Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of Banking & Finance 29 (2005) 2883–2907

Journal of BANKING & FINANCE

www.elsevier.com/locate/jbf

Risk management implications of time-inconsistency: Model updating and recalibration of no-arbitrage models

Andrea Buraschi [a,b,c,*], Francesco Corielli [d,*]

[a] London Business School, Institute of Finance, Sussex Place, Regents Park, London NWI 4SA, UK
[b] Columbia University, New York, United States
[c] CEPR, London, UK
[d] Universita' Bocconi, Istituto di Metodi Quantitativi, Viale Isonzo 25, 20100 Milan, Italy

Available online 27 April 2005

Abstract

A widespread approach in the implementation of asset pricing models is based on the periodic recalibration of its parameters and initial conditions to eliminate any conflict between model-implied and market prices. Modern no-arbitrage market models facilitate this procedure since their solution can usually be written in terms of the entire initial yield curve. As a result, the model fits (by construction) the interest rate term structure. This procedure is, however, generally time inconsistent since the model at time $t = 0$ completely specifies the set of possible term structures for any $t > 0$. In this paper, we analyze the pros and cons of this widespread approach in pricing and hedging, both theoretically and empirically. The theoretical section of the paper shows (a) under which conditions recalibration improves the hedging errors by limiting the propagation of an initial error, (b) that recalibration introduces time-inconsistent errors that violate the self-financing argument of the standard replication strategy. The empirical section of the paper quantifies the trade-off between (a) and (b) under several scenarios. First, we compare this trade-off for two economies with and without model specification error. Then, we discuss the trade-off when the underlying economy is not Markovian.

* Corresponding authors.
E-mail addresses: aburaschi@london.edu (A. Buraschi), francesco.corielli@uni-bocconi.it (F. Corielli).

0378-4266/$ - see front matter
doi:10.1016/j.jbankfin.2005.02.002

JEL classification: D9; E3; E4; G12

Keywords: Time-consistency; No-arbitrage; Term structure; Options; Bonds; Risk management

1. Introduction

The derivative market has grown considerably over the last thirty years, both in terms of volume and product complexity. Increased trading activity has led to an increase in financial institutions' exposure to derivatives. This, in turn, has caused increasing interest in model risk. Broadly speaking, model risk can be attributed to either an incorrect model or to an incorrect implementation of a model.[1] The Basel Committee on Banking Supervision has defined precise rules to monitor certain risks. Specifically, banks must post reserves in proportion to their market and credit risk exposure. In 1999, the Basel Committee declared its intention to subject all banks to capital requirements for operational risk, placing particular emphasis on the component of model risk (Basel Committee Banking Supervision, 1999). In this paper, we study the consequences of one of the most widespread operational aspect of risk management: the recalibration of the model.

In the last fifteen years, more accurate models of the term structure have developed along two distinct paths. The first and more traditional approach consists of specifying a relatively parsimonious equilibrium factor model of the spot interest rate. No-arbitrage conditions are then used to derive implications for the yield curve given an initial value of the spot interest rate. Finally, historical observations are used to match a desirable set of moment conditions. In this framework, the econometrician requires the model to be correct, at least "on average". Important contributions to this literature include Vasiček (1977), Cox et al. (1985), Constantinides (1992), Longstaff and Schwartz (1992), Duffie and Kan (1996), Duffie and Singleton (1997), Dai and Singleton (2000), Backus et al. (1999) and Buraschi and Jiltsov (2003, 2005).

The second, more modern approach consists of specifying a stochastic process of the forward rate curve. No-arbitrage conditions are then used to derive implication for the yield curve, given an initial value of the (infinite dimensional) yield curve. Since Ho and Lee (1986), the "modern" no-arbitrage approach has gained increasing popularity. This includes the models by Black et al. (1990), Black and Karasinski (1991), Cooley et al. (1992), and Hull and White (1990, 1993). It is claimed that a crucial advantage of such an approach is that the model can fit the term structure exactly (both the yield curve and the term structure of volatilities). By comparison, parsimonious equilibrium models, it is argued, give rise to pricing errors that are usually considered unacceptable by practitioners for marking-to-model purposes (unless the parameters are continuously cross-sectionally recalibrated). For instance,

[1] See Derman (1996) for a discussion on the impact of model risk from an investment banking perspective.

the four parameter Vasicek model can at most fit two points on the yield curve and two points on the volatility curve given the spot interest rate on a specific date. At a later date, without recalibration, the same specification will typically miss those four points. No trader would be willing (or allowed) to value his book and design dynamic hedging strategies, with a model that already grossly misprices the benchmarks. For instance, a 1% pricing error for the underlying bond may lead to a 25% error in the value of an option written on the same bond. This motivated a widespread use of recalibration in the implementation of a model.

Hull and Suo (2000) describe the difference between no-arbitrage and equilibrium models as follows: "A no-arbitrage model is a model designed to be exactly consistent with today's term structure of interest rates. In an equilibrium model, today's term structure of interest rates is an *output*. In a no-arbitrage model, today's term structure of interest rates is an *input*." The key difference that makes this possible, as summarized by Hull and Suo (2000), is that "in an equilibrium model, the drift of the short rate is not a function of time. Modern no-arbitrage models are becoming very popular. In Bloomberg, the theoretical values of interest rate derivatives are based on a modified version of the Ho and Lee (1986) model that restricts interest rates to being positive. The same model is used as the default option for computing option-adjusted yield spreads in callable bonds. Bloomberg gives theoretical prices based on two alternatives: Black et al. (1990) and Hull and White (1990) models.

An overlooked drawback of recalibration is that it violates the self-financing condition of the replication strategy since it implies a change in the conditional distribution of the process with respect to which the replicating portfolio weights are computed. The risk manager that forces any "modern" no-arbitrage model to exactly fit the time series of the yield curve is deemed to generate hedging errors if the continuously updated model is used to hedge a book of derivatives. At any point in time, all primitive securities are correctly priced. When the risk manager unwinds his positions, however, he discovers a non-zero cost of carry in his hedging strategy which, over time, can be seen to generate unexpected losses (or gains).

Academics have expressed concern that recalibrating or simply re-initializing a term structure model to a newly observed yield curve may contradict the initial assumptions of the model. Dybvig (1989) argues that the changes in parameter values should be interpreted as direct evidence of the model mis-specification. Black and Karasinski (1991) are more explicit: "When we value the option, we are assuming that its volatility is known and constant. But a minute later, we start using a new volatility." The fixed income model approach is similar. "We can value fixed income securities by assuming we know the one-factor short rate process. A minute later we start using a new process that is not consistent with the old one." Backus et al. (1999) summarize their results by saying that "the misrepresentation of fundamentals, whatever their form, cannot generally be overcome by adding arrays of time-dependent parameters."

Our paper formalizes the above statements of Black and Karasinski (1991) and Backus et al. (1999) and quantifies this particular form of model risk, which we call time-inconsistency. Time-inconsistency arises from attempts to fix model inadequacies as the model is being used. This can be a day-by-day updating of the parameters (which should be constant) or a daily restarting of the model from a new set of initial conditions (which would be incompatible with the dynamics of the model itself). We concentrate on the latter behavior as it is widespread in the implementation of "modern" no-arbitrage term structure models.

In this paper, we address the following three questions of both practical and theoretical interest. First, although it is clear that model updating trivially reduces the pricing errors of a model at the time of the updating, is this procedure correct from the perspective of the risk manager? We show that when the pricing model is based on diffusion processes, model updating reduces the upper bound of pricing and hedging errors due to model mis-specification. Thus, model updating is indeed theoretically justifiable, even if time-inconsistent, when the underlying model is mis-specified.

Second, what is the extent of potential pricing and hedging improvement gained from model updating? In other words, how much can model updating overcome model specification errors? We simulate an interest rate model and then use slightly mis-specified versions of the correct model for pricing and hedging. We find that although model updating can improve the hedging performance of an incorrect model, the gains of model updating are small. When hedging a cross-section of options in a Hull and White economy, model updating reduces the hedging errors due to model mis-specification by less than 10%. Since the model is continuously recalibrated to the benchmark securities, it always seems correct when evaluated with respect to the benchmarks. Model updating, however, simply masks model mis-specification.

Third, what are the features of a model that make model-updating more or less effective? For instance, do model updating procedures imply different results if applied to Markov versus non-Markov economies? We simulate the hedging performance of a continuously updated Markov model when the underlying process is not Markovian. The extent of non-Markovianity is indexed by a parameter. We find that while model updating can still improve the hedging performance of the model, the improvement is substantially smaller than for Markov models. For at-the-money options, the hedging errors of a continuously recalibrated model in a non-Markov economy are three times larger than the errors in the Markov specification of the same model.

This paper is organized as follows: the remainder of this Section discusses the paper's contribution in light of the current literature. Section 2 introduces the notation and describes the characteristics of the problem we want to study. Section 3 discusses the theoretical implications of initial conditions updating and gives a formal definition of time-inconsistency. Section 4 quantifies the effect of initial conditions updating both in Markov and non-Markov settings using Monte Carlo simulations. Section 5 briefly discusses suggested approaches to circumvent the problem of initial conditions updating, namely market models and stochastic string models. Section 6 concludes.

1.1. Related literature

Our work related to Backus et al. (1999) shows that the ability of no-arbitrage models to correctly price a subset of primitive securities may not extend to more general forms of state contingent claims. For instance, the authors show that when interest rates exhibit mean reversion, the popular Black, Derman and Toy model overprices call options on long-term bonds relative to short-term bonds. Backus, Foresi and Zin warn against the use of additional parameters in arbitrage-free models unless they are complemented with an analysis of the fundamentals. The latter might include mean reversion, multiple factors, stochastic volatility and non-normal interest rate distributions. We follow the spirit of Backus et al. (1999), but differ in our approach to the material. Empirically, we quantify the costs and benefits of the (time-inconsistent) re-calibration procedure in the presence of both a correctly specified model and an incorrectly specified model. In contrast, Backus et al. (1999), focus their analysis on models which are mis-specified specifically because of mean reversion in the data.

Our work is also related to Brandt and Yaron (2001). They propose a two step empirical approach to obtain accurate cross-sectional fits of the term structure that is consistent with the dynamics of a set of underlying factors. They achieve this by extracting preference parameters and by allowing them to be time dependent. Assuming that the underlying factors are perfectly observable and that they do not depend on the preference parameters, Brand and Yaron use semi non-parametric methods to reverse engineer the term structure sequence and recover a series of preference parameters that is consistent with a sequence of observed yield curves.

The empirical analysis in this paper quantifies the hedging costs of time-inconsistent models. Our results emphasize the importance of studying time-consistent models. The contribution of Brandt and Yaron (2001) provides one example. Other contributions include Hull and Suo (2000) with regards to using implied volatility function models in option pricing, and Björk and Christensen (1999) who study the relationship between cross-sectional and time series properties of forward interest rate models. Santa Clara and Sornette (2001) suggest an approach that generalizes the Heath et al. (1992) framework. They model the forward curve with stochastic string shocks. This approach allows an exact fit of the yield curve, though the hedging strategy requires an infinite number of instruments. A similar approach is explored by Kennedy (1994, 1997) and Goldstein (2000). Our analysis considers practical, finite dimensional implementations of these models, where hedging is performed with respect to a limited number of liquid instruments.

2. The HJM family of models

The Heath et al. (1992) family of interest rate models is based on a description of the pricing dynamics of fixed maturity zero coupon bonds, which is generated by the instantaneous forward rates evolution. This can be represented as

$$\mathrm{d}f(t,T) = a(t,T)\mathrm{d}t + \sum_{j=1}^{q} b_j(t,T)\mathrm{d}W_t^j \tag{2.1}$$

where $\mathrm{d}W_t$ is a q-dimensional vector of orthogonalized standard Brownian motions. The price at time t of one unit of a non-defaultable bond maturing at time T is defined as $P(t,T)=\exp(-\int_t^T f(t,v)\,\mathrm{d}v)$. Under the usual regularity conditions, the zero coupon bond prices evolution is given by

$$\mathrm{d}P(t,T)/P(t,T)=[r(t)+\alpha(t,T)]\mathrm{d}t+\sum_{j=1}^{q}\beta_j(t,T)\mathrm{d}W_t^j \tag{2.2}$$

where $\alpha(t,T)=-\int_t^T a(t,v)\,\mathrm{d}v+\frac{1}{2}\sum_{j=1}^{q}(\int_t^T b_j(t,v)\,\mathrm{d}v)^2$ and $\beta_j(t,T)=-\int_t^T b_j(t,v)\,\mathrm{d}v$. The drift in excess of the risk free rate, $\alpha(t,T)$, is the reward for the risky position in the discount bond $P(t,T)$. It is well known that the absence of arbitrage opportunities on $P(t,T)$ is equivalent to the existence of a q-dimensional vector of predictable processes, $\gamma_j(t)$, such that

$$\alpha(t,T)=\sum_{j=1}^{q}\gamma_j(t)\beta_j(t,T),\quad \forall T. \tag{2.3}$$

If this no arbitrage condition is satisfied, the Fundamental Theorem of Asset Pricing implies (i) that there exists a probability measure $\widetilde{Q}$ with respect to which $\widetilde{W}_t^j=W_t^j+\gamma_j(t)t$ are q Brownian motions and (ii) that the drift of $\mathrm{d}P(t,T)$ is equal to $P(t,T)r(t)$. It then follows that (a) the price of any contingent claim paying a function of $P(t,T)$ can be represented as the conditional expected value of the future cash flows, discounted at the rate $r(t)$; (b) the drift of the forward rate process must be a function of the forward volatilities term structure. Under the objective probability measure,[2] this condition is given by

$$a(t,T)=-\sum_{j=1}^{q}\gamma_j(t)\int_t^T b_j(t,v)\,\mathrm{d}v+b(t,T)'\int_t^T b(t,v)\,\mathrm{d}v. \tag{2.4}$$

Implementing the model requires two crucial specifications:

1. An initial condition for $P(t,T)$, i.e. the choice of $P(0,T)$.
2. The term structure of volatility $\beta(t,T)$.

The literature on the Heath, Jarrow and Morton approach discusses specification (2) in detail. This is done in the context of pricing a given set of quoted options. It is usually argued that $P(0,T)$ can be arbitrary so that it provides for an exact fit of the *initial* term structure at time 0 as well as at any future time $t>0$. This assertion is incorrect. Any (Heath et al., 1992) model allows *only one* degree of freedom: the choice of one (and only one) yield curve initial condition at time 0. From this, the model completely determines the future movement of the yield curve, and therefore, its possible shapes at any future date. Thus, "modern" no-arbitrage models can

[2] Under the risk neutral probability measure $\widetilde{Q}$, by construction $\gamma_j(t)=0$. Thus, the no-arbitrage process of the forward rate under the risk neutral measure $\widetilde{Q}$ is:

$$\mathrm{d}f(t,T)=b(t,T)\int_t^T b(t,v)\,\mathrm{d}v+\sum_{j=1}^{q}b_j(t,T)\,\mathrm{d}\widetilde{W}_t.$$

provide an excuse for an "inconsistent" user to run a continuous recalibration by inputting the new "initial" term structure. However, given a specific functional form $P(0,T)$ and $\beta(t,T)$, HJM risk management strategies based on the re-initialization of $P(t,T)$ at each new date t would violate the dynamic no-arbitrage constraint.

2.1. Model-updating and time-inconsistency

We begin by giving a general definition of a time inconsistent updating behavior. A model $\{X_t\}_{0<t<T}$ is a stochastic process defined on some probability space. The process defines a family of conditional probabilities $Q(X_{t_{i+1}},\dots,X_{t_k}|X_{t_1},\dots,X_{t_i})$ for any sequences of dates $t_1 < t_2 < \dots < t_k$. Suppose now that at some generic date $i+\tau$ the risk manager computes prices and hedge factors using a probability $\widetilde{Q}(X_{t_{i+1+\tau}},\dots,X_{t_{k+\tau}}|X_{t_{1+\tau}},\dots,X_{t_{i+\tau}})$ on future observations.

Definition 1. A pricing, or hedging, behavior is time inconsistent (simply: inconsistent) with the model $\{X_t\}_{0<t<T}$ when

$$\widetilde{Q}(X_{t_{i+1+\tau}},\dots,X_{t_{k+\tau}}|X_{t_{1+\tau}},\dots,X_{t_{i+\tau}}) \neq Q(X_{t_{i+1+\tau}},\dots,X_{t_{k+\tau}}|X_{t_{1+\tau}},\dots,X_{t_{i+\tau}}).$$

A frequent case of model inconsistency arises when the model cannot fit (if kept consistent) observed bond prices or observed derivative securities prices and the model is subsequently recalibrated. This implies a change in the original probability distribution from Q to $\widetilde{Q}$.

Example 1. Consider the Ho and Lee model. This is obtained as a special case of Heath et al. (1991) when (2.1) can be written as $f(t,T) = f(0,T) + \sigma^2(Tt - t^2/2) + \sigma W_t$. In this case, the price $P(t,T)$ of a zero coupon bond with expiry $T > t$ and terminal value 1 is given by

$$P(t,T) = \frac{P(0,T)}{P(0,t)} \mathrm{e}^{-[\sigma(T-t)W_t + \frac{1}{2}\sigma^2 T(T-t)t]}$$

where $P(0,T)$ and $P(0,T)$ are the price of the zero coupon bonds observed at time 0. Suppose that at time 0 the risk manager observes that $P(0,s) = \exp(-cs)$, then

$$P(t,T) = \mathrm{e}^{-c(T-t)}\mathrm{e}^{-[\sigma(T-t)W_t + \frac{1}{2}\sigma^2 T(T-t)t]}.$$

At time t the risk manager can only choose W_t. Therefore, at best, he can fit exactly one bond for given parameters values. The usual market practice, as discussed in the introduction, is to restart the model at the observed $P(t,T)'$ so that at time s the risk manager has

$$P(s,T)' = \frac{P(t,T)'}{P(t,s)'} \mathrm{e}^{-[\sigma(T-s)(W_s - W_t) + \frac{1}{2}\sigma^2(T(s-t))(T-(s+t)) + ts(s-t)]}.$$

The original model, however, implies that

$$P(s,T) = \mathrm{e}^{-c(T-s)}\mathrm{e}^{-[\sigma(T-s)W_s + \frac{1}{2}\sigma^2 T(T-s)s]}.$$

$P(s,T)'$ and $P(s,T)$ have different conditional (lognormal) distributions. If the model is meant to price and hedge derivatives, at time 0 the risk manager uses the lognormal distribution of $P(s,T)$ while at time t he changes the distribution to the new (lognormal) distribution for $P(s,T)'$. This violates the self-financing property of the replication strategy.

Inconsistent model re-initialization is very common in stock option pricing as well, especially in the case of generalized deterministic volatility models and/or implied trees in which the volatility surface or deterministic parameters are recalibrated to fit quoted option prices.

Our definition of time-consistency is related to, though different from, the one given by Björk and Christensen (1999). They define the class of cross sectional specifications for the forward rate curve which are compatible with a given HJM interest rate model. They give the following definition: "Assume that a given interest rate model $\mathscr{M}(\ldots)$ is an exact picture of the financial market. Now consider a particular family $\mathscr{G}$ of forward rate curves, ..., and assume that the interest rate model is calibrated using this family. We then say that the pair $(\mathscr{M};\mathscr{G})$ is consistent (or, that $\mathscr{M}$ and $\mathscr{G}$ are consistent) if all forward curves which may be produced by the interest rate model $\mathscr{M}$ are contained within the family $\mathscr{G}$, provided that the initial forward curve is in $\mathscr{G}$. Otherwise, the pair $(\mathscr{M};\mathscr{G})$ is inconsistent."

Suppose a model $\mathscr{M}$ is initialized at time t_0 with an element of $\mathscr{G}$, where $\mathscr{G}$ is consistent in the Bjork–Christensen sense with $\mathscr{M}$. At time t_i, the model is re-initialized with a forward curve outside the set $\mathscr{G}$, as in example 1. This is an inconsistent behavior under both our and their definition. Notice, however, that for a given initial condition at t_0, re-initializing the model at time t_i with an element of $\mathscr{G}$ is a consistent behavior under their definition. It could, however, be inconsistent under ours if the model under the new initial condition implies different conditional probabilities on the same future random events.

An inconsistent behaviour implies a violation of the self financing condition of a replication strategy. In Section 3, we find that the hedging errors implied by a time inconsistent recalibration of the model can be very large with respect to the initial value of the position to be hedged. The intuition is simple. Consider a risk manager who wants to hedge an option on a Treasury Bond but he observes that the model-implied prices on both the underlying asset P and a reference set of options V deviates from the market values. He recalibrates the initial value of the Treasury curve and the implied volatility surface to match the level of the option model price function V to observed market values. In order to satisfy these conditions, the shape and therefore the greeks of this function have to change. The dynamic financing properties of the hedging strategy cum recalibration are therefore different.

2.2. A theoretical justification of model updating

Optimal pricing and hedging policies are defined with respect to the set of conditional Q-distributions of $\{X_t\}_{0<t<T}$. Inconsistency, therefore, violates the zero cost of carry of a Q-based hedging strategy. Thus, model updating could have serious con-

sequences on risk management strategies. This raises the following natural question: *Why is time-inconsistent model-updating so common when implementing modern no-arbitrage models*? In this section we try to find theoretical justifications for the widespread practice of model inconsistency.

To simplify notation, consider the forward price process $P(t,T) = P(\tau)$ of a zero coupon bond with generic maturity T. From (2.2), after adjusting for the riskless rate $r(t)$, the canonical representation of the *forward* price is

$$P(\tau) = P(t) + \int_t^\tau \alpha(s,P)\,\mathrm{d}s + \int_t^\tau \beta(s,P)\,\mathrm{d}W_s.$$

We drop the time of maturity to simplify notation. We also assume that usual conditions sufficient for the existence of a strong solution to the above stochastic differential equations are satisfied.

To describe the tradeoff between model mis-specification error and model updating, we consider three cases:

Case 1. (Correct α,β and incorrect $P(t)$)
The risk manager uses the correct model (i.e. the true α and β of the process) while the initial conditions, $\widehat{P}(t)$, are mis-specified, possibly because of bid–ask spreads or observation errors:

$$\widehat{P}(\tau) = \widehat{P}(t) + \int_t^\tau \alpha(s,\widehat{P})\,\mathrm{d}s + \int_t^\tau \beta(s,\widehat{P})\,\mathrm{d}W_s.$$

Case 2. (Incorrect α,β and incorrect, but model-consistent, $P(t)$)
The risk manager uses a mis-specified model $(\tilde{\alpha},\tilde{\beta})$. The initial conditions $\widetilde{P}(t)$ are consistent with the model, but incorrect with respect to observed market prices

$$\widetilde{P}(\tau) = \widetilde{P}(t) + \int_t^\tau \tilde{\alpha}(s,\widetilde{P})\,\mathrm{d}s + \int_t^\tau \tilde{\beta}(s,\widetilde{P})\,\mathrm{d}W_s.$$

Case 3. (Incorrect α,β and marked-to-market, but model-inconsistent, $P(t)$)
The risk manager uses a mis-specified model $(\tilde{\alpha},\tilde{\beta})$. The initial conditions $P'(t)$ are continuously updated and marked-to-market so that $P'(t) = P(t)$:

$$P'(\tau) = P(t) + \int_t^\tau \tilde{\alpha}(s,P')\,\mathrm{d}s + \int_t^\tau \tilde{\beta}(s,P')\,\mathrm{d}W_s.$$

Notice that, in general, this strategy is inconsistent with the model $(\tilde{\alpha},\tilde{\beta})$.

The following proposition proves the first simple intuition of risk managers using model recalibration: the dynamic consequences of errors on the initial conditions are monotonic.

Proposition 1. *Suppose that on a certain probability space we have a standard, one dimensional Brownian motion W_s driving the diffusion process of $P(.)$, $P_l(.)$ and $P_u(.)$, which share the same α and β but with initial conditions*

$$P_l(t) \leqslant P(t) \leqslant P_u(t),$$

then we shall have:

$$\Pr[P_l(\tau) \leqslant P(\tau) \leqslant P_u(\tau), \forall t \leqslant \tau < \infty] = 1.$$

Proof. Under the hypothesis of Proposition 1, it is possible to show that (Prop. 2.18, Ch. 5, Karatzas and Shreve, 1991):

$$\Pr[P_l(\tau) \leqslant P(\tau), \forall t \leqslant \tau < \infty] = 1,$$

$$\Pr[P(\tau) \leqslant P_u(\tau), \forall t \leqslant \tau < \infty] = 1.$$

Suppose A is the event $[P_l(\tau) \leqslant P(\tau),\ \forall t \leqslant \tau < \infty]$ and B is the event $[P(\tau) \leqslant P_u(\tau),\ \forall t \leqslant \tau < \infty]$ Since $P(A \cup B) = P(A) + P(B) - P(A \cap B)$ we have $P(A \cap B) = 1$ (otherwise $P(A \cup B)$ would be greater than 1). □

Suppose that the initial value at time t is only known to lie in an interval $[P_-(t); P_+(t)]$, possibly because of bid–ask spreads or observation errors. The previous result shows that at any time $\tau \geqslant t$, the "true" process, started at t with $P(t)$, shall (almost surely) stay in the interval $[P_-(\tau); P_+(\tau)]$.

The result has implications in terms of the dynamic impact of model recalibration, which are summarized in the following proposition.

Proposition 2. *Let $V(P(\tau))$ be the payoff of a derivative security, expressed as a function of the underlying asset price. Let $V(\cdot)$ be a bounded function with bounded first and second derivatives with respect to its arguments and let P, $\widetilde{P}$ and P' be the above defined processes. Suppose that there exist finite A and B such that $\sup_{s,z,h}|\alpha(s,z) - \tilde{\alpha}(s,h)| = A$ and $\sup_{s,z,h}|\beta(s,z) - \tilde{\beta}(s,h)| = B$. Then, when computing $E_t\ {}_V(P(\tau))$, the conditional expected (quadratic) error generated by the model-consistent strategy is larger than when using the model-inconsistent marking-to-market strategy. Formally: there exist non-negative constants $C1$ and $C2$ such that*

$$E_t|V(P(\tau)) - V(P'_t(\tau))|^2 \leqslant 2(\tau - t)^2 C1 + 2(\tau - t)C2,$$

$$E_t|V(P(\tau)) - V(\widetilde{P}(\tau))|^2 \leqslant 3|V(P(t)) - V(\widetilde{P}(t))|^2 + 3(\tau - t)^2 C1 + 3(\tau - t)C2.$$

Proof. The stochastic differentials for $V(P)$ and $V(\widetilde{P})$ are:

$$\mathrm{d}V(P) = (V_P\alpha + \frac{1}{2}V_{PP}\beta^2)\mathrm{d}t + V_P\beta\mathrm{d}w,$$

$$\mathrm{d}V(\widetilde{P}) = (V_P\tilde{\alpha} + \frac{1}{2}V_{PP}\widetilde{\beta}^2)\mathrm{d}t + V_{PP}\tilde{\beta}\mathrm{d}w.$$

Set $V_P\alpha + \frac{1}{2}V_{PP}\beta^2 = a$, $V_P\beta = b$, $V_P\tilde{\alpha} + \frac{1}{2}V_{PP}\widetilde{\beta}^2 = \widetilde{a}$ and $V_P\tilde{\beta} = \tilde{b}$. Since V is bounded with bounded first and second derivatives and $|\alpha - \tilde{\alpha}|$ and $|\beta - \tilde{\beta}|$ are bounded, then there exist constants $C1$ and $C2$ such that $\sup|a - \widetilde{a}| = C1$ and $\sup|b - \tilde{b}| = C2$. Using the inequality $(a + b + c)^2 \leqslant 3a^2 + 3b^2 + 3c^2$ we have:

$$|V(P(\tau)) - V(\widetilde{P}(\tau))|^2 \leqslant 3|V(P(t)) - V(\widetilde{P}(t))|^2.$$

$$+3(\tau - t)\int_t^\tau |a - \widetilde{a}|^2\,\mathrm{d}s + 3\left|\int_t^\tau (b - \tilde{b})\,\mathrm{d}w\right|^2.$$

Using Ito isometry we have $E|\int_t^\tau (b - \tilde{b})\,\mathrm{d}w|^2 = \int_t^\tau E(b - \tilde{b})^2\,\mathrm{d}s$ so that

$$E|V(P(\tau)) - V(\widetilde{P}(\tau))|^2 \leqslant 3|V(P(t)) - V(\widetilde{P}(t))|^2 + 3(\tau - t)^2 C1 + 3(\tau - t)C2.$$

On the other hand, using the inequality $(a+b)^2 \leqslant 2a^2 + 2b^2$,

$$|V(P(\tau)) - V(P'_t(\tau))|^2 \leqslant 2(\tau - t)\int_t^\tau |a - \tilde{a}|^2\,\mathrm{d}s + 2\left|\int_t^\tau (b - \tilde{b})\,\mathrm{d}w\right|^2,$$

$$E|V(P(\tau)) - V(P'_t(\tau))|^2 \leqslant 2(\tau - t)^2 C1 + 2(\tau - t)C2.$$

From a direct comparison of $E|V(P(\tau)) - V(\widetilde{P}(\tau))|^2$ and $E|V(P(\tau)) - V(P'_t(\tau))|^2$, the result follows. □

The first result of Proposition 2 shows that if the difference between the parameters of the correct and incorrect models is bounded, and one continuously marks-to-market the initial conditions, then the mean square difference between the implied V processes is bounded. A small (in sup norm) model mis-specification error implies a small upper bound in the mean square error.

The second result shows that marking-to-market the initial conditions of the model reduces the mean square error upper bound when evaluating a function V of the price process. This result is somewhat surprising since re-initializing the model is a time-inconsistent procedure so that one could expect the errors due to time-inconsistency to accumulate over time.

Consider the case of a derivative security with terminal payoff $\varphi(P)$. The value $V(s,P)$ is given by the expected value of the payoff under the *forward* risk neutral process, i.e. $\alpha = 0$, of P. The solution must satisfy the following p.d.e.:

$$V_s + \frac{1}{2}b^2(s,P)P^2 V_{PP} = 0,$$

$$V(T,P) = \varphi(P),$$

where $b^2(s,P)P^2 = \beta^2(s,P)$. Suppose, however, that the risk manager does not know the true process and uses a mis-specified model $\tilde{b}^2(s,P)P^2 = \tilde{\beta}^2(s,P)$.

If the initial conditions are continuously marked-to-market, the true process P enters in both $b(s,P)$ and $\tilde{b}(s,P)$. Thus, if we call V and $\widetilde{V}$ the two option values and $e = V - \widetilde{V}$ the valuation error, we have:

$$e_t + \frac{1}{2}b^2P^2 e_{PP} = -\frac{1}{2}(b^2 - \tilde{b}^2)P^2\widetilde{\Gamma},$$

$$e(T,P) = 0,$$

where $\widetilde{\Gamma} = \widetilde{V}_{PP}$ is the gamma of the derivative security computed using the mis-specified model driven by $\tilde{b}$. This equation is similar to the standard option pricing p.d.e.

but with a non-zero source term $\frac{1}{2}(b^2 - \tilde{b}^2)P^2\widetilde{\Gamma}$ and an homogeneous boundary condition. We can now prove the following result.

Proposition 3. *When the model is continuously re-initialized by updating the initial conditions to market prices and the model mis-specification errors* $|b^2(s, P(s)) - \tilde{b}^2(s, P(s))|$ *and* $|\frac{\partial}{\partial P} b^2(s, P(s)) - \frac{\partial}{\partial P}\tilde{b}^2(s, P(s))|$, *are bounded, then* (*a*) *the pricing errors* $|e(t, P)|$*are bounded, and* (*b*) *the error in delta* $|\Delta_e(t, P)|$ *is bounded.*

Proof. The solution of the error p.d.e. can be written as the expected value of a function of P when P is driven by the β process. From Theorem 7.6, Ch. 5, of Karatzas and Shreve (1991), we have

$$e(t,P) = E^{t,P}\left[\int_t^T \frac{1}{2}\left(b^2(s,P(s)) - \tilde{b}^2(s,P(s))\right)P^2(s)\widetilde{\Gamma}(s,P(s))\,\mathrm{d}s\right].$$

The first result follows.

With regards to the delta of the recalibrated model, let Δ_e be the error in delta of the mis-specified model, i.e. $\Delta_e = \widetilde{V}_P - V_P$. Then, Δ_e satisfies a parabolic p.d.e. with homogeneous boundary conditions and source term given by

$$P\left[(bb_P - \tilde{b}\tilde{b}_P)\widetilde{\Gamma}P + \frac{1}{2}(b^2 - \tilde{b}^2)(\widetilde{\Gamma}_P P + 2\widetilde{\Gamma})\right].$$

Thus, Δ_e satisfies the following probabilistic problem:

$$\Delta_e(t,P) = E^{t,P}\left[-\int_t^T P\left[(bb_P - \tilde{b}\tilde{b}_P)\widetilde{\Gamma}(s,P(s))P(s) + \frac{1}{2}(b^2 - \tilde{b}^2)(\widetilde{\Gamma}_P(s,P(s))P(s) + 2\widetilde{\Gamma}(s,P(s)))\right]\mathrm{d}s\right].$$

The second result follows. □

These results assert that when the price process follows a Markovian Ito process: (a) small observation and/or model mis-specification errors imply bounded pricing as well as hedging errors if the model is updated; (b) model updating decreases the hedging errors due to model mis-specification even though the updating is inconsistent; (c) given a (mis-specified) model, inconsistently marking-to-market its initial conditions yields better results than alternative model consistent strategies.

While these results could be seen as a theoretical justification for the widespread approach of marking-to-market the initial conditions, we suggest a more careful interpretation. Model updating is theoretically justifiable under some (strong) conditions. What is still unknown is whether it can induce large hedging errors when the risk managers do not know the correct data generating process. In the following sections, we quantify the trade-off between the hedging errors generated by time-inconsistent marking-to-market of the initial conditions as well as the gains from model updating discussed in this section.

3. Quantifying updating effects: Empirical experiments

We have qualitatively shown that updating initial conditions can reduce pricing and hedging errors of a mis-specified model. Now we explore the magnitude of error reduction by focusing on commonly used models. First, we discuss two simulation settings based on a Markov model. Then we consider examples of non-Markov models.

3.1. The Hull and White model

In this section, we ask the following question: To what extent are hedging errors, which are generated by inconsistency, counterbalanced by the potentially positive effect of incorporating new information? We consider two settings in which a trader is aware that he is behaving in an inconsistent manner. The second scenario differs from the first in that the trader faces an additional unknown specification error.

Assume the risk manager wants to hedge a short call option position. In the first setting the trader uses a correctly specified model. The inconsistency derives solely from periodic re-initialization by marking-to-market the initial conditions. The valuation model coincides with the data generating process, though the trader can only observe a finite set of "liquid benchmarks". Here we ask to what extent inconsistent behavior increases hedging errors in the absence of specification error. The discount factor for the option's expiration date is computed from a set of "liquid" bonds in one of two ways: (a) by deriving the discount function from the model or (b) by interpolating the liquid bonds and using the interpolated vector in the hedging procedure. The latter method is typically motivated by measurement error and model mis-specification. In both methods, the trader uses the model delta which is fully coherent with the dynamics of the bond prices.

In case (a), hedging errors derive solely from discretization as the model is consistent with the data generating process. This can be viewed as the benchmark case in our simulations as it allows us to identify the order of magnitude of the errors that result from discrete time hedging.

In case (b), the formulae for the computation of the option value and its delta are correct and the true model is known to the trader. The discount function used, however, is interpolated in a way that is "inconsistent" with the model. Thus, Case (b) is equivalent to a continuous re-initialization of the model. As mentioned earlier, this procedure violates the dynamic no-arbitrage restrictions.

In the second setting, we introduce specification error. The trader does not know the true data generating process. His dynamic model is not coherent with the zero coupon bond prices that he observes in the market. Instead of changing the specification of his model, he chooses to re-initialize it using current term structure data. In this case, we ask whether inconsistent behavior simply leads to an ex post book fitting, or whether it can indeed reduce the hedging errors that would arise from a mis-specified model by incorporating new information.

Case 1: Inconsistency with a valid model

What is the economic size of a time-inconsistent error when the trader's model is correct? To answer this we simulate a 1-factor Hull and White model of the form

$$\mathrm{d}r_t = [\theta(t) - \alpha r_t]\mathrm{d}t + \sigma \mathrm{d}W_t.$$

The initial parameters, known to the trader, are: a 35 bp/year upward sloping forward rate term structure with a time 0 instantaneous rate level of 0.07; the mean reversion factor, α, is set to 0.0506; the instantaneous standard deviation is set to 0.01.

Specifically, consider the case in which daily yield curve re-initialization is performed using a cubic spline (the most widely used term structure fitting technique). The hedging performance simulation is carried out with respect to a vanilla call option, expiring in 3 years, written on a zero-coupon bond with maturity of 8 years. We consider option moneyness levels (forward price/strike price) ranging from 80% to 120% so as to analyze hedging error dependency on the leverage level implicit in the option.

Table 1 presents our results. We define the "time inconsistency error" as the square root of the average square difference between the cost of replicating the option using the true model and the cost of replicating using the continuously updated (time inconsistent) model. For the 90% out-of-the-money option, the error due to inconsistency is 13.88% of the option value. The error increases to 45.96% for the 80% forward out-of-the-money option. Considering that the risk manager is using the right model these are very large errors! For deep out-of-the-money options, the magnitude of this error resembles that produced by discretely hedging at a daily frequency, independent of the underlying asset shocks. Increasing the hedging frequency necessarily reduces the discretization error. This may, however, increase time inconsistency hedging errors. Furthermore, the simulations show that for

Table 1
Time inconsistency with no specification error

Moneyness (forward/strike)	Option value	Time inconsistency error
80%	0.01	45.96% (5.45)
90%	0.24	13.88% (0.22)
95%	0.90	7.01% (0.03)
100%	2.28	4.18% (0.01)
105%	4.21	2.36% (0.001)
110%	6.43	1.67% (0.0003)
120%	10.67	0.96% (0.0000)

This table presents the results of the cost of time inconsistency for the 1 factor Hull and White model. The "Time-Inconsistency Error" is defined as the square root of the average square difference between the cost of replicating the option using the true model and the cost of the replicating portfolio using the continuously updated (time inconsistent) model. Moneyness is defined as the forward value of the underlying strike price. The underlying asset is a 8 year zero coupon bond, while the expiration of the option is 3 years. The replication portfolio is rebalanced daily. The unit of measure is in percentages.

at-the-money and deep in-the-money options, the inconsistency error becomes substantially larger than that from discrete time hedging. For example, consider a 120% in-the-money option. In this case, the error due to time-inconsistency (equal to 0.96% in absolute terms) is 3200% of the error due to daily discrete time hedging.

Model inconsistency can be exploited by other traders. Suppose a trader is using an inconsistent hedging strategy and that a second trader knows both the model and the inconsistent procedure (recalibration) practiced by the first trader. The second trader can design arbitrage strategies against the first trader. Suppose, for simplicity, a one-factor model in which each possible future bond price is a non-random function of any other price. The inconsistent trader at time t_1 re-initializes his model based on $P_i(t_1, T)$. Thus in t_1, the two traders disagree on the possible prices of the bonds at some future date $t > t_1$. For instance, suppose that the law connecting prices $P(t, T')$ and $P(t, T'')$ at a given time t is $P(t, T') = C(P(t, T''))$ for the consistent trader and $P(t, T') = I(P(t, T''))$ for the inconsistent trader. This means that the event $P(t, T') \neq I(P(t, T''))$ has zero probability for the inconsistent trader and probability 1 for the consistent trader. The first best strategy for the consistent trader is then to buy, from the inconsistent trader, a digital option which expires in the money conditional on the realization of this event. This option will have no value for the inconsistent trader.

These examples show that when inconsistent behavior exists, it is theoretically always possible to profit through arbitrage. In practice, however, the available information is not so accurate (e.g. traders do not know the parameter values). This justifies bid–ask spreads along with the common practice of assigning a non-zero value to securities which are worthless (according to our models). Accounting for these real world practices, limiting our trading to securities (which are reasonably continuous with respect to initial condition variations) and implementing an accurate model, can sufficiently avoid significant arbitrage losses, even when coupled with inconsistent behavior. This is not the case when the model dynamics are incorrect.

Case 2: Inconsistency with an invalid model

What happens when an inconsistent strategy is adopted in the presence of a mis-specified model? We mimic a real world situation in which the term structure evolution generated by the model does not fit the observed data. The trader responds by re-initializing the model rather than modifying it. We explore the impact of this reaction.

Consider a 2-factor Hull and White model in which the original 1-factor model is based on the following particular case:

$$dr_t = [\theta(t) + u_t - \alpha r_t]dt + \sigma_1 dW_{1t},$$

$$du_t = -bu_t dt + \sigma_2 dW_{2t}.$$

The new parameters' values are $b = 0.01$ and $\sigma_2 = 0.002$. The trader is assumed to commit a model specification error by incorrectly using a one factor Hull and White model. We need to establish a rule for selecting the model's parameter values. We estimate these parameters using an exactly identified GMM approach based on the parameter values known to us of the full model.[3] This is intended to mimic the trader's presumed behavior of choosing to fit the model to the data so as to best replicate its statistical properties, e.g. by matching the moments of the observed series.

The trader is aware that the model-implied prices differ from market prices. He tries to fix this by re-calibrating the model, i.e. by introducing inconsistency. As before, re-calibration is based on one of two methods:

1. By deriving the discount function which is consistent with the (mis-specified) model used to compute the hedging strategy. This strategy entails only model specification error.
2. By deriving the discount function by interpolation with some flexible functional form (e.g. cubic spline) using a set of liquid bonds and then applying this discount function to implement the hedging strategy. Here the trader commits both model specification and time-inconsistency error.

Specification errors generate hedging errors. Does model re-calibration reduce these errors? What is the extent of the reduction, if any? Table 2 reports the results.

First, the replication error due to mis-specification is significant. For the at-the-money forward option, the size of the error is a quarter of the option value. When the option is far out-of-the-money, the errors are small in absolute value but large with respect to the option value. Obviously, the sampling standard deviations, measured in percentage points, of the option value are quite large.

Second, we find that time-inconsistent model updating does reduce the error arising from model mis-specification. However, in the most relevant cases, i.e. when the option is in-the-money, the reduction is small. Thus, inconsistency does not solve, or even significantly reduce, model mis-specification problems.

We checked the robustness of our results against changes in the initial shape of the yield curve, e.g. increasing the slope by 20%. We find that the results are not sensitive to changes in the level or slope of the initial term structure.

We repeat our analysis by changing the ratio of the instantaneous standard deviations of the two factors. A relative decrease in the standard deviation of the second factor (not included in the trader's model) decreases the relative value of the model specification error. For instance, if the standard deviation of the second factor is halved with respect to that of the first factor, the value of the specification error for the at-the-money forward option goes from 27.29% to 15.82%. At the same time, the replication error due to the time inconsistence behaviour falls from 22.98% to 13.51%.

[3] This strategy is motivated by our willingness to include, in our analysis, the potential influence of sampling error, which would have been ignored had we used the true theoretical moments. In simulation exercises where an "incorrect" model is used to describe the data, it is important to put the "wrong" model in the best possible position for replicating the data generating process. This implies that the value of the parameters of the "wrong" model should be chosen in such a way that the theoretical distribution of the observables given by the "wrong" model matches the distribution induced by the data generating process.

Table 2
Specification errors and time inconsistency

Moneyness (forward/strike)	Option value	Specification error	Time inconsistency error
80%	0.017	87.51% (19.66)	78.03% (28.24)
90%	0.37	68.38% (8.08)	47.27% (2.85)
95%	1.01	42.57% (2.26)	31.94% (1.09)
100%	2.14	27.29% (0.67)	22.98% (0.41)
105%	3.73	18.76% (0.25)	16.88% (0.18)
110%	5.61	13.67% (0.12)	12.63% (0.09)
120%	9.58	8.20% (0.04)	7.66% (0.03)

This table presents the results of the cost of time inconsistency when the true economy is a 2 factors Hull and White model. Risk management is done using a mis-specified 1 factor Hull and White model. We study the extent of reduction in the replication errors by re-initializing the model at each time t using a cubic spline routing applied to a set of observed bond yields. The "specification error" is defined as the square root of the average square difference between the true value of the option and the cost of the replicating portfolio (in percentage points of the initial option value) using a mis-specified 1 factor model that matches the observed second moment of the spot interest rate. This moment matching is done once based on a out-of-sample simulated process. The "time inconsistency" error is defined as the square root of the average square difference between the true value of the option and the cost of the replicating portfolio (in percentage points of the initial option value) based on a *continuously* re-inizialized (time inconsistent) mis-specified model. Moneyness is defined as the forward value of the underlying strike price. The underlying asset is a 8 year zero-coupon bond, while the expiration of the option is 3 years. The replication portfolio is rebalanced daily. Sampling standard errors are in parentheses.

We verify the robustness of the results with respect to different values of the mean-reversion parameter. We find that reduced mean-reversion in any of the two factor dynamics produces similar results to an increase in the standard deviation of the same factor. The higher the speed of mean-reversion the higher the relative cost of time-inconsistency. Moreover, we find that when the speed of the mean-reversion parameter is doubled the error due to inconsistency is about twice the order of magnitude than the discretization error.

To summarize, marking-to-market the model does not incorporate new information well, nor does it address the market prices deviations from model predictions. A time-inconsistent strategy based on model updating correctly prices a book of options but is unable to replicate them since the replicating strategy is not self-financing. Correct pricing at each trading date of a set of liquid benchmarks is almost irrelevant for risk management purposes. Hedging errors arise because the risk manager is not considering the "model's delta," i.e. his position's sensitivity to a change in the structure of the model. The results obtained clearly show that this sensitivity can add-up to a large hedging error.

3.2. A non-Markov BDT model

The results considered in the previous section are based on Markovian economies. Two otherwise identical Markov models have the property that even if they were to start from different initial conditions, they would still imply the same pricing conditional distribution if their values were to coincide at some stage. Errors in initial conditions can therefore easily be corrected with subsequent observed values of the process.[4] This is not the case for non-Markov models. In a non-Markov economy, asset prices are path-dependent. Updating or restarting a model using current observations could imply very different hedging errors.

While most continuous time models applied to security valuation are Markov in some set of variables, several HJM models imply a non-Markov short rate process. The empirical literature shows strong evidence of a Garch component in the interest rate process. In this case, interest rates are not Markov in a one-dimensional state-space. They can be written as Markov processes only if the state vector is supplemented by some set of unobservable variables.[5] In the following analysis we quantify how model updating affects option hedging performance in a non-Markov economy. We assume that the short rate is a conditionally lognormal discrete time process:

$$\ln r_{t+1} - \ln r_t = A[B - \ln r_t] + \sigma_{t+1}\varepsilon_{t+1}, \quad \varepsilon_{t+1} \sim \mathcal{N}(0,1),$$

where A and B are the mean reversion coefficient and long term value of the log short rate. Models using this specification of the interest rate process include Black, Derman and Toy. When σ is constant, the process is log-Normal with constant volatility. Let us consider a perturbation of the previous specification in which the volatility is a function of the history of interest rate innovations:[6]

$$\sigma_{t+1} = \frac{k}{k+t-1}\bar{\sigma} + \frac{t-1}{k+t-1}s_t,$$

$$s_t = \sqrt{\sum_{i=2}^{t}\left(\ln\frac{r_i}{r_{t-1}}\right)^2 \Big/ (t-1)}.$$

The current volatility is a weighted average of a constant value $\bar{\sigma}$ and an average of the history of interest rate quadratic innovations. Different values of k give different weights to the non-Markov part of the short rate volatility. For $k \to \infty$, the economy is Markov and lognormal. For $k < \infty$ the process is state dependent and non-Markov. The smaller k is, the larger the dependence is on the history of interest rate innovations.

We suppose that the true model structure is known to the trader, who, despite the information available to him, decides to restart the model afresh each day. This results in a hedging errror due to time-inconsistency as the trader forgets the past error history, unintentionally influencing the evolution of the volatility parameter. Our purpose is to evaluate the consequences of this error.

We compare the behavior of short and long term memory traders. Long term memory traders are time-consistent and account for the full evolution of s_t by cor-

[4] An example of such a conversion is seen through the empirical implementations of Bliss and Ritchken (1996) and de Jong and Santa-Clara (1999).

[5] Brenner et al. (1996).

[6] A similar specification is discussed in Bliss and Ritchken (1996).

rectly implementing the model. They value and hedge options correctly apart from model discretization error and residual sampling error in the delta valuations. Short memory traders are time-inconsistent and follow one of the two strategies outlined below:

- Strategy A: *Evaluate bonds and options by recalibrating the Treasury curve daily to fit bond prices exactly.*
- Strategy B: *Evaluate bonds and options by recalibrating both the Treasury curve and σ_{t+1} daily to fit both bond prices exactly and the ATM option.*

The two strategies represent realistic behaviors of risk managers who must hedge a set of options. We compare the results from these two strategies with the "exact" strategy of a long term memory risk manager. Notice that we do not consider any strategical interaction between the two trader types, as this is not an exercise in market microstructure.

We perform Monte Carlo experiments, where the different strategies are used to hedge five different call options with a range of strikes. Experiments are run for different values of A, k and σ. Under each scenario, the set of five options must be hedged for 20 time intervals. The options are all written on the same bond, which expires 5 time intervals after the options expire.

Delta values for each option are estimated, at each time interval, using a Monte Carlo loop. Thus, we use as variance reduction techniques both (a) control variates based on the quantiles of the empirical distribution of ε_{t+1} to stabilize the delta estimate for each time interval, and (b) antithetical variates both in the estimate of deltas and in the evaluation of delta hedging error. Each strategy is evaluated using 1000 replications and, for each step of each replication, deltas are computed using 500 replications.

We consider the following parameter configurations: $B = \ln(0.03)$, $A = [0.5; 0.01]$ implying a high or low level of mean reversion respectively, $\sigma = [0.01; 0.05; 0.1]$, and $k = [100; 0.10; 0.01]$ implying a small or large degree of non-Markov behavior. When $k = 0.01$, the constant value of the volatility $\bar{\sigma}$ has a negligible weight and the evolution of σ_{t+1} is highly path dependent. When $k = 100$, the past history weight in s_t is less that 20%. When $k = 1000$, the results of the three different trader types converge within the sampling error. For $k \to \infty$, the three different behaviors are equivalent (i.e. the model becomes a constant volatility Markov model). We consider five strike prices centered around the at-the-money forward value and ranging from this value by plus and minus 1% and 2%.

Tables 3 and 4 summarize the hedging errors for the two strategies as a percentage of the true value of the option and the hedging errors due to model discretization. The hedging errors are computed as the average Monte Carlo difference between the terminal value of the option and the value of the replicating portfolio. A positive value for the hedging error implies a loss for the option seller who hedges the position. Values in parentheses are Monte Carlo standard errors. In most scenarios the hedging errors due to discretization, and their Monte Carlo standard errors, are small with respect to the true value of the option. The only exceptions are the two

Table 3
Pricing and hedging errors

Hedging errors				
Strike	Correct model	Trader B	Trader A	Option value
$A = 0.5$, $k = 100$, $\sigma = 0.01$				
−2%	0.0000 (0.0000)	0.0000 (0.0000)	0.0001 (0.0000)	2.502 (0.0080)
−1%	0.0000 (0.0000)	0.0000 (0.0000)	0.0001 (0.0000)	1.511 (0.0048)
atmf	0.0000 (0.0000)	0.0002 (0.0000)	0.0003 (0.0001)	0.5210 (0.0018)
+1%	0.0002 (0.0000)	0.0006 (0.0001)	0.0007 (0.0001)	0.0010 (0.0000)
+2%	0.0000 (0.0000)	0.0000 (0.0000)	0.0001 (0.0000)	0.0000 (0.0000)
$A = 0.5$, $k = 100$, $\sigma = 0.05$				
−2%	0.0000 (0.0000)	0.0011 (0.0002)	0.0048 (0.0001)	2.577 (0.0081)
−1%	0.0076 (0.0007)	0.0202 (0.0007)	0.0447 (0.0006)	1.558 (0.0050)
atmf	0.0049 (0.0016)	0.0576 (0.0016)	0.0879 (0.0014)	0.5610 (0.0022)
1%	0.0037 (0.0012)	0.0505 (0.0015)	0.0930 (0.0016)	0.0600 (0.0005)
2%	0.0018 (0.0003)	0.0092 (0.0005)	0.0255 (0.0006)	0.0010 (0.0001)
$A = 0.5$, $k = 100$, $\sigma = 0.1$				
−2%	0.0251 (0.0023)	0.0611 (0.0028)	0.0872 (0.0033)	2.701 (0.0086)
−1%	0.0626 (0.0031)	0.1028 (0.0035)	0.1211 (0.0032)	1.602 (0.0057)
atmf	0.1068 (0.0033)	0.1292 (0.0038)	0.1395 (0.0036)	0.698 (0.0034)
+1%	0.1032 (0.0031)	0.1185 (0.0033)	0.1291 (0.0035)	0.214 (0.0015)
+2%	0.0568 (0.0026)	0.0778 (0.0027)	0.0881 (0.0031)	0.034 (0.0005)
$A = 0.5$, $k = 0.1$, $\sigma = 0.05$				
−2%	−0.0032 (0.0022)	0.0255 (0.0031)	0.0842 (0.0070)	2.676 (0.0081)
−1%	−0.0067 (0.0041)	0.0529 (0.0044)	0.1294 (0.0093)	1.664 (0.0051)
atmf	−0.0051 (0.0063)	0.1486 (0.0067)	0.2238 (0.0117)	0.676 (0.0024)
+1%	−0.0037 (0.0050)	0.1120 (0.0056)	0.1689 (0.0089)	0.092 (0.0009)
+2%	−0.0025 (0.0019)	0.0242 (0.0025)	0.0542 (0.0044)	0.013 (0.0004)

This table summarizes simulation results for different non-Markov settings. The options evaluated in these examples are call options with 5 different strikes: at the money forward (atmf) and at the money forward −2%, −1%, +1%, and +2%. The parameters of the model are: A – measuring the mean reversion intensity of the short interest rate; k – measuring the degree of "volatility memory" of the model, and σ – measuring the initial value of volatility. Hedging errors are computed as Monte Carlo estimates of the differences between the option terminal value and the hedging portfolio. "Correct model" means that the correct model is used for hedging. Errors are due to discrete time and Monte Carlo error. Trader B indicates a short term memory trader who calibrates σ in order to fit the price of the atmf option as quoted with the correct model. Trader A is a short term memory trader who does not fit σ in order to replicate the price of observed options.

economies with parameter values [$A = 0.5$; $k = 100$; $\sigma = 0.1$;] and [$A = 0.01$; $k = 0.01$; $\sigma = 0.05$].

3.2.1. Quasi-Markov (k = 100) model

When [$A = 0.5$; $k = 100$; $\sigma = 0.01$] type A and type B hedging errors are very similar, but somewhat larger than the errors due to discretization. The at-the-money option hedging error is almost negligible for both type A and B strategies. As volatility increases to $\sigma = 0.05$, both type A and type B recalibration strategies yield bad results. Type B errors are 4 to 15 times bigger than the errors in hedging with the

Table 4
Pricing and hedging errors

Hedging errors				
Strike	Correct model	Trader B	Trader A	Option value
$A = 0.5$, $k = 0.01$, $\sigma = 0.01$				
−2%	0.0000 (0.0000)	0.0001 (0.0000)	0.0004 (0.0001)	2.505 (0.0082)
−1%	0.0000 (0.0000)	0.0003 (0.0001)	0.0008 (0.0003)	1.510 (0.0048)
atmf	−0.0006 (0.0001)	0.0051 (0.0005)	0.0084 (0.0007)	0.5200 (0.0017)
+1%	−0.0008 (0.0001)	0.0044 (0.0004)	0.0076 (0.0006)	0.0020 (0.0001)
+2%	0.0000 (0.0000)	0.0002 (0.0000)	0.0003 (0.0000)	0.0000 (0.0000)
$A = 0.01$, $k = 0.01$, $\sigma = 0.01$				
−2%	−0.0016 (0.0003)	0.0031 (0.0000)	0.0049 (0.0005)	2.495 (0.0082)
−1%	0.0020 (0.0011)	0.0101 (0.0023)	0.0181 (0.0029)	1.512 (0.0052)
atmf	0.0240 (0.0033)	0.0746 (0.0049)	0.1039 (0.0052)	0.577 (0.0024)
+1%	0.0174 (0.0031)	0.0418 (0.0030)	0.0761 (0.0041)	0.0624 (0.0007)
+2%	0.0054 (0.0012)	0.0069 (0.0016)	0.0182 (0.0020)	0.005 (0.0002)
$A = 0.01$, $k = 0.01$, $\sigma = 0.05$				
−2%	0.0942 (0.0162)	0.2715 (0.0207)	0.4012 (0.0289)	2.768 (0.0114)
−1%	0.1194 (0.0183)	0.4880 (0.0322)	0.8267 (0.0418)	1.892 (0.0090)
atmf	0.1527 (0.0192)	0.7959 (0.0315)	1.1292 (0.0395)	1.235 (0.0069)
+1%	0.1648 (0.0183)	0.8331 (0.0252)	1.1142 (0.0321)	0.739 (0.0051)
+2%	0.1362 (0.0153)	0.4872 (0.0258)	0.8114 (0.0398)	0.413 (0.0037)
$A = 0.5$, $k = 0.01$, $\sigma = 0.1$				
−2%	0.0690 (0.0114)	0.3903 (0.0262)	0.7944 (0.0355)	2.755 (0.0085)
−1%	0.0773 (0.0133)	0.4768 (0.0278)	0.9897 (0.0378)	1.701 (0.0058)
atmf	0.0850 (0.0143)	0.6209 (0.0279)	1.023 (0.0362)	0.7240 (0.0035)
+1%	0.0758 (0.0114)	0.4500 (0.0208)	0.9163 (0.0379)	0.2490 (0.0010)
+2%	0.0469 (0.0074)	0.2008 (0.0137)	0.4120 (0.0324)	0.078 (0.0011)

This table summarizes simulation results for different non-Markov settings. The options evaluated in these examples are call options with 5 different strikes: at the money forward (atmf) and at the money forward −2%, −1%, +1%, and +2%. The parameters of the model are: A – measuring the mean reversion intensity of the short interest rate; k – measuring the degree of "volatility memory" of the model, and σ – measuring the initial value of volatility. Hedging errors are computed as Monte Carlo estimates of the differences between the option terminal value and the hedging portfolio. "Correct model" means that the correct model is used for hedging. Errors are due to discrete time and Monte Carlo error. Trader B indicates a short term memory trader who calibrates σ in order to fit the price of the atmf option as quoted with the correct model. Trader A is a short term memory trader who does not fit σ in order to replicate the price of observed options.

correct model and they are large with respect to the true option values. For instance the at-the-money option error for a type B trader is 1/10 of the at-the-money option value. The A trader shows an at-the-money hedging error which is about 1/6 of the corresponding option value.

3.2.2. Non-Markov ($k = 0.1$) model

In the case of the less Markov model, i.e. $k = 0.1$, the results are worse. When $[A = 0.5;\ k = 0.1;\ \sigma = 0.05]$ hedging errors for at-the-money options are 30 times the correspondent errors for the correct model and are 1/4 of the value of the option.

These are the results for the trader who fares better. The type A trader shows an hedging error almost twice that of a type B trader.

3.2.3. Highly non-Markov ($k = 0.01$) model

Both trader A and B hedging errors are extremely significant. When [$A = 0.5$; $k = 0.01$; $\sigma = 0.1$] the hedging error of Trader B is 7.29 times the hedging errors implied by a discrete time hedging strategy based on the correct model. Moreover the absolute value of the error is about 85% of the value of the option. This shows the danger of marking-to-market a model when the underlying economy is non-Markov.

3.2.4. Effect of mean reversion

The effect of a decrease in A, the mean reversion parameter, is very similar to the effect of an increase in σ. Compare for instance the case [$A = 0.5$; $k = 0.01$; $\sigma = 0.01$] (Table 3, third panel) to the case [$A = 0.01$; $k = 0.01$; $\sigma = 0.01$] (Table 4, third panel). Hedging errors on at-the-money options for both trader A and B increase roughly tenfold and are of the same order of magnitude as the errors in the case $A = 0.5$, $k = 100$ and $\sigma = 0.05$.

3.2.5. Effect of volatility

We find that when σ is small (0.01), the hedging error with the correct model is negligible and the difference between strategy A and B is small, independent of k. The one exception is the case in which the mean reversion parameter A is very small. A big value of σ implies large hedging errors even when the risk manager uses the correct model. This is the case with $\sigma = 0.1$ and $A = 0.5$ (Table 4, fourth panel) and the case of $\sigma = 0.05$ and $A = 0.01$ (Table 4, panel three). The effect of σ is slightly enhanced by the value of k. Compare the results of Table 3 first panel and Table 3 third panel with the results of Table 4 first panel and Table 4 fourth panel. In both cases, σ changes from 0.01 to 0.1, but in the first setting $k = 100$ while in the second $k = 0.01$. In both cases, the hedging errors with $\sigma = 0.01$ are small (but more significant when $k = 0.01$ as noted above). However, the errors when $\sigma = 0.1$, are about five times bigger for the $k = 0.01$ case than for the $k = 100$ case.

3.2.6. Monte Carlo standard errors

Monte Carlo standard errors are in general larger (by an average from 5 to 10 times) in the highly non-Markov case ($k = 0.1$). However, the size of type A and B hedging errors are always several orders of magnitude larger than the corresponding standard errors. Note that the value of k greatly influences the Monte Carlo standard errors of the hedging errors, while the Monte Carlo standard errors are almost independent of k. This indicates that the more non-Markov the rate process is, the more difficult it is to estimate the deltas. Even with $k = 0.01$ the hedging error differences between the three types of traders are well beyond those of the Monte Carlo standard errors.

In summary, we find that even in a (mildly) non-Markov economy, such as $k = 100$, the errors for incorrect model updating are not negligible except in very extreme cases of low volatility. These results quantify the hedging errors and warn against model updating, especially when the risk manager has reason to believe that the underlying process is not Markov.

4. Conclusion

The practice of re-initializing interest rate models daily with observed prices is so widespread that it has become an accepted market tool. This practice, however, is inconsistent in the sense that it violates the dynamic assumptions underlying the model and from which derivative prices and hedge factors are computed. This behavior is similar in spirit to the practice of using current "implied" parameters, such as the implied volatility, from quoted market option or cap prices.

Both behaviors grew from the model inability to reproduce relevant aspects of price dynamics, along with the need for trading books to be marked to market prices. While searching for more statistically accomplished models, it is important to determine whether initial conditions updating is a helpful way to deal with model error. Our findings are as follows:

For diffusion models, theoretical results about well posedness of stochastic differential equations show that a "small" error in computing initial conditions implies "small" pricing and hedging errors. Moreover, the updating of initial conditions corrects part of the model mis-specification errors. However, we show that these improvements are insufficient in significantly reducing the model mis-specification errors.

We studied two modelling instances in detail:

The first model we studied was the Hull and White one and two factor model. We used simulation to show that, consistent with theory, an error in the initial conditions has bounded consequences on pricing and hedging. When the one factor model was used for pricing and hedging in a two factor economy, however, simply updating initial conditions, was not sufficient to correct for the model error consequences in a significant way.

The second scenario considered an underlying non-Markov economy (and so lay outside the general well posedness results derived above). Although most asset pricing models used in practice are Markov, mainly for ease of computation, there is growing empirical evidence of non-Markovianity in interest rates. We showed that even for a small degree of non-Markovianity in the underlying economy, model updating implies large hedging errors.

The widespread use of model updating in interest rate models, while understandable for bookkeeping rules, is by no means effective in correcting the consequences of model mis-specification.

Acknowledgement

The authors thank George Constantinides, Ian Cooper, Anthony Neuberger and Stephen Schaefer for helpful comments and participants to seminars at the

University of Chicago and London Business School. Purnendu Nath provided excellent research assistance.

References

Basle Committee on Banking Supervision, 1999. A New Capital Adequacy Framework, Bank of International Settlements, Basle.

Backus, D.K., Telmer, C., Wu, L., 1999. Design and estimation of affine yield models. Working Paper, Carnegie Mellon University.

Björk, T., Christensen, B.J., 1999. Interest rate dynamics and consistent forward rate curves. Mathematical Finance 9 (4), 323–348.

Black, F., Derman, E., Toy, W., 1990. A one factor model of interest rates and its implication to treasury bond options. Financial Analysts Journal 46, 33–39.

Black, F., Karasinski, P., 1991. Bond and option pricing when short rates are lognormal. Financial Analysts Journal 47, 52–59.

Bliss, R.R., Ritchken, P., 1996. Empirical tests of two state-variables Heath–Jarrow–Morton models. Journal of Money, Credit, and Banking 28, 452–466.

Brandt, M.W., Yaron, A., 2001. Time-consistent no-arbitrage models of the term structure, Working Paper, The Wharton School, University of Pennsylvania.

Brenner, R.J., Harjes, R.H., Kroner, K.F., 1996. Another look at models of the short term interest rate. Journal of Financial and Quantitative Analysis 31, 85–107.

Buraschi, A., Jiltsov, A., 2003. Habit formation and the term structure of interest rates. Working Paper, London Business School.

Buraschi, A., Jiltsov, A., 2005. Inflation risk premia and the expectations hypothesis: A monetary model of the treasury yield curve. Journal of Financial Economics 75 (2), 429–490.

Constantinides, G., 1992. A theory of the nominal term structure of interest rates. Review of Financial Studies 5, 531–552.

Cooley, T., LeRoy, S., Parke, W., 1992. Pricing interest-sensitive claims when interest rates have stationary components. Journal of Fixed Income 2, 64–73.

Cox, J.C., Ingersoll Jr., J.E., Ross, S.A., 1985. A theory of the term structure of interest rates. Econometrica 53, 385–407.

Dai, Q., Singleton, K.J., 2000. Specification analysis of affine term structure models. Journal of Finance 55, 1943–1978.

de Jong, F., Santa-Clara, P., 1999. The dynamics of the forward interest rate curve: Formulation with state variables. Journal of Financial and Quantitative Analysis 34, 131–157.

Derman, E., 1996. Model Risk, Quantitative Strategies Research Notes. Goldman Sachs, New York, NY.

Duffie, D., Kan, R., 1996. A yield factor model of interest rates. Mathematical Finance 6 (4), 379–406.

Duffie, D., Singleton, K., 1997. An econometric model of the term structure of interest rate swap yields. Journal of Finance, 1287–1381.

Dybvig, P., 1989. Bond and bond option pricing based on the current term structure. In: Dempster, M., Pliska, S. (Eds.), Mathematics of Derivative Securities. Cambridge University Press, Cambridge.

Goldstein, R.S., 2000. The term structure of interest rates as a random field. Review of Financial Studies 13, 365–384.

Ho, T.S.Y., Lee, S.-B., 1986. Term structure movements and pricing interest rate contingent claims. Journal of Finance 41, 1011–1029.

Heath, D., Jarrow, R., Morton, A., 1991. Contingent claims valuation with a random evolution of interest rates. Journal of Futures Markets (January), 55–78.

Heath, D., Jarrow, R.A., Morton, A., 1992. Bond pricing and the term structure of interest rates. Econometrica 60, 225–262.

Hull, J., Suo, W, 2000. A methodology for assessing model risk and its application to the implied volatility function model, Working Paper, University of Toronto.

Hull, J., White, A., 1990. Pricing interest rate derivative securities. Review of Financial Studies 3, 573–592.

Hull, J., White, A., 1993. One-factor interest-rate models and the valuation of interest rate derivative securities. Journal of Financial and Quantitative Analysis 28, 235–254.

Karatzas, I., Shreve, S.E., 1991. Brownian Motion and Stochastic Calculus. Springer-Verlag.

Kennedy, D.P., 1994. The term structure of interest rates as a Gaussian random field. Mathematical Finance 4, 247–258.

Kennedy, D.P., 1997. Characterizing Gaussian models of the term structure of interest rates. Mathematical Finance 7, 107–118.

Longstaff, F., Schwartz, E.S., 1992. Interest rate volatility and the term structure: A two-factor general equilibrium model. The Journal of Finance 47, 1259–1282.

Santa Clara, P., Sornette, D., 2001. The dynamics of the forward interest curve with stochastic string shocks. Review of Financial Studies 14, 149–185.

Vasiček, O.K., 1977. An equilibrium characterization of the term structure. Journal of Financial Economics 5, 177–188.