North American Actuarial Journal

ISSN: 1092-0277 (Print) 2325-0453 (Online) Journal homepage: https://www.tandfonline.com/loi/uaaj20

The Theory of Optimal Stochastic Control as Applied to Insurance Underwriting Cycles

David L. Eckles, David G. McCarthy & Xudong Zeng

To cite this article: David L. Eckles, David G. McCarthy & Xudong Zeng (2016) The Theory of Optimal Stochastic Control as Applied to Insurance Underwriting Cycles, North American Actuarial Journal, 20:4, 327-340, DOI: 10.1080/10920277.2016.1179122

To link to this article: https://doi.org/10.1080/10920277.2016.1179122

Published online: 20 Jun 2016.

North American Actuarial Journal, 20(4), 327–340, 2016

ISSN: 1092-0277 print / 2325-0453 online
DOI: 10.1080/10920277.2016.1179122

The Theory of Optimal Stochastic Control as Applied to Insurance Underwriting Cycles

David L. Eckles,[1] David G. McCarthy,[1] and Xudong Zeng[2]
[1]*Department of Insurance, Legal Studies, and Real Estate, Terry College of Business, University of Georgia, Athens, Georgia*
[2]*School of Finance, Shanghai University of Finance and Economics, Shanghai, China*

We use the theories of optimal stochastic control and engineering process control to analyze the well-known phenomenon of insurance underwriting cycles in continuous time. We show in a continuous time framework that underwriting cycles can be explained with a model where premiums are set rationally, but where there are various reporting and regulatory lags. We find that the observed cycle length depends on the length of these underlying lags. Our result can be seen as consistent with previous empirical work showing underwriting cycles varying across countries and lines of insurance. In the event that no lags exist, our result is also consistent with more recent literature suggesting that insurance cycles may not exist.

1. INTRODUCTION

Until recently, the existence of underwriting cycles, alternating periods of high premiums/high profits and low premiums/low profits for insurers, had been a (presumably) well-established result in the academic literature. From recent papers such as Lazar and Denuit (2011), Meier and Outreville (2006), and Venezian and Leng (2006) to the enduring papers of Cummins and Outreville (1987), Winter (1994), and Gron (1994a, 1994b), academics have (mostly) argued for the existence of underwriting cycles.[1] Indeed, the first sentence of Higgins and Thistle (2000) reads "[i]t is well known that the property-liability insurance industry is subject to cycles in premiums and underwriting profits" (p. 442).

However, two recent papers, Boyer et al. (2012) and Boyer and Owadally (2015), have cast doubt on the actual existence of an underwriting cycle. In particular, Boyer et al. (2012) show that while a potential cycle-suggestive time series result is indeed observed, this does not prove the actual *existence* of a cycle (not to mention the length of the cycle). Noting the difficulty of forecasting cycles, the authors suggest that traditional "tests are biased in favor of finding" a cycle (p. 1010).[2] Boyer and Owadally (2015) continue with this theme of bias and revisit previously published empirical papers on the existence of underwriting cycles. Boyer and Owadally (2015) note that "[t]he tests that have been used in previous research on underwriting cycles were biased in favour of finding such cycles" and "[o]nce these biases are corrected, significant underwriting cycles become the exception rather than the norm" (p. 2).

Turning our attention back, for a moment, to the presumption of the existence of underwriting cycles, we note that even here there is debate. Specifically, the mechanism by which the cycles manifest themselves is contested by academics. While authors such as Cummins and Outreville (1987) and Cummins and Danzon (1997) argue that the cycles are a rational response to reporting and regulatory lags in the structure of the insurance industry, others, such as Winter (1994) and Gron (1994a, 1994b), argue that the cycles are the result of supply-side capacity constraints and opportunity cost differences between internal and external capital.

Again, assuming the existence of underwriting cycles, there are three primary explanations for underwriting profit cycles in insurance. The first, and most naïve, explanation is that the industry brings the cycles on themselves by irrational behavior. Under this hypothesis, a cycle can be generated if, for instance, the actual investment returns earned are less than those assumed in the

Address correspondence to David L. Eckles, Department of Insurance, Legal Studies, and Real Estate, Terry College of Business, University of Georgia, 206 Brooks Hall, Athens, GA 30602. E-mail: deckles@uga.edu
David G. McCarthy was previously professor of actuarial science at the University of the Witwatersrand, Johannesburg.
Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/uaaj.
[1]These are by no means the only empirical papers studying insurance cycles.
[2]Boyer et al. (2012) also suggest that a lack of power in out-of-sample tests hampers the ability to predict cycles.

premium calculation and the premiums cannot be retroactively changed, and/or the realized losses are greater than the expected losses. In the next period, insurers may then opt to charge higher premiums (perhaps in an effort to avoid regulatory scrutiny for a lack of capital), which will lead to fewer people purchasing insurance, which will eventually lead to a hard market. In a similar fashion, better interest rate experience can lead to a soft market or a market characterized by high availability and low prices. Venezian (1985) provides a more sophisticated version of this explanation where the cycles are caused by poor forecasting techniques.

The second explanation given for underwriting cycles is the capacity constraint hypothesis put forth by Winter (1994) and Gron (1994a, 1994b). They argue that because externally generated capital is more expensive than internally generated capital, insurers are capacity constrained. Hence, when there is a loss shock, supply of insurance falls and buyers of insurance will chase prices upwards. These periods of high prices and low quantity are subsequently followed by periods of low prices and high quantity. Although Gron (1994a, 1994b) finds empirical support for these models, they are not without several criticisms, as described in Cummins and Danzon (1997). In particular, the model ignores that insurance quality can differ between companies, and empirical support can be found only for short-tailed lines where the theory suggests that cycles should be found in all lines. Finally, Gron's results have not been consistently found by other researchers. Like Gron (1994a, 1994b), Niehaus and Terry (1993) and Froot and O'Connell (1997) show evidence of an inverse relationship between capital and prices, where Cummins and Danzon (1997), utilizing an alternative methodology, find a positive relationship between capital and prices.

The final theory of underwriting cycles was suggested by Cummins and Outreville (1987). They suggest that underwriting profit cycles are a result of institutional, regulatory, and accounting factors. Specifically, because the loss data emerge slowly, there is a delay between the time insurers become aware of an accident and the time that new rates can go into effect; the underwriting profits can then be shown to follow an autoregressive process of order one (AR(1)). Additionally, if this informational lag is combined with regulatory lags and/or renewal lags, the underwriting profits can be shown to follow an AR(2) process. This AR(2) process is necessary for cycles to exist. The Cummins-Outreville model does not depend on capacity constraints, and their model allows price to differ according to insurer quality.

In the spirit of Winter (1994), Gron (1994a, 1994b), and Cummins and Outreville (1987), we believe that if insurance cycles exist, they are a result of a rational response to underlying economic circumstances. We propose that insurers adjust prices, underwriting procedures, and policy conditions in response to changes in their profit stream (including changes in claims patterns), and the effect of these changes is stochastic due to a changing economic environment (including supply-side constraints). Also, there is a random delay between implementing these changes and observing their effect on the profit stream (e.g., reporting lags, regulatory lags). We also recognize that insurers can only survive losses for a certain length of time and cannot leave the market without incurring a cost. We investigate this model using the mathematics of stochastic control theory.

Our model can be considered as a generalized version of the Cummins-Outreville model in continuous time. Ultimately we show that when an insure optimizes future profits by changing a control variable (prices), which has an unknown stochastic effect, and an unknown stochastic delay, this rational price-setting behavior results in underwriting profit cycles, independent of the loss process. We further show that the length of underwriting cycles depends on the length of the regulatory/reporting lags. However, we note that the aforementioned delay may (1) not exist, (2) vary over time, (3) vary across firms, (4) vary across countries, and/or (5) vary across lines of insurance. To that end, *if* cycles exist, they may not be permanent. In this way, our model can be thought of as a "unifying model" of the existence (or lack thereof) of cycles. Considering short-term market/regulatory conditions, a cycle may be apparent, as previously documented in the literature. However, considering a long-term perspective, one where these delays may change (or disappear altogether), a result similar to that argued by Boyer et al. (2012) can also be argued. In fact, the idea of changing (or disappearing) delays would be consistent with the lack of cycle predictability documented in Boyer et al. (2012).

The remainder of this article is divided into seven sections. Section 2 concretely presents our theory and model basics. Sections 3 and 4 explain our model in the context of stochastic control theory and engineering process control, respectively. Section 5 solves our basic model using Brownian motion as the loss process. Section 6 shows how our model might be generalized to include capital constraints. Section 7 shows that the results from the model result in underwriting profit cycles. Section 8 concludes and provides ideas for future research.

2. FUNDAMENTAL INSURANCE MODEL

We imagine an insurance company in a perfectly competitive market. Losses in this market follow some continuous Itô process, which we denote $(y_t)_{t\geq 0}$. For now we will not restrict $(y_t)_{t\geq 0}$ at all; it can easily be imagined that $(y_t)_{t\geq 0}$ could be constrained to be positive. We further assume that the insurer maximizes the expected value of some concave function of profits. Though we later allow for a broader class of objective functions, we initially consider the insurer to minimize the squared deviations of premiums from losses for four reasons. First, minimizing a convex function is equivalent to maximizing a concave function for our choice of function (e.g., Beck 2014; Cohen and Megiddo 1990). Second, this objective function is consistent with the notion of a firm operating in a competitive environment with zero expected profits: Choosing premiums to minimize the variance of profits is mathematically equivalent to setting premiums equal to the expected value of losses and therefore setting expected profits equal to zero.[3,4] Third, our objective function is consistent with corporate finance theory explaining why firms might appear to be risk averse (see Froot et al. [1993] for a review of rationales). For example, if earnings are uncertain, a firm maximizing expected revenues less a progressive tax (a concave function) would appear to be similar to a firm minimizing the variability of earnings around their expected value (a convex function) (again, note that minimizing a convex function is equivalent to maximizing a concave function). Fourth, the use of this objective function allows us to provide a tractable solution to our model. Ultimately we will show that our model yields the same results with a very broad class of objective functions. We further idealize that the managers of this company are permitted to alter premiums continuously, and further that this alteration is instantaneous and costless.

Insurance underwriting profits may follow a variable pattern, which could be a rational response to underlying economic circumstances. Insurers adjust prices, underwriting procedures, and policy conditions in response to changes in their profit stream (perhaps due to supply-side constraints). The effect of these changes can also be considered stochastic due to a changing economic environment. Then there is a random delay between implementing these changes and observing their effect on the profit stream (e.g., reporting lags, regulatory lags, and the time that naturally elapses between when policies are written and when claims in respect of those policies arise). Considering the initial objective function, the insurer's objective is to minimize the deviations from a goal of zero profits over time.[5] This could be written mathematically as follows:

$$\begin{aligned} &\arg\max_{\{P_s\}} E_\alpha \left[-\int_{t+\alpha(t)}^{T} o\left(x_s\right)\, ds \,|\Im_t \right] \\ &\text{s.t. } x_s = P_{s-\alpha(s)} + y_s, \text{ for} t + \alpha(t) \leq s \leq T, \end{aligned} \tag{1}$$

where x_s is the insurer's profit made up of the premium, $P_{s-\alpha}$, and the underwriting losses, y_s(the loss process $(y_s)_{s\geq 0}$can be any continuous-time Itô process). $o\left(x_s\right)$ is the insurer's objective function with respect to profits, and $\alpha(t)$ is the delay random variable with an independent and identical known distribution in the control for each t.[6] We assume that the realization of $\alpha(t)$is observed at t. For notational simplification, we suppress the time dependence of $\alpha(t)$and refer to the delay as α going forward.

We assume α is independent of the loss process. The filtration $(\Im_t)_{t\geq 0}$ satisfies the usual conditions. $\Im_t$ is the set of all information available at time t. The premium process $(P_t)_{t\geq 0}$is $(\Im_t)_{t\geq 0}$predictable. We use the convention $E_\alpha$to emphasize that the conditional expectation is taken with respect to a randomly determined α. Since the effect of the premium (control) is not seen until α later (the delay), we consider the cumulative effects from time $t + \alpha$ to the terminal time horizon T. We solve this model for the quadratic objective function:

$$o\left(x_s\right) = x_s^2$$

and a loss process $(y_t)_{t\geq 0}$following a Brownian motion process (it is straightforward to extend our results to the case where the loss process is a Brownian motion with drift).

The use of the quadratic objective function is equivalent to assuming rational expectations, since our quadratic objective function is mathematically equivalent to setting premiums to ensure that the expected value of profit is zero; that is, our objective function is consistent with long-run equilibrium profits in a competitive environment. Finally, we note that we consider only continuous processes in our model. The use of continuous processes makes the model solution easier to find. Thus, our model (as we solve it) is probably best considered in the context of relatively simple risks (e.g., automobile liability, automobile property damage), where underwriting cycles are also observed (Cummins and Outreville 1987).

[3]Our maximization problem is rooted in the foundations of insurance economics. Borrowing from Einav and Finkelstein (2011), we note that in "the textbook case of insurance demand and cost, in which perfectly competitive, risk-neutral firms offer a single insurance contract that covers some probabilistic loss; risk-averse individuals differ only in their (privately-known) probability of incurring that loss; and there are no other frictions," then "Consumers in this market make a binary choice of whether or not to purchase this contract and firms in this market compete only over what price to charge for the contract" (p. 116).

[4]In using this objective function, we do *not* mean to imply that the squared deviation of profits is the only possible objective function. We are simply trying to model a firm setting a premium where the underlying risk is unknown: Such a firm would be penalized both if they set the premium too low (because of underwriting losses) and if they set the premium too high (because of low sales). We use this particular function simply to allow for tractability of the results. Importantly, the results of this article are *not* dependent on this assumption, as shown in the Corollary to Lemma 1. All we require is that the objective function be concave. For instance, we could add an arbitrary positive constant economic rent to the analysis in the article. This positive rent would guarantee positive profits but would not change the cycles observed in the model.

[5]Again, the zero-profit goal is only meant to be an example of an insurer's objective function.

[6]For further discussion of a random delay in a continuous-time setting, see recent papers such as Gozzi et al. (2009) and Federico (2011).

Beyond the mathematics, the model is fairly intuitive. The insurer is attempting to minimize their deviations from zero profits. Without a delay in observing random losses and setting premiums, this would be trivial. However, our model includes an unknown delay (α) between setting premiums (P) and observing unknown losses (y). Since the delay is unknown, the insurer must consider the distribution of the delay. We will also point out that our model does not specify the exact mechanism for the delay. That is, the delay could be due to reporting, regulatory, random claims, simple accounting-related issues, or a combination of any or all of these.

Last, the delay variable, α, is important in our model. Since we are using a continuous-time framework, this variable can be thought about as virtually any time period (e.g., instantaneous, three seconds, six hours, two days, five years). However, because we are modeling an insurer setting premiums and realizing losses, it seems appropriate to consider the delay variable in terms of months or even years; that is, for a short-tailed line of business, the delay may well be relatively quick (e.g., half of a year) while for a long-tailed line of business, the delay may be much longer (e.g., five years). Ultimately, though, as we do not posit a mechanism for the delay, the scale of the delay is also not really material.[7]

3. OPTIMAL STOCHASTIC CONTROL

This section will give a nonrigorous introduction to the theory of stochastic control and explain its applicability to this problem. A further introduction to stochastic control can be found in Astrom (1970).

Dynamic programming was developed by Bellman (1962). It can be regarded as a method to solve the problem of minimizing an objective function of the path over time of a state variable by changing a control variable, where the dynamics of the state variable are specified by an ordinary differential equation, which may contain the control variable as a parameter.

Stochastic control theory can be regarded as a generalization of dynamic programming, where the state variable follows a stochastic differential equation rather than an ordinary differential equation, and the control variable is chosen so as to minimize the expected value of the objective function. To our knowledge, stochastic control theory was first applied in a finance setting by Merton (1969, 1971). Stochastic control theory has also been applied to insurance by those investigating optimal insurance choices of individuals (Moore and Young 2006), optimal investment strategies of insurers (Hipp and Plum 2000, 2003; Schmidli 2005; Browne 1995), new business investment strategies of insurers (Hipp and Taksar 2000), optimal reinsurance purchases (Asmussen et al. 2000), and optimal dividend distribution (Asmussen and Taksar 1997; Asmussen et al. 2000; Albrecher and Hartinger 2007), to name a few (see also Schmidli [2008] for a overview on the use of stochastic control in insurance). Additionally, Goovaerts et al. (1992) incorporate insurance cycles into a stochastic model when studying insurer solvency. However, stochastic control theory has yet to be used to examine the *existence* of underwriting cycles.

A general formulation of a stochastic control problem (of the Bolza type) is as follows:

$$\max_{c\in C} E\left[\int_0^T f\left(c_{t,}, X_t, t\right) dt\right] \text{s.t. } dX_t = \mu\left(c_t, X_t, t\right) dt + \sigma\left(c_t, X_t, t\right) dw_t. \tag{2}$$

(Note that this restricts $(X_t)_{t\geq 0}$to follow an Itô process; since Itô processes are Markov, we need consider only controls c_t of the feedback [or Markov] type.) Defining

$$J\left(X_t, t\right) = E\left[\int_0^T f\left(c_s^*, X_s, t\right) ds \middle| \Im_t\right] \tag{3}$$

where c_t^* is the optimal control, we can invoke Bellman's principle of optimality (see Kamien and Schwarz 1991), which states that any subpath of the optimal solution must be optimal for that subproblem. This implies that along the optimal (expected) path, the change in $J(X_t,t)$ must be given exactly by the negative of the value of the objective function evaluated at t, after changes in t and the drift in $J(X_t,t)$ caused by (expected) changes in $(X_t)_{t\geq 0}$ are taken into account. Since $J(X_t,t)$ is a function of an Itô process, we can invoke Itô's lemma to write Bellman's principle of optimality in this case as

$$-f\left(c_t^*, X_t, t\right) = J_t + \sup_{c\in C}\left[J_X \mu\left(c_t, X_t, t\right) + \frac{1}{2}\cdot J_{XX}\sigma^2\left(c_t, X_t, t\right)\right]. \tag{4}$$

[7]It is possible that the delay is shorter than insurer reporting periods. That is, if the delay is three months, then an insurer reporting annually may not necessarily exhibit a cycle since the year-over-year accounting results would have "smoothed" the cycle. Of course, this does not mean the cycle does not exist. Rather, it simply means the cycle is harder to observe.

A further condition for c_t^* to be optimal is that it satisfies the first-order condition obtained by differentiating the expression inside the brackets in (4) with respect to c_t. These two conditions together form the Hamilton-Jacobi-Bellman (HJB) equation of stochastic control. Note that we require only that μ and σ, (and eventually our y process) be sufficiently smooth and integrable to ensure that J exists.

The HJB equation is a partial differential equation and is often nonlinear. It needs to be solved using a boundary condition specified in the original problem (in the specification given above, the boundary condition is $J(X_T, T) = 0 \forall X_T$). Needless to say, this is often difficult, making analytical solutions hard to come by, although numerical methods can easily be applied to many problems. A further complication is that the HJB equation is a necessary but not sufficient condition for optimality, meaning that any candidate solution needs to be tested further.

A further drawback of the method of stochastic control is that it is impossible to introduce constraints on the state variable explicitly. Intuitively, it can be seen that imposing a constraint on the state variable when the state variable moves randomly could lead to nonsensical results if the problem has not been set up to naturally incorporate these constraints. Merton (1969, 1971) dealt with this problem by showing that under his optimal solution the non-negative constraint on wealth could never be breached.

With these warnings in mind, we hope to apply the method of stochastic control theory to our insurance company. However, before we do so, we would like to briefly examine a discrete time alternative to stochastic control theory, called engineering process control.

4. ENGINEERING PROCESS CONTROL

Another branch of control theory, called engineering process control, contains results that we believe are applicable to our insurance company setting. This discussion will roughly follow Box and Jenkins (1976).

Box and Jenkins (1976) describe engineering process control in a discrete-time setting based on ARIMA models. In their framework, a discrete version of our insurer underwriting loss model would be written as

$$x_t = P_{t-\alpha} + \phi^{-1}(B)\,\theta(B)\,\varepsilon_t, \tag{5}$$

where B is the backshift operator, ε_t is white noise, and φ and θ are deterministic polynomial functions of order p and q, respectively. Note that we now require the loss process $(y_t)_{t\geq 0}$ to be an ARIMA(p,d,q) process.

We can then derive what Box and Jenkins (1976) call a minimum mean squared error (MMSE) control, which is a feedback control designed to minimize the mean squared error of the process from a target value (in our case, zero). This gives

$$P_t^* = -E\left[\phi^{-1}(B)\,\theta(B)\,\varepsilon_{t+\alpha}\middle|\,\phi^{-1}(B)\,\theta(B)\,\varepsilon_t\right]. \tag{6}$$

Using the Wiener-Kolmogorov prediction formula (see Hamilton 1994), we can write this as

$$P_t^* = -E\left[\frac{\phi^{-1}(B)\,\theta(B)}{B^{\alpha}}\right]_{+}\varepsilon_t \tag{7}$$

or alternatively as

$$P_t^* = -E\left[\frac{\phi^{-1}(B)\,\theta(B)}{B^{\alpha}}\right]_{+}\left[\phi^{-1}(B)\,\theta(B)\right]^{-1}P_t^*, \tag{8}$$

which is a restatement of our later results in an ARIMA (p,d,q) setting.

Our formulation is somewhat more general than Box and Jenkins (1976). Our processes are continuous in time and are more general than stationary ARIMA processes (general Itô processes need not necessarily be mean stationary or homoskedastic).

However, Box and Jenkins (1976) highlight one interesting result: They note that under certain circumstances, MMSE solutions take on more and more of an "alternating" character, the control at time t reversing a substantial portion of the control at time $t-1$. In other words, under certain circumstances, MMSE controls lead to cyclical solutions.[8]

They then use this idea to introduce what they call constrained control, which in our paradigm could be regarded as introducing a cost to adjusting the control. They note that large reductions in the variance of the optimal control can be achieved by relatively

[8] Box and Jenkins seem to imply that their MMSE *control* takes on an alternating character; we find that the stochastically controlled *output variable* alternates, while the control follows some (not necessarily cyclical) Itô process. We have not been able to resolve this apparent contradiction.

small sacrifices in the mean square error of the process—in other words that even a small cost of adjusting the control variable might significantly reduce the cyclicality of the optimal control.

Demonstrating the equivalences of optimal stochastic control and engineering process control is beyond the scope of this article; however, with this Box-Jenkins result in mind it is intriguing to examine the circumstances under which optimal stochastic control will produce cyclical profits.

5. RATIONAL EXPECTATIONS (BASIC) MODEL

Our basic model will be as described above, but here we set the general Itô process $(y_t)_{t \geq 0}$, to be standard Brownian motion, $(w_t)_{t \geq 0}$. As discussed above, the insurance firm will be assumed to minimize the mean squared error. Thus, (1) becomes

$$\arg\max_{\{P_s\}} E_{\alpha,w}\left[-\int_{t+\alpha}^{T} x_s^2 ds \,|\Im_t\right] \text{s.t. } x_s = P_{s-\alpha} + w_s, \text{ for } t+\alpha \leq s \leq T. \tag{9}$$

As before, in this program the insurer is attempting to minimize the deviations of profit from zero while adjusting premiums (*P*) in response to losses (now, w, representing Brownian motion) that are delayed by α. Again, since both the delay and losses are unknown to the insurer, the insurer must consider the distributions of both.

The optimal premium at time *t* can then be written as (see Appendix A for proof)

$$P_t^* = P_{t-\alpha}^* - x_t. \tag{10}$$

Here the insurer can be thought of as setting the premium today with explicit consideration of the realized delay. Now, recalling the profit function, the optimal profit can be written as

$$x_t^* = P_{t-\alpha}^* + w_t. \tag{11}$$

Equation (11) shows that the profit observed at time *t* is the premium received at time $t - \alpha$ in excess of the underwriting loss at time *t*. This would be a fairly reasonable outcome for an insurance company that collects premium early (e.g., at time $t - \alpha$) for losses and subsequent profits observed at a later time (*t*). Note that since α is random we do not know exactly how "late" the losses will be *before* time *t*. When we calculate the expectation of the premium (e.g., Eq. [10]) or profits (e.g., Eq. [11]) *at* time *t*, we know the realization of α. We also note here that we have solved for the premiums and profits at one point in time. When solving for the premiums and profits in other periods, the mathematics would remain the same, but the realizations of both losses and the delay could change. Finally, although there is no explicit learning of the delay function in the model, the distribution of the delay is considered to be known by the insurer.

Substituting the optimal $P_{t-\alpha}^*$ into (11) we can solve for the optimal profit at time *t*, and (11) becomes

$$x_t^* = P_{t-2\alpha}^* - x_{t-\alpha}^* + w_t. \tag{12}$$

Using (11) again, (12) becomes

$$x_t^* = P_{t-2\alpha}^* - (P_{t-2\alpha}^* + w_{t-\alpha}) + w_t. \tag{13}$$

We can now see that the optimal profit at time *t* is

$$x_t^* = w_t - w_{t-\alpha}. \tag{14}$$

Equation (14) suggests that insurer profits in time *t* will be a function of the realized losses in time *t* as well as the losses with an α delay. Intuitively one can consider the profit as being a function of current losses and the losses that were considered in setting the premium. Again, this results from both random losses and a random delay process.

If we slightly reformulate our model, we have a continuous time version of Cummins and Outreville (1997). Writing (1) as

$$\arg\max_{\{P_s\}} E_{\alpha}\left[-\int_{t}^{T} (Q_s + w_s)^2 \, ds \,|\Im_{t-\alpha}\right] \text{s.t. } dw_t = dB_t, \tag{15}$$

where $Q_s = P_{s-\alpha}$ is $\Im_s$measurable , $(B_t)_{t\geq 0}$ is standard Brownian motion, reconciles our framework with Cummins and Outreville (1997).[9]

From this model, we can show the optimal premium path can be written as (see Appendix B for the proof)

$$Q_t = -w_{t-\alpha}. \tag{16}$$

Recall that profits at time t are written as $x_t = Q_t + w_t$. Substituting (16) into this equation yields the optimal profits at time t:

$$x_t = -w_{t-\alpha} + w_t. \tag{17}$$

Notice that this is exactly the same as the result of our model.

When testing the cyclicality of (14) and (17), it is necessary to examine the autocovariance function for $(x_t)_{t\geq 0}$. Appendix C shows that the autocovariance function of order j for our optimal profit path can be written as

$$\gamma = \max(\alpha - j, 0). \tag{18}$$

We will discuss this term and its implications more in the following sections. Note that the profit under the optimal control is the expected value of the loss process $(y_t)_{t\geq 0}$ at time $t-\alpha$, when $(P_t)_{t\geq 0}$ is chosen, less the actual loss process experience at time t. This can be shown to be true in the general case.

6. ANALYSIS OF CYCLICALITY

Having solved for the optimal profit path, we now demonstrate that it is cyclical. We first show that in spite of the fact that the unconditional expectation of profits is consistently 0, the optimal path of profits is not a martingale. We also show that the optimal profits follow an autoregressive process of order α. We derive the theoretical autocorrelation function of our differenced optimal profit stream in the case of a Brownian motion loss process. We also demonstrate our model and calculate sample autocorrelation functions by simulating a Brownian motion loss process. We ultimately find that this model results in underwriting cycles.

Although we report results for one Itô process, our model guarantees cyclical behavior for *any* nondeterministic Itô process.

Lemma 1

For all nondeterministic Itô loss processes, the insurance profits will not follow a martingale; in particular, they will follow some autoregressive process of order α.

Proof

$$\begin{aligned} E\left[x_s^* \middle| \Im_t\right] &= P_{s-\alpha}^* + E\left[w_s \middle| \Im_t\right] \text{ if} t < s < t+\alpha \\ &= E\left[w_s \middle| \Im_{s-\alpha}\right] + E\left[w_s \middle| \Im_t\right]. \end{aligned}$$

Now, $x_t^* = P_{t-\alpha}^* + w_t = E\left[w_t \middle| \Im_{t-\alpha}\right] + w_t$, and hence, for $\left(x_t^*\right)_{t\geq 0}$ to be a martingale, we require

$$\begin{aligned} E\left[w_t \middle| \Im_{t-\alpha}\right] + w_t &= E\left[w_s \middle| \Im_{s-\alpha}\right] + E\left[w_s \middle| \Im_t\right] \text{ or} \\ E\left[w_t \middle| \Im_{t-\alpha}\right] - E\left[w_s \middle| \Im_{s-\alpha}\right] &= E\left[w_s \middle| \Im_t\right] - w_t. \end{aligned}$$

Now, since all Itô processes are Markov, the left-hand side depends only on the development of $(w_t)_{t\geq 0}$ up to time $s-\alpha$, while the right-hand side depends on the development of $(w_t)_{t\geq 0}$ up to time t. Since the process is not deterministic, the left-hand side cannot equal the right-hand side with probability 1. Hence, $\left(x_t^*\right)_{t\geq 0}$ cannot be a martingale.

Note that $x_t^* = P_{t-\alpha}^* + w_t = E\left[w_t \middle| \Im_{t-\alpha}\right] + w_t$, and since $(w_t)_{t\geq 0}$ is Markov, this can be written as

$$f\left(w_{t-\alpha}, t-\alpha\right) + w_t$$

for some deterministic function f. Hence, $\left(x_t^*\right)_{t\geq 0}$ is autoregressive of order α. ♦

[9]The slightly odd-looking SDE in (15) is simply to emphasize that the state variable follows a standard Brownian motion.

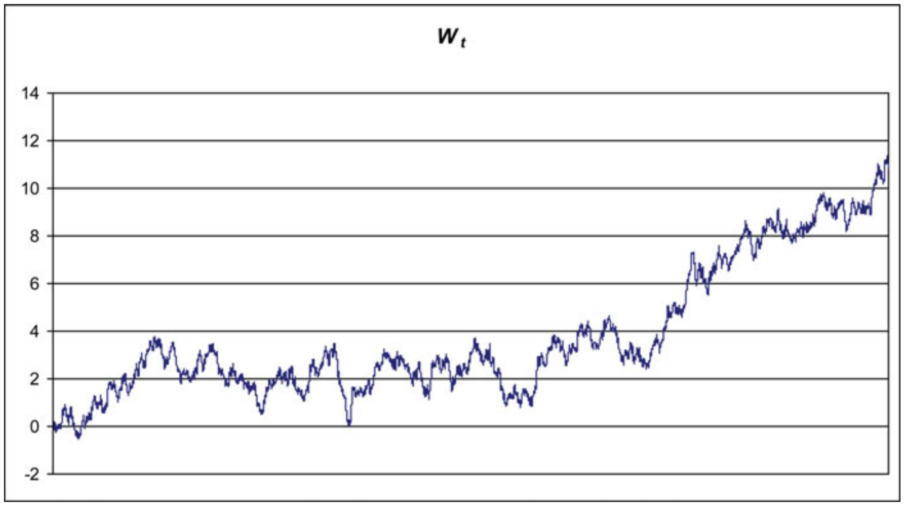

FIGURE 1. Example Loss Process.

Note that this lemma states that the optimal profit stream does not follow a martingale, even though the unconditional expectation of profits is zero for all t.

Furthermore, while we have hitherto provided justification for seeking to minimize the mean squared error of underwriting profits, our results are not dependent on this assumption. In fact, we find that our result will hold for any strictly concave function of current profits.[10]

Corollary

Lemma 1 holds for any concave objective function of current profits.

Proof

Note that the quadratic function, by definition, minimizes the difference between the observed value $(w_t)_{t\geq 0}$ and the predicted value of $(w_t)_{t\geq 0}$ at time $t-\alpha$. Hence, any other objective function will increase the distance, implying that the argument used in Lemma 1 still applies. ⧫

We continue this section by analyzing the result from our basic model for evidence of cyclicality. In this section we treat the delay (α) as deterministic to facilitate the simulation.

7. BASIC MODEL CYCLICALITY

Recall the profit process from our first model:

$$x_t^* = w_t - w_{t-\alpha}, \tag{19}$$

where $(w_t)_{t\geq 0}$ is standard Brownian motion. Also recall the autocovariance function for this process:

$$\gamma = \max(\alpha - j, 0). \tag{20}$$

From this we can derive the autocovariance function of the differenced profits.

Discretizing $(w_t)_{t\geq 0}$ by writing it as a sum of independently identically distributed normal random variables, and taking first-order differences, we can see that

$$x_t^*(1) = \varepsilon_t - \varepsilon_{t-\alpha}, \tag{21}$$

[10]As an example, we can show that the solution to the problem with capital constraints can be written as the solution without capital constraints plus some perturbation term due to capital constraints.

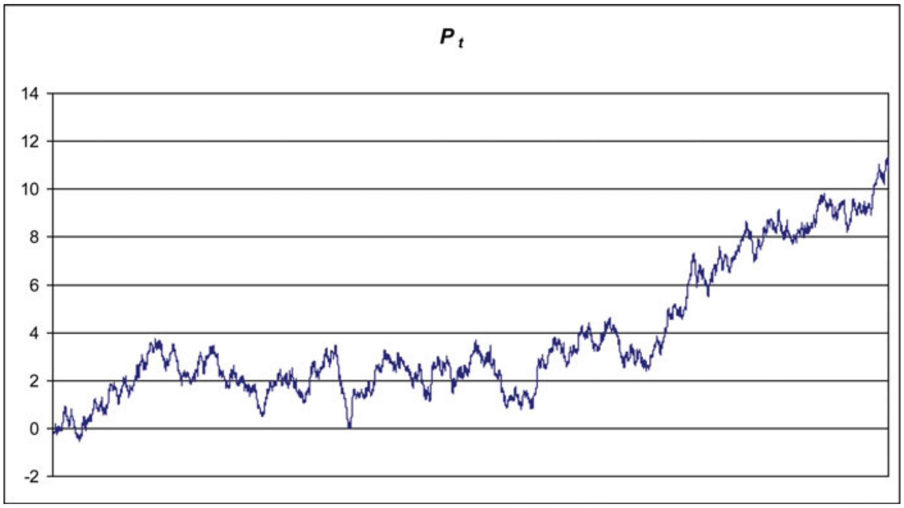

FIGURE 2. Example Premium Process.

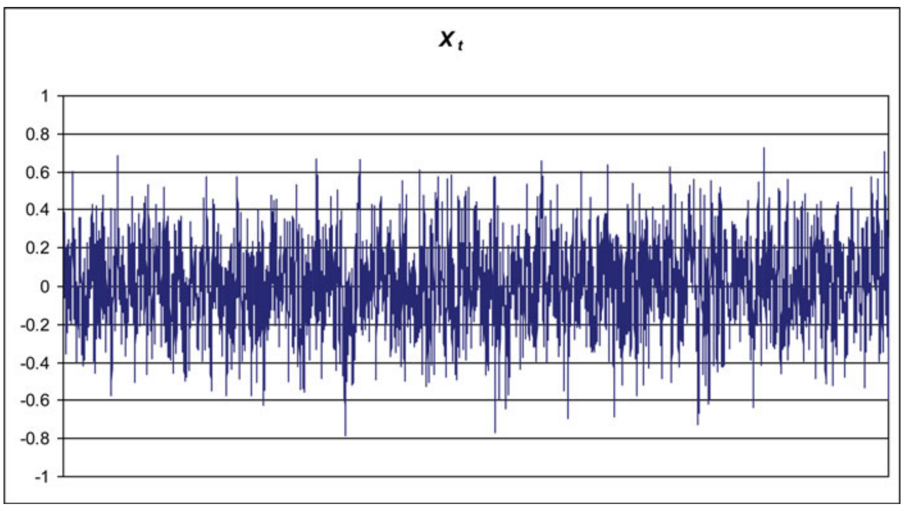

FIGURE 3. Example Profit Process.

where ε_t has a $N(0,1)$ distribution. This allows the autocorrelation function to be easily deduced as

$$\rho_j(1) = \begin{cases} 1 & \text{if } j = 0 \\ -0.5 & \text{if } j = \alpha \\ 0 & \text{otherwise} \end{cases}. \tag{22}$$

Hence, we have shown that profits follow an integrated autoregressive process of order α.

As a simple example, we generated our loss process $(w_t)_{t\geq 0}$ with a lag equal to five. Figure 1 shows one such Brownian motion process. With this simulated loss process and using the results from our first model (Eq. [10]), we can now determine the premium process, $(P_t)_{t\geq 0}$ shown in Figure 2.

Since the premiums are a function of losses, Figure 2 is very similar to Figure 1. The difference between the two series (as suggested by the optimal profit process shown in Eq. (14)) results in the profit process shown in Figure 3.

Although it may not be immediately clear from Figure 3 that the profit process is cyclical, Figures 4 and 5 show the partial autocorrelation function of the original profit process and the autocorrelation function of the differenced profit process, respectively. While we can use the partial autocorrelation function to get an idea as to cyclicality, we specifically notice that the autocorrelation

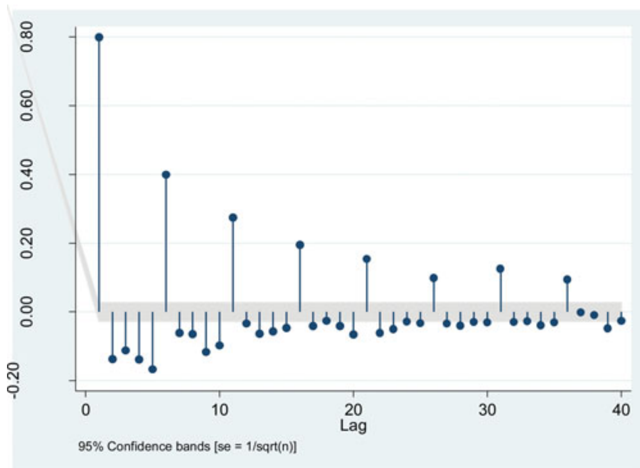

FIGURE 4. Partial Autocorrelation Function (Profit Process).

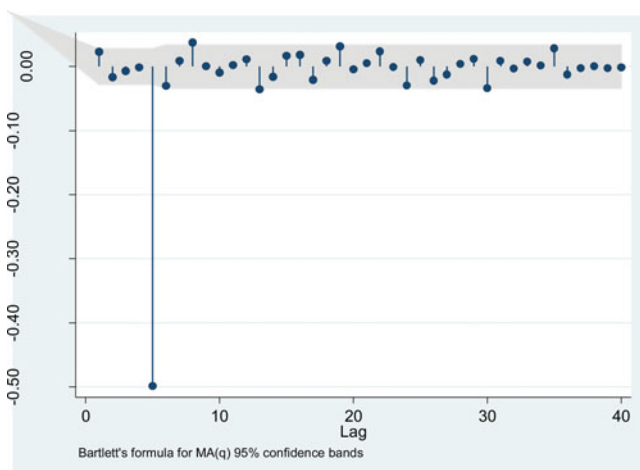

FIGURE 5. Autocorelation Function (Differenced Profit Process).

function of the differenced process has a coefficient of -0.5 on the α-lag term. According to Box and Jenkins (1976), these parameter values are sufficient for the profits process to display *pseudo-periodic* behavior.

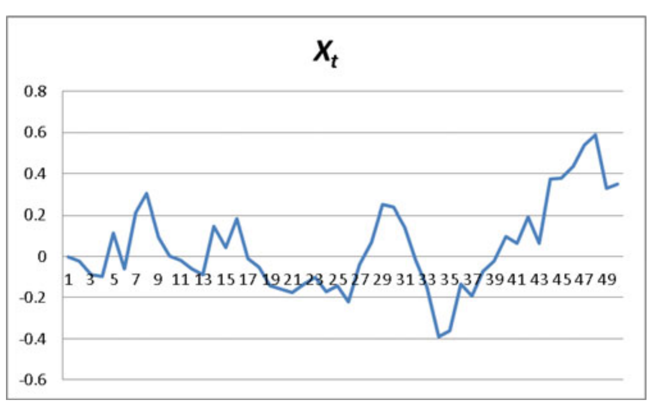

FIGURE 6. Example Profit Process (First 50 Realizations).

Figure 6 shows the same profit function for the first fifty realizations of the profit in the simulated process. With fewer observations, the cyclicality of the profit process is much more apparent.

8. DISCUSSION AND CONCLUSION

Underwriting cycles in insurance profits have historically been well studied. Until recently, most research has attempted to explain the existence of the cycles. In particular, Cummins and Outreville (1987) put forward a rational expectations model explaining cycles. We have generalized Cummins and Outreville (1987) in continuous time. Additionally, we have allowed for a stochastic delay in regulatory/accounting lags that account for cycles in Cummins and Outreville. We have found, as do those authors, that underwriting cycles can be explained with a rational expectations model. However, we also find that the cycle length is entirely dependent on the length of the lag.

Although our results are consistent with those theoretical results of Cummins and Outreville (1987), they are also consistent with empirical studies on underwriting cycles. Cummins and Outreville themselves show underwriting cycles of differing lengths in different countries. Further, Lamm-Tennant and Weiss (1997) find variation in underwriting cycle length across countries and across insurance lines. Our theoretical results showing cycles varying with the various lags due to the regulatory and reporting environments are consistent with both of these empirical studies; that is, the stochastic delay we introduce in our model can be easily imagined to be different across countries and lines of insurance, thereby giving rising to different cycle lengths.

This result, consistent with extant research on underwriting cycles can also be viewed as comporting with recent research of Boyer et al. (2012) and Boyer and Owadally (2015), who have called into question the existence of the cycles. In particular, if there is indeed no lag, there is no cycle.

Much work can still be done for insurance cycles in continuous time. In particular, our model with capital constraints can be examined further. Additionally, it may be interesting to introduce two lags, one for premiums and another for losses. The stochastic model we develop in this article is still rather rudimentary. As a result, we can only solve the control problem by a point-to-point optimization method. In fact, it may be possible to develop a more general model. However, it turns out that solving the general model would involve irregular stochastic control techniques. We hope that the model presented here provides a starting point for a more thorough investigation into underwriting cycles using stochastic control.

FUNDING

Xudong Zeng is supported by NSFC Grant No. 71271127.

REFERENCES

Akerlof, G. 1970. The Market for “Lemons”: Quality Uncertainty and the Market Mechanism. *Quarterly Journal of Economics* 84(3): 488–500.

Albrecher, H., and J. Hartinger. 2007. A Risk Model with Multi-layer Dividend Strategy. *North American Actuarial Journal* 11(2): 43–64.

Asmussen, S., B. Hojgaard, and M. Taksar. 2000. Optimal Risk Control and Dividend Distribution Policies. Example of Excess-of-Loss Reinsurance for an Insurance Corporation. *Finance and Stochastics* 4: 299–324.

Asmussen, S., and M. Taksar. 1997. Controlled Diffusion Models for Optimal Dividend Pay-out. *Insurance: Mathematics and Economics* 20: 1–15.

Astrom, K. J. 1970. *Introduction to Stochastic Control Theory.* New York: Academic Press.

Beck, A. 2014. *Introduction to Nonlinear Optimization*. Philadelphia: Society for Industrial and Applied Mathematics.

Bellman, R. E., and S. E. Dreyfus. 1962. *Applied Dynamic Programming.* Princeton: Princeton University Press.

Box, G. E. P., and G. M. Jenkins. 1976. *Time Series Analysis—Forecasting and Control*. Oakland: Holden-Day.

Boyer, M. M., E. Jacquier, and S. Van Norden. 2012. Are Underwriting Cycles Real and Forecastable? *Journal of Risk and Insurance* 79(4): 995–1015.

Boyer, M. M., and I. Owadally. 2015. Underwriting Apophenia and Cryptids: Are Cycles Statistical Figments of our Imagination? *Geneva Papers on Risk and Insurance—Issues and Practice* 40: 232–255.

Browne, S. 1995. Optimal Investment Policies for a Firm with a Random Risk Process: Exponential Utility and Minimizing the Probability of Ruin. *Mathematics of Operations Research* 20(4): 937–958.

Cohen, E., and N. Megiddo. 1990. Maximizing Concave Functions in Fixed Dimension. Research Report RJ 7656 (71103), IBM Almaden Research Center, San Jose.

Cummins, J. D., and P. Danzon. 1997. Price, Financial Quality and Capital Flows in Insurance Markets. *Journal of Financial Intermediation* 6: 3–38.

Cummins, J. D., and J. F. Outreville. 1987. An International Analysis of Underwriting Cycles in Property-Liability Insurance. *Journal of Risk and Insurance* 54(2): 246–262.

Einav, L., and A. Finkelstein. 2011. Selection in Insurance Markets: Theory and Empirics in Pictures. *Journal of Economic Perspectives* 25(1): 115–138.

Federico, S. 2011. A Stochastic Control Problem with Delay Arising in a Pension Fund Model. *Finance and Stochastics* 15: 421–459.

Froot, K. A., and P. G. J. O’Connell. 1997. The Pricing of U.S. Catastrophe Reinsurance. NBER Working Paper No. 6043. Cambridge, MA: NBER.

Froot, K. A., D. S. Scharfstein, and J. C. Stein. 1993. Risk Management: Coordinating Corporate Investment and Financing Policies. *Journal of Finance* 48(5): 1629–1658.

Goovaerts, M. J., F.DeVylder, and R. Kaas. 1992. A Stochastic Approach to Insurance Cycles. *Insurance: Mathematics and Economics* 11: 97–107.

Gozzi, F., C. Marinelli, and S. Savin. 2009. On Controlled Linear Diffusions with Delay in a Model of Optimal Advertising under Uncertainty with Memory Effects. *Journal of Optimization Theory and Applications* 142: 291–321.

Gron, A. 1994a. Capacity Constraints and Cycles in Property-Casualty Insurance Markets. *Rand Journal of Economics* 25: 110–127.
Gron, A. 1994b. Evidence of Capacity Constraints in Insurance Markets. *Journal of Law and Economics* 37: 349–377.
Hamilton, J. D. 1994. *Time Series Analysis*. Princeton: Princeton University Press.
Higgins, M. L., and P. D. Thistle. 2000. Capacity Constraints and the Dynamics of Underwriting Profits. *Economic Inquiry* 38(3): 442–457.
Hipp, C., and M. Plum. 2000. Optimal Investment for Insurers. *Insurance: Mathematics and Economics* 27: 215–228.
Hipp, C., and M. Plum. 2003. Optimal Investment for Investors with State Dependent Income, and for Insurers. *Finance and Stochastics* 7: 299–321.
Hipp, C., and M. Taksar. 2000. Stochastic Control for Optimal New Business. *Insurance: Mathematics and Economics* 26: 185–192.
Kamien, M. I., and N. L. Schwartz. 1991. *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management.* Boston: Elsevier Science.
Lamm-Tennant, J., and M. Weiss. 1997. An Investigation of International Underwriting Cycles. *Journal of Risk and Insurance* 64: 415–439.
Lazar, D., and M. Denuit. 2011. New Evidence for Underwriting Cycles in US Property-liability Insurance. *Journal of Risk Finance* 13(1): 4–12.
Meier, U. B., and J. F. Outreville. 2006. Business Cycles in Insurance and Reinsurance: The Case of France, Germany, and Switzerland. *Journal of Risk Finance* 7(2): 160–176.
Merton, R. C. 1969. Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case. *Review of Economics and Statistics* 51(3): 247–257.
Merton, R. C. 1971. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory* 3: 373–413.
Moore, K. S., and V. R. Young. 2006. Optimal Insurance in a Continuous-Time Model. *Insurance: Mathematics and Economics* 39: 47–68.
Neihaus, G., and A. Terry. 1993. Evidence on the Time Series Properties of Insurance Premiums and Causes of the Underwriting Cycle: New Support for the Capital Market Imperfection Hypothesis. *Journal of Risk and Insurance* 60(3): 46–479.
Schmidli, H. 2005. On Optimal Investment and Subexponential Claims. *Insurance: Mathematics and Economics* 36: 25–35.
Schmidli, H. 2008. *Stochastic Control in Insurance*. London: Springer.
Venezien, E. C. 1985. Ratemaking Methods and Profit Cycles in Property and Liability Insurance. *Journal of Risk and Insurance* 52: 477–500.
Venezian, E. C., and C.-C. Leng. 2006. Application of Spectral and ARIMA Analysis to Combined-Ratio Patterns. *Journal of Risk Finance* 7(2): 189–214.
Winter, R. 1994. The Dynamics of Competitive Insurance Markets. *Journal of Financial Intermediation* 3: 379–415.

Discussions on this article can be submitted until July 1, 2017. The authors reserve the right to reply to any discussion. Please see the Instructions for Authors found online at http://www.tandfonline.com/uaaj for submission instructions.

APPENDIX A. PROOF FOR OPTIMAL PREMIUM

From our first model, the optimization problem is as follows:

$$\underset{\{P_s\}}{\arg\max}\, E_\alpha\left[-\int_{t+\alpha}^{T} x_s^2 ds\,|\Im_t\right] \text{s.t.} x_s = P_{s-\alpha} + w_s \text{ for } t+\alpha \le s \le T. \tag{A.1}$$

Substituting in for x_s yields

$$\underset{\{P_s\}}{\arg\max}\, E_\alpha\left[-\int_{t+\alpha}^{T} (P_{s-\alpha} + w_s)^2\, ds\,|\Im_t\right]. \tag{A.2}$$

Using the law of iterated expectations, we can easily verify that (A.2) can be written as

$$\underset{\{P_s\}}{\arg\max}\, E_\alpha\left[E\left[-\int_{t+\alpha}^{T} (P_{s-\alpha} + w_s)^2\, ds\,\middle|\Im_t, \alpha\right]\right]. \tag{A.3}$$

Similarly, we can interchange the integration and expectation operators to yield

$$\underset{\{P_s\}}{\arg\max} - \int_{t+\alpha}^{T} E_\alpha\left[E\left[(P_{s-\alpha} + w_s)^2\,\middle|\Im_t, \alpha\right]\right]. \tag{A.4}$$

To further simplify (A.4), it will be useful to evaluate $E_\alpha\left[E\left[(P_{s-\alpha} + w_s)^2\middle|\,\Im_v, \alpha\right]\right]$ for some v. Expanding the binomial term yields

$$E_\alpha\left[E\left[\left(P_{s-\alpha}^2 + 2w_s P_{s-\alpha} + w_s^2\right)\middle|\,\alpha, \Im_v\right]\right]. \tag{A.5}$$

When $s - \alpha \le v$, distributing the expectation operator across each term, (A.5) becomes

$$P_{s-\alpha}^2 + 2P_{s-\alpha} E\left[w_s|\,\alpha, \Im_v\right] + E\left[w_s^2\middle|\,\alpha, \Im_v\right]. \tag{A.6}$$

Manipulation of (A.6) yields the following:

$$P_{s-\alpha}^2 + 2P_{s-\alpha}E\left[w_t + (w_s - w_v)\middle|\, \Im_v, \alpha\right] + E\left[(w_s - w_v)^2\middle|\, \Im_v, \beta\right]$$
$$+2E\left[w_v\,(w_s - w_v) + w_v^2\middle|\, \Im_v, \alpha\right]. \tag{A.7}$$

When $s \geq v$, substituting $w_t = x_t - P_{t-\alpha}$, recognizing that Brownian motion has mean of zero, and that $E_w\left[(w_s - w_t)^2\middle|\, \Im_t, \alpha\right] = E[s - t|\alpha]$, we can rewrite (A.5) as

$$E_\alpha\left[P_{s-\alpha}^2 + 2P_{s-\alpha}(x_v - P_{v-\alpha}) + (s - v) + (x_v - P_{v-\alpha})^2\right]. \tag{A.8}$$

To solve (A.1), it is sufficient that we minimize the term inside of the expectation in (A.8) for all $s$and v.

Let $s = t + \alpha$, $v = t$in (A.8). We take the first-order condition with respect to our decision variable P_t, which can easily be seen to be

$$P_t + (x_t - P_{t-\alpha}) = 0. \tag{A.9}$$

Hence, the optimal premium satisfies the following (recursive) equation:

$$P_t^* = P_{t-\alpha}^* - x_t.$$

APPENDIX B. PROOF FOR OPTIMAL PREMIUM PATH

Recall our model within the Cummins and Outreville framework is written as

$$\arg\max_{\{Q_s\}} E_\alpha\left[-\int_t^T (Q_s + w_s)^2\, ds\, |\Im_{t-\alpha}\right] \text{s.t.wtisastandardBrownianmotion.} \tag{B.1}$$

Expanding the binomial term allows us to rewrite (B.1) as

$$\max_{\{Q_s\}} E_\alpha\left[-\int_t^T \left(Q_s^2 + 2Q_s w_s + w_s^2\right)\, ds\, |\Im_{t-\alpha}\right]. \tag{B.2}$$

By the law of iterated expectations, (B.2) can be written as

$$\max_{\{Q_s\}} E_\alpha\left[-\int_t^T E\left[\left(Q_s^2 + 2Q_s w_s + w_s^2\right)\, ds\, |\Im_{t-\alpha}, \alpha\right] ds\right]. \tag{B.3}$$

Distributing the expectation operator yields

$$\max_{\{Q_t\}} -E_\alpha\left[\int_t^T \left(Q_s^2 + 2Q_s E\left[w_s|\, \Im_{t-\alpha}, \alpha\right] + E\left[w_s^2\, |\Im_{t-\alpha}, \alpha\right]\right)\, ds\right]. \tag{B.4}$$

The first-order condition of the integrand of (B.4) with respect to Q_t is

$$Q_t = -E\left[w_{t-\alpha}\middle|\, \Im_{t-\alpha}, \alpha\right] = -w_{t-\alpha}. \tag{B.5}$$

APPENDIX C. AUTOCOVARIANCE FUNCTION

Consider the generic autocovariance equation of order j as

$$\gamma_j = E\left[x_t x_{t-j}\right]. \tag{C.1}$$

Substituting our optimal profit path, from Equation (14), into (C.1) yields

$$E\left[(w_t - w_{t-\alpha})\left(w_{t-j} - w_{t-j-\alpha}\right)\right]. \tag{C.2}$$

Expanding this term results in

$$E\left[w_t w_{t-j}\right] - E\left[w_{t-\alpha} w_{t-j}\right] + E\left[w_{t-\alpha} w_{t-j-\alpha}\right] - E\left[w_t w_{t-j-\alpha}\right]. \tag{C.3}$$

Using the rules for Brownian motion covariates (discussed in Appendix A) (C.3) simplifies to

$$(t - j) - \min(t - \alpha, t - j). \tag{C.4}$$

Rearranging terms and simplifying ultimately results in the following autocovariance formula:

$$\gamma = \max(\alpha - j, 0). \tag{C.5}$$