t+k} \\
& \qquad + \pi_{0-} V_{0-}^{t+k} + \pi_{+-} V_{+-}^{t+k} + \pi_{-+} V_{-+}^{t+k}\Big) \qquad (19)
\end{aligned}
$$

where

$$
\begin{aligned}
\pi_{00} &= 1 - 4(a\tfrac{h}{g} + b\tfrac{g}{h} + c\tfrac{1}{2})E & \pi_{0+} &= (2b\tfrac{g}{h} + eg + \begin{Bmatrix} c \\ 0 \\ 2c \end{Bmatrix})E \\
\pi_{+0} &= (2a\tfrac{h}{g} + dh + \begin{Bmatrix} c \\ 2c \\ 0 \end{Bmatrix})E & \pi_{0-} &= (2b\tfrac{g}{h} - eg + \begin{Bmatrix} c \\ 2c \\ 0 \end{Bmatrix})E \\
\pi_{-0} &= (2a\tfrac{h}{g} - dh + \begin{Bmatrix} c \\ 0 \\ 2c \end{Bmatrix})E & \pi_{+-} &= -\begin{Bmatrix} c \\ 2c \\ 0 \end{Bmatrix} E \\
\pi_{-+} &= -\begin{Bmatrix} c \\ 0 \\ 2c \end{Bmatrix} E & & \text{and } E \equiv k/gh(2+kf)
\end{aligned}
\tag{20}
$$

The entries in braces correspond to cases 1, 2 and 3 respectively of (17) being substituted for $U_{\mu r}$. Notice that the π's sum to 1 by construction. Since $4ab \geq c^2$ the π's are also non-negative for sufficiently small h and k and appropriate selection of h/g. Interpreting the π's as probabilities of moving to adjacent gridpoints in the next time step, equation (19) can be viewed as follows: V_{00}^t is the discounted expected value of V^{t+k}, where the discount factor is $(1 + kf/2)/(1 - kf/2)$ and the probabilities of being at the various neighbouring gridpoints at time $t + k$, conditional on being at i, j at time t, are given by the π's. Were the values V^{t+k} known, one could obtain $V_{i,j}^t$ by determining this expectation. Of course they are generally not known except on a boundary. However they may be represented in turn as conditionally expected values of V^{t+2k}, et cetera, until the known values V^T are encountered at the maturity boundary. The solution value $V_{i,j}^t$ is the expected value over all possible paths starting at $\mu = ih$, $r = jh$, t and ending at time T, of the attained maturity value multiplied by the cumulative discount factor

$$
\prod_{\tau=t}^{T}\left(\frac{1 + f(\mu_\tau, r_\tau)k/2}{1 - f(\mu_\tau, r_\tau)k/2}\right) \tag{21}
$$

The Monte Carlo method determines this expected value by sampling. A random walk is taken through the μ, r grid, starting at i, j and using the π's for the current grid location as transition probabilities. The discount factor, which depends on the grid location through f, is accumulated. Upon reaching maturity T, the factor is multiplied by V^T at the terminal grid point to get a final score. Repeating this procedure many times, the average score obtained is the estimated $V_{i,j}^t$.

For securities making payments before maturity, such as coupon bonds and rate caps, each payment is multiplied by the discount factor accumulated by that date and added to the score for that particular walk. If the payment is history dependent, as with average caps, the history variable $s(t)$ must be updated as the walk proceeds. However this adds little computational burden. The virtue of the Monte Carlo method in such contexts should be apparent. Since it proceeds by moving forward from the initial state to maturity, a particular history is generated as part of each walk. This is easily incorporated into the periodic payments. In contrast, direct methods start from the known maturity values and work backwards to the initial state, requiring all possible values of the history variable to be considered at each time step. [4]

Equation (19) was derived from (16) without reference to the context giving rise to the pde. Consideration of the process that gave us (2) shows the Monte Carlo solution to in fact be a simulation of possible interest rate scenarios. The coefficient function $f(\mu, r)$ equals $-r$ in the valuation equation. The factor $(1+fk/2)/(1-fk/2)$ is a second order approximation to e^{-rk}. Thus the Monte Carlo solution is the expected value of payments discounted by the average instantaneous interest rate prevailing between time 0 and the payments. From the transition probabilities π given in (19), one can calculate the expected rate of change in the state variables r and μ at time t:

$$\begin{aligned} E[\mu(t_{i+1}) - \mu(t_i)]/k &= \frac{d(r_i, \mu_i)}{1 + kf/2} \\ E[r(t_{i+1}) - r(t_i)]/k &= \frac{e(r_i, \mu_i)}{1 + kf/2} \end{aligned} \tag{22}$$

The implied process for the state variables used in the simulation is not the objective process given by (1), but rather the 'risk-adjusted' process

$$\begin{aligned} dr &= (\kappa_1 r \ln(\mu/r) - \lambda_1 r^{3/2})\, dt + \sigma_1 r\, dz_1 \\ d\mu &= (\kappa_2 \mu \ln(\gamma/\mu) - \lambda_2 \mu r^{1/2})\, dt + \sigma_2 \mu\, dz_2 \end{aligned} \tag{23}$$

The Monte Carlo approach to solving pde's in effect values securities at the expected present value of their cash flows, using a suitably altered probability distribution for those cash flows and future interest rates. This is the same valuation principle arrived at through a different route by Cox, Ingersoll and Ross [3, p.380].

4 Implementation Details

4.1 Parameter Estimates

The parameters of the interest rate model were estimated from U.S. Government Treasury Bill and Note data following the procedure described in [7]. Weekly

[4] The reverse side of the coin, however, is that it is relatively to easy to incorporate into backwards solving procedures any forward looking decisions affecting the value of the security — for example, the decision whether to exercise an American option. This is cumbersome to implement in a Monte Carlo setup since the future, in a manner of speaking, has not yet unfolded as part of the solution procedure.

Figure 5: Theoretical yield curves

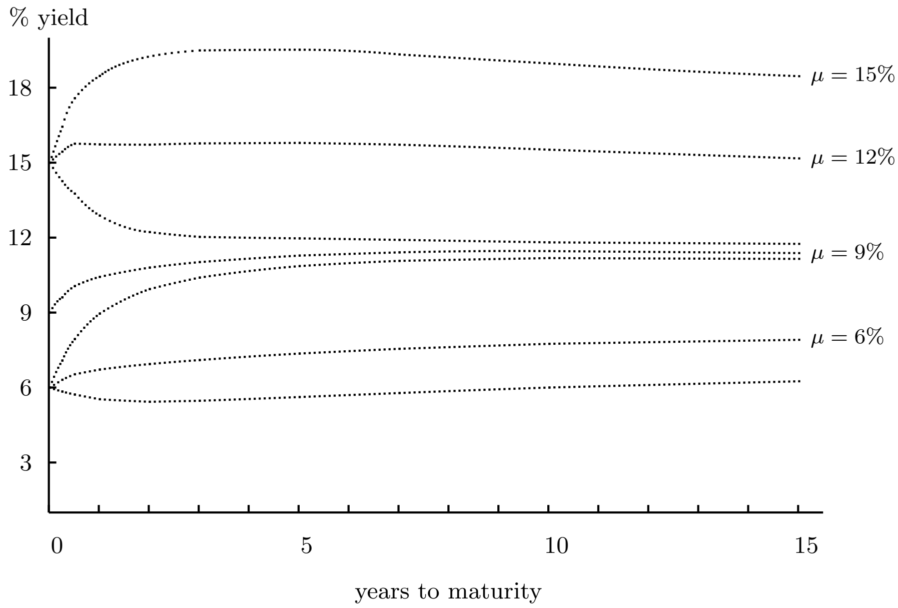

observations of the prices of 4, 13, 26 and 52 week Bills and yields on 3, 5 and 7 year Notes were used for the period January 1978 to April 1986. The yield on 1 week Treasury Bills was used for the instantaneous interest rate $r(t)$. The unobservable $\mu(t)$ was treated as a latent variable, its time series estimated along with the fixed parameters. The results were as follows:[5]

$$\begin{array}{lllllllll} \kappa_1 & = & .0507 & \sigma_1 & = & .1055 & \lambda_1 & = & \text{-}.2629 \\ \kappa_2 & = & .00080 & \sigma_2 & = & .0274 & \lambda_2 & = & \text{-}.0119 \\ \gamma & = & .00163 & \rho & = & \text{-}.281 & & & \end{array}$$

The above estimates take the time unit to be one week. Figure 5 depicts the theoretical yield curves for a variety of (μ, r) states.

The simulations and illustrations that follow use these parameter values. Their most notable feature is the very large value of κ_1 relative to κ_2. The value of κ_1 implies of half-life of deviations of r from μ of about thirteen weeks;

[5]The parameter estimates differ from those of [7] for several reasons. First, the data set is extended by almost two years. Second, the one week bill rate was used instead of the Federal Funds rate as the instantaneous interest rate. Third, serial correlation in the pricing residuals was handled in a slightly different fashion. The result is somewhat lower estimated volatility of r and a lower value of κ_1.

the value of κ_2 implies a half-life of deviations of μ from γ (which is equivalent to 8.48%/year) of about sixteen years. We now briefly discuss some more specifically numerical issues in implementing the Monte Carlo solution.

4.2 At the Boundary

Although the interest rate process has natural boundaries at $r = 0$ and $\mu = 0$, it has no natural upper boundaries. Clearly a grid of finite size must be used in the simulations. We somewhat arbitrarily set the maximum levels of r and μ to be 30%/year. This is rationalized by noting that the stochastic process has a type of stochastic stability given by the tendency for r to be drawn toward μ and μ to in turn be drawn toward γ.[6] This means that the probability of reaching the boundary is extremely small on any given realization. Consequently the treatment of the boundary should have little effect on the current valuation of the security, which is an expectation. Nevertheless, allowance must be made in setting up the transition probabilities for this occurrance. For this application we forced all boundaries to be reflecting by setting to 0 the probabilities of moves outside the region of the grid and reallocating the remaining probability mass to the remaining adjacent gridpoints equiproportionately.

4.3 The Starting Point

For practical application one wishes to be able to value a security at an arbitrary initial state, not just one that lies at a gridpoint on a preselected grid. To allow for this, the actual largest and smallest values of r and μ used for the grid are selected so that the initial state lies exactly at a gridpoint. This must be done prior to computing the arrays of transition probabilities to be used in the random walks of course. The lower and upper levels of 0% and 30% may thus be increased by up to one step — 1.5% for r and .75% for μ).

4.4 Negative Transition Probabilities

There are dramatic differences in the coefficients of the valuation pde, or equivalently in the expected drift and volatility of the state variables, across the r, μ space. This leads to some problems in implementing the Monte Carlo scheme. Equation (20) may stipulate negative values for some transition probabilities. The situation arises from trying to represent the continuous time process by moves to adjacent gridpoints only. For example, with a time step of one week, the expected value of r at the next step may be more than one step away from

[6] See [9] for different concepts of stochastic stability. We suspect that stability in the sense of bounding at a low level the probability of ever reaching the artificial boundaries could be formally demonstrated using his stochastic Lyapunov function approach. However we have not done so.

its current value at some locations in plane. There is no way that moves to adjacent gridpoints only could accurately replicate the expected drift. Similarly, for a given mesh size, there are limits on how much volatility can be simulated.

There are several ways that one might deal with this problem. The tradeoff is between computation time and accuracy of the representation of the continuous time process. One route is to shrink the time step k in the random walk.[7] However an infeasibly small stepsize would be required to eliminate the problem for all grid locations. Moreover the problem areas tend to be for states that have very low probabilities of ever being reached, and hence, one suspects, have little influence on the ultimate value.

We handle the problem in a number of ways. First, the relative sizes of g and h were selected to give more similar probabilities of vertical and horizontal moves at the most relevant levels of the term structure. A step size of 1.5% was used in the r direction and .75% in the μ direction to recognize the higher drift and volatility of r.[8] Second, when setting up the transition probabilities (which only need be done once for a given interest rate process), case 2 or case 3 of the difference expression for $U_{r\mu}$ was used instead of case 1 if that gave all positive computed probabilities. And third, the remaining negative probabilities were simply set to 0 and the other probabilities increased to sum to 1. To keep our representation of the true process as faithful as possible, this was done in a manner that maintained the correct marginal distribution of movements in the μ direction. Thus the distribution of movements that had more long lasting implications were distorted the least.

4.5 Numerical Derivatives

For hedging and risk measurement purposes, it is desirable to know the sensitivity of a contract's value to changes in interest rates — that is, the partial derivatives of P with respect to the state variables r and μ. When using a direct method, the whole grid of contingent security values is available at the end of the algorithm. Numerical derivatives may readily be calculated. The Monte Carlo method, however, only generates a value for the chosen initial state. Calculating numerical derivatives would require conducting random walks from adjacent initial states. The additional computations required were held to a minimum by the following procedure. When walking from the initial state, the process is halted after some number of steps (prior to the first payout) and the current value of the discount factor stored. The walk is then restarted. The final

[7] This is analogous to the instability problem encountered when using explicit (as opposed to implicit) finite difference methods to solve pde's. Stability is assured only for sufficiently small values of k relative to g and h. This is not surprising in light of Brennan and Schwartz's [2] demonstration that explicit schemes can be viewed as computing expected discounted values assuming only moves to adjacent gridpoints are possible, while the more complex implicit schemes can be viewed as permitting moves to *any* gridpoint.

[8] These values were settled upon after some experimenting with valuing non-history dependent claims and comparing the results to those of the direct implicit method.

outcome from that point onward is both stored in an array of average values from that starting point, and also multiplied by the stored discount factor to get a score for that random walk. Upon completing the valuation of the security for the chosen initial state, one proceeds to value it starting from adjacent states. These walks are also paused after the preset number of steps. A check is made whether previous walks have gone through the same gridpoint at that same time. If so, then the average outcome of those previous walks is substituted for completing of the current one. If not, then the walk continues to the security's maturity. This procedure accomplishes two objectives. It reduces the number of walks conducted.[9] It also reduces the random error introduced into the numerical derivatives since the valuation errors associated with the initial and adjacent states will be positively correlated.

4.6 The Floating Rate

The state contingent floating rate $\phi(r, \mu)$ must be specified for the cap contracts. For illustration we suppose it is the yield on 13 week discount bonds. The ADI method was used to calculate a 41 by 41 matrix of theoretical values of \$1 maturity value 13 week bills for r and μ values from 0 to 30%. Each entry was converted to an annualized simple interest rate. A matrix of contingent floating rates appropriate for the grid spacing of the Monte Carlo procedure was extracted from this array using second order Taylor Series interpolation when necessary.

The effect of using different instruments as the basic for the floating loan rate is examined by substituting different bill maturities in the first step above.

5 Testing the Monte Carlo Procedure

To test the Monte Carlo procedure we compare its valuation results with those obtained using the ADI method for history independent claims: discount bonds, European options on same, and instantaneous rate cap contracts. This exercise has a dual purpose. First, it helps in selecting a mesh size for the Monte Carlo procedure that minimizes the bias introduced by *ad hoc* treatment of the problems discussed in the preceding section. Second, it allows one to verify whether the estimated standard error of the Monte Carlo value, generated as a byproduct of the sampling, is a reliable guide to the actual random error inherent in the procedure.[10]

[9] Typically less than 20% of these secondary walks had to proceed to maturity. Moreover those stored averages that were less reliable because they were based on smaller numbers of realizations were also those of less consequence for the valuation in the secondary walk because they represented infrequently encountered states.

[10] Pseudorandom numbers for the simulations were generated using the IBM Scientific Subroutine algorithm $x_n = 65539 x_{n-1} \bmod 2^{31}$, with a seed of any odd integer between 0 and 2^{31}. Since this algorithm has been subject to criticism, runs were also made using the algo-

Table 1: Comparison of Monte Carlo and ADI pricing of discount bonds

Maturity	r	$\mu = 6\%$		$\mu = 9\%$		$\mu = 12\%$	
	6%	98.45	(.006)	97.90	(-.005)	97.38	(.016)*
13 weeks	9%	98.26	(.017)**	97.65	(.000)	97.09	(.025)**
	12%	98.11	(.014)*	97.46	(.011)*	96.85	(.013)*
	6%	93.60	(.045)*	92.56	(-.028)	91.66	(-.023)
52 weeks	9%	91.59	(.056)*	90.32	(.012)	89.25	(.030)
	12%	89.73	(-.055)*	88.27	(-.037)	87.04	(-.028)
	6%	69.54	(.337)**	68.65	(.322)**	67.90	(.064)
260 weeks	9%	58.53	(-.222)*	57.56	(-.103)	56.76	(-.067)
	12%	49.25	(.093)	48.26	(-.014)	47.45	(.051)
	6%	46.21	(.387)*	45.59	(.509)*	45.07	(.424)*
520 weeks	9%	33.10	(-.020)	32.52	(-.065)	32.04	(-.211)
	12%	24.16	(-.143)	23.64	(-.284)*	23.22	(-.246)*

All prices are for bills with $100.00 maturity value. Main entry is the equilibrium bill price calculated by the Alternating Direction Implicit (ADI) method. In parentheses is Monte Carlo method estimate minus the ADI value. * indicates that this difference is more than one estimated standard deviation from 0 but less than two. ** indicates that it is more than two standard deviations from 0 but less than three.

Based on these tests, a grid was selected with r and μ values from 0 to 30% in steps of 1.5% for r and .75% for μ. The time step k was set at 0.5 weeks. Any finer grid mesh, or longer time step without coarsening the grid mesh, increased the problem of negative transition probabilities. The number of random walks to maturity was set at 2000 for each test.

Table 1 provides the ADI valuations of $100 face amount pure discount bonds for various interest rate states, and the amounts by which the Monte Carlo valuations differed from them. The maturities range from 13 weeks to 10 years. The Monte Carlo valuations were always within three estimated standard errors from the direct solutions, and were within one standard error in slightly more than half of the cases. The estimated standard pricing error ranged form about .01% for 13 week bonds to about 1% for 10 year bonds. Casual inspection does not reveal any consistent bias in the Monte Carlo pricing of these instruments.

rithm URN11 of Dudewicz and Rolley [5]. The results were basically similar to those reported here, suggesting that the particular shortcomings of these algorithms were not influencing our results.

Table 2: Monte Carlo and ADI pricing of instantaneous rate caps

	Cap level relative to current rates[a]	Term of Rate Cap 52 weeks		260 weeks	
Current μ, r equal 6%	0%	.638	(-.008)	6.192	(-.225)**
	3%	.068	(.005)	1.958	(-.140)**
Current μ, r equal 12%	0%	1.519	(-.001)	10.828	(-.166)
	3%	.612	(-.003)	6.208	(-.225)*

[a]The absolute cap rates are thus 6.31 and 9.31%/year when the initial state is $\mu = r = 6\%$ and 12.99 and 15.99%/year when the initial state is $\mu = r = 12\%$.

All caps are on a $100 constant balance loan with quarterly interest payments and quarterly readjustment of the loan rate to the prevailing simple interest rate on 13 week discount bonds.

Tables 2 and 3 provide analogous results for European put and call options on 13 week discount bonds. The striking prices were selected so that the diagonal entries for each time to expiry represented values of options that were currently 'at the money'. Entries above the diagonal are out of the money puts or in the money calls; entries below the diagonal are in the money puts or out of the money calls. These results check the Monte Carlo method's simulation of the tails of the interest rate distribution. The pricing errors tend to be smaller in absolute terms than the errors for discount bonds, but somewhat larger in proportionate terms, especially for low value out of the money options. There appears to be some overpricing of the long term put options at high levels of interest rates (last row of Table 2). This is consistent with the fact that the negative transition probability problem becomes greater at the more volatile upper tail.

Finally, Table 4 compares the Monte Carlo and ADI valuation of instantaneous rate caps. This test is somewhat redundant since an instantaneous rate cap is equivalent to a portfolio of European put options on discount bills. The fact that the Monte Carlo value is too low in seven out of the eight cases is thus probably coincidental, since the put option values (Table 2) were evenly split between being too high and too low. Again the absolute size of valuation errors is not obviously related to the value of the cap, implying larger proportionate errors for cap levels above current floating rates.

These comparisons suggest that the Monte Carlo procedure is a reliable alternative to direct solutions of the valuation pde. Moreover the estimated standard valuation error is a fair indicator of the potential true error. Let us now use the procedure to examine the history dependent average and hybrid

Table 3: Cap values for different terms to maturity

Cap term	Cap level relative to current rates	Cap Type Instantaneous	Hybrid	Average
1 year	0%	1.02	.94	.88
	3%	.26	.08	.07
3 year	0%	4.78	4.24	3.91
	3%	1.92	.86	.75
5 year	0%	8.71	7.77	7.18
	3%	4.06	2.17	1.95

rate caps.

6 Rate Cap Properties

Having verified the accuracy of the Monte Carlo procedure, we now use it explore the properties of the three types of interest rate caps. Specifically, we examine the relation between cap value and the following factors: time to maturity, level of interest rates, cap level, rollover frequency, floating rate maturity and interestrate history.

Our benchmark contract is a three year (156 week) cap on a \$100 constant balance loan. The floating rate is taken to be the simple interest yield on 13 week discount bills. The loan rate is reset every 13 weeks. The beginning of the loan is the first rollover date. The initial state is taken to be $(\mu, r) = (9\%, 9\%)$ unless otherwise indicated. This implies a floating rate of 9.61%/year. Each valuation uses a sample of 3000 random walks.

Table 5 lists values for the three cap types for maturities of 1, 3 and 5 years, and cap rates of 0% and 3% above the initial floating rate. The entries are dollar values per \$100 of loan. Thus the equilibrium value of a contract that caps the average rate on a 3 year floating rate loan at 3% above current rates, or 12.61%, would be \$.75 . As expected, longer maturity caps are more valuable, with instantaneous caps more valuable than hybrid caps and hybrid caps more valuable than average caps. If the cap rate is at the money, the three cap types do not differ greatly in value, with the average cap worth 80–90% of an instantaneous cap. However there is considerable difference for out of the money caps, with the difference being proportionately the greatest for shorter term caps. This occurs because it will on average take a fair amount of time for rates to reach the cap rate, during which time a positive balance will build up

Table 4: Cap values at different interest rate levels

State μ / r	Cap level relative to current rates	Cap Type Instantaneous	Hybrid	Average
6%/6%	0%	3.04	2.66	2.41
	3%	.69	.15	.14
9%/9%	0%	4.78	4.24	3.91
	3%	1.92	.86	.75
12%/12%	0%	6.41	5.76	5.30
	3%	3.27	1.82	1.57

Values are for 3 year caps on $100 constant balance loans. Loan interest rate is readjusted every 13 weeks to the yield on 13 week discount bills. This rate is 6.31, 9.61 and 12.99% respectively for the three interest rate states listed.

in the 'reserve account' $s(t)$. The average and hybrid caps do not pay out until this reserve is exhausted. Thus the probability of any payout being made over the relatively short term is much lower than for instantaneous caps.

Table 6 displays the effect of the level of interest rates on cap values. The main point to observe here is that the cap values all rise with higher interest rates. For 0% caps, the increase is slightly more than in proportion to the level of rates. This results from the stochastic process chosen to model rate movements: The volatility of the state variables μ and r was assumed proportional to their levels. For out of the money caps the increase is substantially more than in proportion to the level of rates. This results from the fact that the fixed absolute cap level of 3% represents a proportionately smaller rise in rates if rates are high to begin with.

Let us focus for a moment on the effect of the cap level for a given initial state. Table 7 varies the cap level from 2% in the money to 3% out of the money for three year caps. Once in the money, a 1% drop in cap rate causes the cap value to rise by almost the value of a certain 1% payment flow over the cap term. More interestingly, as the cap level moves into the money the values of the three cap types converge. This convergence is in both absolute and proportionate terms. In fact, for absolute cap rates of zero or infinity, the three cap contracts are equivalent. With a cap rate of zero, each contract pays all the interest on the loan; with a cap rate of infinity, each contract will pay out nothing with probability one.[11]

[11] With infinite cap rates, all caps would clearly have zero values. With zero cap rates, their values would be the market value of the stream of floating interest rate payments on the loan. This would be the market value of the loan with interest payments made, which would be

Table 5: Cap values for different cap rates

State μ / r	Cap level relative to current rates	Cap Type Instantaneous	Hybrid	Average
	-2%	8.62	8.45	8.26
9%/9%	0%	4.78	4.24	3.91
	3%	1.92	.86	.75

A seemingly minor aspect of a rate cap contract is the frequency with which the loan rate is reset. Table 8 values at the money 3 year caps with rollover periods from one to six months. The floating rate can either have the maturity corresponding to the rollover period, or be can be kept at the benchmark 13 week rate. It is apparent that cap values increase with the frequency of rate rollover, and decrease with the maturity of the instrument used to define the floating rate. The frequency effect derives from two sources. First, the more frequently is the rate reset, the larger is the number of options represented by the cap contract (although each is for a shorter term); second, the more frequent is the rate reset, the sooner it is to exercise of the first option, which occurs at the first rollover date after the contract's inception. The floating rate maturity influences the cap value through its impact on the volatility of the loan rate. The structure of the interest rate process implies that short term end of the yield curve is more volatile than he long term end. Thus a cap based on 4 week rates is an option on a more volatile asset than one based on 13 week rates, giving it higher value. What is surprising is the quantitative importance of this aspect of the contract. Decreasing the rollover period and floating rate maturity from 26 to 4 weeks raises the cap values by almost 40%.

One must know the sensitivity of cap values to changes in the interest rate state variables in order to hedge a position in interest rate cap contracts, or to properly assess one's exposure to interest rate risk if not hedged. Table 9 gives estimated partial derivatives of cap values with respect to μ and r. First note that values are more sensitive to μ than to r, particularly for longer term caps. This would be countered to some extent by the fact that r is the more volatile factor. However fluctuations in μ have a longer term impact on the general level of rates. The striking feature of these results is that at the money average and hybrid caps are actually more sensitive to interest rate fluctuations than are instantaneous caps — despite their lower values. Thus larger hedge positions would have to be taken in some circumstances to cover a writing of average

$100 if the rollover frequency is the same as the floating rate maturity, minus the market value ofthe loan with no interest payments made, which would be the value of a $100 maturity value discount bond.

Table 6: Cap values for different rollover frequencies

Weeks between rate changes	Floating rate maturity in weeks	Cap Type Instantaneous	Average
4 weeks	4 weeks	5.74	4.58
4 weeks	13 weeks	5.15	4.18
13 weeks	13 weeks	4.78	3.91
26 weeks	13 weeks	4.27	3.57
26 weeks	26 weeks	3.81	3.28

Values are for 0% caps on 3 year constant balance loans of \$100. The current state is assumed to be $\mu = r = 9\%$.

caps than a writing of instantaneous caps on the same loan. The situation is consistent with the observation from Table 7 that the caps converge in value as they become more in the money. For this to occur, average caps must rise faster in value with 'in the moneyness' than instantaneous caps. For well out of the money caps the absolute risk exposure of average caps is less than for the others, though the exposure relative to value is still greater.

Finally, let us look at the influence of interest rate history on cap values. This is only relevant for caps that have been in place for some time, but is necessary for dynamic hedging strategies. Table 10 provides values for caps have one year and three years remaining in their term for possible current levels of the history state variable — the hypothetical reserve account. A value of \$3 for $s(t)$ on a \$100 loan could arise, for example, from floating rates having averaged 1% below the cap rate for the preceding three years, or 3% below the cap rate for just one year. This is of no consequence for the value of instantaneous caps, of course. A higher positive value for s reduces the value of average and hybrid caps since it represents the payouts that will not be made if interest rates rise. This effect is more pronounced the closer the caps are to being in the money. For well into the money caps the effect would be one-for-one since a payout would be received at the first opportunity and it's amount would be reduced by the amount s.

Negative values for s are meaningless for hybrid caps. By construction, it has a minimum value of 0. However increased negative values of s reduce the value of average caps. The effect is greater the more out of the money is the cap rate. This occurs because the borrower is obliged to actually make payments to the issuer of the cap until the reserve account is built up to 0. His total payment is limited by the cap rate. The more out of the money is the cap, the

Table 7: Sensitivity of cap values to changes in interest rates

Cap term	Cap level relative to current rates	Cap Type Instantaneous		Hybrid		Average	
		$\partial U/\partial\mu$	$\partial U/\partial r$	$\partial U/\partial\mu$	$\partial U/\partial r$	$\partial U/\partial\mu$	$\partial U/\partial r$
1 year	0%	.44	.14	.45	.16	.46	.15
	3%	.19	.05	.09	.03	.08	.02
3 year	0%	1.94	.14	2.07	.19	2.11	.15
	3%	1.06	.05	.76	.06	.69	.04
5 year	0%	2.99	.07	3.25	.13	3.21	.08
	3%	1.81	.00	1.57	.04	1.43	.01

Results are for caps on \$100 constant balance loans, with rate readjusted every 13 weeks, in state $\mu = r = 9\%$. Entries are changes in \$ value of cap per 1%/year change in the state variables μ and r respectively. A 1%/year increase in both state variables together is associated with a roughly 1.1%/year increase in the equilibrium yield on 13 week bills.

Table 8: Cap values for different rate histories

Cap term	Accumulated reserve $s(t)$	Cap Type Instantaneous		Hybrid		Average	
		0% cap	3% cap	0% cap	3% cap	0% cap	3% cap
1 year	3	1.02	.26	.07	.00	.07	.00
	0	1.02	.26	.94	.08	.88	.07
	-3	1.02	.26	—	—	.59	-1.95
3 year	3	4.78	1.92	2.40	.40	2.31	.38
	0	4.78	1.92	4.24	.86	3.91	.75
	-3	4.78	1.92	—	—	3.35	-1.35

Results are for caps on \$100 constant balance loans, with rate readjusted every 13 weeks, in state $\mu = r = 9\%$. The history state variable $s(t)$ has a minimum value of 0 for the hybrid rate cap.

more immediate is this reimbursement likely to be. For very high cap rates the effect would be one-for-one. The effect of s on the value of average caps is thus not monotonic.

The second point to note is that average caps can take negative values if current interest rates leave it out of the money and past rates have averaged more than the cap rate ($s < 0$). Notice that the value is less negative for the cap with three years remaining. There are better prospects of receiving future benefits from this cap, despite the likelihood of negative near term payouts, than for the one year cap. There is a bound on how negative the cap value can become, as s becomes more negative, that is determined by the cap rate, that limits the flow of negative payouts, and the remaining term. For practical purposes, the potential negative value of the average cap implies there is two way default risk associated with the contract. Situations can arise where the cap purchaser has an incentive to renege. Such is not the case with instantaneous or hybrid caps, which carry no default risk for the issuer.

7 Conclusion

We have, we hope, demonstrated the viability of Monte Carlo methods for valuing complex contingent claims. The virtue of the procedure lies in the ease with which deterministic functions of the time path of exogenous factors can be introduced into the analysis. Incorporating such variables adds little computational difficulty or expense. The procedure is ideally suited to handling history dependent claims such as those that have recently emerged in financial markets.

The procedure as given is general. It can be adapted to other two factor interest rate processes, or joint interest rate/stock price processes, by simply changing the coefficient functions used to define the transition probabilities, and the rules that determine the contract payouts.

Some lessons can be learned from the analysis of the interest rate cap contracts. The cap values were more sensitive to some contract features than might have been anticipated *ex ante*. In particular, they were influenced greatly by the level of interest rates and the rollover frequency. Moreover the significance of having an average cap specification as opposed to an instantaneous cap specification depended critically on whether rates were currently above or below the cap rate.

Some of these results depend heavily, no doubt, on the particular specification chosen for the interest rate process. To the extent that this is the case, it serves as a caution against relying exclusively — either explicitly or implicitly — on any one specification of those aspects of the financial environment that are stochastic.

References

[1] Boyle, Phelim P., "Options: A Monte Carlo Approach," *Journal of Financial Economics*, 4 (1977), 323–38.

[2] Brennan, Michael J., and Eduardo S. Schwartz, "Finite difference methods and jump processes arising in the pricing of contingent claims: a synthesis," *Journal of Financial and Quantitative Analysis*, 13 (September 1978), 461–474.

[3] Cox, John C., Jonathan E. Ingersoll, Jr., and Stephen A. Ross, "An Intertemporal General Equilibrium Model of Asset Prices," *Econometrica*, 53 (March 1985), 363-84.

[4] Curtiss, J. H., "Sampling Methods Applied to Differential and Difference Equations," *Proceedings of Seminar on Scientific Computation*, IBM Corporation, November 1949.

[5] Dudewicz, Edward J., and Thomas G. Rolley, *The Handbook of Random Number Generators and Testing*, Columbus, Ohio: American Sciences Press, 1981.

[6] Hammersley, J. M., and D. C. Handscomb, *Monte Carlo Methods*, London: Methuen, 1964.

[7] Jacobs, Rodney L., and Robert A. Jones, "A Two Factor Latent Variable Model of the Term Structure of Interest Rates," mimeo, Simon Fraser University, February 1986.

[8] Klahr, Carl N., "A Monte Carlo method for the solution of elliptic partial differential equations," in Anthony Ralston and Herbert S. Wilf (eds.), *Mathematical Methods for Digital Computers*, vol. 1, pp. 157–64, New York: John Wiley & Sons, 1960.

[9] Kushner, Harold J., *Stochastic Stability and Control*, New York: Academic Press, 1967.

[10] Lattes, Robert, *Boundary Value Problems Occurring in Mathematical Physics*, 1969.

[11] McKee, S., and A. R. Mitchell, "Alternating direction methods for parabolic equations in two space dimensions with a mixed derivative," *The Computer Journal*, 13 (February 1970), 81–86.

[12] Ramaswamy, Krishna, and Suresh M. Sundaresan, "The Valuation of Floating Rate Instruments: Theory and Evidence," mimeo, University of Pennsylvania, December 1985.

[13] Vemuri, V., and Walter J. Karplus, *Digital Computer Treatment of Partial Differential Equations*, Englewood Cliffs: Prentice-Hall, 1981.

List of Figures

List of Tables