Research Study

Insurance Against Misinformation in the Securities Market: Actuarial Aspects

Harry Panjer

May 23, 2006

Commissioned by the
Task Force to Modernize Securities Legislation in Canada

Harry Panjer

Dr. Harry Panjer is a Professor for the Department of Statistics and Actuarial Science at the University of Waterloo.

Dr Panjer has taught actuarial science in universities in the US and Canada for over 30 years. He has published many articles in insurance related scientific publications. His books include; Actuarial Mathematics (1986), Insurance Risk Models (1992), Loss Models (1998, 2nd Ed. 2004), Financial Economics: with Applications in Investments, Insurance and Pensions (1998) and Operational Risk: Modeling Analytics (2006).

He holds BA, MA and PhD degrees in Mathematics, Statistics and Actuarial Science, and is a Fellow of the Society of Actuaries and the Casualty Actuarial Society and was named an Honorary Fellow of the Institute of Actuaries in the UK.

Dr Panjer has served on many professional committees at the national and international level and as President of two major professional actuarial associations, the Society of Actuaries (US) and the Canadian Institute of Actuaries. He is currently Chief Editor of the /North American Actuarial Journal/ and an Associate Editor of several other insurance related publications.
He is a frequent speaker at professional and industry meetings and has written many articles and editorials in the professional press. In addition, he provides consulting services to the insurance sector.

Table of Contents

1. Executive Summary

This article addresses some technical aspects of creating a program that increases competitiveness of Canadian markets by protecting investors from losses due to misinformation provided by issuers of publicly traded securities traded on Canadian exchanges. The paper examines modelling, pricing and funding questions for three scenarios, following the work of Tom Baker in his related paper. In particular, the parameters that are required are discussed. The paper makes no attempt to quantify any specific model parameters.

In order to evaluate the cost of the misinformation protection program under each of the three scenarios, it will be necessary to understand both the frequency and the severity of losses to investors. The portion of losses compensable varies under each scenario. If one uses historical data to estimate losses it may be necessary to risk-rate issuers of securities.

2. Summary of Recommendations

This paper explores a number of issues but does not come up with clear cut proposal for a specific program to protect investors. The following recommendations provide some guidance regarding the next steps in the process of focusing attention on specific outstanding issues.

Recommendation #1: Determine the appetite for such a program by governments; investors; dealers; and issuers. Without some support for a program, at least in principle, any further work may be futile.

Recommendation #2: Determine the form the program should take. I recommend pursuing the fund model with or without government support. There is likely little appetite for insurers and reinsurers to assume this risk without significant rewards, due to the high degree of uncertainty surrounding the risk.

Recommendation #3: Determine an appropriate size for the fund and level of aggregate inflow needed. This requires some serious study of past misinformation events.

Recommendation #4: Determine who the payer should be. There are good arguments for several options. A flat percentage per-trade fee is likely the best place to start. It can be refined using a risk-rating process at a later stage if desired.

Recommendation #5: Determine an appropriate financing mechanism that is able to absorb large swings in losses from the fund. This may be as simple as retroactive assessment of all issuers or as complex as a financially engineered instrument that allows investors to assume misinformation risk.

3. Introduction

In making investment decisions, such as buying of selling a particular stock through a stock exchange, investors rely on information that is in the public domain. Investors' confidence in financial markets relies on the accuracy and timeliness of that information. Insider trading results from asymmetric information in markets. This asymmetry occurs if insiders know information that is not revealed to other investors or revealed to investors at a later time. It also results when the information that is revealed by issuers of securities to market participants is incorrect. Incorrect information can be provided to market participants intentionally or unintentionally. Information that is identified as being incorrect can be subsequently corrected. For example, quarterly earnings can be restated by an issuer after the initial report of earnings is made. This restatement has a direct effect on investors' views of the appropriate value of the security and the prices at which the security subsequently trades.

In this paper, we address the question of whether a misinformation protection program can be developed to protect investors from misinformation provided by the issuer (or its representatives) of a security. There are numerous difficult issues that need to be addressed in designing and financing a misinformation protection program. These include issues of defining the specific "misinformation events" that trigger losses; the size of the losses suffered; the appropriate amount recoverable by investors for such losses; and the form such a program would take. It could be designed as an insurance program operated by the insurance industry, as an insurance fund operated by the stock exchange, or as a government-backed guaranty fund. There are also several ways to finance such a program. For example, premiums could be paid by investors, or by issuers or the program could be funded by the stock exchange.

It is reasonable to assume that the risk of misinformation varies dramatically across different issuers of securities. Thus if such a program is to be actuarially fair, it will be necessary to assess the risk of misinformation for each issuer, so that the payer of the cost of the protection is charged for the appropriate level of risk exposure.

Misinformation can occur in three circumstances: (i) the purchase of a security offered by a prospectus during the period of distribution; (ii) the acquisition or disposition of an issuer's security in connection with or pursuant to a take-over bid or issuer bid; and (iii) the acquisition or disposition of an issuer's security at other times. This paper does not address the first two of these circumstances. It addresses only the case of misinformation in the case of the "normal" operations of the issuer after the initial offering, but excluding takeover bids.

4. What Legislation is Relevant?

For the purpose of this paper, we shall assume that Ontario is the relevant jurisdiction and that Ontario legislation[1] is applicable[2]. In Ontario, the determination of the occurrence of misinformation, the loss to investors and the limits of liability are described in the Ontario Securities Act, Part XXIII.1 (Civil Liability for Secondary Market Disclosure), Section 138. The Securities Act uses the term "misrepresentation" rather than term "misinformation" that has been used for this project.[3] Actions for misrepresentation can be made against the issuer as well as its directors, officers or other "influential persons", such as a majority shareholder, or an "expert", such as an accountant. For the purpose of this paper, we consider only actions against the issuer. Losses made by investors can be partially recovered by actions against directors and other parties and their insurers. Such recoveries may reduce the net losses after recovery. Thus, such recoveries may reduce the amounts recoverable from the misinformation protection programs contemplated in this paper.

Section 138 makes it clear that "misrepresentation" refers to both written and oral statements made by the responsible issuer (or persons or companies acting, with authority, on behalf of the responsible issuer). For both written and oral statements, there is no requirement for plaintiffs to prove that the provided misinformation was relied upon in making investment decisions. However, there is an additional burden of proof placed on the plaintiff when the misrepresentation is made orally rather than in written statements. For information presented orally, the plaintiff is required to prove that the provider of the information (except for experts) either (i) knew or should have known that the information was not correct; or (ii) acted in a grossly negligent fashion in providing the misinformation.

The Act also deals with timing issues. Section 138 refers to two key time points: the time of the misrepresentation or failure to disclose, and the time of corrective action. The time of misrepresentation is the time at which the public or oral misrepresentation is made (or the time at which a disclosure should have been made). The time of corrective action is the time at which the misrepresentation (or a failure to

[1] This section is based on a layman's interpretation of the legislation and should not be relied upon as legally sound without obtaining a legal opinion.
[2] This may present a problem for securities that are traded in more than one exchange. The recent case of Nortel involving New York courts illustrates this problem. We do not attempt resolve this issue in this paper.
[3] The terminology "misrepresentation" seems to imply that the information provided was deliberately or mistakenly in error based on the information available to the issuer (or its representatives) at the time the information was provided. It does not refer to forward-looking information where appropriately cautionary language is used. For example, forecasts of sales of products are always incorrect, in the sense that actual future sales will almost never be exactly the same as forecasts. With appropriate wording added to forecast information, users of the information do

disclose) is identified and publicly corrected (or disclosed). Multiple misrepresentations (or failures to disclose) with common subject matter or content may be treated as a single misrepresentation. Similarly multiple instances of failure to make timely disclosure of a material change or material changes with common subject matter may be treated as one misrepresentation.

Section 138 identifies the extent to which securities acquired or disposed of are relevant to the determination of damages. For securities acquired after the time of misrepresentation but before the time of corrective action, the damages are limited to the difference (including commissions) between the actual purchase and actual sale prices of the security (if sold during the 10 days following the time of corrective action), or an estimate of the market price if not sold. For securities disposed of after the time of misrepresentation but before the time of corrective action, the damages are limited to the difference between the actual sale price and actual subsequent repurchase price of the security (if purchased during the 10 days following the time of corrective action), or an estimate of the market price if not purchased in that period.

From this, one can infer that there are no damages to persons who held the security during the entire period from the time of misrepresentation to the time of corrective action. Thus the loss to investors relates only to the securities that are traded during the period.[4]

---

so at their own risk and do not have a basis for an action against the provider of the information.

[4] It should also be noted that because there is a buyer and a seller for each trade, there will be a loss associated with every trade, either to the buyer or to the seller. It is interesting to note that because there is a buyer and a seller for each trade, if there is a "misinformation loser", there must also a "misinformation winner". In the hypothetical situation where the "loser" is fully compensated by an insurance mechanism, the investor's downside risk is protected, while the upside gain of the "winner" remains unchallenged. This extreme situation of full protection at a sufficiently attractive price could induce investors to take misinformation risk or to try to arbitrage this risk. For example, if an investor suspects misinformation, the investor could trade between accounts that s/he controls, thus creating the opportunity for gain from misinformation at the cost of only the transaction costs involved in trading. In most insurance situations, the "loser" does not have a comparable "winner". In legal actions by "losers" against "winners", the winners are also the alleged wrongdoers.

5. Risk Modelling Elements

We begin examining the technical aspect of developing a misinformation protection program by first attempting to understand the nature of losses suffered by investors. Once that is complete, we can examine how specific program designs would provide partial protection against such losses.

In order to model the risk associated with misinformation in the marketplace, it is necessary to develop a mathematical model and determine the values of various quantities in the model. We will focus in misinformation for equities only, leaving other securities for later studies. We consider a single exchange, say the Toronto Stock Exchange (TSX), in which many companies are listed. These companies vary vastly in terms of size as measured by market capitalization.

We also choose arbitrarily a one-year time horizon for any "misinformation protection program". This choice of time period is only for the purpose of discussing probabilities and losses on a time-consistent basis - in this case an annual basis - and should not imply any particular form of the "misinformation protection program".

It is also necessary to distinguish the amount of losses suffered by the investor as a result of misinformation, from the amount of such losses that are compensable by the misinformation protection program. Compensable losses will necessarily be less than actual losses, partly to not induce inappropriate actions by the investor, as described above. Before venturing into more technical pricing issues of the program, it is imperative to understand the aggregate losses to investors in order to appreciate the possible magnitude of the program.

With a clear understanding of the detailed nature of losses, we can then evaluate specific forms of programs that are designed to cover a portion of those losses.

In general terms, we will need to assess the annual frequency of losses as well as the severity of losses for the exchange as a whole. This can be done in at least two different ways. In a "top-down" approach, we start at the aggregate level and estimate the frequency of misinformation events per year for the exchange as a whole. This can be done by examining the actual history of the stock exchange and making adjustments to reflect the changing size of the exchange. We would estimate the size or "severity" of losses suffered by investors by examining historical misinformation events and by building a model of the

size of losses given the current characteristics of the exchange.

In a “bottom-up” approach, we would develop a risk model that applies to each equity security listed by the exchange. This becomes a very large-scale exercise due to the number of listed securities in the exchange. The dimensionality of the problem can be reduced by grouping various issuers into categories based on some risk characteristic or by building a model of risk that captures key common variables such as industry or size. In order to fairly price the risk associated with each issuer, it will be necessary to use some version of a “bottom-up” approach; either one that models each security, or one that models the aggregate losses and then “allocates” the risk exposure back to the individual issuers. Such allocation will be based on the relative levels of risk exposure for each issuer. Whatever approach is used, the results of the two approaches should be consistent.

Finally, it is important to note that Section 138 distinguishes between two types of losses: (i) those due to misinformation that is subsequently corrected, and (ii) those due to delayed information that is subsequently released. It also distinguishes between “release of documents” and “public oral statements”. It will be necessary to consider all four combinations of these categories separately since the risk characteristics may vary.

The expected loss for the program as a whole in one year is the product of the expected frequency (average number of losses), per year for the program and the expected severity (average loss size), per year for the program. We will also obtain the distribution of the total loss (aggregate loss distribution), per year for the program. The expected loss is the mean of this distribution. The aggregate loss distribution provides us with the probability of all possible total loss amounts.

Calibration of the model frequency and severity presents a major challenge. If the past is used as a basis for estimating frequency and severity of losses for the future, it will be necessary to identify historical events that would be now identified as misinformation events. Although it might seem rather straightforward to look at misinformation events associated with the release of documents, especially regularly scheduled documents, it will likely to be very difficult to reasonably comprehensively identify public oral statements that led to significant share price change and that were subsequently corrected.

We now discuss each of the frequency and severity in more detail.

i. Frequency of Losses

The frequency of losses by investors due to misinformation can be broken down into elementary components: the frequency of released documents; and the probability of misinformation in the document that is released.

a) Frequency of Release of Documents

Corporations listed on a stock exchange release quarterly report documents and annual statements to the public. They may also release additional documents to the exchange or to the public depending on specific circumstances. Any misrepresentation in these documents constitutes the basis for an action for damages by any party that purchased or sold securities (see Securities Act Section 138.3 (3)), between the time of release of the document containing the misinformation and the time of public correction of that information. Issuers that release documents more frequently are exposed to greater risk of misinformation as a result of greater frequency of releases.[5]

b) Probability of Misinformation in the Released Document

Any document that is released may contain misinformation. The investor's risk of misinformation will be greater for some issuers than for others. In particular, the risk of deliberate misinformation made by the issuer may be greater during times where the issuer is experiencing financial stress. If this is the case, then for companies that are rated by a rating agency, the rating of a company may be useful for determining the risk of misinformation. The hypothesis that the risk of misinformation is related to the rating of the issuer needs to be established before using company ratings in assessing risk of misinformation. Misinformation may also arise from unintentional errors that are made in documents that are released to the public but that affect subsequent prices at which the issuer's security trades.

In any case, it will be necessary to assign a risk factor to each issuer in a modelling exercise. This is analogous to assigning a rating to each issuer. Development of ratings is a complex exercise. Development of a misinformation risk factor will be equally challenging if there is little historical data on

[5] Since the "release of documents" refers mainly to formal reports such as quarterly and annual reports, there is likely not a great deal of variation between issuers in this frequency. However, this will need to be investigated.

misinformation events.[6]

c) Frequency of Public Oral Statements

Public oral statements may be made from time to time by persons with authority to speak on behalf of the issuer. For the purpose of determining a right-of-action for damages, Section 138 treats such public oral statements in a manner similar to the release of documents. As with the release of documents, the frequency of public oral statements is a critical factor in determining risk of misinformation. In developing a detailed bottom-up model the frequency of public oral statements by each issuer will need to be well understood.

d) Probability of Misinformation in a Public Oral Statement

As with the release of documents, any public oral statement may contain misinformation. The probability of misinformation in a public oral statement may differ significantly from that in the release of a document, since public oral statements may have a different level of internal control than documents that are subject to audit and other internal controls. Alternatively, the existence of risk factors may induce issuers to make fewer public oral statements.

e) Frequency of Delayed Information

In Section 138, information that is delayed is handled in the same way as misinformation in documents or public oral statements. From our point of view, delayed information can be treated exactly like misinformation, where the time at which information should have been released is treated as the time of misinformation and the time of (delayed) release of the information is treated as the time of corrective action.

ii. Severity of Losses

Section 138 defines the amount of loss to investors arising from misinformation or delayed information. Purchasers of securities suffer losses following a misinformation event from prices that were higher than

[6] The development of misinformation risk factors may be a contentious issue. In fact, it could undermine confidence in the market by identifying issuers that are more risky. It may be necessary to use only one, or very few risk categories.

they would have been with full and timely information. The loss for investors who purchased shares following a misinformation event is the number of shares purchased at prices that were excessive, multiplied by the average excess price. The average excess price is determined from the actual price of purchase and subsequent sale price if the sale occurred during the misinformation period (up to 10 trading days following the time of corrective action), or an imputed sale price for shares that were not sold by the end of the misinformation period.

Those who sell shares following a misinformation event may also suffer losses. For shares that were sold at prices that were lower than they should have been with full and timely information, the loss is defined analogously.

From this, it can be seen that the severity of losses associated with a single case of misinformation depends on the daily volumes of shares traded following the release of a document or public oral statement, as well as the length of time between the time of misinformation (including the failure to release information) to the time of corrective action. As noted earlier, there is a potential loss associated with every trade that takes place during the relevant period.

a) Daily Trading Volumes Following the Release of a Document or Public Oral Statement

Since losses can arise from each trade during the relevant period, the potential for large aggregate losses is greater for actively traded securities than for thinly traded securities. Furthermore, the level of trading activity may be related to the time since the release of a document such as a quarterly report, or to the time since a public oral statement. The extent to which the release of documents and public oral statements induce trading activity immediately following the release or statement will need to be understood and captured in a model.

b) Length of Time from Misinformation to Corrective Action

Because losses are based on the volume of trades between these times, the longer the time until corrective action, the greater will be the potential aggregate loss. Longer delays until corrective action mean a larger number of trades since the time of misinformation.

c) Length of the Delay of Information

For delayed information, since losses are based on the volume of trades during the delay period, the longer the time until corrective action, the greater the potential loss.

d) Magnitude of Change in Price Following Corrective Action

Some misinformation may have a greater impact on the price of a security than others. Thus, for each security, it is necessary to understand the distribution of changes in price resulting from the misinformation (or delay in information) events. If one is able to classify misinformation into various categories, then one can refine this by developing distributions of price change for each misinformation type.

iii. Calibrating the Risk Elements

If one is to develop a full-blown bottom-up model, all of the above listed items need to be calibrated to obtain estimates for aggregate losses for the future. This can be based on historical information within the market or on external information. In the absence of such detailed data, aggregate historical losses and their volatilities may be used to provide rough estimates of aggregate losses.

One concern that needs to be considered is how representative the past data is of future losses. In particular, does the existence of a misinformation protection program or other recent actions (e.g. Sarbanes Oxley in the U.S.) have an effect on the behaviour of issuers in ensuring that information is correct and timely?

a) Aggregate Annual Loss

The probability distribution of the aggregate annual loss to investors of misinformation (or delayed information), occurring during the year and that is subsequently corrected (or released, in the case of delayed information), can be calculated based on models of the frequency and severity of losses if these models are correctly calibrated. The technical details of this calculation are described in the appendix of this report.

The probability distribution of the aggregate annual loss describes the chances that aggregate market

losses exceed any prescribed level. This model can be used for designing the details of a misinformation protection program. In particular, this model provides the expected losses over a one-year horizon. The expected losses represent the net amount of premium income that needs to flow into the program (before addition of risk-loading, and various expense items). The aggregate risk-loading can be determined from this probability distribution by setting the premium at an appropriate safety level. The aggregate risk-loading can then be allocated to individual issuers. Similarly all expenses of operating the program will need to be allocated to the individual issuers.

b) Annual Frequency of Loss

As discussed above, the frequency of misinformation events for an individual issuer depends on the frequency of release of documents, public statements and delayed information. For the release of documents and public statements, there are associated probabilities of misinformation in the document or public statement. These probabilities may be related to various other variables (covariates). These covariates could include the size (market capitalization) of the issuer, the capital structure of the firm, the nature of industry or many other factors.

With very low frequencies of misinformation events, the probabilities for individual issuers are very small and extremely difficulty to estimate. However, this situation is not uncommon. Other guaranty funds have developed risk-based assessments for the member organizations. These guaranty funds usually have a homogeneous group of members, typically from the same industry and the same geographical region. For a particular industry, such as banking or insurance, it is much easier to build a risk-based framework than it is for a heterogeneous group of industries as would be the case of an entire exchange. Furthermore, such guaranty funds of banks and insurers protect customers from the financial failure of the firm, which is a very well-defined event. Misinformation is a less well-defined loss event.

As with other rare events, the number of misinformation events for a single issuer or for the market as a whole should be modelled by a Poisson distribution with the Poisson parameter being set to the expected number of misinformation events in a year. This is the standard model used in insurance studies and can be applied to individual firms, industry sectors or all issuers as a group. If there is more observed year-to-year variation in observed misinformation events, additional uncertainty can be added by generalizing the Poisson to a negative binomial distribution. This is standard in insurance and in some models of credit risk; e.g. CreditRisk+, developed by Credit Suisse First Boston.

c) Severity of Losses

The size associated with a specific loss is much easier to calibrate than frequency, as it is better defined: at least the maximum loss is fixed. First, consider shares that are purchased following the misinformation event at a price that is higher that it should have been. Section 138 defines the loss as the difference between the purchase price and an actual or imputed price. The imputed price is based on actual trading prices during the ten-day period following the time of corrective action. In efficient markets, price adjustments occur quickly following the time of corrective action. In the absence of other forces affecting prices, the actual and imputed price should be very close together. The actual number of shares traded during the misinformation period is also known precisely.[7]

In order to calibrate the severity of losses, it will be necessary to model the aggregate level of trading levels of each issuer's shares, since more trades mean higher risk exposure. Trading levels can be determined from historical studies for each security listed in the exchange, and may be different immediately after an information event than at other times. It will also be necessary to model the change in security values resulting from each misinformation event. For most securities, there is no, or little, history of such events, and security-specific loss levels will be difficult to calibrate with a high degree of confidence. It may be best to use consistent percentage loss amounts for all securities. When expressed as percentages of pre-misinformation event share prices, historical data for all securities can be combined to develop a single probability of distribution of percentage loss for a single misinformation event. Alternatively, if there is sufficient data, one could use a separate distribution for each of several classes of securities.[8]

The natural probability for the loss rate (the proportion lost) is a beta distribution. The beta distribution can be calibrated using the mean and variance of historically observed losses. Technical details appear in the appendix.

---

[7] The issue of shares being traded on more than one exchange and where the investor protection program applied to only one exchange is not resolved. The purchase (sale) price of the share during the relevant period is known, and the imputed price following the time of corrective action is known since it should be the same in all markets. The unknown is to which shares the actual sale (purchase) price is applied for shares sold (bought) during the relevant period following a purchase (sale) also during the relevant period. Relatively straightforward rules could be developed for this situation.

[8] Similar issues arise in credit risk where one needs to model the proportion of bond value that is recovered (the recovery rate) by the bondholder following bankruptcy of the issuer of the bond. The loss to the bondholder is the balance of the bond's value. The recovery rate for bonds is often based on the seniority of the bond since bonds with highest seniority have first call on the remaining assets following bankruptcy. The recovery rate is often assumed to be fixed within each seniority class for the purpose of pricing credit risk.

6. Level of Risk per Issuer

We can examine the level of risk exposure for each issuer of securities using the expected (average) loss-level per year. For misinformation in the release of documents, the expected loss per event is proportional to each of: the daily average number of shares traded during the misinformation period (after the release of misinformation and before corrective action); the average number of releases of documents per year; the probability that a single release contains misinformation; and the average change in price of the security from before the misinformation to after it is corrected. For delayed information, similar quantities are required.

Obtaining reasonable estimates for each of these quantities requires a very large study of historical data for each issuer. This kind of study is well beyond the scope of this paper, but will need to be done at a future date in order to develop hard numbers if risk-based measurements are going to be credible.

It has been suggested that a simple per-trade fee would be appropriate to fund this program. An issuer's misinformation risk exposure increases with the size of the share ($2 vs. $70), the daily volume traded, the length of the misinformation period, and the individual issuer's information risk propensity. A per-trade fee that is a flat amount per share would capture variation due to the second item (volume) only. A per-trade fee that is a percentage of the share price captures the first two items (size, volume). If we assume no bias in the third item (length of the misinformation period), then we fail to capture factors related to the issuers' individual differences in risk propensity. Using a percentage per-trade fee is analogous to assigning all firms the same credit rating. If an objective, among others, of the program is to affect behaviour of the issuer, using a percentage per-trade fee offers no reward for good behaviour. Under the percentage per-trade fee, the cost of the program will likely be borne by the trading of actively traded large-cap issuers, such as the major banks and insurers. My guess is that these organizations are probably the least likely to have misinformation events. Thus investors in riskier shares are being subsidized by investors in less risky stocks.

On the other hand, one could argue that a flat per-trade fee per share reflects the higher risk for small-cap shares, an unproven hypothesis. Under a system of a flat fee per share, there would be an incentive to not have stock splits.

The point of the above two paragraphs at this juncture is point out that the risk-rating of issuers may be necessary to have a fair allocation of costs.

7. Modeling Losses for Specific Scenarios

Professor Tom Baker has identified three possible program forms (called "scenarios") in his paper, “Insurance Against Misinformation in the Securities Market”. In this section, we use these scenarios as a basis for developing corresponding probability models of annual losses to the insurance program under that scenario. Imposing limits of liability may change the distributions significantly.

Section 138 limits the liability to the issuer of the security to the greater of $1 million or 5% of the market capitalization in the case of misinformation or delayed information. In insurance terminology, this cap is a "per occurrence" limit. The severity distributions described above is the distribution on a per occurrence basis. The cap is either $1 million or 5%. However, since virtually all companies have a market capitalization over $20 million, the 5% cap applies to virtually all the issuers in the market.

i. Scenario 1: Government-Provided Primary Insurance

Prof. Baker develops this scenario by assuming a crown corporation modelled on the Canadian Deposit Insurance Corporation (CDIC). He calls this the Canadian Investor Insurance Corporation (CIIC). Under this model, the CIIC acts as primary insurer for investors. Investors do not need to first proceed with civil action against violators since violations will be determined through a process of the CIIC. The CIIC will determine the occurrence of a violation that is covered by CIIC insurance and then determine, based on market trading data, the identity of eligible investors and the amount of loss recoverable from CIIC for each investor.

Under Ontario securities law, Section 138 limits the liability of the issuer to the greater of $1 million or 5% of market capitalization. Presumably in cases where the aggregate loss exceeds 5%, there would be a pro rata adjustment to the amount recovered by each investor. For example, if losses are determined to be 10% of market capitalization, investors would each recover 50% of their losses.

a) Funding

In terms of actuarial pricing of losses, the expected losses for the fund as a whole can be calculated using the general procedures described earlier. In addition to these losses, the CIIC also needs to have way of covering losses that exceed this level. This means that a risk reserve is built up using annual "premium" income and other assessments; and that the CIIC has a mechanism to assess listed companies, as funds are

required to cover losses that exceed the amounts available from the risk reserve. When losses are limited to 5% of violator's market capitalization and the number of violations per year is small, any assessments should be easily manageable by the market as a whole.

There are several possibilities for funding the program:

- Investors could, in effect, buy insurance coverage though a per-trade fee. The expected loss for a single issuer is roughly proportional to the price of the security (a $10 share can only drop at most $10, a $100 share can drop much more), the chance of misinformation in a single released document for that issuer, the number of documents released per year, etc. This suggests that the per-trade fee, although small, should vary by issuer unless it can be shown that the resulting overall risk level per trade is quite uniform for all issuers or it becomes impractical to do so.[9] Part 6 discussed this in more detail in the context of the possible risk-rating of issuers.
- Issuers could pay a risk-based premium as an additional normal listing fee to the listing exchange(s). The premium considerations would be similar to those considered under the per-trade option. This approach is used in the life insurance industry's guaranty fund in Canada, in which a risk-based assessment is made of each member company.
- When reserves that are built up are insufficient to cover a loss, an assessment can be made of all listing issuers to cover losses. The assessment for losses in a single year could be spread over a number of future years to reduce the impact. However, in order for the CIIC to pay losses immediately, it would need to borrow funds in the marketplace or delay payment of covered losses if it does not use direct government support.

Subrogation presents an interesting challenge from a pricing perspective. Under the CIIC model, the CIIC will obtain subrogation rights. There is no history on such subrogation. From the perspective of CIIC management, it would be sensible to set a level of recovery through subrogation, initially. This level could be selected based on a consensus of expert opinions. Alternatively, one might begin by assuming zero subrogation initially. As experience develops, the CIIC could adjust premiums to reflect actual subrogation recoveries. Presumably subrogation plays an important role, allowing the CIIC to recover its losses from all of the at-fault parties, including the issuer.

[9] It is not unreasonable for investors to pay this insurance premium. Efficient market theory would suggest that we should expect the market price of the security to increase as a result of the existence of the program, and that the aggregate amount of increase equals the value (or approximately the price) of the insurance. An ancillary benefit of is that the cost of capital to the issuer is reduced because of lower volatility to investors as a result to the program.

Finally, the optimal size of the CIIC fund also needs to be determined. The success of the CIIC depends on its ability to make retroactive assessments to all issuers. The size of the fund can be determined in a number of ways. I assume that the fund will begin by charging premiums that reflect the risk level for each issuer and have a "loading" for expenses and risk. By balancing the desire for low premiums and the desire for a low level of volatility in assessments, some funds will need to be built up. The life insurance industry in Canada, for example, has built up a fund of about $100 million and can assess members for losses that exceed this amount. It maintains this level of funds as its ultimate level.

For our situation, the ultimate fund level can be determined through a process of negotiation. From the distribution of annual loss, one can determine the likelihood of exceeding the fund level in one, two, three or more years. It is possible to determine, using classical actuarial ruin-theory arguments, an upper limit on the amount necessary to never have an assessment for losses exceeding the fund. However, this requires that the fund be allowed to grow indefinitely. With a size-limited (in constant dollars) fund, this theory shows that there will be assessments at some time in the future.

An alternate way to manage loss is for the CIIC itself to securitize total losses or the portion of losses beyond the capacity of the fund; in effect, obtaining insurance against excess losses. It may be able to do this by issuing a security (e.g. a high-yield bond or other credit derivative-like instrument), whose return depends on the losses incurred by CIIC each year. Good loss experience results in high yields to investors. Poor loss experience means low yields, possibly even negative yields.

This approach has the effect of allowing those investors in the marketplace who wish to act as "insurers" of misinformation risk to assume this risk.[10] This type of program may be attractive to some investors who feel that misinformation risk is partially diversifiable with other market risk, or who feel that the returns provide adequate compensation for the additional risk. With sufficiently high-yielding bonds (or other securities), the investor may be able to increase expected return for a minimal increase in the aggregate risk for the investor.[11]

In any case, whatever form CIIC financing takes, there needs to be an appropriate risk charge made to the party that is paying to cover the risk. The risk charge is necessary to cover, in part, the cost of capital

[10] If this type of program is deemed to be "insurance", a licensed insurer will need to be involved. If it is viewed simply as a product guarantee, an appropriate security will need to be developed by an investment banker.
[11] For example, consider a conservative investor holding only high-grade bonds. Adding a high-yield (but credit-worthy) bond whose only risk is misinformation (not credit) can add some yield by assuming a portion of the market misinformation risk.

necessary to support the risk. This risk charge flows into the fund and accumulates over time.

ii. Scenario 2: Industry-Provided Excess Insurance

Prof. Baker develops this scenario by assuming that the exchange (TSX) requires that all issuers on the exchange be required to participate in a misinformation insurance program called the Canadian Securities Fraud Protection Fund (CSFPF). Under this model, the CSFPF provides protection against eligible losses under Section 138 that are not recovered from private-party civil action against the responsible party. Prof. Baker suggests specific claims procedures to determine eligibility of investors. Unlike Scenario 1, in order for an investor to claim for losses, the investor must first obtain a civil judgement or an approved settlement from the responsible defendant. Hence, the determination of the occurrence of a violation would be made by civil courts. CSFPF would be a defendant in such actions and be involved in any settlement process.

Actuarial considerations are similar as under Scenario 1, but there are further complications. If losses payable by the CSFPF depend upon the amounts recovered through civil actions against other parties, the amount recovered from the CSFPF depends upon the forms of D&O and other insurance in place. The issuer will be a defendant in any civil action designed to recover investors' losses. It is possible that the compensable loss from the issuer exceeds the capacity of the issuer. If this is the case, the issuer will be bankrupted and the residual loss will be the responsibility of the CSFPF. In this case, the ability of the issuer to withstand misinformation claims against it is a critical piece of information in determining an appropriate premium. This information is similar to the credit rating assigned to issuers. Thus the credit rating of the issuer (if available) is an important risk variable. Finally, if the otherwise compensable loss exceeds the issuer's financial capacity, the court may recognize that it is not in the best interests of all parties to bankrupt the issuer and direct the parties to a settlement that allows the issuer to survive. In this case, the exposure to risk by the CSFPF increases, as it may be called on to cover additional losses.

a) Funding

In terms of actuarial pricing of losses, the expected losses for the fund as a whole can be calculated using the general procedures described earlier for the CIIC. However, rather than the fund recovering losses through subrogation, the fund reduces losses by the amount paid directly by the issuer and various related insurance companies.

Although Prof. Baker suggests only one way, there are several possibilities for funding the CSFPF program:

- Investors could, in effect, buy insurance coverage through a per-trade fee. Considerations in determining the size of the fee are similar to those under the CSFPF model. In general, the per-trade fee is rather unattractive to investors who feel that they should be able to rely on information of the issuer and should therefore not be required to pay an additional fee.
- Prof. Baker's suggestion is that issuers could pay a risk-based premium as an additional normal listing fee to the listing exchange(s). The premium considerations would be similar to those considered under the per-trade option. From the investor's perspective, by having the program in place, the issuer is warranting the information provided.
- Finally, as with the CIIC, when funds that are built up are insufficient to cover a loss, an assessment can be made to all listing issuers to cover losses. Similarly, the assessment for losses in a single year could be spread over a number of future years to reduce the impact. However, in order for the CSFPF to pay losses immediately, it would need to borrow funds in the marketplace or delay payments to beneficiaries.

Consideration of the ideal ultimate level follows the same arguments as for the CIIC.

iii. Scenario 3: Private Market Insurance

In this scenario, the exchange or legislation requires each issuer to obtain insurance coverage that compensates investors for misinformation losses. Alternatively, the exchange could require that issuers disclose the lack of such insurance to investors in every released document or public oral statement. This would certainly be a discouragement to self-insuring losses by issuers. This insurance would be over and above all applicable D&O and liability coverage.

One major challenge in this scenario is ensuring that investors are indeed protected. Since the insurance is purchased by the issuer but the beneficiaries are investors, investors rely on the issuer to select insurers that are creditworthy and to apply for insurance in a truthful manner, so that claims are not denied due to fraudulent information provided the insurer by the issuer. This is in contrast to D&O and other liability insurance obtained by the issuer to protect itself from claims.

Claims would be processed using the procedures of the insurer. As with Scenario 1, civil litigation may be

involved in settling claims against insurers. In this scenario, Prof. Baker has designed specific limits of liability for the individual investor of 75% of the compensable loss up to $250,000 per investor, per loss event. Such limits may be useful from the insurer's point of view; but fail to provide protection for the institutional investor.

The insurance market would set prices. If the insurance is voluntary, there is a potential for adverse selection. If the insurance is mandatory, there would need to be sufficient insurance capacity in the marketplace and enough insurers to create a competitive market of companies willing to write this type of coverage.

One major difficulty in writing this line of insurance is that losses across beneficiaries (investors) are not diversifiable. One claim of loss by an investor could trigger claims by all investors who purchased or sold the security during the misinformation period. In effect, the loss associated with a single misinformation event is the sum of all losses.

Another major difficulty is that the risk is not diversifiable with D&O and other liability coverage obtained by the issuer. An insurer exposed to an issuer's D&O risk may also be exposed to the issuer's misinformation risk. If D&O coverage is difficult to obtain due to market capacity, misinformation coverage will be even more difficult to obtain unless the prices for such coverage are very attractive for insurers (but not for issuers).

While the non-diversifiability of coverage for a single issuer presents difficulties for a single insurer providing that coverage, the losses for the entire market may be spreadable across insurers if there are a sufficiently large number of insurers and reinsurers around the world willing to assume a piece of the entire risk. This would suggest a pooling arrangement in which issuers purchase insurance and participating insurers each provide a fraction of the coverage for each issuer.

Finally, insurers that provide this coverage and are also issuers in the same market will have to obtain coverage for their own misinformation risk.

a) Funding

In the private insurance market, insurers will have to price this coverage for each and every issuer in the exchange. There will be an even greater need to accurately rate the risk level of each issuer. In the

absence of historical data for each issuer, the insurer will need to build a rating model that reflects the specific covariates of the issuer. The degree of loading depends upon the "price of risk" in the marketplace at a given time. As insurance capacity is absorbed by misinformation risk, the price of risk can increase, producing higher returns to insurance company investors to reflect the changed supply-demand relationship.

The "loading" in an insurance premium consists of the direct expenses and the amount of return (based on the price of risk) to investors for assuming risk. In the long run, if pricing models are accurate, this loading for risk flows to the insurer (and its investors).

The calibration of the elements in the overall risk model is similar to the other two scenarios. However, the limit of $250,000 suggested by Prof. Baker puts an upper boundary on the loss per investor. In previous scenarios, there was no need to identify the number of investors who traded during the misinformation period - from the time of misinformation to the time of corrective action. In the Scenario 3 model, there is the need to measure how many investors are eligible. This may result in investors taking actions to minimize the impact of this limit. For example, an investor who places a portion of his/her investments in the names of family members may be able to increase the effective limit substantially.

There will be further elements in a bottom-up model. For instance, we will need to understand the distribution of shares traded per investor. The average number is the total shares traded divided by the number of investors engaged in trading during the misinformation period.

Finally, because of limits and coinsurance, investor claimants will never cover actual losses and may still take action against the issuer for excess losses.

8. The Role of Diversification for Investors

Under the current model - with no investor protection program - the risk of misinformation is shared by all investors. For a portfolio of securities, holding a diversified portfolio can reduce the volatility in a portfolio for a given level of return. The optimal relationship between risk and return is a topic in the field of modern portfolio theory. Using mean-variance analysis, the efficient frontier giving the optimal portfolio can be obtained for any collection of securities. The calibration of stock returns and their covariances is done by using historical data. For stocks that are publicly traded, this information is easy to obtain and use.

Misinformation risk is one element of risk that is reflected in historical prices. Hence, it is already embedded in any analysis of optimal portfolios. Misinformation may be identified by a sharp rise or fall in stock price at the time of corrective action. If one were to be able to separate out the volatility due to misinformation, the remaining volatility would reflect other sources.

For actively traded securities, the impact of misinformation risk on estimated total volatility is probably very small in general. It is only large if (i) there are misinformation events affecting the share price in the historical data; (ii) the change in the price of the security due to correct information is large; and (iii) the period from the time of misinformation to the time of corrective action is large.

Because the misinformation period is typically quite short, when volatility is measured on an annual or monthly basis, a misinformation period has no impact if it occurs during the year or if the month in which it occurred is missed in the data.

9. Observations and Conclusions

Misinformation risk is faced by all investors, both individual investors and institutional investors such as pension plans. The loss occurs when securities are bought (sold) at prices that are too high (low). It occurs only for shares that are traded during the misinformation period. Thus, misinformation causes a temporary aberration in the price of a security; misinformation that is corrected is likely to have little long-term impact on markets. However, the cost of a misinformation event to a single investor can be large. The idea of an investor protection program is to protect those investors in order to develop greater confidence in the securities market. This greater confidence in an insured Canadian market enhances its competitiveness by differentiating it from uninsured markets.

It would appear to this author that of the three scenarios, Scenario 2 has the greatest likelihood of success. The major difference between Scenario 1 and Scenario 2 is the involvement of government in Scenario 1. It was suggested at the Roundtable in Toronto that it is unlikely that governments would be interested in being the insurer of last resort for the small segment of the population that trades securities sufficiently actively to have a significant risk of misinformation. In particular, misinformation only occurs when documents are release, or public oral statements are made.[12] Long-term investors using buy-and-hold strategies would not suffer losses unless there was misinformation at the time of purchase or sale of the security. However, it may be argued that government support for the enhanced competitiveness of Canadian markets is in the national interest: it would increase Canada's attractiveness for placing capital, resulting in higher foreign investment in the country, creating jobs, etc..

Scenario 3 requires the private insurance sector to assume the risk of misinformation at a price established by the insurance marketplace. The risk for an insurer is not diversifiable by increasing the number of shares of a single issuer that are traded, or by increasing the number of trades. Since the results of all trades in a single security are all perfectly correlated, increasing the number of trades increases the risk exposure. The risk is also highly correlated with related D&O and liability insurance covering the same issuer or related parties, and insured by the same insurer.

However, the risk for an insurer is diversifiable by increasing the number of issuers that it insures, and reinsuring a portion of each risk in order to manage the total exposure. If there is enough insurance capacity and there are enough insurers and reinsurers that are willing to venture into this type of

[12] Therefore, not trading a security would largely remove the risk of misinformation resulting from the release of document.

insurance, there is some chance that this approach could work, at least temporarily. However, it is well-documented that immediately following major, highly-correlated insurance losses,,(e.g. hurricane), prices for insurance coverage rises, which, in turn, often causes a loss of coverage as prices become unaffordable. Until we know how large losses could be, we don't have a good handle on the viability of a private market solution.

Scenario 2 is similar to other industry-based, self-governing guaranty funds. Premiums flow into the fund and are invested; benefit payments flow out of the fund. When available funds are insufficient to cover a loss, the member organizations are assessed additional amounts to cover the loss. This can work if the exchange has the power to "tax" its member listing organizations. A fund can be built up over time if losses are less than the premiums and assessments flowing into the fund. The ideal ultimate level of the fund and the loading-in of the premiums can be determined based upon some agreed-upon criteria; e.g. the expected time until an additional assessment is made is ten years, or the probability of an additional assessment is 1% within year. When the fund reaches the ideal ultimate level, a premium holiday could be declared, or the fund could pay dividends.

Scenario 2 can also work if the capital market itself absorbs the risk associated with the fund being depleted. For example, any losses in excess of what is available from the fund could be securitized in some way; such as with a high-yield bond whose yield is reduced to reflect losses. Investors could participate by buying this security. In effect, willing market participants become the insurers.

In Scenario 2, the demand for accurate information to operate the fund successfully is less if the fund is operated using the assessment mechanism. As long as premium charges are correctly ordered according to the level of risk (e.g. greater risk means a larger premium), and the absolute level is reasonably accurate, the premium structure should be seen to be fair to the issuers. The assessment devise allows for "errors" to be corrected each year. Furthermore, as experience is gathered, premiums can be adjusted annually to reflect new information.

If it is agreed to pursue further study in this or any other option, significant work will be needed to analyzing historical information. I look forward to further discussion and participation.

Appendix: Modelling of Loss Distributions

In this appendix, we describe some of the mathematical tools used to evaluate the annual aggregate loss distribution.

1. Frequency of Misinformation Events

Consider an issuer labelled *i*, the *i*-th issuer of a total of *n* issuers. We consider only losses arising from released documents causing the share price to fall. Analogous treatment can be developed for price increases, for public statements, and for delayed information.

We use the following notation:

- λ_i: Expected annual number of misinformation events for issuer *i*.
- λ : Expected annual number of misinformation events for all issuers.
- r_i: Expected annual number of released documents for issuer *i*.
- p_i: Probability that a released document for issuer *i* contains misinformation.

From this, we have:

$$\lambda_i = r_i \cdot p_i, \quad i = 1,2,...,n.$$

For all issuers combined, we have:

$$\lambda = \sum_{i=1}^{n} \lambda_i \ .$$

2. Severity of Losses

We use the following notation for a single issuer:

- s_i: Expected daily number of shares of issuer *i* traded following release of a document.
- d_i: Expected number of days until corrective action for issuer *i*.
- P_i: Share price of issuer *i* immediately prior to a misinformation event.
- δ_i: Expected fraction of share value lost as a result of misinformation.
- μ_i: Expected loss from a single misinformation event for issuer *i*.
- c_i: Market cap of issuer *i* immediately prior to a misinformation event.

From this, we have:

$$\mu_i = s_i \cdot d_i \cdot P_i \cdot \delta_i, \quad i = 1,2,...,n.$$

For issuer i, the probability distribution of the size of a loss can be modelled as a fraction of the market cap. This distribution has mean μ_i. it can be modelled as a beta distribution with mean μ_i, lower end point zero, upper end point c_i and variance determined by the specific variance assumptions made about each of the factors s_i, d_i, p_i, and δ_i. These variance assumptions can be calibrated from historical data.

We denote the resulting distribution of the amount of a single loss from issuer i by $f_i(x)$.

3. Aggregate Losses

For all issuers combined, the expected loss for a single misinformation event is

$$\mu = \sum_{i=1}^{n} \frac{\lambda_i}{\lambda} \mu_i.$$

Using the Poisson assumption for the annual number of misinformation events for all issuers combined, the distribution of a single loss from a single misinformation event is:

$$f(x) = \sum_{i=1}^{n} \frac{\lambda_i}{\lambda} f_i(x)$$

and the distribution of aggregate losses can be written as:

$$f_{S_1}(x) = \sum_{j=1}^{\infty} e^{-\lambda} \frac{\lambda^j}{j!} f^{*j}(x).$$

We can obtain similar distributions for price increases, for public statements, and for delayed information (a total of five more distributions).

Assuming stochastic independence between events, the distribution of total loss for the year is obtained a the convolution of all these distributions:

$$f_S(x) = f_{S_1} * f_{S_2} * f_{S_3} f_{S_4} * f_{S_5} * f_{S_6}(x).$$

The distributions described above can be evaluated using a number of different possible tools including fast Fourier transform, Panjer recursion or simulation.

It should be noted that more complex modelling must be done if the simplifying assumptions do not hold. The empirical analysis of historical data can be used for this more sophisticated exploration. The model presented above is relatively simple and is computationally manageable.

Harry Panjer
Professor, University of Waterloo
Department of Statistics and Actuarial Science
Waterloo, Ontario N2L 3G1

hpanjer@uwaterloo.ca