

The CAS *E-Forum*, Fall 2021

The Fall 2021 edition of the CAS *E-Forum* is a cooperative effort between the CAS *E-Forum* Committee and various CAS committees, task forces, working parties and special interest sections. This *E-Forum* contains the winner of the CAS Reinsurance Call Paper Prize.

Reinsurance Working Group

Zeeshan Ali	Linda Grand	Chris E. Nelson
Fernando Alberto Alvarado	Wei Juan Han	Kagabo E. Ngiruwonsanga
Diego Fernando Antonio	Sara J. Hemmingson	Vincent Philip Pomo
Glenn R. Balling	Weiming Hong	Jayson James Ramdany
Jon Beaver	Ya Jia	Jiandong Ren
Karen B. Buchbinder	Kara Dawn Kemsley	Terence Richard Robinson
James Chang	Richard Scott Krivo	Viqas Ahmed Shah
Aritra Chatterjee	Clifford Kin Lok Lau	Cunhua Shi
Wei Chen	Hao Li	Karen Sonnet
Jordan Paul Comacchio	Wenyi Li	Mou Jian Teo
David Core	Ruide Lin	Darren Choon Boon Toh
Matthew John Engelbert	Henry Ding Liu	Tracy Leslie Valentine
Isaac R. Espinoza	Brian E. MacMahon	Phillip Anthony Victory
Katrina Evans, <i>Staff Liaison</i>	Stephanie T. Magnuson	Chengwei Wang
Brian A. Fannin, <i>Staff Liaison</i>	Alexander James Martin	Hao Wang
Daniel Enrique Fernandez	Gregory F. McNulty	Ming Yi
Thomas Patrick Fiorillo	Xingyi Meng	
	Brian J. Mullen	

CAS *E-Forum*, Fall 2021

Table of Contents

2021 Reinsurance Call Paper Prize Winner

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

Neil M. Bodoff, FCAS, MAAA. 1-33

***E-Forum* Committee**

Derek A. Jones, *Chairperson*

Michael Li Cao

Mark M. Goldburd

Karl Goring

Gregory F. McNulty

Bryant Edward Russell

Shayan Sen

Brandon S. Smith

Elizabeth A. Smith, *Staff Liaison/Director of Publications*

Zongli Sun

Yingjie Zhang

For information on submitting a paper to the *E-Forum*, visit <https://www.casact.org/publications-research/guide-submission-e-forum>.

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

Neil M. Bodoff, FCAS, MAAA

Abstract

Motivation. Current paradigms for insurance risk derive primarily from financial and mathematical models of risk and ignore the role of human motivations.

Method. Investigate key ideas of Principal-Agent Theory.

Results. Applying to insurance the framework of Principal-Agent Theory unlocks new insights that differ from prior paradigms and better explains observable phenomena in the insurance market.

Conclusions. Non-systemic risk and non-tail risk nevertheless manifest principal-agent costs that ought to be incorporated into the cost of insurance risk.

Keywords. Insurance. Reinsurance. Risk. Pricing. ERM. Principal-Agent Theory.

1. INTRODUCTION

One of the foundational challenges in insurance and reinsurance is how to price risk and how to manage risk. Historically, answers have been developed by applying the contexts of Portfolio Theory, Enterprise Risk Management (ERM), Rate of Return, Cost of Capital, and Capital Allocation to insurance and reinsurance. However, the insurance industry has mostly skipped over the framework of Principal-Agent Theory (PAT), which provides a different approach to these questions of how to price insurance risk and how to manage insurance risk.

1.1 Research Context

Insurance risk pricing and insurance risk management have been heavily influenced by mathematical models developed in asset pricing theory; an example is Kreps's "Investment-Equivalent Reinsurance Pricing." I have not found significant usage of Principal-Agent Theory.

1.2 Objective

The objective of this paper is to supplement the current frameworks and to leaven the current discussion by citing Principal-Agent Theory as a key component of any discussion of insurance risk pricing.

1.3 Outline

The remainder of the paper proceeds as follows. Section 2 will describe the importance of

choosing a suitable “mental model” or conceptual framework for analyzing a topic, Section 3 will describe Principal-Agent Theory and why it is relevant to insurance risk, Section 4 will provide examples of its application to insurance, and Section 5 will provide initial suggestions for how to apply it quantitatively to insurance risk going forward.

2. MENTAL MODELS AND CONCEPTUAL FRAMEWORKS: THE METAPHORS WE SELECT SHAPE OUR THINKING

“The limits of my language are the limits of my world” – Wittgenstein

The conceptual framework or “mental model” that we select to represent a situation has great influence on our thinking. Choosing such metaphors can uncover great insight for us but can also limit and even skew our thinking. Thus, we must exert great care to scrutinize our selected analogies and metaphors to make sure they are suitable for the task at hand.

2.1 The Roulette Wheel and the Drunken Sailor: Process versus Outcome, Luck versus Skill, Randomness versus Determinism, and Exoneration versus Blame

One situation in which choice of metaphor is crucial arises whenever one must choose to judge based on process versus outcome. When someone follows a process which then results in an unexpected or unfavorable outcome, one has a few options about how to evaluate the situation:

1. Option #1: all that matters is the outcome. If the outcome is undesirable, then this fact indicates someone must have done something incorrect or inappropriate in the process.
2. Option #2: all that matters is the process. So long as someone follows the right process and the appropriate decision-making procedures, then an unexpected or undesirable outcome is essentially irrelevant.

Take an oversimplified example: I hand some money to my friend and instruct him (or her) to go place a bet on the roulette wheel in the local casino. When he returns, he indicates that all my money has evaporated. Now I imagine two possibilities: the friend correctly followed my instructions, placed a bet on the roulette wheel, and had an unlucky outcome arising from the roulette wheel’s randomness. Or he really acted like an “untrustworthy drunken sailor” and got drunk, spent my money on his own enjoyment in town, and did not follow my instructions. The more I adopt the framework of “roulette wheel,” the more I chalk up the loss to randomness and bad luck and thus find my friend blameless. But the more I adopt the mental model of “untrustworthy drunken sailor”

the more I view the outcome as nothing to do with randomness and luck, everything to do with skill, and thus find my friend to be blameworthy.

2.2 Uncertainty versus Certainty in Probabilities

"The mistake I'd say 98 percent of money managers and individuals make is they feel like they've got to be playing with a bunch of stuff," Druckenmiller said during his speech "Lost Tree Club" speech. "If you really see it, put all your eggs in one basket and watch the basket very carefully."

One of the hidden assumptions in adopting a “roulette wheel” framework is the problem of certainty of probabilities. Specifically, when one spins the roulette wheel, one knows for certain what all the possible outcomes are and what the probabilities associated with these outcomes are. There is no uncertainty in the underlying probabilities, only randomness or “process variance” in what the outcome will be.

In that situation, it does indeed make sense to denigrate “resulting” or “outcoming” i.e., judging based upon outcome.¹ And because all the probabilities are known going into a situation and because we would not revise our views of those probabilities based upon the observed outcome, it would make sense to “think in bets.”

But is such a framework suitable, appropriate, or useful for the real world? Examples from betting are stylized and might be inaccurate metaphors for the real world.² In the real world, probabilities are often uncertain, unknown, and murky; observing an unexpected, undesirable outcome leads us to revise our views of estimated probabilities.

2.3 Bayes' Theorem

"It seemed like a good idea at the time..." - Anonymous

In the real world, when we take action and then experience an adverse outcome, it often prompts us to re-evaluate the estimates and the information that we relied upon in the first place. We can use the analytical framework of Bayes' Theorem to give some partial weight to the outcome to revise our prior views of initial probabilities.

$$P(A|B) = P(A) * P(B|A) / P(B) \quad (2.1)$$

¹ See Duke [10]

² Taleb [27] and [28] refers to the “Ludic Fallacy”: incorrectly comparing real life situations to games

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

Where:

$P(A)$ = initial probability of A

$P(B)$ = probability of B

$P(B | A)$ = probability of B given A

$P(A | B)$ = revised probability of A given observed B

Let's revert back to our example in which I gave money to a friend to place a bet on a roulette wheel. Say that I am initially 99% confident that my friend is trustworthy and that I estimate only a 1% chance that he'll squander the money on his own enjoyment like a drunken sailor. I give him \$100 to place a bet on a roulette wheel's even numbers and he returns to me with zero dollars. This outcome can be attributed either to his being completely trustworthy and getting unlucky at the roulette table or rather it can be attributed to him squandering the money as an untrustworthy spendthrift.

$P(A)$ = initial probability that my friend is trustworthy = 99%

$1-P(A)$ = initial probability that my friend is untrustworthy = 1%

$P(B | A)$ = Probability of losing my money given that my friend is trustworthy = $1-18/38$ {i.e., there are 18 even numbered slots out of total of 38 slots on American roulette wheel} = 52.63%

$P(B)$ = Probability of all the ways of losing my money = $[99\% * 52.63\% + 1\% * 100\%]$ = 53.1%

$P(A | B)$ = revised probability that my friend is trustworthy = $P(A) * P(B | A) / P(B)$ = $99\% * 52.63\% / 53.1\%$ = 98.1%

$1-P(A | B)$ = revised probability that my friend is untrustworthy = $1-98.1\%$ = 1.9%

After sending my trusted friend to do a task that then resulted in an unfavorable outcome, I cannot judge entirely based upon outcome nor can I completely dismiss the outcome; rather, the outcome leads me to dramatically revise my estimated probability of his untrustworthiness from 1% to roughly 2%, which might exceed my level of comfort in the future.

In the real world, probabilities are uncertain and subject to revision. A superficial appeal to the metaphor of the roulette wheel and its randomness, luck, stochasticity, and process variance is inappropriate in the real world. And yet, the paradigm of fixed, known probabilities leading to pure random outcomes has persisted even as it has been transplanted from its initial natural habitat into the new domains of investment risk and insurance risk. But if the paradigm is the wrong, then the implications of this paradigm are also wrong. In the real world, outcomes matter.

3. INTRODUCTION TO PRINCIPAL-AGENT THEORY

“He doesn’t have to worry about the interim consequences of error. Most people do.” - Howard Marks, talking about Warren Buffett versus others

Principal-Agent Theory seeks to explain phenomena in the real world in light of the “principal versus agent problem”: that the principal (the person with authority) and the agent to whom he or she assigns a task cannot be perfectly aligned in their interests. The principal desires the agent to act solely for the benefit of the principal, but the agent also seeks his or her own agenda. In many situations, not only does this divergence of interests exist, but also the principal and agent might themselves be acutely aware of such a conflict of interest. Moreover, in a world in which information is murky, opaque, and asymmetrical, it may be difficult or impractical or impossible to eliminate such conflicts of interest. Similarly, it might be impractical to accurately measure the performance of the agent. In such a world both the principal and the agent will seek out or agree to various arrangements in response to this fundamental problem of non-alignment of interests and the fundamental problem of the inability to know what is really happening in order to measure performance.³

There are levels to this phenomenon: an individual or group can be either a principal or an agent depending on the situation. For example, in the relationship between shareholders and the CEO, shareholders are the principal (owners) and the CEO is the agent (manages the company). But in the relationship between a CEO and senior executives, the CEO is the principal and the senior executives are the agent. Similarly, senior managers can be principals (who delegate and assign tasks) to junior employees who then act in the role of agents (doers of tasks). Whenever someone with authority delegates tasks to someone else, the principal-agent problem arises.

3.1 Implications of Principal-Agent Theory

Given the divergence of interest between a principal and an agent, we can anticipate several types of costs that arise:

1. Cost of monitoring: the principal adopts activities to monitor the agent’s compliance with the wishes of the principal; this choice creates direct costs of monitoring but also generates

³ See Jensen [14] and Jensen [15]. An example: both an employer and an employee might agree to monitor “clocking in” and clocking out” of the workplace to monitor time on the job. It’s an imperfect measure of effort but it is something that can be measured concretely versus other less concrete measures. Similarly, an employer and an employee salesperson might agree to base compensation on sales, which can be measured concretely, despite its imperfect barometer of the efforts of the salesperson.

indirect costs as the agent responds and reacts to this monitoring.

2. Cost of imperfectly measuring outcomes: when monitoring is not practical or feasible, the principal may choose to monitor the agent's performance by measuring outcomes. But when information is opaque, the measurement will be imperfect. An example is when a principal cannot measure the effort expended by traveling sales representatives and thus chooses to compensate them via commission on sales. The sales reps who work harder and visit more clients might achieve more sales, but luck and other factors also contribute and so the correlation between diligence and sales outcomes is imperfect.
3. Cost of excessive risk-seeking: when a principal hires an agent to do an activity that involves risk-taking, there is a cost that the agent might be too risk-seeking and will take more risk than the principal desires. The agent has a flawed incentive to take on extra risk with the principal's financial resources. In such a situation, we might anticipate that the principal would impose risk limits on the agent.
4. Cost of excessive risk aversion: there is also a risk that the agent will be too risk-averse and take on less risk than the principal desires. Why would the agent take less risk than optimal? As demonstrated previously in the discussion of Bayes' Theorem, if the agent takes a risk and the outcome is unfavorable, the principal will revise his assessment of the skill of the agent, which could create a loss of position and loss of income to the agent. The agent's interest in self-preservation could incentivize the agent to create a cost for the principal by foregoing business opportunities that might otherwise have paid off for the principal. In such a situation, we might anticipate that the principal would construct inducements for taking risk. Moreover, the principal might agree to the agent's desire to spend money to provide risk hedging mechanisms to reduce the risk of adverse outcomes, thereby reducing the threat to the agent and thereby facilitating more optimal risk-taking by the agent.

We can imagine an implied dialogue between principal and agent as follows:

Principal: "Agent, I need to monitor and measure your performance."

Agent: "But any such system will be imperfect."

Principal: "Still, I need to measure you despite system imperfections."

Agent: "Then I will take certain actions as a result."

Principal: "Such as?"

Agent: "For example, I will spend money to hedge my risk on risk-taking business."

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

Principal: “But hedging your risk will be a needless cost to me that only benefits you!”

Agent: “Very well, consider this an ‘agency cost’.”

Principal: “Very well, but I will monitor and constrain your ability to spend money on risk hedging so that it is not an unlimited expenditure.”

Agent: “Very well. But keep in mind that if you are too onerous then in a competitive labor market, I can switch to another employer...”

The discussion above highlights that principal-agent costs can be both direct and indirect and can flow not only from the principal to the agent but also can flow in a continuous circle from the agent back to the principal.

3.2 Relevance of Principal-Agent Theory to Insurance and Reinsurance

The framework of principal-agent theory is relevant to property-casualty insurance because results depend upon a mixture of underwriting skill combined with random outcomes. This commingling of skill and luck means that neither the “roulette wheel” nor the “drunken sailor” is a perfect model for insurance underwriting activities. Any time there is a discrepancy in process versus outcome, the principal-agent problem will become part of the risk management landscape.

4. APPLICATIONS OF PRINCIPAL-AGENT THEORY TO INSURANCE RISK MANAGEMENT

Principal-Agent Theory allows us to understand various aspects of insurance risk management.

4.1 Risk Limits for Individual Underwriters on a Gross of Reinsurance and Net of Reinsurance Basis

Many insurers impose risk limits for individual underwriters on both a gross of reinsurance basis and net of reinsurance basis. Why? And what should the limit be? Under the theory of Enterprise Risk Management, risk is relevant only at the total company level. According to this theory, individual risk limits only make sense to the extent that they can be connected to the firm’s total risk, but this connection is often quite tenuous. Moreover, because the contribution of risk from each underwriter to the total company’s risk is so small, the maximum amount of risk that each individual underwriter could take would be much higher than what we actually observe in the real world. But if the rationale for risk limits flows from principal-agent theory, then we can understand the need for risk limits on individual underwriters: limiting risk-taking is a cost incurred by the senior managers to

reduce the possibility that an underwriter will improperly underprice risk at the company's expense. Hence the need to impose risk limits. At the same time, individual underwriters might desire to purchase excess-of-loss (XOL) reinsurance to reduce their own risk in a way that hedges risk more than what the company might optimally prefer. This is a cost imposed by the underwriters upon the company because of the underwriters' awareness of flawed measurement arising out of principal-agent dynamics. Ultimately, the direct and indirect costs of individual risk limits can be attributed to principal-agent costs.

4.2 Reinsurance for a New, Small Line of Business

Why do insurers purchase reinsurance when starting up a new book of business whose volume is quite small? From the perspective of managing the total risk of the company, the new book of business is too small to materially change the overall risk. But from the perspective of principal-agent theory, the senior managers of the firm (agent) do not want to have an adverse outcome on the new business that could influence the shareholders (principal) to doubt their talent and skill.

4.3 Parameter Risk

Parameter risk refers to a situation in which it is difficult to estimate the underlying loss cost for an insurance or reinsurance cover. But the parameter risk of an individual deal or product or line of business wouldn't likely alter the risk profile of an overall company, so why does it matter? Some thinkers have offered various explanations:

1. Statistical bias in loss pick: Marc Shamula⁴ notes that if an actuary is pricing an XOL reinsurance layer and the actuary is modeling the underlying loss severity with a single parameter pareto curve with "shape parameter" alpha, then parameter risk could introduce statistical bias. For example, if alpha could equal 2.0 or 3.0 and the actuary selects 2.5 as an average, then the layer loss cost using alpha 2.5 will underprice the layer loss cost relative to an average of layer loss cost via alpha 2.0 and 3.0 via Jensen's inequality $f(E[X]) \neq E[f(x)]$.
2. Winner's curse: Russ Wenitsky notes that parameter risk, when combined with a competitive auction process of bidding in the market, could lead to the winner's curse.⁵ Parameter risk per se does not create a problem but combining with open auction bidding leads to mis-estimation of the loss pick because the winner of the prize to write the insurance/reinsurance cover likely won the auction by underpricing.

⁴ Private communication

⁵ See Wenitsky [30]

3. Principal-agent: adopting principal-agent theory unlocks a new explanation. Parameter risk is important because it could lead to situation in which the underwriter has badly mis-priced the business and gets blamed for a bad outcome. In a framework that acknowledges the cost of the principal-agent problem, this cost needs to be incorporated into pricing.

4.4 Non-Systematic and Non-Systemic Insurance Risk Pricing, Business Unit Insurance Risk Pricing, and the Hole-in-one Insurance Problem

What should be the price of insurance risk? What type of risk matters as a driver of compensable insurance risk? Should only systematic risk that cannot be diversified in an investor's portfolio matter? What about diversifiable, idiosyncratic, non-systematic risk in an investor's portfolio that nevertheless affects the holistic downside risk of an individual insurance company? What about the risk of a small division or business unit within an insurer that does not meaningfully contribute to the all-in downside risk of the insurer?

Jon Major⁶ writes in 2009 about an in-depth actuarial discussion on the cost of risk. The catalyst was a paper taking a view based on the investment world view of risk: what matters is systematic downside risk that correlates with the overall market, but not much else. The problem is that according to this view, much of insurance, which doesn't correlate to overall financial markets, ought not require much of a loading for the cost of risk. Much of the subsequent analysis by Major focuses on whole-company franchise value to illuminate why the severe downside risk of insurance, despite being uncorrelated with financial markets, matters (and thus should be priced for). The idea of pricing for whole-company downside is widely accepted in insurance and is not controversial (for the most part). Yet Major's discussion and conceptual analysis, which briefly raises agency costs and then proceeds onward to whole-company franchise value, does not seem to find a reason to assign cost of risk for a business unit or component of an overall company. As a *gedanken* thought experiment, he discusses a hypothetical situation of an insurer who writes "golf hole-in-one" insurance. Yet in that discussion, he doesn't allude to the possibility that this hypothetical insurance operation might be just a business unit of a larger division or larger company. In such a case, downside risk to the hole-in-one insurance operation would not create risk to the franchise value of the entire company.

The missing piece is agency costs: the executives and underwriters who write the hole-in-one insurance face a principal-agent problem. Namely, if a golfer scores a hole-in-one and collects a large payout, the principal (CEO of the overall company) shines a harsh spotlight on the hole-in-one

⁶ Major [22]

insurance business unit, loses confidence in the senior underwriting team, and perhaps terminates their employment. Thus, the executives and underwriters who provide hole-in-one insurance do need to charge policyholders for the cost emanating from principal-agent problems; the same executives likely would desire to hedge risk by spending on the cost of reinsurance and would need to charge policyholders sufficient margin upfront to pay for such risk management. Tellingly, in the thought experiment, Major writes out hypothetical numbers for the hole-in-one insurance case but assumes the probability of a hole in one is knowable, known, and has no uncertainty; there is no Bayesian updating after an adverse result. In a world of uncertainty around underlying probabilities, and in a world where underwriters are supposed to be skilled at estimating these unknown probabilities, and in a world of principal-agent problems, then an adverse outcome (such as a hole-in-one) would likely have adverse outcomes for the underwriters. This principal-agent problem is likely to be a reason for significant price of risk even for downside insurance outcomes that only affect a single business unit and do not correlate with the broader insurance company or broader financial market.

4.5 Basis Risk

Basis risk arises if a company purchases a risk hedge that responds imperfectly to the company's risk: for example, purchasing reinsurance protection triggered by industry loss metrics even though the company loss metrics might align imperfectly with the industry loss. Why do insurers recoil so dramatically at low likelihoods and small amounts of basis risk? Although the risk to the overall company-wide downside is small, the risk within principal-agent theory is large. If a downside event occurs and it turns out that the basis risk within a hedge cause the insurer to not recover fully, then this mis-step raises the antennae of the shareholders (principal) towards the firm's senior management (agent) and could impose large costs to the agent.

4.6 Minimum Prices

Why do insurers impose minimum prices such as minimum price per million of increased limit? Why do reinsurers impose minimums such as minimum Rate on Line (RoL)? Some possibilities relate to fixed costs and expense ratios, but another contributor is likely principal-agent considerations. When an individual underwriter takes on the risk of a policy or reinsurance contract with a highly skewed ratio of limit to premium, it creates unwanted agency costs on the individual dealmaker. In a world of principal-agent costs, it makes sense that the individual risk-bearer would impose minimum premiums on the purchaser to cope with principal-agent costs. Similarly, in a world of principal-agent costs, it makes sense that the re(insurer) would impose minimum premium

guidelines on its own underwriters.

In summary, there are likely several candidate explanations for minimum prices:

1. Fixed expenses
2. Principal-agent costs
3. Buyer's asymmetric information: if the risk-bearer is correct in estimating, for example, that the loss cost is "1-in-a million" and therefore the premium should be very low, then why is the buyer purchasing this insurance? The desire of the buyer to purchase this (re)insurance indicates that there may be an asymmetry of information: the buyer is essentially revealing that the loss cost is much higher than just "1-in-a million". In such a situation, minimum prices might be useful.⁷ Perhaps this concept of asymmetric information is what is being alluded to in the old maxim that exhorts underwriters to ensure that diversification does not turn into "di-worse-ification."
4. "Fair share": the idea proposed is that the (re)insurer seeks to recoup some of the overall value it creates for the buyer, rather than simply price the product at the lowest possible technical breakeven price. At first blush this explanation seems quite appealing, building upon economic concepts such as *consumer surplus* and *price differentiation*. Yet this explanation falls short because it falls into the trap of *competitor neglect*: in a competitive market, one would expect to find that competitors are willing to provide insurance cover for a lower yet profitable price irrespective of "fair share".⁸

4.7 Pricing Models Based Upon Arbitrage-Free Pricing Assumptions: Desirable or Undesirable?

In a competitive market, no-arbitrage assumptions make sense and we should see arbitrage-free pricing.⁹ But in a world of principal-agent costs, where the agent can get blamed when things go wrong, we should expect the complete opposite: we should absolutely anticipate pricing that seems to "violate" principles of no-arbitrage because of the phenomenon of principal-agent costs.

⁷ One could view this situation as also relating to principal-agent: it relates to the principal-agent problem of the insurer not knowing if the individual underwriter will estimate the loss pick too optimistically and price the deal too low.

⁸ Perhaps one ought to distinguish between situations of *perfect competition* and *imperfect competition*. A situation of *imperfect competition* would undermine the critique of *competitor neglect* and reinstate the rationale of *fair share*.

⁹ Venter [29]

4.8 Packaging of Risk

Moreover, we should expect to see the emergence of activities that take granular, individual insurance risks and “package them” in order to make the principal-agent costs more manageable for counterparties such as insurers and reinsurers. These “packaging activities” can be performed either inside an insurance company by ceded reinsurance staff, outside the insurance company via intermediaries such as wholesalers, brokers and reinsurance brokers¹⁰, or outside the company via arms-length transactions¹¹. The added cost of intermediation might be outweighed by the benefit of reduced principal-agent costs.

Mathematically, we can investigate this phenomenon of combining granular risks into larger packages by noting, based on Venter[29]:

$$E[X+Y] = E[X] + E[Y] \tag{4.1}$$

$$\text{Covariance}(X+Y, Z) = \text{Covariance}(X, Z) + \text{Covariance}(Y, Z) \tag{4.2}$$

$$\sigma_{x+y} \leq \sigma_x + \sigma_y \tag{4.3}$$

In a stylized, arbitrage-free pricing world, prices should be additive: the price of a sum should equal the sum of prices. Thus, pricing based upon expected value (equation 4.1) or covariance (equation 4.2) satisfies this desirable criterion. Pricing based upon volatility (equation 4.3), in contrast, manifests the supposedly undesirable trait of sub-additivity. However, once we identify that principal-agent costs are important in re(insurance) pricing and many other real world situations, we can acknowledge that principal-agent costs ought not be additive. Thus, the sub-additivity of volatility (standard deviation) is a feature, not a bug.¹² Principal-Agent Theory, together with equation 4.3, bolsters our intuition and validates our everyday lived experience that combining granular individual risks into larger packages of risk is a valuable, unsurprising, and normative aspect

¹⁰ Reinsurance brokers might set up a “facility” through which reinsurers participate automatically or semi-automatically on various deals, thus reducing the principal-agent cost emanating from any one individual deal.

¹¹ Example of packaging via arms-length transactions: a reinsurer can be viewed as packaging risk from individual insurers before ceding risk to the insurance-linked securities (ILS) aka “cat bond” market

¹² Venter [29] similarly notes that Utility Theory shows the same phenomenon of non-additivity and thus concludes its non-desirability in a no-arbitrage pricing environment; our discussion shows why non-additivity of Utility Theory is, contra Venter [29], not disqualifying but actually desirable.

of the insurance risk management landscape.

4.9 Cat Bonds Versus Traditional Reinsurance

Both cat bonds and traditional reinsurance cover the risk of property losses arising from natural catastrophes, yet sometimes the price of cat bonds are cheaper than equivalent traditional reinsurance. Why? Historically, one explanation has been to point to the different types of investors in each instrument: to the extent that large, more diversified pools of investment portfolios such as pension funds could support cat bonds, goes the theory, then that more diversified portfolio could support taking cat risk at a lower price. While this fact may be true, it overlooks another critical difference between cat bonds and reinsurance: principal-agent costs. Investors who supply equity capital to insurance/reinsurance company executives must worry about what the executives will do with the money: will they use the funds appropriately or not? This concern causes significant agency costs when investors supply equity capital to company executives to support writing cover for natural catastrophes. But when the capital is supplied in the format of a cat bond, agency costs are minimized because the funds cannot be used for any purpose other than supporting the cat bond. In the parlance of behavioral economics: by having the writers of nat cat cover choose to forfeit the flexibility of equity capital to do anything they please, and instead constrain themselves in the “commitment device” or “Odysseus contract” of cat bond capital which can be used for no other purpose other than supporting the nat cat bond, they have reduced agency costs.¹³ Thus, one reason why cat bonds can sometimes be cheaper than traditional reinsurance cover is that they generate lower principal-agent costs.

4.10 The Curious Case of “Private Cat”

A curious phenomenon occurs when executives in the insurance and reinsurance market care about downside risk for their company, but instead of measuring risk relative to balance sheet, earnings, or other risk metrics that relate to the company’s financial resources, they focus on the risk of an adverse outcome *relative to peer companies*. An adverse outcome relative to peers is often called a private catastrophe or “private cat”. In any classical risk management framework, this view relative to peers makes no sense; after all, if a loss is small relative to balance sheet or earnings why should it matter? And yet in a principal-agent framework, outcomes relative to peers are even more crucial than absolute risk because an outcome worse than peers could cause a principal to revise its views of the competence of the agent. In contrast, if the company suffers a loss but is in line with

¹³ Odysseus contract named after the story of Odysseus/Ulysses tying himself to the mast of his ship in Homer’s *The Odyssey*

the loss of peers then the pain of such a loss is partially diminished.

Turning the capital asset pricing model (CAPM) and Markowitz¹⁴ on their heads, principal-agent theory indicates that systematic risk that correlates with the broader market is actually desirable (to some extent) whereas idiosyncratic risk that does not correlate with the market is undesirable (to some extent). We might say that “being uncorrelated with the broader market is favorable for risk cost, but unfavorable for principal-agent cost” and that “being correlated with the broader market is unfavorable for risk cost, yet favorable for principal-agent cost”.

4.11 Narrative: Why Risk is More Than Just a Quantum

In a traditional insurance risk framework, only the quantitative attributes of risk should matter, not the “backstory” of where the risk came from. And yet, sometimes the risk aversion and cost of risk associated with Casualty business is different than a numerically equivalent amount of risk emanating from Property (*ceteris paribus*; all else equal; setting aside time duration effects and other differences). Why? In a principal-agent framework, we identify the fraught relationship between principal and agent as a key source of risk and a key source of cost. If adverse outcomes are more easily “explained away” in one arena versus another, then that business will have lower agency costs. Indeed, in property catastrophe, we can say that agency costs are lower than otherwise would be because after any adverse outcome the agent can point to a clear geo-physical phenomenon and say, “You can’t blame me; the wind blew, we had a hurricane, and we suffered a financial loss!”. In contrast, when adverse outcomes arise from Casualty business, there is no clear natural phenomenon to blame; all the blame adheres to the underwriter who wrote the insured business. Thus, because of the principal-agent problem, the “storyline” or “narrative” associated with downside risk matters, not just the numerical quantity of risk.

Could a similar dynamic apply not only to distinguishing Casualty risk from Property Risk, but to distinguishing Reserve Risk from Underwriting Risk? It’s possible that the pain of principal-agent costs could be higher for reserve risk than underwriting risk, in the sense that a “miss” on reserves could trigger more severe loss of confidence than a miss on underwriting. Counterintuitively, this could imply that some categories of risk loads, for example principal-agent costs, could increase over the development pattern of an insurance contract rather than decrease monotonically from policy inception to ultimate maturity.

¹⁴ Bodie, Kane, Marcus [2] *Investments*

4.12 On the Optimal Corporate Structure of Insurers and Reinsurers

Acknowledging the reality of principal-agent costs also illuminates some advantages and disadvantages of various forms of corporate structure that are adopted by insurers and reinsurers. If an insurer is publicly-traded in the stock market, then the principal-agent problem is exacerbated: the separation of ownership versus control means that senior executives manage the business whereas the equity shareholders, who are diffuse and removed from day-to-day operations, retain ownership. What happens if the company has bad results – is this adverse outcome attributable to bad luck or to bad skill? The owners cannot always discern if the managers are doing a good job because the information is opaque. The risk of adverse outcomes matters not only because of a threat to capital and not only because of a threat to the franchise value of the firm and not only because volatility of earnings but also – maybe even primarily – because adverse outcomes threaten to alter the principal’s (shareholders’) perception of the skill of the agent (CEO and senior management).

In contrast, an alternative corporate form might bring a more favorable situation regarding the principal-agent problem. For example, if an insurer or reinsurer were housed within a larger conglomerate¹⁵ or family business¹⁶ then these corporate arrangements might generate significant improvements in reducing the costs of the principal versus agent problem.

Any given form of corporate structure has many different advantages and disadvantages; given the ambiguity in quantifying the exact numerical benefit of reduced principal-agent costs, it would be difficult to definitively say that one form of corporate structure is superior to another. Given this uncertainty in tabulating costs versus benefits, it is no surprise to see various insurers and reinsurers taking on all sorts of different corporate forms. No single option is a dominant choice, and the cost of the principal-agent problem is a major contributor to this ambivalence.

What would happen if one could observe two “twin” insurers with the exact same risk profile and the exact same financial resources and the exact same management, with the only difference that one insurer is owned by shareholders in the publicly-traded stock market whereas the second insurer is owned within a family business or other conglomerate? If the analytical framework developed above is correct, then it would predict that *ceteris paribus*, all else equal, these two “twin” insurers would have different risk tolerances: it is likely that the publicly traded insurer would have lower risk tolerance and thus purchase more reinsurance than the equivalent “twin” insurer housed within a private family business or conglomerate. This surprising implication is predicted by exactly zero of the existing known risk frameworks such as ERM or CAPM or earnings volatility; only principal-

¹⁵ For example, Berkshire Hathaway

¹⁶ For example, Partner Re housed within the Exor conglomerate controlled by the Agnelli family

agent theory would explain such an observation.

4.13 The Purpose of Publishing Risk Metrics: To Transmit Information to the Audience or to Inoculate the Speaker?

Senior insurance executives often disclose risk metrics to key stakeholders. For example, a company might use its annual report or investor day presentation to indicate the numerical value of its prospective “1-in-250 year” return period loss from Windstorm, Earthquake, Wildfire, Cyber Event, or Casualty Catastrophe Scenario. Why? According to a more traditional framework, the disclosure intends to inform: to convey useful information to investors so that they understand the potential risks and capital calls from future possible events. When viewed through the proposed prism of principal-agent theory, however, these communications take on an additional element: to inoculate the speaker. The more that management discloses potential future downside risks to the shareholders ahead of time, the more that the senior managers can defend themselves after an adverse event against accusations of mismanagement. Thus, the agent anticipates the future scrutiny of the principal and proactively lays the communication groundwork *ex ante* in order to enhance the agent’s position *ex post*. The implication is that senior executives might be able to achieve significant benefit for themselves (not just the company) by investing in broader, deeper and more granular risk disclosures in the future.

4.14 Enterprise Risk Management (ERM) versus Group Risk Management (GRM)

Enterprise Risk Management (ERM) is a popular and common framework for managing insurance risk. One of the central ideas of ERM is that the company is an integrated “enterprise” and that various risks should be measured on the basis of the company *in toto*. But acknowledging the framework of principal-agent theory is to acknowledge that within any company there are different business units who all have their own agendas, subject to monitoring by home office. The entire company is not really an integrated holistic enterprise but rather a Group: a collection, a confederation. Group (or *Gesellschaft*) Risk Management (GRM) is a modern risk management framework alternative to ERM that recognizes principal-agent costs and better aligns with the real world.¹⁷

¹⁷ Wilson [31] differentiates between a “syndicate,” a loose collection, and a “team,” an integrated whole.

4.15 Size, Scale, and Scope of the Company: Economies of Scale Versus Diseconomies of Scale

In a traditional insurance risk framework, each business unit in a company has some idiosyncratic risk that can be partially diversified away as part of a broader, larger, more diversified insurance company. In this framework, larger, broader, more diversified conglomerate insurers should be preferred to smaller specialist companies. But in light of principal-agent theory, we can see that having a large, broad, diversified conglomerate generates significant agency costs to home office of monitoring all the various business units within the company. Meanwhile, a smaller company that specializes in a more targeted range of business will have lower principal-agent costs, albeit unable to diversify risk as well as a larger conglomerate. Thus, the benefits of scale and diversification conflict with the costs of the principal-agent problem; the optimal company size is murky. As a result, it is no surprise that no single strategy dominates the market as companies scale-up, scale down, evolve, and grapple with the economies of scale and the diseconomies of scale.¹⁸

4.16 Summary of Key Differences Between Portfolio Theory and Principal-Agent Theory

The discussion above highlighted some key differences between a traditional risk management framework rooted in Portfolio Theory versus the proposed framework rooted in Principal-Agent Theory. We can summarize some of these difference in the table below:

	Portfolio Theory	Principal-Agent Theory
Central theme	Rate of Return on Capital	Asymmetric Information
Performance Measurement	Relative to Capital	Relative to Peers
Risk Management	ERM	GRM
Risk Measurement Scope	Whole Company	Whole Company and also Individual Business Units
Risk Metrics	Probability of Shareholder Ruin Probability of large loss relative to balance sheet Probability that earnings are highly volatile Loss of Firm's Franchise Value	Probability of Executive Ruin Probability of large loss relative to peers Probability that earnings fall short of business plan Loss of Executives' Reputations and Careers

¹⁸ Coase [8] grapples with the question of the size of the firm and identifies “transaction costs” as the rationale for differentiating between market-based activities outside the firm versus non-market-based activities inside the firm; Cowen [9] pushes back against this orthodox view and instead thinks of a firm as “a carrier of reputation”.

4.17 Blending Together Price Loadings for Risk Cost and Principal-Agent Cost

The discussion above highlights the importance of including principal-agent costs in any insurance pricing or insurance risk measurement situation, but the intent is not to exclude all other sources of risk and cost. Rather, the intent is to broaden the discussion beyond the traditional portfolio theory view of risk to also include principal-agent considerations. Ultimately the final view of risk and cost will include elements of both; some insurance pricing situations will have low Risk Costs yet have high Principal-Agent Costs, and vice-versa.¹⁹ Some suggested **guesses** can be found in the following table that attempts to directionally map out the landscape of lines of business in a reinsurance pricing context:

	Risk Cost: Low	Risk Cost: Medium	Risk Cost: High
Principal-Agent Cost: Low	Personal Auto	Workers' Compensation	Florida Hurricane
Principal-Agent Cost: Medium	Commercial Property	Casualty	California Earthquake
Principal-Agent Cost: High	Personal Umbrella	Excess Casualty	Cyber

We can acknowledge that the total risk load for insurance risk ought to blend both traditional risk cost and principal-agent cost. We can say:

$$\text{Premium} = \text{Loss Cost} + \text{Risk Load} + \text{Expenses} \quad (4.4)$$

$$\text{Risk Load} = \text{Risk Cost} + \text{Frictional Cost} \quad (4.5)$$

$$\text{Frictional Cost} = \text{Principal Agent Cost} + \text{Other Frictional Costs (e.g., tax effects, etc.)} \quad (4.6)$$

¹⁹ Pricing for “risk cost” often focuses on contribution to holistic downside. One approach is to allocate the company’s cost of capital; this process can be done via various methods, among them Capital Allocation by Percentile Layer (see Bodoff [3]). Others recommend alternative frameworks for calculating risk cost; this topic is discussed extensively in the literature and is outside the scope of the current discussion. Here we focus on principal agent frictional cost.

$$\text{Premium} = \text{Loss Cost} + \text{Risk Cost} + \text{Principal Agent Frictional Cost} + \text{Other Frictional Costs} \quad (4.7) \\ (\text{e.g., tax effects, etc.}) + \text{Expenses}$$

Equations 4.5 and 4.6 and 4.7 help us to understand the observed pricing of some lines of business. Some lines of business (for example: Personal Umbrella insurance, Hole-in-one insurance, etc.) tend not to contribute too much risk to the overall company's risk; the cost of risk is quite low, whether measured as allocated cost of downside capital or whether measured as the line's contribution to the company's earnings volatility. Rather, the observed price loading for such types of insurance likely derive significantly from principal-agent costs.

5. QUANTITATIVE MODELS AND ACTIONABLE DECISION CRITERIA

Despite the evidence marshalled above that principal-agent theory supplements, supersedes, and supplants the traditional view rooted solely in portfolio theory; that it is more nuanced and more accurate than portfolio theory; and that it is more relevant to insurance risk management than portfolio theory, these attributes alone are not sufficient for principal-agent theory to prevail.

Why? Because the usefulness and widespread adoption of a theory is only partially related to the accuracy of its depiction of the real world. Crucial is its ability to provide clear, actionable guidance for decision making.

Scholars in a subject sometimes hold onto a model despite its flaws if the model is sufficiently "useful." One example arises in Economics. Some of the drawbacks of the traditional model of "rational human" were known but were elided because the underlying model was so "useful" for economists to write down equations that could be solved for optimization. Only later, when the critiques of Behavioral Economics became too problematic to ignore, did economists finally acknowledge the limitations of the traditional model.²⁰

The widespread adoption of a model in a field of study despite being flawed is further entrenched because of the phenomenon of "intersubjectivity".²¹ Neither objective nor subjective, an idea is intersubjective if it derives its potency from the beliefs of many individuals in a large group; in such a case, an individual is constrained by and cannot overcome or dislodge the beliefs of the group.

²⁰ Kahneman [16], Gilbert [12], and Ariely [1]

²¹ Yuval Noah Harari [13]

Conclusion: an idea can be inaccurate and yet become entrenched because it is “useful,” and especially so if this useful model, then becomes the *lingua franca* of how individuals interact and transact.

In the case of insurance risk, the framework of portfolio theory, while incomplete, has a major advantage: its numerical output provides clear decision-making guidance. For example, in Kreps’s framework of “investment-equivalent reinsurance pricing,” the risk margin for a reinsurance contract is evaluated in the language of investments. If the “yield” or “return on capital” on insurance and/or reinsurance risk taking exceeds a target hurdle rate of return then the risk bearer should take on the risk, but if the rate of return is too low then the risk bearer should decline. Because this framework focuses on rate of return on downside capital and/or volatility of earnings, it does not consider any pricing mechanisms related to principal-agent dynamics. The model skips over principal-agent costs and yet the model is still very useful and very widely used because it provides a clear “go versus no-go” criterion for risk pricing and risk taking.²²

Thus, if principal-agent theory is to take hold, it likely will require some quantitative framework for estimating agency costs and for evaluating insurance risk pricing. For example, if two different situations of insurance risk-taking display more versus less principal-agent cost, how will we know how much more or less to charge for (re)insurance in such a situation? How will a purchaser of reinsurance evaluate whether the benefit of the reinsurance cover is worth the price? Principal-agent theory might be accurate about the reality of our world, but we will likely need to begin exploring and developing quantitative models of this framework if we want it to become part of the conversation.

5.1 Principal-Agent Costs, Actionable Decision Criteria, and Quantitative Models

Developing a framework for evaluating principal-agent costs is a challenge that continues beyond this paper. As a first draft, I present one possible set of examples of how to move forward.

5.1.1 Insurance Risk Pricing and Risk Hedging as a Form of Compensation

The insight that various insurance and reinsurance structures, transactions, price mechanisms, and strategic moves all relate to principal-agent costs suggests a new prism through which to refract these costs: the prism of labor costs. More specifically, we can think of principal-agent costs as a form of executive compensation.

²² Kreps [19] and [18]

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

This approach would be a significant change in perception from the current framework. In the current traditional view, the risk margin that is embedded in insurance risk pricing relates to the financial risk for the company, as measured by downside risk to capital or volatility of earnings. But in this new framework, the cost of risk would also encompass agency costs: the cost of risk to the underwriters and executives agreeing to the deal and taking on the potential reputational hit from an adverse outcome.

In one way, this approach simply resuscitates the framework of Utility Theory: for underwriters and executives, taking on additional insurance risk or reducing insurance risk is the equivalent of a decrease or increase in their utility, which is a form of compensation. As such, the price of risk can be estimated using concepts from Utility Theory such as utility functions, indifference curves, or “pain functions”.²³

Yet there are two differences now being proposed. First, although the actuarial literature has references to utility, aversion, preferences, and pain, these are usually invoked at the total company level. Here, we advance the novel concept of a new plane of risk pricing: the utility preferences of the underwriters and executives as distinct and non-identical to those of the company.²⁴ This new way of thinking about the cost of risk has not appeared meaningfully in the actuarial literature for pricing insurance and reinsurance; it represents a new mindset and new framework.²⁵

The second novel idea we propose is that principal-agent theory provides justification for **why** the risk and disutility to underwriters and executives ought to be compensated. If we could guarantee that the underwriters and executives taking on the insurance risk would be viewed as a “roulette wheel” then there plausibly would be no principal-agent costs, only unsystematic, diversifiable randomness of outcomes that wouldn’t necessarily be compensable. Yet because of the principal-agent problem, all counterparties foresee that adverse financial outcomes will likely not be consequence-free. Bad outcomes will cause investors, boards of directors, CEOs, and other principals to view the business unit executives as “drunken sailors” and cause them loss of compensation; this risk of loss is a principal-agent cost and ought to be compensable.

If one adopts this conceptual framework, then it could provide guidance for evaluating decisions in this domain. Namely, if the costs are viewed as labor compensation costs, then one could make decisions about transactions and pricing in light of labor compensation costs being acceptably low

²³ Kreps [20]

²⁴ Ross [26] is a source in the economics literature of dis-entangling the two separate utility functions of the principal and the agent.

²⁵ Major [22] refers to principal-agent costs as “mysterious” and “psychological”.

or unacceptably high; the actionable decision-making criterion on risk pricing would relate to principal-agent cost as a component of overall labor compensation costs.

5.1.2 Purchasing Reinsurance to Hedge Risk for a Business Unit of a Larger Organization

Let's return to golf hole-in-one insurance. Downside risk is not correlated to the overall financial markets and thus based on portfolio theory and CAPM the risk charge should be zero or minimal. Jon Major [22] suggests that nevertheless there is a rationale for such a company to incorporate risk pricing for the non-systematic risk of hole-in-one insurance: namely, downside risk that is specific to this company could threaten the overall firm and thus reduce Franchise Value.

But what if the operation providing hole-in-one insurance is not a standalone company but rather is a division or business unit of a larger insurance company? The downside risk of the hole-in-one division washes out in the broader company. Now, not only is there no systematic risk to justify risk load but there is also no threat to this company's franchise value. So then on what basis would there be justification for adding risk load to the pricing of this insurance product? The second side of the same coin would be: can a business unit that is a division of a larger insurer choose to spend money on reducing risk via reinsurance? This is more than just a hypothetical example because in the real world, business units and divisions of larger companies do indeed seek to hedge risk by purchasing reinsurance; if purchasing the reinsurance costs money, and there is benefit in reducing the downside risk of the business unit but there is no concomitant benefit to reducing the risk of the overall company, then in what sense can we justify the transaction if there is cost but no company-wide benefit? Indeed, from the point of view of any traditional framework the validity of the purchase falls short of the decision criterion: the benefit does not exceed the required rate of return. But from an agency cost perspective, the framework is quite different. If we view the expenditure as an agency cost, then the suitable prism through which to evaluate the situation might be executive compensation. The purchase reduces risk at the business unit level, which creates a utility benefit primarily to the unit's executives but creates cost and no benefit to the overall company: that is compensation cost.

5.1.3 Towards a Mathematical Model of Principal-Agent Costs

How should we think about the benefits and costs of reinsurance for a business unit that desires to hedge risk by purchasing reinsurance? According to Major [22], agency costs are "mysterious" and "psychological," but if we want to gain traction with principal-agent theory, then we ought to develop a mathematical framework for applying it. What are the costs and benefits of purchasing reinsurance within the framework of principal-agent theory? How should we evaluate if the cost is

worth the benefit? At what price is the act of hedging risk worthwhile or not?

We can propose the following set of quantitative approaches:

1. Framework #1: Volatility
2. Framework #2: Utility
3. Framework #3: Sustainability of Earnings (Franchise Value)

We note that these three concepts generally appear in the literature, but with a focus on the pricing of risk cost at the total portfolio risk level; now, for the first time, we adopt these frameworks at the individual underwriter and/or individual business unit level in order to capture the pricing of principal agent costs.

5.1.4 Framework #1: Volatility

One approach would be to acknowledge that principal agent costs often emerge from the variability of outcomes of individual standalone deals, which suggests the possible usefulness of setting principal agent costs as a function of volatility (standard deviation):

$$\text{Principal Agent Frictional Cost} = k\% * \sigma_{\text{standalone}} \tag{5.1}$$

One advantage of such an approach is that it's relatively straightforward to adopt. Another advantage, serendipitously, is that some market participants already use such an approach.²⁶ Another advantage is that the formula responds well to directional changes in standalone volatility: all else equal, if the standalone volatility is larger, then the loading for principal agent frictional cost ought to be larger.²⁷

The framework of equation 5.1 could help defend the usefulness of loss-sensitive features in reinsurance contracts covering non-systemic lines of business; some examples of such loss-sensitive features are swing-rating, sliding scale commissions, profit commissions, and paid reinstatements. To the extent that these loss-sensitive features reduce the standalone volatility of the profit-and-loss distribution of a reinsurance treaty, they might have a large effect on principal-agent cost, despite

²⁶ Some market participants calculate total risk load on the basis of standalone volatility; this approach is incorrect for "risk cost" but here we acknowledge that this approach could be resuscitated for use in "principal agent frictional cost". See equation 4.7 for reference.

²⁷ A further refinement to equation 5.1 could be to bifurcate standalone volatility into two buckets, parameter risk and process risk, and to apply different loadings to each.

likely having only a small effect on the total downside risk of the company.

One major disadvantage of Equation 5.1 is that it does not give guidance about what the parameter k ought to be; in other words, it does not provide a framework for evaluating if the price loading is too low or too high, thus limiting its usefulness as an actionable decision criterion. One solution to this problem would be to invoke the conceptual framework developed in section 5.1.1: labor compensation costs. One might be able to use market-based peer benchmarking of compensation costs to approximate whether the price loading of equation 5.1 is too high or too low, whether from the perspective of the buyer of protection or the seller of protection.

5.1.5 Framework #2: Utility

A second approach could be to adopt the framework of Utility Theory. Specifically, what is the premium that is required to make the transaction worthwhile to the individual underwriter taking on the risk? Following Venter [29], we seek to determine the *certainty equivalent* premium such that the utility function $u(\cdot)$ of not accepting the deal is equivalent to accepting the deal, thus making the individual underwriter *indifferent*.²⁸

$$u(\text{wealth}) = E[u(\text{wealth} + \text{premium} - \text{losses})] \quad (5.2)$$

A similar yet more tractable approach in this framework would be to adopt the “pain function” proposed by Kreps [20] and [21]. Similar to Utility Theory, the idea is to flesh out all the various possible outcomes and to assign value and/or pain to each of them. Because extreme bad outcomes are far more painful *to the individual underwriter* than the gain of small, good outcomes, an individual underwriter would prefer not to take on an individual deal unless the price loading for principal agent frictional costs were sufficient to make the deal worth writing. A similar utility analysis applies to the buyer as well.

One advantage of this framework is that it provides an actionable decision-making criterion, at least from the perspective of the individual buyer and seller, for whether or not the compensation for principal agent costs is sufficiently attractive to make it worthwhile to them to do the deal; simultaneously, from the perspective of the broader (re)insurance companies employing the buyer and seller, they would still need to evaluate and benchmark the compensation costs to evaluate the reasonability of the deal’s pricing and the overall expenditure.

²⁸ Venter [29] and others discuss various choices for utility functions, such as Exponential, among others.

A disadvantage of such an approach is the difficulty in estimating the Utility Value or Pain Function of different outcomes.

5.1.6 Framework #3: Sustainability of Earnings

A third approach that we explore derives from Feldblum [11], Panning [24], and Bodoff [5], among others, and is sometimes described as Lifetime Value, Franchise Value, or Sustainability of Earnings.

This approach differs from the Utility Theory approach of section 5.1.5 in three ways:

1. Adopts a multi-period perspective rather than a single period perspective.
2. Focuses more on extreme, terminal bad outcomes rather than general variability.
3. Provides more concrete quantitative guidance on the strategic choice of the value of buying or selling protection rather than the more fungible guidance of an arbitrary utility function.

The framework of “sustainability of earnings” focuses on strategic choices in the here and now, how they alter the sustainability of future earnings and, as a result, how they affect the current **value** of future earnings.

Briefly, we can introduce a framework in which:

E = annual expected earnings

Value = current value of future earnings discounted to reflect the time value of money, the general risky variability of future earnings, and the sustainability risk of a termination event that prevents survival into the future to collect future earnings.

As a general proposition, Bodoff [5], following Panning [24], shows that an activity that increases earnings but decreases sustainability could in fact lower the current value of future earnings; conversely, a choice that decreases earnings but increases sustainability could in fact increase the current value of future earnings.²⁹ More specifically, in our context of (re)insurance risk and principal agent costs, we might view an underwriter who takes on the risk on a specific deal as someone who increases earnings but also increases the risk of a “termination event”; the pricing embedded in the deal must compensate for principal agent cost in such a way that this tradeoff between earnings and termination is sufficiently attractive for the underwriter taking on the risk.

²⁹ Although such a framework is most directly associated with Warren Buffet’s metaphor of a “moat” protecting future earnings, such a framework might also hypothetically be useful for addressing strategic choices relating to a multitude of forces affecting “sustainability”: climate change, supply chain fragility, pandemic preparedness, terrorism, cyber attack, nuclear accident, etc.

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

Conversely, the purchaser of risk protection could be viewed as decreasing earnings but also decreasing the risk of a “termination event”; the pricing embedded in the deal must compensate for principal agent cost in such a way that this tradeoff between earnings and termination is sufficiently attractive for the purchaser who is buying protection.

Now let’s return to the case of a business unit that desires to hedge risk by purchasing reinsurance.

First, we can view the risk position of the business unit as a standalone unit:

Exhibit 1

	Business Unit View Without Expenditure on Business Unit Reinsurance	Business Unit View With Expenditure on Business Unit Reinsurance
(1) E = expected earnings	100	90
(2) r = discount rate	10.0%	10.0%
(3) Discount Factor = 1/(1+r)	90.9%	90.9%
(4) g = growth rate	4.0%	4.0%
(5) Growth Factor = 1+g	1.04	1.04
(6) p = annual survival probability	96.0%	99.0%
(7) 1-p	4.0%	1.0%
(8) Value (V)	944.9	1,265.6
(9) P/E multiple	9.4	14.1

Notes

$$(8) = (1) * \{ (6) * (3) \} / \{ 1 - (5) * (6) * (3) \}$$

$$(9) = (8) / (1)$$

The formula (8) for Value (V) derives from the formula for an infinite geometric series, as shown by Panning [24]. Namely:

$$V = E * DF * p + E * GF * DF^2 * p^2 + E * GF^2 * DF^3 * p^3 + \dots \tag{5.3}$$

Multiply both sides of equation 5.2 by GF/E and let $r = GF * DF * p$, then solving for an infinite geometric series:

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

$$GF / E * V = r + r^2 + r^3 + \dots = r/(1-r) \quad (5.4)$$

Then replacing r yields:

$$V = E * DF * p / (1 - GF * DF * p) \quad (5.5)$$

Returning to Exhibit 1: for the business unit, reducing risk (for example by purchasing reinsurance) decreases earnings, yet this strategic move nevertheless **increases** the value of the business; reducing the risk of business failure increases the sustainability of earnings, increases the Price-to-Earnings multiple, and increases the franchise value of the business unit. Thus, reducing risk by purchasing reinsurance will lead to an increase in value, suggesting that the purchase should be approved.

However, this business unit is not actually a standalone unit but is part of a much larger whole; redoing the analysis at the broader total company level produces the following:

Exhibit 2

	Total Company View Without Expenditure on Business Unit Reinsurance	Total Company View With Expenditure on Business Unit Reinsurance
(1) E = expected earnings	1,000	990
(2) r = discount rate	10.0%	10.0%
(3) Discount Factor = 1/(1+r)	90.9%	90.9%
(4) g = growth rate	4.0%	4.0%
(5) Growth Factor = 1+g	1.04	1.04
(6) p = annual survival probability	99.00%	99.01%
(7) 1-p	1.000%	0.990%
(8) Value	14,063	13,944
(9) P/E multiple	14.063	14.085

Notes

$$(8) = (1) * \{ (6) * (3) \} / \{ 1 - (5) * (6) * (3) \}$$

$$(9) = (8) / (1)$$

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

By adopting the view of the total company, we see that reducing risk by purchasing reinsurance for a business unit decreases overall earnings and **decreases** the value of the company. After all, reducing risk at the business unit level does not lead (in this numerical example) to a meaningful reduction of risk at the company level, leading to a decrease in franchise value. Reducing risk by purchasing reinsurance at the business unit level decreases value and suggests that the purchase should be disapproved.

What about the perspective of the business unit executives? Let's simplify as follows:

1. Let's arbitrarily assume that executive compensation is roughly 10% of the business unit's earnings. This is a stylized choice designed for ease of discussion.
2. Moreover, if a severe adverse outcome were to occur, let's assume that the company would terminate the executives and that these executives would earn zero future earnings. Again, this is a stylized choice designed for ease of discussion.

In essence, the principal-agent problem creates a probability that the agents (senior executives of the business unit) who preside over an adverse outcome will be terminated and will lose out on future compensation. This creates an undesirable risk to the business unit executives who will then seek to purchase reinsurance to reduce their risk.

Let's examine the view of the executives using the same framework as above:

Exhibit 3

	Business Unit Executives' View Without Expenditure on Business Unit Reinsurance	Business Unit Executives' View With Expenditure on Business Unit Reinsurance
(1) E = expected compensation	10.0	9.0
(2) r = discount rate	10.0%	10.0%
(3) Discount Factor = $1/(1+r)$	90.9%	90.9%
(4) g = growth rate	4.0%	4.0%
(5) Growth Factor = $1+g$	1.04	1.04
(6) p = annual survival probability	96.0%	99.0%
(7) $1-p$	4.0%	1.0%

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

(8) Value (aka "wealth" aka "utility") | 94 | 127 |

Notes

$$(8) = (1) * \{ (6) * (3) \} / \{ 1 - (5) * (6) * (3) \}$$

For the executives, reducing risk by purchasing reinsurance reduces annual compensation but increases wealth by increasing the probability of surviving into the future; for them, purchasing reinsurance increases their welfare and thus it is “rational” for them to desire to do so. Exhibit 3 provides the quantitative analysis that indicates that the purchase, despite giving up income, is worthwhile for the executives; but if the price were higher, then this conclusion would at some point no longer be true.³⁰ The framework as shown in Exhibit 3 thus provides an actionable decision making criterion, at least for the executives, to evaluate whether or not to purchase the reinsurance.

The risk of the business unit is small relative to the larger company and is not highly correlated with the outcomes of the overall firm; the cost of purchasing reinsurance essentially provides zero or very limited benefit to the overall company.³¹ And yet, the company might prefer to allow the business unit executives to move forward with hedging the business unit’s risk even at significant cost: it increases the “utility” or “wealth” of the executives and thus is a form of executive compensation, and the firm wants to retain its business unit executives in a competitive labor market.

The actionable decision criterion about whether the company should allow the business unit executives to purchase this reinsurance for non-systematic, non-company-wide risk would relate to whether the cost makes sense within the context of compensation costs, unrelated to whether the purchase sufficiently improved the company’s overall downside risk or achieved a hurdle rate of return against downside risk capital.

The discussion above focuses specifically on a situation of a business unit seeking to spend money on a reinsurance transaction that reduces risk for the business unit but does not meaningfully

³⁰ A similar “mirror-image” analysis would need to be done by the reinsurance underwriters: does the price of the deal include sufficient price loading to compensate them for increasing their own principal-agent risk? If and when there’s no price that can address the principal agent costs of both the reinsurance purchaser and the reinsurance seller, the deal will fail to clear the market.

³¹ The temptation in such a situation, in which purchasing risk protection for the individual business unit creates value for the business unit but is a cost for the overall company, would be to implement “internal reinsurance”. The drawback to this approach is that the business unit executives intuit that when bad outcomes remain within the overall company, rather than being ceding externally, the overall company lacks a strong enough “commitment device” to confidently restrain the overall company from exacting negative consequences upon the business unit executives. For this reason, “internal reinsurance” has severe drawbacks in relation to principal agent frictional costs.

reduce risk for the overall company. Yet a similar analytical framework rooted in principal agent frictional costs could be adopted in the inverse case and in the converse case, among others:

1. If the business unit wanted to cede less reinsurance and thus increase its risk, with virtually zero impact on the overall company, it would seek to demonstrate that the added income for the business unit sufficiently offset the additional principal agent frictional costs of retaining more risk (i.e., increasing the exposure of business unit's executives to a termination event).
2. Similarly, if a reinsurance company's business unit wanted to take on the risk of writing a new treaty, then that unit would need to load the pricing sufficiently for the principal agent frictional costs involved, even in a situation when writing such a deal would create zero risk to the overall reinsurance company. Although the customer would lobby for a low price because of low systemic risk, and although the overall reinsurance company could theoretically prefer a lower technical price to grow market share for such non-systemic risk, the reinsurance business unit executives' perspective would be that they need to obtain sufficient premium to compensate for the specific principal agent costs that accrue to them.

The discussion above indicates that whether we are analyzing buying or selling of insurance or reinsurance, one ought to incorporate principal agent frictional costs into one's pricing methodology and decision analysis.

In summary, this section has shown three quantitative approaches to estimating principal agent frictional costs: Volatility, Utility, and Sustainability of Earnings. There are levels to this: principal agent frictional costs and associated pricing and commensurate decision making ought to be evaluated co-instantaneously on both the individual underwriter level and also the total company level. The proposed quantitative frameworks are initial suggestions for how to quantify principal-agent frictional costs; further research is needed in this area.

5.1.7 Strategies to Reduce the Frictions of Principal Agent Frictional Cost

The discussion above stresses the importance of acknowledging and pricing for the realities of principal agent frictional cost. But "frictional cost," especially in pricing, is seldom a good thing. In a competitive market, participants should seek out strategies to reduce frictional cost. What strategies and techniques can be used by competitive market participants to reduce principal agent frictional cost?

1. Packaging: combining granular individual risks into larger buckets.
2. Adding ballast: for non-proportional reinsurance structures, attaching less "out of the

money” than strictly needed. This would mean that an Excess of Loss (XOL) treaty or an Aggregate Stop Loss (ASL) or Adverse Development Cover (ADC) might have an attachment point closer to the mean loss than initially viewed as optimal. Although doing so risks adding cost via unneeded “dollar trading,” doing so adds “ballast” to the contract and reduces principal agent frictional cost; the benefit often outweighs the cost.

3. Introducing loss-sensitive features into the reinsurance contract.
4. Supplementing the underwriter’s judgment by adopting external sources of guidance to estimate likelihood of claim (vendor cat models³², actuarial pricing tools, computerized AI algorithms, etc.). Adopting robust, dispassionate, partially automated systems for evaluating risk could reduce the sting of principal agent costs that occur for the underwriter when outcomes are unfavorable.
5. Adopting linear, rather than non-linear, compensation incentives. An example of a non-linear incentive would be an “all or nothing” approach to “making plan” or “hitting target.” Because missing target profit by 1% could lead to a loss of executive compensation much larger than 1%, this arrangement tends to heighten, not reduce, principal agent cost. Another example of non-linear incentive structure occurs in the inverse case, when exceeding profit targets does not lead to a linear increase in executive compensation; by suppressing the upside, such an arrangement reduces the overall incentive for executives to take on deals that have idiosyncratic, non-systemic risk. These two examples of limiting upside compensation and amplifying downside compensation both tend to increase principal agent frictional cost. The opposite approach of a more linear incentive compensation structure tends to reduce principal agent frictional cost.
6. Innovating new approaches to organizational structure and performance measurement. Those companies that innovate new approaches to organizational structure in such a way as to reduce principal agent frictional cost could have an advantage in the marketplace. One example would be to experiment with different approaches to smoothing individual underwriters’ performance, such as smoothing across multiple time periods.

Any strategic move to reduce principal agent frictional cost also runs the risk of increasing other costs; one should consider a broader analysis of cost versus benefit before adopting any of the

³² Some risks are more reliably modelled than others, suggesting that principal agent frictional costs embedded in market pricing could vary across different modelled risks. See Bodoff [4] for application to natural catastrophe bonds and also to the credit market’s “credit spread puzzle.”

strategic moves listed above. In a competitive market, various players will experiment with these strategies as they seek out an optimal approach.

6. CONCLUSIONS

If we view insurance risk taking and insurance risk hedging in light of principal-agent theory, then a component of risk pricing will relate to principal-agent costs. Any participant in the insurance market, whether as buyer or seller, should anticipate that principal-agent theory will influence the utility considerations of taking risk or hedging risk. Any participant in the insurance market should expect that principal-agent costs, as a form of risk-taking executive compensation, will contribute to the price of risk. Any model such as CAPM or arbitrage-free pricing or allocated-cost-of-capital or capital allocation or other models that rely only on a hyper rational world or only on a total company view with no considerations of principal-agent costs inside an insurance company “group” is likely to be flawed.

Principal-agent theory is an important component of the overall risk pricing and risk management landscape for insurance and reinsurance and ought to be incorporated more fully in real world analyses.

7. REFERENCES

- [1] Ariely, *Predictably Irrational*.
- [2] Bodie, Kane, Marcus, *Investments*
- [3] Bodoff, “Capital Allocation by Percentile Layer,” *Variance* **2009**
- [4] Bodoff and Gan, “An Analysis of the Market Price of Cat Bonds,” *Variance* **2012**
- [5] Bodoff, “Sustainability of Earnings”
- [6] Borch, *ASTIN Bulletin: The Journal of the IAA* , Volume 1 , Issue 5 , July 1961 , pp. 245 – 255, DOI: <https://doi.org/10.1017/S0515036100009685>
- [7] Buchanan and Tullock, *The Calculus of Consent: Logical Foundations of Constitutional Democracy*
- [8] Coase, “Nature of the Firm,” <https://doi.org/10.1111/j.1468-0335.1937.tb00002.x>
- [9] Cowen, *Big Business*
- [10] Duke, *Thinking in Bets*
- [11] Feldblum, “Personal Automobile Premiums: An Asset Share Pricing Approach for Property/Casualty Insurance”, *Proceedings of the CAS*, 1996, Vol. LXXXIII, Part 2, 190-296, <http://www.casact.org/pubs/proceed/proceed96/96190.pdf>
- [12] Gilbert, *Stumbling on Happiness*
- [13] Harari, *Sapiens*
- [14] Jensen, *Foundations of Organizational Strategy*
- [15] Jensen, *A Theory of the Firm*
- [16] Kay and King, *Radical Uncertainty*
- [17] Kahneman, *Thinking Fast and Slow*
- [18] Kreps, “Reinsurer Risk Loads from Marginal Surplus Requirements”
- [19] Kreps, “Investment Equivalent Reinsurance Pricing” 1995
- [20] Kreps, “Management’s Best Estimate of Reserves”

The Roulette Wheel and the Drunken Sailor: Principal-Agent Theory and its Ramifications for Insurance and Reinsurance Risk Management

- [21] Kreps, “Riskiness Leverage Models,” Proceedings of CAS 2005
- [22] Major, “The Cost of Risk: A COTOR-VALCON1 Discussion,” CAS E-Forum 2009
- [23] Mango, “Capital Consumption. An Alternative Methodology for Pricing Reinsurance,” 2003
- [24] Panning, William H., “Managing the Invisible: Measuring Risk, Managing Capital, Maximizing Value”, 2006 ERM Symposium, <http://www.erm-symposium.org/2006/pdf/papers/Panning%204-10-06.pdf>
- [25] Patrik, “Reinsurance,” Foundations of Casualty Actuarial Science, 2001
- [26] Ross, “The Economic Theory of Agency: The Principal’s Problem”
- [27] Taleb, *Fooled by Randomness*
- [28] Taleb, *The Black Swan*
- [29] Venter, “Premium Calculation Implications of Reinsurance Without Arbitrage.” ASTIN Bulletin, 21(2), 223-230. doi:10.2143/AST.21.2.2005365
- [30] Wenitsky, “The Reinsurer’s Conundrum: Loser’s Curse,” Best Review
- [31] Wilson, R. (1968). The Theory of Syndicates. *Econometrica*, 36(1), 119-132. doi:10.2307/1909607

Acknowledgment

The author acknowledges the CAS Reinsurance Research Committee for reviewing and providing constructive feedback on an earlier draft of this paper.

Biography of the Author

Neil Bodoff is Executive Vice President at Willis Re. He is a Fellow of the CAS and a Member of the American Academy of Actuaries. He can be contacted at neil.bodoff@willistowerswatson.com and neil_bodoff@yahoo.com