BANK OF CANADA
BANQUE DU CANADA

Staff Discussion Paper/Document d’analyse du personnel 2018-2

Tail Risk in a Retail Payment System: An Extreme-Value Approach

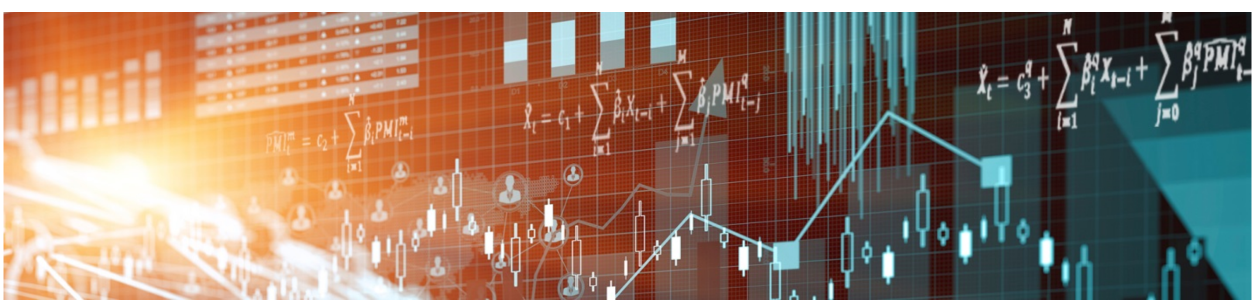

by Hector Perez-Saiz, Blair Williams and Gabriel Xerri

www.bank-banque-canada.ca

Bank of Canada Staff Discussion Paper 2018-2

January 2018

Tail Risk in a Retail Payment System: An Extreme-Value Approach

by

Hector Perez-Saiz,[1] Blair Williams[2] and Gabriel Xerri[3]

[1] International Monetary Fund
hpereze-saiz@imf.org
[2] Department of Economics
University of Toronto
blair.williams@mail.utoronto.ca
[3] Financial Stability Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
gxerri@bankofcanada.ca

ISSN 1914-0568

Acknowledgements

We want to thank James Chapman, Paul Miller, Miguel Molico and particpants of the presentations at the Bank of Canada, Payments Canada and the Bank of Finland for useful comments and discussions.

Abstract

The increasing importance of risk management in payment systems has led to the development of an array of sophisticated tools designed to mitigate tail risk in these systems. In this paper, we use extreme value theory methods to quantify the level of tail risk in the Canadian retail payment system (ACSS) for the period from 2002 to 2015. Our analysis shows that tail risk has been increasing over the years, but the pace of growth has been reduced towards the end of our data sample, which suggests a slower rate of growth of collateral required to cover that risk.

Bank topics: Econometric and statistical methods; Financial stability; Payment clearing and settlement systems
JEL codes: G21, G23, C58

Résumé

L'importance croissante que revêt la gestion des risques dans les systèmes de paiement a mené à l'élaboration d'un éventail d'outils destinés à atténuer le risque extrême dans ces systèmes. Nous utilisons des méthodes fondées sur la théorie des valeurs extrêmes pour quantifier l'ampleur du risque extrême dans le Système automatisé de compensation et de règlement (SACR) pour la période allant de 2002 à 2015. Selon notre analyse, le risque extrême s'est accru au fil des ans, mais à un rythme qui a ralenti vers la fin de la période étudiée, ce qui dénote une diminution du rythme de croissance de la valeur des garanties nécessaires pour couvrir ce risque.

Sujets : Méthodes économétriques et statistiques; Stabilité financière; Systèmes de compensation et de règlement des paiements
Codes JEL : G21, G23, C58

1 Introduction

Payment systems and other financial market infrastructures (FMIs) are crucial for the smooth functioning of a financial system. Because of their importance, FMIs must comply with risk management standards that ensure the financial system operates without any disruptions and that any potential loss from a default is covered with a high degree of confidence. As such, it is critical that regulators and policy-makers have confidence when applying risk-management tools.

Currently, central banks and regulators are implementing a set of international risk-management standards being implemented by central banks and regulators that are designed to improve the safety and soundness of systemically important payment, clearing and settlement systems. The Principles of Financial Market Infrastructures (PFMIs) from the Committee on Payments and Market Infrastructures (CPMI) of the Bank for International Settlements (BIS 2012) serve as important risk management principles for systemically important FMIs that are being used in numerous countries. In Canada, an additional set of standards has been developed specifically for prominent payment systems (PPS).[1] These standards adopt aspects from the PFMI but are appropriately calibrated for PPS. Canada's retail payment system, the Automated Clearing Settlement System (ACSS) has been designated by the Bank of Canada as a PPS and, as such, will need to meet the standards for PPS. For this study, we focus on the PPS standard for credit risk, which states ″A PPS should effectively measure, monitor and manage its credit exposures to participants and those arising from its payment, clearing and settlement processes. A PPS should maintain sufficient financial resources to cover its credit exposure arising from the default of the participant and its affiliates that would generate the largest aggregate credit exposure for the PPS.″ [2]

Our goal is to provide a forward-looking tool that supplies an estimate of how large exposures can become in extreme but plausible circumstances so that policy-makers or system operators can choose the appropriate level of collateral for the system with a high degree of confidence. Extreme events may not be observed historically but may arise and therefore should be considered when designing a safe and sound policy on collateral.

In addition, Canada is currently undergoing an initiative aimed to modernize the Canadian payments ecosystem. This paper will provide important insight on potential future exposures that may arise in Canada's retail system in extreme but plausible events to properly inform system designers of the potential credit risks that may be exhibited in the future.

Our analysis follows the literature that studies the extreme value methods, which has gained impetus among commercial banks, FMIs and other financial institutions to comply with various regulations (Christoffersen 2012). Our approach also complements the collateral schemes shown

[1] PPS, while not systemically important, are critical for economic activity. In these systems, disruptions or failures could have the potential to pose risks to the economic activity and affect general confidence in the payment system.

[2] For more information about criteria for PPS, see http://www.bankofcanada.ca/wp-content/uploads/2016/02/criteria-risk-management-standards.pdf.

in Perez Saiz and Xerri (2016) that use historical data to calculate the optimal level of required collateral by estimating the largest exposure within a time window. For our analysis, we use a Peak Over Threshold method to fit the tail using a Generalized Pareto distribution.

Using the entire sample of exposures for the period 2002-15, we are able to model the tail distribution and show that tail risk has been increasing over the years, but the rate of growth has been diminishing in the end of our data sample, which suggests a slower rate of growth of collateral required to cover tail risk.

2 The Automated Clearing Settlement System

2.1 Description

The Automated Clearing Settlement System (ACSS) is a deferred net settlement system for retail payments in Canada and was designated by the Bank of Canada to be overseen as a PPS on May 2, 2016. The ACSS was introduced in 1984 and is owned and operated by Payments Canada. The majority of retail payment items in Canada are cleared through the ACSS (approximately 73 per cent of retail payment value on average each business day at the end of 2016). The core of the ACSS is an information system used to track the volume and value of payment items exchanged between participants of ACSS and determine the final balances. The ACSS is used to process a high volume of lower value, less time-sensitive payments that do not require intra-day finality provided by Canada's Large Value Transfer System (LVTS). Settlement for ACSS takes place on the settlement accounts of direct participants on the books of the Bank of Canada using LVTS payments, on a deferred (next day) multilateral net settlement basis after final positions are determined.

2.2 Regulatory environment

Oversight and regulation responsibilities for payment systems in Canada are shared between the Bank of Canada and the Department of Finance. The Department of Finance has broad responsibility for the financial system in Canada, including payment systems, and also has authority to make regulations. The Bank of Canada has responsibility for the oversight of payment, clearing and settlement systems it has designated as having the potential to pose systemic or payments system risk to the Canadian financial system.

3 Data

3.1 Data sources and patterns

In this section we describe the nature of the data we use in our analysis. The data used consists of daily payments received and sent between participants in the ACSS. Using total received and sent payments, we construct the final daily multilateral net settlement obligation of each participant with ACSS at the end of each cycle (day) for 2002–15. This is the amount that each participant owes to the system (or is owed by the system) at the end of each cycle, and it can be described as the exposure for the ACSS if a participant defaults and cannot meet its end-of-day payment obligations.

Participant-specific patterns The final net obligation for each participant exhibits a bell-shaped distribution (see figures 1 and 2). In these figures, we have the following sign convention: A positive sign means a debit settlement obligation for the participant, and a negative sign means a credit settlement obligation for the participant. Generally, the mean/mode of the distribution is not too far away from zero. However, this depends on whether the participant tends to be a net receiver or a net sender in the system. If the participant tends to be a net sender, the distribution will be skewed to the right (meaning more often there is a net debit settlement obligation) and if the participant tends to be a net receiver then the distribution will be skewed to the left (meaning more often there is a net credit net settlement obligation). As we can see from Figure 2, the distribution of payments of every participant has a relatively small mean and median, on the order of few million. The participants can have relatively large extreme values, with the largest observed debit and credit positions for 2002–15 on the order $2 billion.

We test the distribution of exposures for each participant as well as the full sample for normality using an Anderson-Darling test. We find that we can reject the null hypothesis that the data follows a normal distribution (see Table 1). Visually it is also clear that this data is not normally distributed. A histogram with a fitted normal distribution (see Figure 1) shows that the data is not normally distributed but rather that the distribution exhibits many more observations tightly centered around the mean as well as fatter tails than the normal distribution. Because the data is not normally distributed, certain extreme value approaches, such as the variance-covariance method, that rely on assumptions of normality, would underestimate the tails of this distribution. These methods would thus not be appropriate for our estimation.

Time patterns We also analyze the evolution of the distribution of net settlement obligations over the sample period. We consider all net obligations in a given month for every participant, and we obtain several key statistics. Figure 3 provides a broad overview of the evolution of payment patterns in ACSS over the years. Figures 3a and 3b show the mean and median of the absolute value of net positions across quarters for 2002–15. The median shows a constant growth over the years, and the mean shows a slightly curved and increasing growth. Figures 3c and 3d show the standard deviation of the absolute value of net positions and net positions across quarters for 2002–15.

3.2 Credit risk in the Automated Clearing Settlement System

The motivation for our analysis stems from the existence of credit risk exposures in the system. With the current configuration of the ACSS, net settlement obligations that are incurred at time T are not settled until $T + 1$. This delay in settlement increases the potential for participants to default on their net settlement obligation and exposes the system and its participants to credit risk. In this article, we propose a methodology to estimate future exposures that have not been observed in the past. The goal will be to provide more confidence that the chosen level of required collateral for the system is sufficient.

We define a net settlement obligation, $np_{b,t}$, as the end-of-day net settlement obligation of participant b in period t with ACSS, with the following sign convention: $np_{b,t} > 0$ if participant b owes funds to the system (debit position), and $np_{b,t} < 0$ if ACSS owes funds to participant b. Because we are concerned with the exposure of the ACSS to the default of a participant, we define the exposure of the ACSS to the default of a participant b in period t, $e_{b,t}$, as

$$e_{b,t} \equiv \max(np_{b,t}, 0). \tag{1}$$

4 Methodology

4.1 Motivation

The methodology presented in this paper complements the collateral schemes shown in Perez Saiz and Xerri (2016) to account for future possible exposures that are not observed in the data.Perez Saiz and Xerri (2016) propose a collateral scheme that is based on the creation of a collateral pool with contributions from all participants. Following risk management standards provided in the Bank of Canada's standards for PPS,[3] the size of the collateral pool is determined by the single largest debit settlement obligation (exposure) within the window across all ACSS participants,

$$\overline{K}^{W}_{pool} = \max_{b,t \in W}(e_{b,t}), \tag{2}$$

where $e_{b,t}$ is the exposure of participant b in given day t, and W is the rolling time window considered. Perez Saiz and Xerri (2016) show that the size of the window determines the size of the collateral pool, and a very large window incorporates the largest historical values of $e_{b,t}$ observed

[3] http://www.bankofcanada.ca/wp-content/uploads/2016/02/criteria-risk-management-standards.pdf

in the time period considered.

A disadvantage of this approach is the reliance on historical values to control risk. This methodology uses historical simulation to cover the largest exposure among participants rather than some arbitrary large confidence level (e.g., 99 per cent), as it is usually done by extreme-value methods in the risk-management literature. As discussed in Kuster, Mittnik and Paolella (2006) and others, this historical approach ignores predictions extending beyond the extreme returns observed in the time window considered. In addition, this approach ignores the non-stationary nature of the data. Figure 3 shows a clear increasing trend of $e_{b,t}$ over the years resulting from the increasing rate of substitution of cash payments to other electronic means of payment, population increases, and other factors. We are interested in estimating the probability of having extreme but plausible credit exposures, and what would be the required collateral to cover the risk associated with these tail events would be.

In this paper, we provide an alternative approach to the historical method. We estimate the net settlement obligations that are not present in the historical data and therefore estimate the tail probabilities of these occurrences. This would represent extreme but plausible events that we would like to consider when determining the optimum level of collateral that is required to cover the single largest exposure created in the system on any given day. The estimated tail settlement obligations would represent the single largest settlement obligation in the system and therefore would represent the size of the collateral pool that members will contribute to.

For our analysis and estimation we use a peaks-over-threshold method to fit a GPD to model the tail distribution of our data. We chose this method over the classical extreme-value theory method of using block maxima to model extreme events due to its more efficient use of data. Classical extreme-value theory uses only the maximum value observed over blocks of time to estimate a generalized extreme-value distribution. However, the maximum value in each time period may not be truly indicative of extreme behaviour. We instead use the peaks-over-threshold method, which uses only observations exceeding a certain threshold, defined as extreme, to model extreme events. This allows us to more accurately model the tail of this distribution.

4.2 Threshold selection

We define the random variable exposure, e, as the exposure of an ACSS participant in a given day. The first step in estimating the distribution of the tails would be to determine an optimum threshold $\overline{e}$ to analyze the exceedances over this selected threshold. These exceedances, i.e., values $e > \overline{e}$, will be used to estimate the distribution of the tail and ultimately provide us with estimates for the optimum required collateral. Choosing a threshold $\overline{e}$ involves a trade-off between bias and variance. Choosing too low a threshold will skew the parameters towards lower observations and not accurately reflect the behaviour of extreme events. However, choosing too high of a threshold leaves too few observations to accurately estimate the parameters.

A method discussed in Coles et al. (2001) for choosing an optimum threshold $\overline{e}$ would be to estimate the parameters using maximum likelihood estimation with a series of given thresholds. Once the parameters are estimated for a series of given thresholds, we could seek linearity of the estimated parameter $\widehat{\sigma}$ in $\overline{e}$, and stability in the estimated shape parameter $\widehat{\xi}$ to choose the optimal threshold.

4.3 Tail estimation

As explained by Coles et al. (2001), we assume that random variable Y (exceedances over large threshold $\overline{e}$) is asymptotically distributed following a GPD with cumulative density function equal to

$$H(y;\overline{e},\xi,\sigma) \equiv \Pr(Y \leq y|e > \overline{e};\xi,\sigma)= 1 - (1+\frac{\xi y}{\sigma})^{-1/\xi}, \tag{3}$$

where σ is a scale parameter and ξ is a shape parameter to be estimated with maximum likelihood.

The GPD can be divided in three different types of distributions depending on the value of the shape parameter ξ:

1. Gumbel distribution: If $\xi = 0$, $H(y;\overline{e},\xi,\sigma) = 1- \exp\left(-1\frac{y}{\sigma}\right)$ exhibits an exponential tail
2. Frechet distribution: If $\xi > 0$, $1 - H(y;\overline{e},\xi,\sigma) \approx ay^{-1/\xi}$ exhibits a polynomial tail, where a is some constant
3. Weibull distribution: If $\xi < 0$, distribution $H(y;\overline{e},\xi,\sigma)$ has an upper end point $w_f{=}\frac{\sigma}{|\xi|}$

Once the threshold $\overline{e}$ is selected, we estimate the parameters using maximum likelihood. From our data, we observe a set $\{y_1,...,y_n\}$ of values of exceedances for a given time period and threshold $\overline{e}$. For $\xi \neq 0$ the log likelihood is derived from equation (3) as follows:

$$L(\sigma,\xi;n) = -n\log(\sigma) - (1+\frac{1}{\xi})\sum_{i=1}^{n}\log(1+\frac{\xi y_i}{\sigma}). \tag{4}$$

This expression in equation (4) is maximized with respect to the parameters ξ and σ, so the maximum likelihood estimators of the parameters, $\widehat{\sigma}$ and $\widehat{\xi}$ are obtained by maximizing the likelihood,

$$\max_{\sigma,\xi} L(\sigma,\xi;n). \tag{5}$$

The estimated parameters $\widehat{\sigma}$ and $\widehat{\xi}$ are obtained by solving this optimization problem. Using these estimates, we can model the distribution of exposures e where the tail over the threshold level is distributed following the assumed distribution in equation (3). With this, we can calculate the optimum collateral levels for a given high confidence level c (e.g., 99.9 per cent). The optimum collateral using the this extreme value methodology would be defined as the level $\overline{K}_{EV}$ such that

$$\Pr(e \leq \overline{K}_{EV}; \widehat{\sigma}, \widehat{\xi}) \leq 1 - c. \tag{6}$$

Confidence intervals to $\overline{K}_{EV}$ can be calculated using the profile likelihood method suggested in the literature (see Davidson 2003). The idea is to obtain confidence regions for the parameters by using the profile likelihood that maximizes the likelihood with respect to one parameter conditional on the rest of the parameters being constant.

5 Results

5.1 Threshold selection

We estimate the parameters using maximum likelihood for a series of selected threshold levels. In Table 2, we have values of the estimated parameters $\widehat{\xi}$ and $\widehat{\sigma}$ within a given range of thresholds from the 0.95 quantile to the 0.9999 quantile, i.e., we look at thresholds ranging from the largest 5 per cent of exposures to the top .01 per cent of exposures. As we discussed previously, this provides a secondary method of selecting the optimal threshold $\overline{e}$ for our tail estimation. We are looking to select a threshold where σ shows linearity in $\overline{e}$ and ξ stability. We observe that around a threshold of the 0.99 quantile there is relative stability around $\widehat{\xi}$ and linearity in $\widehat{\sigma}$ around this same point. Since the 0.99 quantile corresponds to a threshold just below 0.6 billion, this threshold falls within the area where the mean residual life plot exhibits near linearity, as desired. Therefore, using these two methods to select $\overline{e}$ as inputs to our selection choice, we can say that the 0.99 quantile is a good choice of threshold for this data set. It is important to note that at thresholds from the 0.9995 to 0.9999 quantiles there are very few exceedences beyond the threshold, and thus the estimate provided can be considered unreliable.

5.2 Tail estimation

Using a threshold of the 0.99 quantile (the largest 1 per cent of exposures in our sample), we use maximum likelihood estimation as shown in equation (5) to estimate the distribution of the tail. Figure 4 shows the probability density function at a threshold of the 0.99 quantile. Also, Figure 5 shows the fitted quantiles of the GPD against the empirical quantiles for exposures beyond a 99 per cent threshold. Figure 5 shows a reasonable fit of the estimated distribution.

We perform an Anderson-Darling test and a Kolmogorov-Smirov test to check if our data came from our fitted GPD. For both tests we find that we cannot reject the null hypothesis that the data came from a GPD with parameters as fitted (Table 3). For this threshold we estimate a value approaching zero for our shape parameter, $\widehat{\xi}$. This indicates our sample follows a Gumbel distribution with an exponential tail.

We express our results in terms of the level of collateral that would provide different levels of daily confidence. A 99 per cent confidence level would represent a level of exposure we would only expect to see surpassed once every 100 business days. When examining Table 2, at the threshold of the 0.99 quantile, we can observe at the 99 per cent level of confidence the required collateral would be approximately \$1.2 billion, at the 99.9 per cent level of confidence (one surpass every 1,000 business days) the required collateral would be approximately \$1.8 billion and at the 99.99 per cent confidence level (one surpass every 10,000 business days) the required collateral would be approximately \$2.3 billion. In Figure 6 we plot the fitted generalized Pareto cumulative distribution function to visually compare with the empirical cumulative distribution function to see how close the fit is.

Note that when comparing the results using a 0.99 quantile threshold to other thresholds in Table 2 we can see that the level of required collateral is relatively stable under most levels of confidence. The level of collateral only becomes unstable around very high thresholds such as the 0.9995 and 0.9999 quantiles that do not have enough exceedences to make reliable estimates of $\widehat{\xi}$ and $\widehat{\sigma}$. We choose to discard these results for any analysis or policy decision, but we use them to show that as the amount of exceedences reduces past a certain point, the robustness and confidence of the Pareto estimation breaks down.

5.3 Projections ahead

The analysis shown in the previous section does not take into account possible structural changes in the payment patterns in the ACSS that create non-stationary changes in the distribution of exposures. As we discussed before, Figure 3 shows a clear pattern in the evolution of net settlement obligations over the years. This figure suggests that the ACSS exposures are constantly growing over the years, but the distribution is getting more compressed because the standard deviation does not increase at the same pace as the mean or median of the distribution from years 2010 or 2011. This compression of the distribution could be due to a more deterministic pattern of payments that causes a significant reduction of the extreme values of the distribution. As the distribution gets compressed, the tails of the distribution decrease, which has relevant consequences in terms of the risk management analysis and optimum collateral.

Motivated by these facts, we extend our previous static analysis to take into account the changes in the distribution of exposures over the years. We use a four-year rolling window to calculate the parameters of our extreme-value model to try to identify any significant changes in the parameter estimates over the years, and the implications for risk management and collateral. The parameters are re-estimated over four-year windows. Since each four-year window has fewer observations than the full sample used previously we use a lower threshold for the rolling window analysis. Following a process similar to that used in our full sample analysis, we use a combination of mean residual life plots and testing a series of thresholds for stability in ξ and linearity of σ in $\overline{e}$ to choose a threshold. We find a threshold of the 0.97 quantile to be a good choice for a four-year window. This choice of threshold leaves 360 observations in each window with which to estimate our parameters $\widehat{\xi}$ and $\widehat{\sigma}$ using maximum likelihood.

We estimate our parameters using maximum likelihood and use the scale parameter $\widehat{\sigma}$ from each window, and the mean shape parameter $\widehat{\xi}$ across all windows to calculate return levels for each window. Coles et al. (2001) note that it is difficult to estimate $\widehat{\xi}$ precisely and recommend against estimating this parameter as a function of time. Accordingly, we do not use the shape parameter $\widehat{\xi}$ from each window when estimating return levels as these difficulties in estimating this parameter precisely introduce high volatility in $\widehat{\xi}$ when re-estimating on each window. The $\widehat{\xi}$ calculated over then entire sample gives a more accurate representation of the shape of this distribution.

Figure 7 shows the results of our analysis using the rolling window analysis for three levels of confidence intervals (99 per cent, 99.9 per cent and 99.99 per cent). Consistent with the patterns shown in Figure 3, we observe that the needs for collateral increased dramatically from 2006 to 2010 before leveling off. Required collateral even tends to decrease in 2015. We plot a linear fit curve for 2010–15 and we observe a negative but very small slope. It is hard to predict what will be the future evolution of the distribution of exposures, but we do not discard that this trend of the compression in the distribution continues over the medium term.

Figure 8 shows the return level plot of our results from the final window in our sample, the period from 2012–2015. Figure 8 shows that if exposures remain distributed as they were over the four years from 2012–2015, a collateral pool of approximately $2 billion should cover exposures with a 99.9 per cent level of daily confidence. In other words, we would expect to see an exposure larger than $2 billion only once every 1000 business days. However, it is important to note that the 95 per cent confidence interval on these estimates is relatively large, ranging from approximately $1.5 billion to $2.5 billion. This suggests that a very conservative collateral scheme that takes into account statistical inaccuracies in our methodology, for a confidence level of 99.95 per cent or above, could be imply an optimum collateral level close to $3 billion.

6 Conclusion

In this paper, we complement previous analysis found in Perez-Saiz and Xerri (2016) to quantify the level of tail risk in the Canadian retail payment system (ACSS) for period from 2002–15. Our analysis allows us to take a forward-looking approach to quantify the optimum level of collateral in the ACSS. We show that tail risk has been increasing over the years, but the pace of growth has been reduced towards the end of our data sample, which suggests a slower rate of growth of collateral required to cover that risk over time.

Our results will help us refine the existing methodologies used to control credit risk in payment systems so we can have a higher level of confidence such that the collateral pledged in the payment system would be sufficient to cover loses in the event of a future default. These results should be useful for a more efficient design of the future generation of payment systems in Canada and in other countries, which should ultimately lead to a safer and sound financial system.

7 References

Arango, C., K. P. Huynh, B. Fung and G. Stuber 2012. "The Changing Landscape for Retail Payments in Canada and the Implications for the Demand for Cash." *Bank of Canada Review* (Autumn): 31-40.

Bank for International Settlements (BIS). 2012. *Principles for Financial Market Infrastructures.*

Christoffersen, P. F. 2012. *Elements of Financial Risk Management.* Academic Press.

Coles, S., J. Bawa, L. Trenner and P. Dorazio. 2001. *An Introduction to Statistical Modeling of Extreme Values*, vol. 208. Springer.

Davison, A. C. 2003. *Statistical Models*, vol. 11. Cambridge University Press.

Kuester, K., S. Mittnik and M. S. Paolella. 2006. "Value-at-Risk Prediction: A Comparison of Alternative Strategies." *Journal of Financial Econometrics* 4 (1): 53-89.

Perez Saiz, H. and G. Xerri. 2016. "Credit Risk and Collateral Demand in a Retail Payment System." Bank of Canada Staff Discussion Paper No. 2016-16.

Wiener, Z. 1999. Introduction to Var (Value-at-Risk). *Risk management and Regulation in Banking*, Springer, 47-63.

Appendices

Figure 1: Histogram of exposures across all participants

This histogram shows the distribution of all end-of-day exposures across all participants for the 2002-15 period.

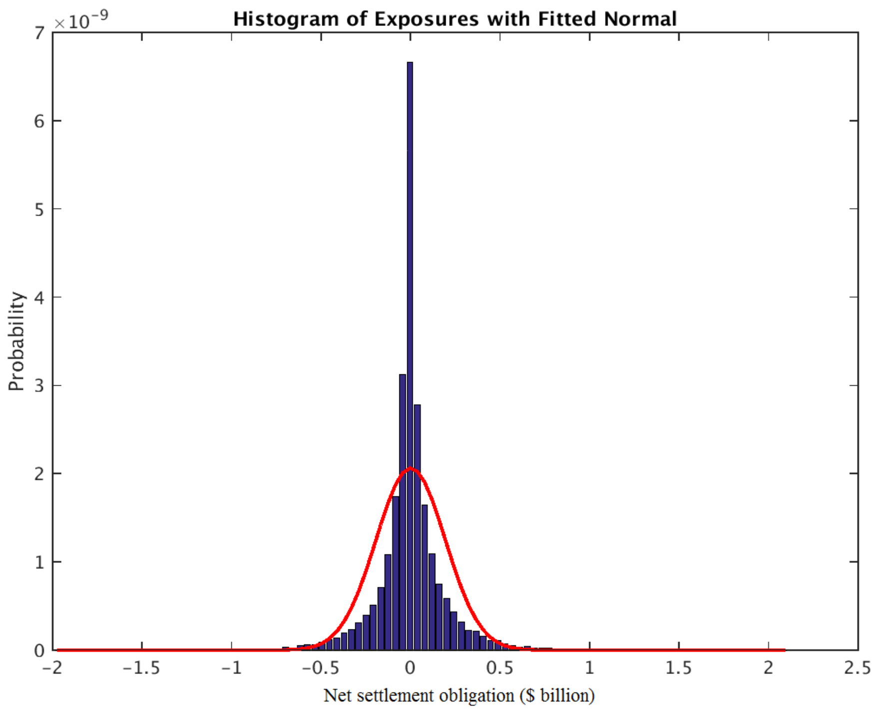

Figure 2: Distribution of net settlement obligations

In the next figure we show the distribution of net settlement obligations for every participant in the ACSS for the 2002-15 period. A positive sign means a net payment obligation of the bank with ACSS (debit position of the bank). A negative sign means a net payment obligation of the ACSS with the bank (credit position of the bank). Source: Bank of Canada calculations using Payments Canada data.

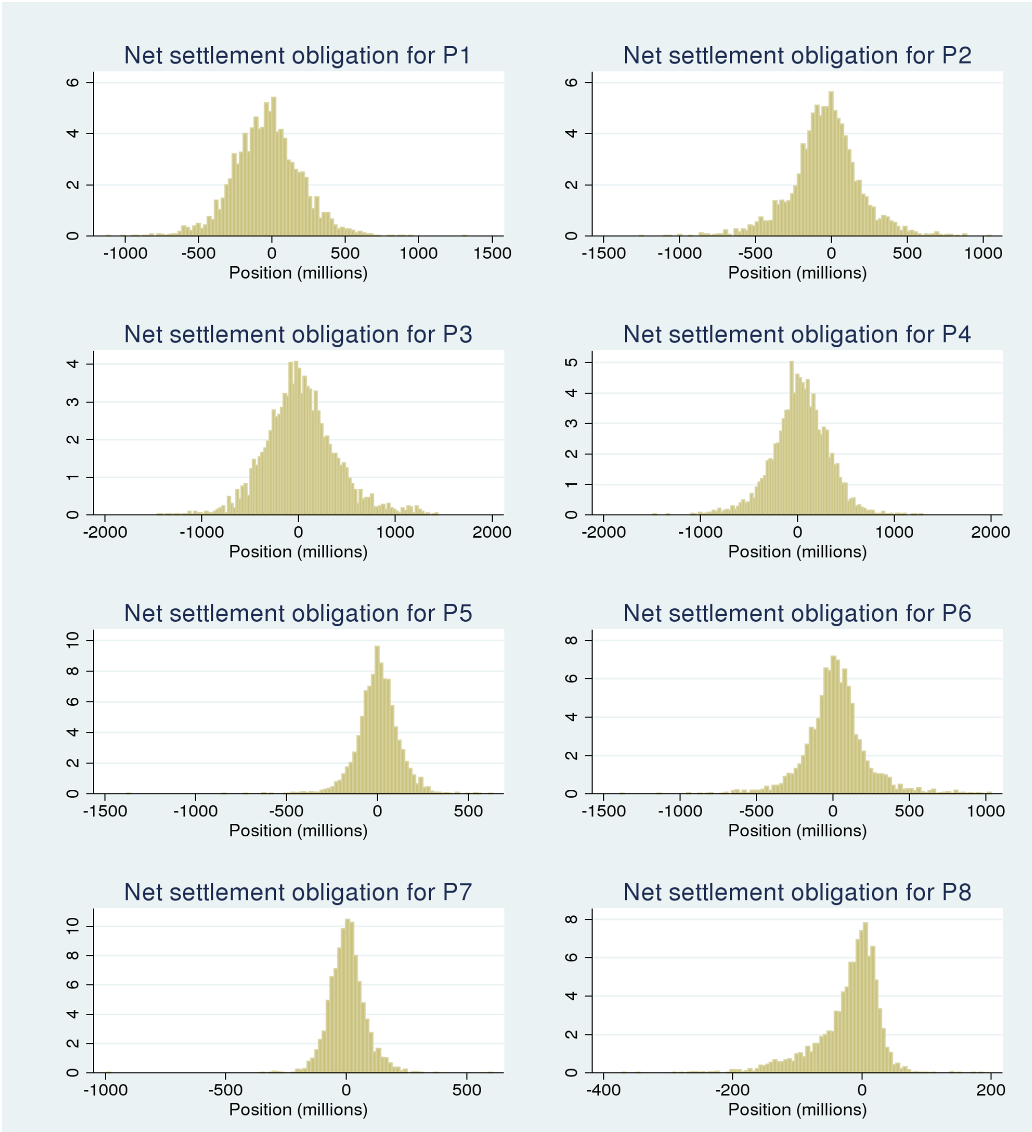

Figure 3: Statistical patterns in ACSS over the years

This figure shows several key statistics of ACSS over the years. Source: Bank of Canada calculations using Payments Canada data.

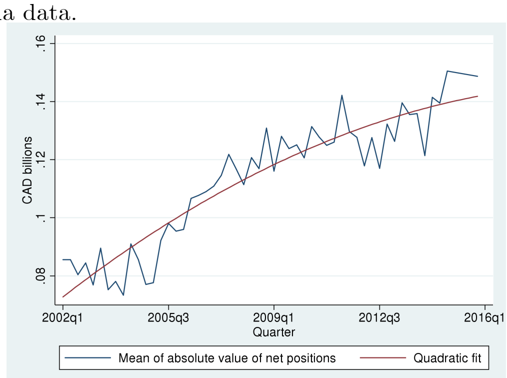

(a) Mean of absolute value of net positions

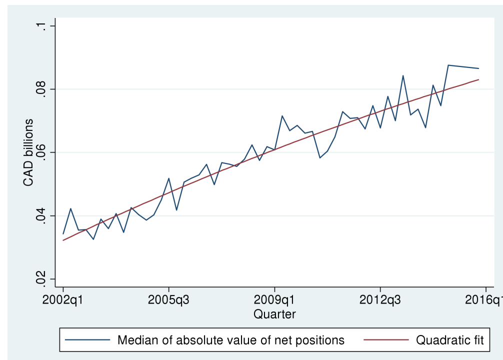

(b) Median of absolute value of net positions

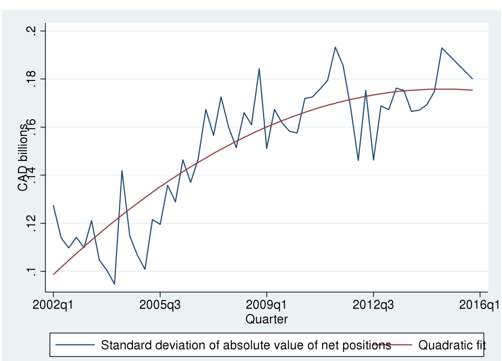

(c) Standard deviation of absolute value of net positions

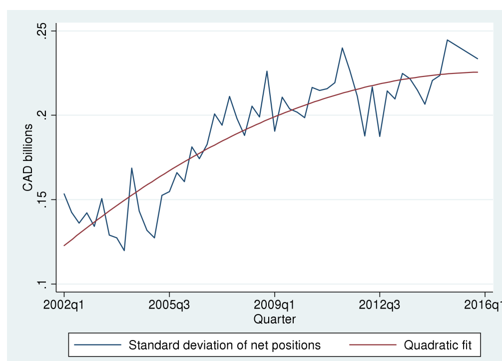

(d) Standard deviation of net positions

Figure 4: 0.99 threshold probability density function

This graph shows fitted generalized Pareto distribution against the empirical PDF for exposures beyond a 99 per cent threshold on the entire sample period (2002 to 2015).

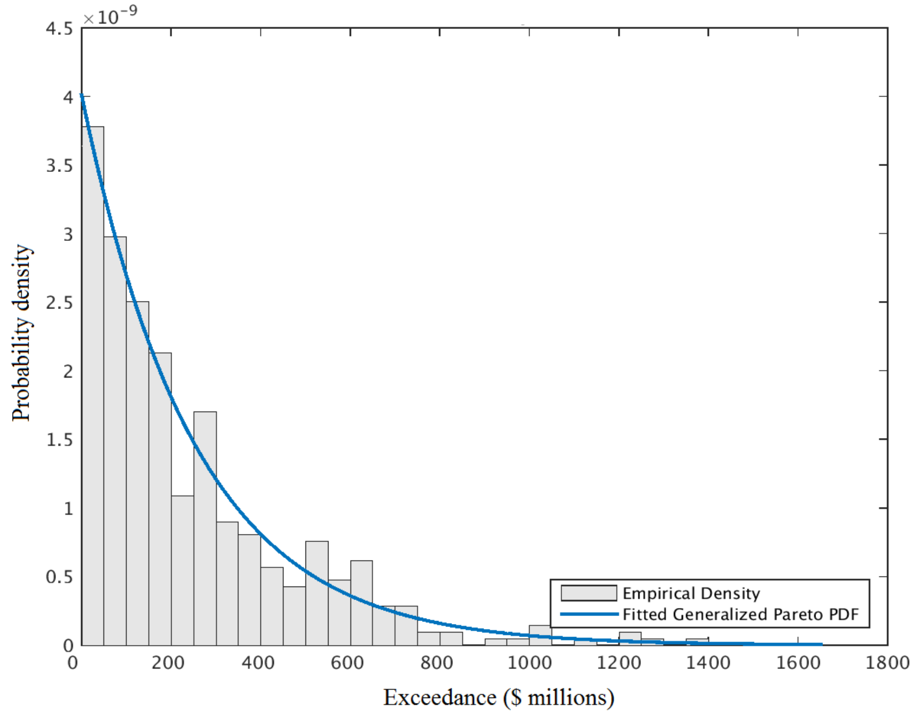

Figure 5: 0.99 threshold quantile-quantile plot

This graph shows fitted quantiles of the generalized Pareto distribution against the empirical quantiles for exposures beyond a 99 per cent threshold over the entire sample period (2002 to 2015).

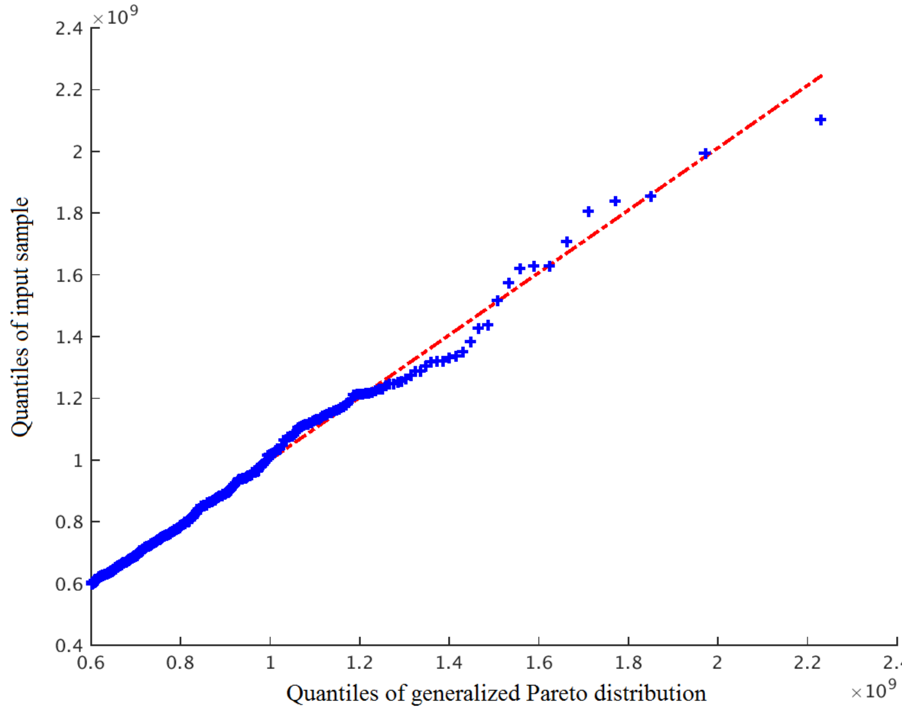

Figure 6: 0.99 threshold cumulative density function

This graph shows fitted generalized Pareto distribution against the empirical CDF for exposures beyond a 99 per cent threshold on the entire sample period (2002 to 2015).

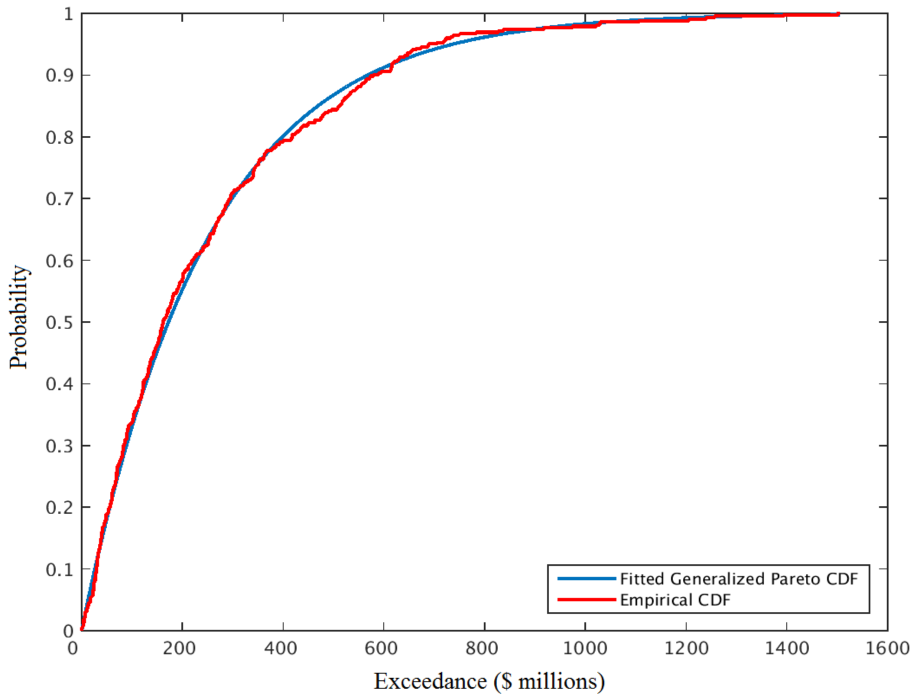

Figure 7: Collateral estimates using rolling window

Collateral required for 99 per cent, 99.9 per cent, and 99.99 per cent daily confidence levels estimated using a 4 year rolling window. Dotted lines show the OLS estimates of the trend from 2009 to 2015.

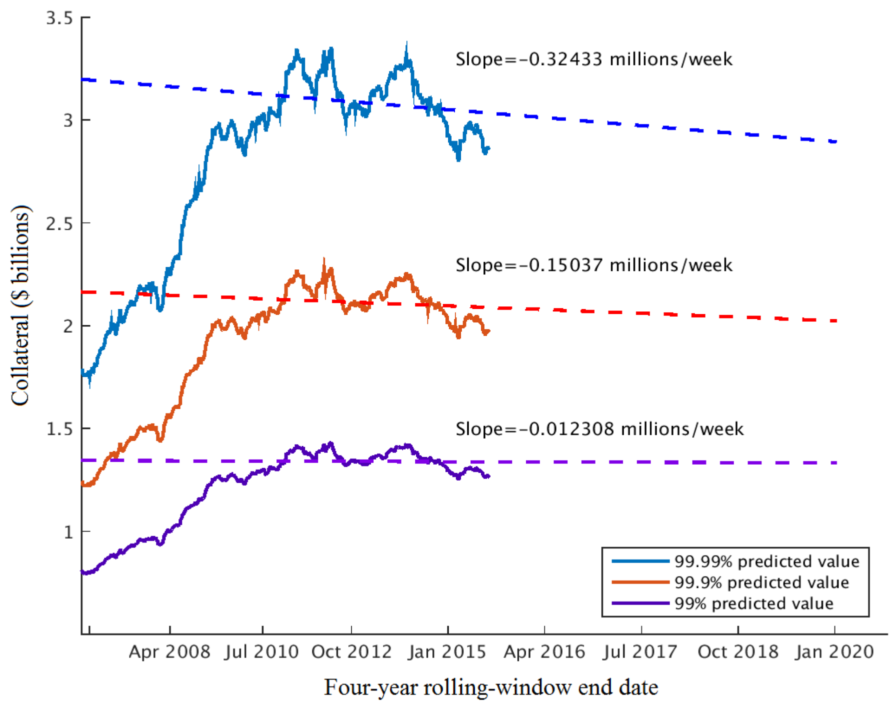

Figure 8: Optimum collateral

Optimum collateral as a function of daily confidence levels using the final window in our sample (2012 to 2015). We include confidence intervals for the collateral.

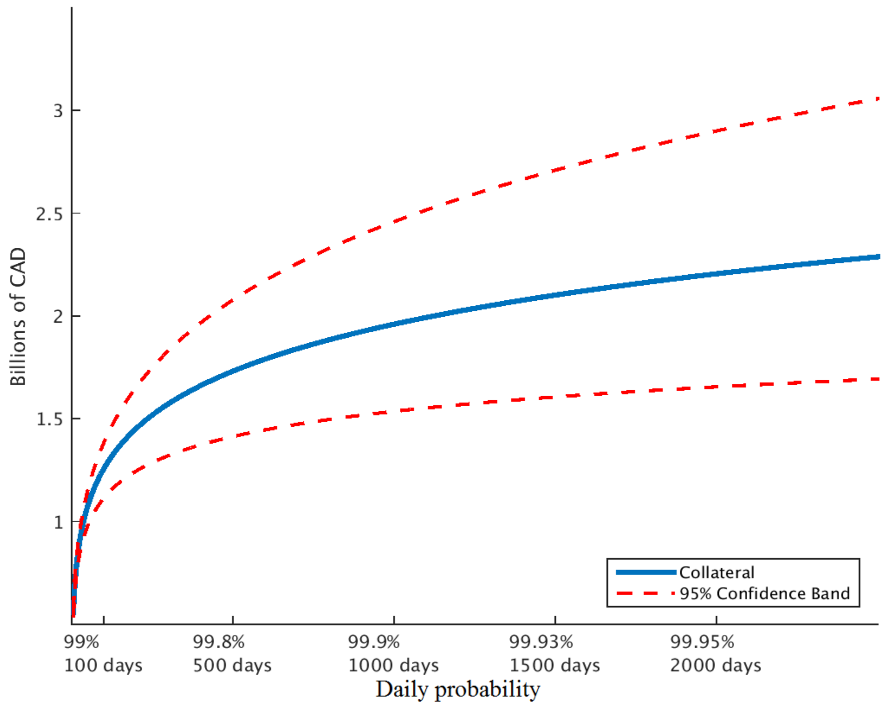

Table 1: Results of Anderson-Darling normality tests

This table shows results of Anderson-Darling tests for normality for the entire sample as well as each individual participant. The null hypothesis is that the data follows a normal distribution. *Value* gives the value of the Anderson-Darling test statistic. Critical Value gives the cutoff critical value for a 1% level of confidence.

Participant	**Value**	**1% Critical Value**	**P-Value**
All Participants	Inf	1.0935	0.0000
P1	6.5733	1.1189	0.0000
P2	60.9279	1.0467	0.0000
P3	55.1032	1.0380	0.0000
P4	35.9552	1.0237	0.0000
P5	21.3282	0.9774	0.0000
P6	19.1270	0.9174	0.0000
P7	11.0131	1.0248	0.0000
P8	30.6748	0.9397	0.0000
P9	46.3911	1.0457	0.0000
P10	17.4861	1.0511	0.0000
P11	99.9256	1.0880	0.0000
P12	Inf	1.0975	0.0000

Table 2: Results of estimation of threshold model

This table shows results of estimation of threshold model. $\widehat{\sigma}$ and collateral are expressed in billions CAD. Collateral is expressed as a daily confidence level (i.e. Collateral 99% represents the value we would expect to see exceeded once in 100 days). Source: Bank of Canada calculations using Payments Canada data .

Threshold $\overline{e}$ (probability)	**Threshold $\overline{e}$ (billions)**	**Number observations**	$\widehat{\xi}$	$\widehat{\sigma}$	**Collateral 99%**	**Collateral 99.9%**	**Collateral 99.99%**
0.9500	0.3015	2116	0.1103	0.1747	1.2167	1.9392	2.8706
0.9600	0.3443	1693	0.1392	0.1703	1.2298	2.0268	3.1249
0.9700	0.3918	1270	0.1147	0.1845	1.2210	1.9578	2.9174
0.9750	0.4249	1058	0.1067	0.1910	1.2192	1.9387	2.8586
0.9800	0.4699	846	0.1262	0.1897	1.2231	1.9840	3.0013
0.9850	0.5186	635	0.0552	0.2196	1.2173	1.8512	2.5711
0.9900	0.6014	423	-0.0097	0.2495	1.2238	1.7782	2.3204
0.9925	0.6709	317	-0.0176	0.2528	1.2258	1.7743	2.3010
0.9950	0.7626	212	-0.0680	0.2752	1.2368	1.7546	2.1974
0.9975	0.9455	106	-0.0841	0.2706	1.2396	1.7544	2.1785
0.9990	1.1820	42	0.1470	0.1873	1.2244	1.7547	2.4986
0.9992	1.2308	32	0.1818	0.1888	1.2186	1.7523	2.5633
0.9995	1.2957	21	-0.3304	0.3651	1.1101	1.7976	2.1189
0.9998	1.5129	11	-0.7029	0.4377	0.7253	1.8561	2.0802
0.9999	1.8165	4	-1.1805	0.3383	-1.2335	1.8828	2.0885

Table 3: Results of tests for generalized Pareto distribution

This table shows results of Anderson-Darling and Kolmogorov-Smirov tests for a generalized Pareto distribution over the entire sample. The null hypothesis is that the data follows a generalized Pareto distribution with our estimated parameters. *Value* gives the value of the test statistic. Critical Value gives the cut-off critical value for a 5 per cent level of confidence.

Test Statistic	**Value**	**5% Critical Value**	**P-Value**
Anderson-Darling	0.4470	2.4928	0.8011
Kolmogorov-Smirov	0.0279	0.0656	0.8875

Table 4: Estimates using .99 quantile threshold

This table shows results of maximum likelihood estimates on the full sample using a .99 quantile threshold. Results are given for our estimated shape and scale parameters as well as the inverse cumulative distribution values representing exposures which we would not expect to be exceeded on 99 per cent, 99.9 per cent and 99.99 per cent of days respectively. $\widehat{\sigma}$ and the collateral values are in billions of Canadian dollars. Lower and Upper confidence interval give the 95 per cent confidence interval for our estimates.

	Lower CI	**Estimated Value**	**Upper CI**
$\widehat{\xi}$	-0.1096	-0.0097	0.0903
$\widehat{\sigma}$	0.2173	0.2495	0.2864
99% Collateral	1.1682	1.2238	1.2924
99.9% Collateral	1.6242	1.7782	2.0437
99.99% Collateral	1.9861	2.3204	3.0015