

Raykov, Radoslav

Working Paper

Systemic risk and collateral adequacy: Evidence from the great crisis

Bank of Canada Staff Working Paper, No. 2019-23

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Raykov, Radoslav (2019) : Systemic risk and collateral adequacy: Evidence from the great crisis, Bank of Canada Staff Working Paper, No. 2019-23, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2019-23>

This Version is available at:

<https://hdl.handle.net/10419/210762>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Staff Working Paper/Document de travail du personnel 2019-23

Systemic Risk and Collateral Adequacy: Evidence from the Great Crisis



by Radoslav Raykov

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Bank of Canada Staff Working Paper 2019-23

June 2019

Systemic Risk and Collateral Adequacy: Evidence from the Great Crisis

by

Radoslav Raykov

Financial Stability Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
rraykov@bankofcanada.ca

Acknowledgements

This paper should not be interpreted as a factual description of, an analysis of or a view on any aspect of the TMX Group or any of its subsidiaries by the Bank of Canada. The author would like to thank Rod Garratt, Charlie Kahn, Miguel Molico, Héctor Pérez Saiz and Maarten van Oordt for valuable advice and comments and the Canadian Derivatives Clearing Corporation for providing the data. Marion Boddy and Nick Kazaka provided excellent research assistance.

Abstract

Conventional collateral requirements are highly conservative but are not explicitly designed to deal with systemic risk. This paper explores the adequacy of conventional collateral levels against systemic risk in the Canadian futures market during the 2008 crisis. Our results show that conventional collateral levels adequately absorb crisis-level systemic risk, even allowing for an implausibly large margin of error. However, this occurs at the expense of unequal buffering of systemic risk across banks. We document that the largest systemic risk contributors are buffered relatively less than the rest and that there is a large cross-country difference in the behavior of US and Canadian institutions. Nonetheless, even this does not result in meaningful risk spillovers. The maximum expected market shortfall in excess of collateral comes up to at most 1% of the banks' market capitalization, and hence the added systemic risk does not exceed the effect of a 1% downward stock price move.

Bank topics: Financial markets; Financial institutions

JEL codes: G10, G20

1 Introduction

Until the 2008 crisis, many financial institutions traded over-the-counter (OTC) derivatives bilaterally, using privately negotiated collateral requirements that were neither transparent nor subject to regulation. Although the financial system has evolved and most OTC derivatives are now traded centrally, collateral requirements in many OTC derivatives markets are still computed with conventional methods from the 1980s, not explicitly designed to deal with systemic risk. Conventional collateral requirements are often conservative in provisioning for every loss as for a potential default; however, they also ignore the comovement of default risks, as they were never designed to protect against correlated defaults. As a result, it is not a priori clear whether conventional collateral levels provide sufficient protection against systemic risk in derivatives markets. This paper aims to answer this question by exploring the history of systemic risk in the Canadian futures market during the 2008 financial crisis with an original risk methodology. It is illustrated with proprietary positions data provided by the Canadian Derivatives Clearing Corporation.

Existing research (e.g. Menkveld, 2016; Huang and Menkveld, 2016) has traditionally framed systemic risk in derivatives markets in terms of correlated losses or correlated loss probabilities derived from bank portfolios (see also Cruz Lopez et al., 2017, and Perez-Saiz and Li, 2018). This loss-based approach comes from existing risk management practice, which conservatively equates a financial loss with the payor’s potential default, i.e. with the worst possible outcome. While this may be a good method for risk-proofing, it is less so for risk measurement purposes, as it misses a key ingredient of systemic risk: the individual bank’s *ability to pay* such correlated losses when they arise. Correlated losses that are honored in full pose no risk to the market, systemic or otherwise; in that case, their correlation should not matter. Due to this, portfolio-based systemic risk measures based solely on loss correlations likely *overstate* the systemic risk of banks that experience contemporaneous losses but remain creditworthy.

At the same time, many derivatives markets rely on conventional collateral requirement models, which treat participants on a separate basis and do not capture their default risk comovements, thus likely *underestimating* the systemic risk involved (Cruz Lopez et al., 2017). Adding these two factors together, it is not a priori clear whether conventional collateral methods provide adequate protection against systemic risk.

To answer this question, we calculate a bank-level systemic risk metric, the *systemic payment shortfall*, and its market aggregate, the *systemic market shortfall*. The bank-level metric is a combination of two components: the bank’s daily payment obligation emanating from its portfolio, and a market-based systemic risk indicator measured on the [0,1] scale, capturing its extent of systemic risk. With this methodology, positions data determines each bank’s total payment obligation at each end-of-day, but a positive payment obligation only counts towards the bank’s systemic risk to the extent that it may not be honored conditional on a crisis. Our composite metric therefore combines information from both public stock returns and confidential portfolio positions.

As market-based risk indicators in the above calculation, we use two alternative measures: the marginal expected shortfall or MES (Acharya et al., 2017), and the tail beta, defined as the probability of a bank’s stock crashing below a certain threshold conditional on a banking index crash (Straetmans et al., 2008; De Jonghe, 2010). MES is commonly used in the bank systemic risk literature, and tail betas are common proxies for a bank’s probability of default conditional on a crisis. The two measures complement each other by capturing different aspects of systemicity. MES calculates the percent drop in a bank’s market capitalization conditional on a crisis, and so captures the degree to which a bank is affected by sector-wide distress. By contrast, tail betas estimate the probability that a bank’s return drops below a certain quantile concurrently with a banking index, regardless of the drop’s extent. We use both metrics in order to obtain a more accurate picture. Despite measuring systemic risk differently, both MES and tail beta analysis agree on the paper’s main point: that the inbuilt conservatism of traditional collateral requirements offsets their inability to deal with systemic risk.

Market-based systemic risk measures, such as MES and tail beta, likely reflect better aggregate information about a bank’s comovement with the rest of the sector than a single portfolio in a single derivatives market, because banks have incentives to split their investment portfolios to keep them proprietary and to economize on margin (Glasserman et al., 2015). Market-based risk measures predict bank performance at least as well as regulator supervisory ratings (Flannery, 1998), but in addition, are more forward-looking, higher frequency, and more granular (Cole and Gunther, 1998; Berger and Davies, 1998).

Each bank’s daily payment shortfall, calculated as the product of its payment obligation times either MES or tail beta, is then aggregated across market participants to gauge the

total systemic risk in the Canadian futures market. After calibration where appropriate, we buffer each bank’s systemic shortfall against two tranches of collateral typical of a centrally cleared market: transaction-specific collateral (margin), computed according to the industry standard SPAN method, and non-transaction-specific collateral (default fund deposit). We aggregate the result to obtain the market’s net expected shortfall, which indicates risk not absorbed by collateral.

We decompose the systemic market shortfall along several dimensions to gain an idea of how the 2008 financial crisis affected systemic risk in the Canadian futures market. We perform decompositions by bank, by trade type (proprietary or submitted on behalf of a customer), prior to the application of collateral, and after its application. The decomposition by bank identifies the key firms driving systemic risk during the crisis and the times when systemic risk peaked, therefore serving as a useful tool for risk monitoring. The decomposition by trade type shows the systemic risk brought on by a bank’s own proprietary trades versus those submitted on behalf of customers. Finally, repeating this decomposition before and after the application of each bank’s collateral shows the effectiveness of conventional collateral requirements in mitigating systemic risk. We first perform these decompositions with the MES version of the shortfall metric, calculated at daily frequency, and then validate our collateral adequacy and contagion results with tail betas.

In contrast to some previous studies (e.g. Paddrick and Young, 2017), we find that conventionally computed collateral provided adequate protection against systemic risk during the 2008 crisis, and resulted in minimal shortfalls above pre-pledged collateral, especially when mutualized among the survivors. The main reason for this is that, rather than making assumptions about the mechanics of stress propagation, instead we gauge a bank’s ability to honor obligations directly from its systemic risk indicator. With this method, we find that the 2008 crisis did not generate significant expected spillovers above the affected firms’ pre-pledged resources, despite the presence of two of the worst-affected firms – MF Global and Merrill Lynch – as active participants with significant overall activity. The maximum expected market shortfall in excess of collateral barely reaches 1% of the surviving banks’ market capitalization, and fails to meaningfully increase their systemic risk.

Our evidence thus shows that the conservatism built into conventional collateral models seems to offset their inability to deal with systemic risk. However, this comes at the cost of unequal buffering of systemic risk across banks, with the highest systemic risk contrib-

utors being buffered relatively less than the rest. Nonetheless, even this does not result in meaningful risk spillovers above collateral. This outcome is not obvious, since the Canadian market was not any less affected than the US: on the contrary, we are able to show that the bulk of the unmitigated exposures in the Canadian market came precisely from US firms.

Nonetheless, the expected shortfall’s grand total did not meaningfully breach the participants’ pre-pledged collateral despite the fact that its distribution across banks did not always match their systemic risk contributions. While supporting the case for joint margining systems such as CoMargin (Cruz Lopez et al., 2017), our results nonetheless suggest that much of the perceived systemic risks associated with centralized derivatives clearing under conventional margin models are overstated. The propensity of such methodologies to look at contemporaneous losses, while automatically equating them with defaults, naturally amplifies systemic risk estimates beyond the actual risk involved.

In addition to studying the distribution of systemic risk across banks prior to and after the application of collateral, we also study its variation across different institution types and over time. Our results show that the balance between proprietary bank trades and those submitted on behalf of bank customers is important in shaping systemic risk in the futures market. There are systematic differences in the behavior of bank customers and banks themselves along at least three dimensions: the composition of risk buildup (Canadian vs. American entities); response to the onset of the crisis (flattening out vs. increasing positions); and the timing of market activity, leading to different peaks of systemic risk (earlier onset and later tapering-off for systemic risk brought on by customers).

For example, we find evidence that with the onset of the crisis, American banks on average flattened out or reduced long positions on their proprietary trades, while Canadian banks did not. Customers of American banks, however, not only did not flatten out but increased their activity as the crisis deepened, driving over 80% of all customer-contributed systemic risk in 2009:Q2 and ultimately bringing the overall systemic risk in the market to peak in that quarter. Systemic risk brought on by customer trades also starts earlier and tapers off slower than that brought on by proprietary trades, as banks and their customers react differently to the crisis. This poses important questions about the correct timing of policy actions by authorities during a crisis, as most market members are Canada-regulated, whereas their customers need not be. The tools developed in this paper help monitor the ratios of systemic risk brought on by a bank versus its customers and identify when they

become excessive.¹

The paper contributes to several different literatures. On the one hand, it enriches the discussion of the risks and benefits of central derivatives clearing started by Duffie and Zhou (2011), but unlike these authors, it focuses on systemic rather than individual risk. The contributions of Menkveld (2016) and Huang and Menkveld (2016) are among the first to discuss systemic risk in centrally cleared derivatives markets. Both papers frame systemic risk in terms of “crowded trades” – bets on the same side of a single asset class that, conditional on the right stress event, could generate large simultaneous losses. The former paper studies the socially optimal balance between crowdedness and diversification, and how crowded trades could create spillover effects across market participants; the latter one decomposes the risk from crowded trades into position and price components. Both studies think of systemic risk in terms of simultaneous losses arising from correlated positions. A similar approach is followed by Cruz Lopez et al. (2017), who frame systemic risk in terms of P&L correlations and propose joint margining across banks, and by Perez-Saiz and Li (2018), who use joint default probabilities to measure interconnectedness between financial market infrastructures. The above studies are representative of the risk-proofing approach, which implicitly equates every potential exposure to an actual loss.

By contrast, our approach acknowledges the endogenous nature of systemic risk and its time variation, as evinced in the banks’ ability to pay. Its goal is better risk measurement, rather than simply risk-proofing, and is therefore closer to an adaptation of the systemic risk measurement literature developed by Acharya et al. (2017) and Straetmans et al. (2008) to the case of derivatives markets. Because of this, the paper’s results differ from that of previous studies on systemic risk in centrally cleared markets (e.g. Paddrick and Young, 2017), which rely on network methods with constant stress-transmission factors or fixed loss-given-default ratios assumed *ad hoc*. In spirit, our results are more similar to those of Heath et al. (2016), who find that the risk-propagation potential of properly regulated derivatives markets remains limited. To our knowledge, this is the first paper to explore whether the conservatism built into conventional collateral requirements offsets their inability to deal with systemic risk.

The rest of the paper is organized as follows. Section 2 provides an overview of the

¹Monitoring this balance is also required by international regulatory standards, known as the Principles for Financial Market Infrastructures (CPMI-IOSCO, 2012), especially Principle 19.

Canadian futures market. Section 3 describes the data and our approach to simulating SPAN margins. Section 4 explains the methodology used to compute systemic risk metrics and disaggregate them along different dimensions. Section 5 lays out the results, and Section 6 validates the collateral adequacy analysis with tail betas. Section 7 summarizes the findings and sets forth the paper’s conclusions.

2 Background on the Canadian Futures Market

The futures market in Canada is organized around the Montreal Exchange (MX), owned by the TMX Group. It includes all major Canadian banks and financial institutions as well as some of the largest US institutions trading in Canadian futures; the financial institutions covered in our sample are listed in Table 1. The market is centrally cleared through the Canadian Derivatives Clearing Corporation (CDCC), a TMX Group subsidiary. The average notional amount traded daily in Canadian futures is \$128.9 billion, comprising a total of approximately 15 broad contract types (Canadian bankers’ acceptance futures, Government of Canada bond futures, various share futures, index futures, repo and index swap futures, and sector index futures). Futures trading currently accounts for an average volume of 260,012 contracts daily, with an average open interest of 1.9 million contracts monthly.²

Traditionally, the three most actively traded futures contracts on the Montreal Exchange are Canadian bankers’ acceptance futures (BAX), 10-year Government of Canada bond futures (CGB), and S&P/TSX 60 index standard futures (SXF). Our historical sample contains data on these three contracts, which account for the bulk of the trading during the sample period (January 2, 2003 – March 31, 2011). The summary statistics of the contracts traded in our sample are discussed in Section 3.1.

An institutional feature of the Canadian futures market is that it is centrally cleared. This means that all trades, regardless of the parties transacting, are submitted to a central counterparty, the CDCC, which becomes the buyer to every seller and the seller to every buyer. This process is known as novation. Novation aims to reduce bilateral counterparty risk, since it replaces the original counterparty with a centralized regulated entity that is subject to stringent risk management standards. In the event of nonpayment or nondelivery by one of the original transactors, the central counterparty assumes responsibility to make

²Based on November 2018 regulatory reporting to the Bank of Canada.

good on the trade. To do so, it determines and collects collateral aimed to cover the potential future default exposure with a high degree of statistical certainty (above 99%) according to a rigorous risk analysis procedure overseen by the authorities;³ thus, collateral requirements are determined centrally according to a transparent and publicly accessible rulebook.

There are two types of collateral collected in a centrally cleared market: transaction-specific and non-transaction-specific. Transaction-specific collateral, known as margin, protects against the risk from a specific transaction, and is posted on a per-transaction basis.⁴ There are two types of margin: initial margin and variation margin. Initial margin is posted at the beginning of the transaction, and variation margin calls are used to update the initial margin amount to the level required given how the market has moved in the meantime. In addition, non-transaction-specific collateral, known as a default fund deposit, is required from every market member to protect against losses exceeding the margin posted. In the event that a bank has unpaid obligations so large that they exceed both its margin and default fund deposit, the excess loss is mutualized among surviving market members by consuming a pro rata amount of their respective default fund deposits according to CDCC's rules.⁵ The members have to subsequently replenish their deposits, which creates a cost to them: we calculate this cost in Section 5.2.2. Since such costs could induce risk propagation from unhealthy to healthy financial institutions, we loosely refer to this as contagion and measure its size by comparing the resulting costs to the respective banks' size. Our methodology reflects both the actual position data as well as the market's institutional features described above.

3 Data

To study systemic risk in the Canadian futures market during the crisis, we combine data from several sources. We use proprietary positions and default fund data courtesy of the Canadian Derivatives Clearing Corporation (CDCC), and financial market data on the sampled banks' stock market returns, retrieved from the Bloomberg and Eikon databases. The

³CDCC is overseen jointly by the Bank of Canada and by Canada's provincial securities regulators, and is subject to stringent international risk management standards known as the PFMI.

⁴However, the entirety of a bank's initial margin can be used to protect against a specific defaulted transaction.

⁵If the entire default fund is depleted, the remaining risk can be mutualized by additional methods (cash calls, variation margin haircutting, etc.).

data provided by the CDCC consists of three different components: positions data, default fund data, and margin interval data on futures contracts.

3.1 Positions data

The positions data covers each active CDCC member’s end-of-day position in the three most traded futures contracts at CDCC: Three-Month Canadian Bankers’ Acceptance Futures (BAX), Ten-Year Government of Canada Bond Futures (CGB), and S&P/TSX 60 Index Standard Futures (SXF). The first two types of contract are interest rate derivatives and the last one is an index derivative; all are centrally cleared. Collectively, these three contracts accounted for more than 90% of the futures contracts cleared annually at CDCC, and are highly representative of Canada’s futures market (Cruz Lopez et al., 2017; Campbell and Chung, 2003; TMX Montreal Exchange, 2013a, 2013b, 2013c). Each member’s end-of-day positions are reported for each of three types of accounts: Firm, Client, and Multipurpose. The Firm account registers the bank’s proprietary trades, whereas the Client account registers trades submitted on behalf of customers who use the bank to access the market. There is also a third type of account, a Multipurpose account, which commingles bank and customer trades, but such accounts make up less than 5% of activity in our sample and are therefore only analyzed as part of the Total accounts (the sum across the three account types). The BAX, CGB, and SXF contracts do not constitute the full set of derivatives cleared by CDCC, but nonetheless typically represent around 90% of the futures open interest (TMX, 2017) and 75% of CDCC’s total clearing activity (Cruz Lopez et al., 2017).

The dataset provides information not only on the contract type (BAX, CGB, or SXF), but also on the contract’s expiry month; thus we are able to distinguish not only between contract types, but also between contracts themselves. The data covers a total 113 contracts from the three types traded between January 2, 2003, and December 31, 2011. There are 39 expiry dates for BAX contracts traded by sample members and 34 expiry dates for CGB and SXF contracts. Basic summary statistics on the positions value and size are provided in Tables 3 and 4.

As can be seen from Table 3, BAX futures is the most popular contract type overall, as measured by both open position size and position volume in total accounts, followed by the CGB and SXF contracts. When decomposed by firm and customer accounts in Table 4, we see that bankers’ acceptance futures and government bond futures (BAX and CGB)

were almost equally popular in proprietary trading until 2006, with government bond futures falling out of popularity during the crisis and BAX emerging as the leader in proprietary trading. By contrast, bank customers largely preferred betting on BAX futures throughout, with the largest number of open contracts held around the financial crisis (2006-2008). This suggests that participating financial institutions and their customers seem to have had different investment goals and strategies, something that also comes up in the analysis of systemic risk across these two groups.

Since market-based systemic risk measures are constructible only for publicly traded firms, the statistics in Tables 2 and 3 reflect the sample of market players who are publicly traded, as they are the ones that can be analyzed. Their list is provided in Table 1. While acknowledging that they do not constitute the full market, the data shows that publicly traded financial institutions accounted for over 76% of our contract universe, and included all systemically important banks in Canada (especially “the Big Six” that account for above 80% of the banking market: RBC, TD, BMO, CIBC, Scotia Bank and National Bank) as well as the significant American members Goldman Sachs, J.P. Morgan, Merrill Lynch, and MF Global.

3.2 Default fund data

In addition to positions data, the CDCC has provided default fund data covering the sample period. In centrally cleared derivatives markets, participant-contributed default funds (sometimes also called guaranty funds or clearing funds) are used as a non-transaction-specific collateral buffer on top of the margin requirement. CDCC maintains a rigorously sized pre-funded default fund to buffer losses exceeding the defaulter’s margin deposit in the event of default. Consistent with industry practice, the default fund is sized by stress testing and calibrated to cover the biggest simulated loss over the potential defaulter’s initial margin observed in the past 60 business days, plus an additional 15% buffer. A market member’s contribution to the default fund is determined according to the member’s 60-day share of total initial margin, used as a risk-weighted proxy of its activity. Default fund requirements are calculated on the member-month basis, and are based on each member’s total activity in both proprietary and customer trades. In the event of default, a bank’s entire default fund deposit can be used to absorb losses, independent of whether the default is only in a single contract type or service, on a proprietary trade, or in a trade submitted on behalf

of a customer. This feature is taken into account when simulating expected losses due to systemic risk. We do not simulate CDCC’s stress testing, which determines the default fund size, due to the availability of actual monthly default fund deposit data on the participant level.

3.3 Margin interval data and calculation of simulated SPAN margins

The third type of data provided by CDCC is a key input for the calculation of initial margin, and is referred to as the *margin interval*. The margin interval is a measure of the potential future exposure that margin should cover against and is based on the volatility of the asset and how quickly it can be liquidated. We use it to calculate initial margin requirements according to the methodology of the industry-standard SPAN system.

The Standard Portfolio Analysis of Risk system (SPAN) was introduced by the Chicago Mercantile Exchange in 1988 and remains the core system behind initial margin calculations for derivatives in many important derivatives markets.⁶ SPAN is conventional in the sense that it margins market participants on a separate basis without explicit factoring of systemic risk (although it takes into account risk-offsetting effects within an institution’s portfolio). The technical details of the SPAN methodology and its simulation appear in the Appendix.

The simulated SPAN margin requirement, MR , reflects the Base initial margin on day 1 of the contract, and the updated margin requirement for subsequent days (so initial margin changes due to subsequent market moves are reflected). Margining in this market is done at the Total account level; however, to analyze the unbuffered systemic risk in firm and customer accounts, we additionally simulate a separate margin requirement for each Firm and Client account as a counterfactual. This allows us to draw further insights about the source of systemic risk across participant groups.

3.4 Price data

To compute market-based systemic risk measures (MES and tail betas), we obtain the participants’ daily stock market returns from the Bloomberg database. Since Bloomberg does

⁶Having become the de facto industry standard, it is used in the markets served by the Chicago Mercantile Exchange, Eurex, LCH.Clearnet, Nymex, Options Clearing Corporation, and CDCC, among others.

not provide historical prices for companies that have been delisted, such as Merrill Lynch and MF Global, we obtain the historical returns of Merrill Lynch and MF Global from the Eikon database instead.⁷ To calculate the value of open positions at CDCC, we use the daily market prices of each contract as reported by CDCC. We then merge the newly calculated systemic payment shortfall measure with the above positions data, default fund data, and author-calculated simulated initial margin (discussed in Section 3.3) covering the years 2003-2011. This unique dataset allows us to look not only at actual positions data of both CDCC participants and their customers during the years surrounding the 2008 financial crisis, but also at systemic risk they accumulated at CDCC and how its composition changed over time.

4 Methodology

We seek to construct a bank-level systemic risk indicator that can be aggregated on the market level and captures not only the risk from correlated losses per se, but also the banks' ability to service such losses when they occur. As such a measure, we construct the expected systemic payment shortfall, derived from two components: a bank's daily payment obligation on its derivative contracts, and a bank-level systemic risk measure on the $[0, 1]$ scale: either the marginal expected shortfall (MES) (Acharya et al., 2017), or the tail beta (Straetmans et al., 2008). We use these complementary metrics, which capture different systemic risk aspects and are available at different frequencies, to gauge the unserviced fraction of an outstanding payment obligation. To make use of the daily frequency of our position data, we first calculate systemic payment shortfalls with MES, which is available daily, and then validate the results with tail betas. In each case, the resulting bank-level shortfalls are aggregated to produce a market-level systemic risk indicator that takes into account both payment obligations and the banks' systemic inability to honor them.

We use this indicator to study the composition of systemic risk across banks both prior to and after the application of collateral. The former exercise compares risks on a relative basis and studies the distribution of systemic risk across banks and over time. The latter exercise gauges the adequacy of collateral on an absolute basis by estimating (an upper

⁷For technical reasons, a company may take some time to be delisted from an exchange after declaring bankruptcy or becoming insolvent. We therefore consider MF Global and Merrill Lynch as active CDCC members only until the dates of their respective bankruptcy filing (MF Global) and acquisition by Bank of America (Merrill Lynch).

bound for) the expected losses in excess of collateral and the cost of mutualizing such excess losses per the existing arrangements in these markets. For this exercise, risks are no longer on a relative basis, and the systemic payment shortfall needs to be put on the same scale as posted collateral to generate plausible dollar shortfalls.

As probabilities, tail betas already produce expected shortfalls in dollars when applied to a dollar obligation. To make sure that our higher-frequency MES results are not driven by issues of scale, we conduct our MES analysis based on assumptions that artificially amplify shortfall estimates. Firstly, we conduct the main MES analysis assuming that the resale value of the defaulted contract is zero.⁸ Secondly, we assume that variation margin on physically settled contracts is also physically settled, which is more risky. And thirdly, we also assume that MES, which is a correlate of the true default probability conditional on distress in the sector, understates that probability with a severe downward bias of 2 to 4 times, and show that our results remain intact. Finally, we validate our MES analysis with tail betas, which fully confirm our contagion and collateral adequacy results.

4.1 Systemic risk measures

4.1.1 Marginal expected shortfall

The first market-based systemic risk measure we employ is the marginal expected shortfall, or MES (Acharya et al., 2017). MES is a measure of comovement between an individual bank and the rest of the banking sector based on the drop in the bank’s market capitalization conditional on a sector-wide downturn. It can be measured both in dollar terms and in percentage terms (as a fraction of the bank’s market value). Since a bank’s solvency is directly related to its financial capital, capital shortfalls during a crisis are indicative of the bank’s (systemic) inability to service its obligations. Acharya et al. (2017) show that this measure is both theoretically well-motivated and has significant explanatory power in predicting the firms contributing to a crisis, performing better than traditional risk measures such as CAPM betas, expected tail losses, or volatility. This combination of a solid theoretical foundation, high predictive power and computational simplicity has made MES one of the widely used measures of systemic risk in the literature.

Compared to other systemic risk measures, such as CoVaR (Adrian and Brunnermeier,

⁸Or equivalently, that the default auction has failed.

2016), MES has distinct advantages in capturing tail losses beyond the threshold of VaR-based measures (Acharya et al., 2017); moreover, unlike CoVaR, MES is additive, which is crucial for being able to aggregate systemic risk across contributors or portfolios (Artzner et al. 1999). This factor is important in our selection of MES compared to other systemic risk metrics.

We construct MES for a bank i at a daily frequency as the average of i 's daily returns, conditional on the days where Canada's banking sector returns were within their worst 5% over a rolling window of one business quarter (65 business days).⁹ Although not all such worst days represent a crisis, Acharya et al. (2017) show, using extreme value theory, that this 'regular tail' of the return distribution is strongly linked to its 'extreme tail' (events that will materialize once a decade or less), which is what gives MES its predictive power.

Following Acharya et al. (2017), if $X_{i,d}$ is the return loss (negative return) of bank i on day d ($d \in [t - 65, t]$) and $Index_{i,d}$ is a stock market index of Canada's banking sector, then this bank's MES for day t is defined as

$$MES_{i,t} = \frac{1}{|I|} \sum_{d \in I} X_{i,d}, \quad \text{where } I = \{\text{worst 5\% of days for the returns of } Index_{i,d}\}, \quad (1)$$

where $Index_{i,d}$ is the S&P/TSX Composite Banks Index.¹⁰ As is conventional in this literature, MES is reported as a positive number, and negative MES values are set to zero since negative capital shortfalls are not meaningful in our setting.

As shown by Figures 1 and 2, and consistent with the results of Acharya et al. (2017), MES adequately captures the onset, peak, and abatement of systemic risk caused by the 2008 crisis in a timely and accurate manner, with more severely affected institutions displaying higher risk than the rest, and doing so at the times when they were most distressed. For example, Table 2 and Figure 2 show Bank 5 and Bank 7, two firms heavily affected by the financial crisis, displaying the highest MES values of 23.95% and 23.37%, respectively, just after negative news of their financial condition had emerged. In line with expectation, our

⁹This is done for consistency with the banking literature, where MES is computed quarterly.

¹⁰As alternative candidates for $Index_{i,d}$, we also consider two custom indices: an author-generated value index of CDCC's Canadian participants, and another value index based on all publicly traded CDCC participants. The potential benefit of these indices is that they exclude bank i from the index, thereby avoiding a mechanical relation between the bank and the index comovements. After comparing MES obtained with the three different methods, however, we find that, since Canadian banks are highly correlated, the custom indices did not make any practical difference (the resulting MES series are 99% correlated and their values are extremely similar).

MES calculations also show that US institutions comoved closely during the crisis, and that Canadian banks were less affected, scoring only half of the peak MES of their American counterparts. The summary statistics for MES of the remaining banks appear in Table 2.

4.1.2 Tail betas

The second market-based systemic risk measure we use is the tail beta (Straetmans et al., 2008). The literature has used different definitions of the term “tail beta.” Following De Jonghe (2010) and Perez-Saiz and Li (2018), we define tail beta as the conditional probability that a bank i ’s stock return crashes within its lowest 5% quantile conditional on the S&P/TSX Composite Banks Index return dropping to its worst 5% quantile.¹¹ Tail betas are often used as a proxy for a bank’s probability of default conditional on sector-wide distress, and provide an alternative measure of comovement between the bank and the sector. They also measure systemic risk differently than MES, ranking the banks with most frequent joint drops below the 5% threshold as most systemic.

Tail betas enable estimation of the behavior of the marginal and joint tail without imposing a particular distribution of the underlying returns, which often exhibit heavy tails. Another advantage of tail betas relative to MES is that, as probabilities, they produce expected shortfalls in dollars when applied to position values. We estimate tail betas using standard extreme value theory methods, described in detail in De Haan and Ferreira (2006); here, we provide only a brief overview. For consistency with the literature, we adopt the convention to take the *negative* of the return when describing the methodology.

Mathematically, if X is the return loss of a bank i and Y is the return loss of the S&P/TSX Composite Banks Index, we are interested in the likelihood

$$\Pr(X_i > Q_x(p) \mid Y > Q_y(p)), \quad (2)$$

where $Q_x(p)$ and $Q_y(p)$ are the respective top 5% quantiles of the two return loss distributions, defined implicitly by $p = \Pr(X > Q_x(p)) = \Pr(Y > Q_y(p)) = 0.05$. The conditional

¹¹Or equivalently, that bank i ’s return *loss* exceeds its 95th percentile, conditional on the S&P/TSX Composite Banks Index return *loss* exceeding its own 95th percentile.

probability (2) equals

$$\Pr(X > Q_x(p) | Y > Q_y(p)) = \frac{\Pr(X > Q_x(p), Y > Q_y(p))}{\Pr(Y > Q_y(p))}. \quad (3)$$

The joint probability in the numerator of (3) is determined by the dependence between the assets and their marginal distributions, which need not be identical. To avoid the inconvenience of working with different marginal distributions whose quantiles differ, it is standard to transform the series X and Y to their respective unit Pareto marginals, $X_t = 1/(1 - F_X(X_t))$ and $Y_t = 1/(1 - F_Y(Y_t))$, where F denotes the CDF of each variable. The respective empirical Pareto marginals are commonly calculated as

$$\tilde{X}_t = \frac{n+1}{n+1 - \text{rank}(X_t)} \quad \text{and} \quad \tilde{Y}_t = \frac{n+1}{n+1 - \text{rank}(Y_t)}. \quad (4)$$

\tilde{X} and \tilde{Y} are constructed to have identical marginal distributions, so their quantiles are now the same: $Q_x(p) = Q_y(p) = q = 1/p$. If we further define $Z = \min\{\tilde{X}, \tilde{Y}\}$, then the numerator of (3) can be rewritten as

$$\Pr(X > Q_x(p), Y > Q_y(p)) = \Pr(\tilde{X} > q, \tilde{Y} > q) = \Pr(\min\{\tilde{X}, \tilde{Y}\} > q) = \Pr(Z > q), \quad (5)$$

which is a univariate probability. Univariate probabilities of fat-tailed variables are commonly estimated with De Haan et al.'s (1994) semiparametric estimator

$$\widehat{\Pr}(Z > q) = \frac{m}{n} \left(\frac{Z_{n-m,n}}{q} \right)^{\hat{\alpha}}, \quad (6)$$

where $Z_{n-m,n}$ is the $(n-m)$ -th largest value of Z , and $\hat{\alpha}$ is a parameter called a tail index, which is estimated separately by the Hill (1975) estimator

$$\hat{\alpha}(m) = \left[\frac{1}{m} \sum_{j=0}^{m-1} \ln \left(\frac{Z_{n-j,n}}{Z_{n-m,n}} \right) \right]^{-1}. \quad (7)$$

The integer m sets the number of tail observations that enter the estimation based on which we judge about the tail's shape. It should be neither too small nor too large, as smaller values reduce the estimation's efficiency, while larger ones include non-tail observations leading to

bias. Following De Jonghe (2010), we optimally select m based on Huisman’s weighted least squares procedure, thus determining $\hat{\alpha}$ and m jointly. (For details, see Huisman et al., 2001). Combining equations (3)–(7) allows us to estimate the tail beta as:

$$\beta_{tail} \equiv \Pr(X > Q_x(p) | Y > Q_y(p)) = \frac{\frac{m}{n}(Z_{n-m,n})^\alpha}{p^{1-\alpha}}. \quad (8)$$

For further details on extreme value theory estimation, see, for example, De Haan and Ferreira (2006).

4.2 Daily payment obligation

The second ingredient of the systemic payment shortfall is each bank’s daily payment obligation emanating from its open positions, $PO_{i,k,t}$.¹² Absent any past due margin payments, it equals the member’s two-day loss.

The extent to which an unfulfilled payment obligation to the central counterparty crystallizes as a loss depends on two factors: (1) the ability to auction off the non-paying bank’s position and receive proceeds; and (2) the size of the defaulter’s pledged collateral (initial and variation margin and default fund deposit). With respect to the first factor, we conservatively assume that the resale value of a defaulted contract is zero (i.e. the default auction has failed), so that any unpaid obligation crystallizes in full as a loss. (This assumption is relaxed in Sections 5.2.3 and 6.) With regards to the second factor, we assume that variation margin settles in cash for cash settled contracts, and physically for physically settled contracts. The latter assumption reflects the additional risk that, in a crisis accompanied by a surge in demand for safe assets, the payor may be unable to quickly obtain physical securities, resulting in a rollover of the margin obligation towards a subsequent day.¹³ Padrick, Rajan and Young (2016) identify variation margin calls as one of the largest sources

¹²By payment obligation, we mean a financial obligation. Under CDCC Rules (Section A-804), a failure to deliver a security or an underlying interest does not automatically make a member non-conforming. Instead, CDCC’s reciprocal payment obligation in favor of the clearing member is reduced accordingly and any remaining quantity of undelivered securities is converted to a rolling delivery obligation each subsequent business day until delivered in full or until maturity of the security. If not delivered by maturity, the rolling delivery obligation is converted to a financial obligation. Only if that obligation is also not honored is the member considered to be in default.

¹³That is, we allow for failure-to-deliver risk for margin paid in securities. Failure to deliver is not the same as failure to pay (see footnote 12) and generally results in a rollover of the payment obligation. (This risk is not present in CDCC, which settles all variation margin in cash, but cannot be ruled out in a more generic market.)

of distress for banks trading derivatives.

These assumptions go considerably beyond the actual risk experienced and even beyond the extreme but plausible range considered by CDCC. We consider alternatives in Sections 5.2.3 and 6 and discuss the size of unpaid obligations relative to collateral in Sections 4.4 and 5.2.2. To make this useful for calculating systemic risk on a continuous basis, $PO_{i,k,t}$ is calculated daily, although contracts themselves expire only on specific dates each month. The payment obligation’s value in the remaining days is a counterfactual meant to keep a running counter of evolving systemic risk.

4.3 Constructing systemic payment shortfall

The final step of constructing the systemic payment shortfall indicator consists of putting the above two ingredients together. This produces a bank-level indicator of the systemic risk each bank brings to the futures market that takes into account both the size of its payment obligations and its systemic risk. We compute each bank’s daily expected *systemic payment shortfall* (SPS) on the bank-day level as

$$SPS_{i,t} = SysRisk_{i,t} \sum_k d_k PO_{i,k,t}, \quad (9)$$

where $SysRisk_{i,t}$ equals either $MES_{i,t}$ or bank i ’s tail beta. The contract-specific depreciation rate (haircut) d_k further reflects losses realized due to differences between the previous day and next-day resale value of the position once it is defaulted on and the default becomes known, thus reflecting the possibility that the default or the systemic event triggering it may itself change the contract price. We first set this parameter to hypothetical values ranging from 25% to 100% to compare scenarios, and then estimate it with extreme value theory in Section 6. The payment shortfall SPS is then aggregated across the bank’s contracts k .¹⁴

The systemic payment shortfall measure is a gross measure of risk, because it is not buffered with the potential defaulter’s collateral; however, it can be informative about risk composition in helping determine which banks are systemic *for the particular derivatives market*. This measure need not (and here, indeed does not) coincide with commonly com-

¹⁴For customer-submitted contracts, the payment obligation is computed in the same way and the clearing bank’s MES is assigned, because the risk posed to the market by a customer failure is only as big as that of the clearing bank, which in a legally binding way guarantees its customer’s obligation. I am grateful to the risk management team of LCH.Clearnet (London) for discussing this aspect.

puted systemicity rankings for big banks, since a bank’s activity in a particular market is not always correlated with its size or risk. The *SPS* measure, when aggregated across banks, is therefore useful in indicating the gross amount of systemic risk concentrated in the particular market and its composition. The aggregated version on the market level is called systemic market shortfall and correspondingly defined as

$$SMS_t = \sum_i SPS_{i,t}. \quad (10)$$

4.4 Calculating systemic risk spillovers

In the previous section, we created an indicator for the aggregate systemic risk in the Canadian futures market. In this section, we compare the magnitude of that risk against the collateral available to buffer it. In the case of MES, this requires additional calibration of the risk measure to make sure that risk and collateral are measured on a comparable scale. This is not necessary for the tail beta, which is already a probability and can be applied directly to position values to generate expected payment shortfalls in dollars.

As documented by Acharya et al. (2017), MES is an accurate predictor of distressed institutions during financial crises, and is therefore a strong correlate of their probability of default conditional on financial sector distress; however, since it is not a probability itself, we instead calibrate the calculated shortfalls using assumptions that artificially amplify them beyond the extreme but plausible case to make sure that the peak observed MES (times a large safety factor) bounds the peak of the latter probability from above. As long as this is the case, our MES risk spillover analysis remains valid.

We pursue this strategy by conducting the MES analysis based on two different sets of assumptions. The first set of assumptions assumes a depreciation $d_k = 25\%$ in the resale value of each defaulted contract and further assumes that observed peak MES understates the peak conditional probability of default by an implausibly large factor of 4 times, resulting in default probabilities of 96% and 94% for Bank 5 and Bank 7 at their respective peaks. A 4 times error factor is very close to the maximum possible, as a larger one would cause the conditional default probability to exceed unity.

The second set of assumptions uses the actual calculated MES, but instead assumes that the residual value of the defaulted contract is zero (i.e. that the default auction has failed) – a scenario also considerably beyond the ‘extreme but plausible’ range considered by regulators.

Since this method generates slightly larger, albeit similar shortfalls, we use it for all baseline computations, and discuss the results of the other method in Section 5.2.3.

Having thus addressed potential issues of scale, next we buffer the systemic payment shortfalls with transaction-specific collateral: the conventional industry-standard SPAN margin, simulated as explained in Section 3.

If a bank's gross expected shortfall exceeds its margin requirement, we further buffer the excess loss against non-transaction-specific collateral – the member's actual pre-pledged default fund contribution, as provided by the CDCC. Any spillovers above these resources are then calculated and their composition analyzed by contributing member and over time. The resulting picture is indicative of the systemic risk in the futures market net of conventionally computed margins and default fund deposits. This metric is calculated for each member i as its net systemic payment shortfall, or *NetSPS*:

$$NetSPS_{i,t} = \sum_k SPS_{i,k,t} - MR_{i,t} - DF_{i,t}, \quad (11)$$

where $MR_{i,t}$ is the simulated SPAN margin requirement for member i , based on the its position and aggregated over contracts k , and $DF_{i,t}$ is the member's (actual) default fund deposit at time t . As before, we use the *NetSPS* to explore the composition of the aggregate version of this metric, the net systemic market shortfall, or *NetSMS*:

$$NetSMS_t = \sum_i \left[\sum_k SPS_{i,k,t} - MR_{i,t} - DF_{i,t} \right]. \quad (12)$$

The net market shortfall provides a good measure about whether risk in this market is adequately contained, because it takes into account not only the gross expected risk present, but also the collateral available to buffer it, only after which losses become a threat to remaining market participants. The results from this analysis are discussed in Sections 5 and 6.

To get a better idea of the threat posed by losses exceeding the posted collateral, we take advantage of the specific institutional setup in centrally cleared markets. As discussed in Section 2, centrally cleared markets differ from bilateral markets in that losses in excess of collateral are mutualized across market participants to diffuse the risk. On an operational level, this occurs through consuming the non-defaulting banks' default fund contributions

in a way controlled by the central counterparty serving the market, and prorated according to survivor banks' risk-weighted market activity. This setup implies that risk mutualization might serve to potentially propagate systemic risk if not adequately contained by collateral, resulting in risk-mutualizing assessments (cash calls) as well as default fund replenishment calls that might come at a time of already heightened stress (Brunnermeier and Pedersen, 2009; Heath et al., 2016). From the standpoint of systemic risk, it is important to also look at the potential for systemic risk propagation resulting from such assessments.

To gauge this potential, we construct the variable $Cost$, calculated as each bank's cost of mutualizing such residual risk by replenishing portions of its default fund consumed by other banks' defaults. It is constructed as the bank's share of the uncovered marketwide shortfall, weighted by its share of initial margin. Thus,

$$Cost_{i,t} = a_{i,t} \left(\sum_{j \neq i} \sum_k SPS_{j,k,t} - MR_{j,t} - DF_{j,t} \right), \quad (13)$$

where $a_{i,t}$ is bank i 's 60-day share of total initial margin determining its default fund exposure.

This metric is especially important during a crisis, as excessive assessments can amplify procyclicality and destabilize banks that are already close to becoming unstable (Brunnermeier and Pedersen, 2009), or else become progressively difficult to collect, threatening the centrally cleared market's viability (Raykov, 2018). Such risk-propagation capabilities remain understudied; to our knowledge, this is the first paper in the literature to study spillovers resulting from simultaneous partial defaults representing systemic risk.

5 Expected Shortfall Analysis Using MES

In this section, we analyze the systemic payment shortfalls computed from daily-frequency position and MES data, based on the conservative position depreciation rates ranging from 25% (in Section 5.2.3) to 100% in the baseline (worst-case scenario) analysis. Section 6 validates this collateral adequacy analysis using lower-frequency conditional default probabilities (tail betas) and endogenous haircut rates estimated with extreme value theory.

5.1 Systemic risk prior to the application of collateral and its composition during the crisis

Our first set of results analyzes the composition of systemic risk in the Canadian futures market prior to the application of collateral. To find out which firms persistently drive systemic risk in this market, we focus on quarterly averages for the gross systemic risk metric, systemic market shortfall (SMS). The SMS metric and its composition are presented separately for proprietary bank trades, trades submitted on behalf of bank customers, and total trades in Figures 3, 4, and 5, Panel A. Figures 6, 7, and 8 repeat this decomposition by country. Due to data disclosure requirements, we do not identify the financial institutions involved, but instead assign them with arbitrary codes ranging from Bank 1 to Bank 14. Finally, since the magnitudes represent the expected systemic shortfalls gross of any collateral pledged, we do not focus on the absolute shortfall size, but rather, on its composition; the discussion of magnitudes requires appropriate calibration of MES and takes place in Sections 5.2.2 and 5.2.3.

Figure 3(A) presents the average systemic risk brought on by proprietary trades before collateral. The figure shows that the SMS metric not only accurately captures the 2008 financial crisis, whose worst quarter according to broad consensus was 2008:Q4, but importantly, also shows the buildup of risk leading to it. The pre-crisis period can be divided into two subperiods: from 2003:Q1 to 2005:Q2, and from 2005:Q4 to 2008:Q1. The first subperiod can be taken as representative of “normal times,” and features low expected payment shortfalls consistent with the risk levels immediately after the crisis abated. The second subperiod shows two buildups of systemic risk in the years just prior to the crisis: from 2005:Q4 to 2006:Q2, and another buildup during 2007:Q2 to 2007:Q4. While the earlier buildup (2005:Q4 to 2006:Q2) is driven almost entirely by Canadian institutions, the second buildup is due to both Canadian and US institutions in roughly equal proportion; this can be seen in Figure 6(A). With the onset of the 2008 crisis, however, the composition of systemic risk again shifts back to Canadian banks. Since, during this time, American banks had substantially higher MES (roughly double the size of Canadian banks – see Figure 2), this composition shift implies that US banks responded to the crisis by flattening out their proprietary trade positions in Canadian futures until the worst was over, whereas Canadian banks remained active. Consistent with this, we observe a return of American banks in

2009:Q2, where one of the American institutions contributes most of the systemic risk from proprietary trades.

Comparing this picture to the trades submitted on behalf of bank customers, we observe a different pattern in Figures 4(A) and 7(A). The two risk buildups before the crisis are still visible, but they are both dominated by Canadian bank customers, with only one American member contributing to the second buildup. With the onset of the crisis, however, the main risk contribution increasingly shifts to US customers until they become responsible for over 80% of the systemic risk in 2009:Q2. Since in that quarter, the American institutions' MES is already on the decline (from the 20-25% range down to the 10-15% range), the increase in the marketwide systemic shortfall from that quarter to the next can be attributed only to increasing open positions. For example, the largest US contributor of systemic risk in 2008:Q4 and 2009:Q1 owes its risk increase to large customer positions in bankers' acceptance notes futures (BAX) and Canadian standard index futures (SXF). Thus, while American banks flattened out or reduced long positions in proprietary Canadian futures, their clients not only did not follow the same investment strategy, but sometimes even increased activity, either trying to hedge at the height of the crisis or profit from it.

Due to this expansion in activity, customer-driven systemic risk followed a different time pattern, and peaked in 2009:Q2, a quarter later than the risk from proprietary trades. While proprietary trade risk increases from the beginning to the end of 2008 and then abates, the risk brought on by customers begins with a pre-crisis buildup as early as 2006:Q2. Within the 2007-08 period, systemic risk from customer trades sets in earlier by nearly half a year compared to proprietary trades (in 2007:Q3, compared to 2008:Q2). Thus systemic risk brought on by customers both starts earlier and tapers off more slowly than risk from proprietary trading. Since according to our conservative methodology, customers are assigned their clearing bank's MES,¹⁵ this difference is also entirely due to the size of open positions. Hence, we conclude that there are systematic differences in the behavior of bank customers and banks themselves in the Canadian futures market along at least three dimensions: the composition of risk buildup (Canadian vs. US entities); response to the onset of the crisis (flattening out vs. increasing positions); and the timing of market activity, leading to different peaks of systemic risk (earlier onset and later tapering off for customer trades).

¹⁵Because the risk posed by a customer is only as large as that posed by its clearing bank, which legally guarantees performance on the trade.

To appreciate the extent to which customer trades shape the Canadian futures market, we compare the size of the shortfalls in Figures 3(A) and 4(A) with Figure 5(A), showing the total systemic risk from all trades (both proprietary and submitted on behalf of customers). Figure 5 shows that the time pattern of total systemic risk in the futures market is dominated by that of customer trades rather than proprietary trades. This is consistent with the fact that aggregate customer positions are larger than the proprietary ones of the clearing bank, sometimes by a large factor (on average, 7:1 for BAX contracts, 4:1 for CGB contracts, and 2.5:1 for SXF contracts, based on Table 4 over the sample years). This implies that the patterns associated with customer trading are also essential in shaping the composition and peaks of systemic risk market-wide; moreover, these patterns happen to be opposite to the behavior of the big clearing banks overseen by Canadian regulators.

5.2 Systemic risk after the application of collateral

In our next set of results, we present the net systemic market shortfalls left after the application of the potential defaulter’s collateral. This requires additional calibration of the MES risk indicator to make sure that risk and collateral are measured on a comparable scale. As documented by Acharya et al. (2017), MES is an accurate predictor of distressed institutions during financial crises, and is therefore a strong correlate of a bank’s probability of default conditional on sector-wide distress; however, MES is not a probability itself. To generate expected shortfalls in dollar terms, MES needs to be replaced by this conditional probability. However, since the latter is unobserved, we instead pursue a different approach. Since we are interested in the maximum shortfall after collateral, we instead ascertain that the peak observed MES (times a large safety factor) bounds the peak conditional default probability from above; as long as this is the case, our risk spillover analysis remains valid. To this end, we test out several sets of assumptions in Sections 4 and 5.2.3, including the assumption that MES understates the maximum conditional default probability by a factor of 4 times, which corresponds to a maximum default probability very close to 1 – the theoretical upper bound. Then we conduct all baseline analysis based on the assumptions producing the largest maximum shortfall, \$524.4 million.

Having addressed potential issues of scale, next we buffer the systemic payment shortfalls with two tranches of collateral: the conventional industry-standard SPAN margin, simulated as explained in Section 3.3, and each institution’s default fund deposit. Since we are inter-

ested in the maximum risk present after collateral, in this section we focus on the quarterly maxima of the net systemic market shortfall as computed according to equation (12). We present the results for proprietary, customer, and total pooled trades in Panels B of Figures 3-5. For continuity with the previous section, we first discuss the composition of risk and then its magnitude, as well as the possibility of contagion.

5.2.1 Risk composition

Figure 3(B) shows the maximum net systemic market shortfalls (*NetSMS*) left after the application of transaction-specific and non-transaction-specific collateral. The figure indicates that post-collateral shortfalls arising from proprietary trades still reflect the financial crisis as it occurs, but are less useful than gross SMS in predicting it as a lead indicator. Prior to the second quarter of 2008, there are virtually no signs of an approaching systemic crisis, as the entire expected risk buildup gets netted against posted collateral, which nullifies the exposure. From this, we conclude that out of the two measures considered, the gross payment shortfall is better as a predictor of crises, but less indicative of the actual risk present, while the net payment shortfall is more indicative of actual risk of contagion once the crisis begins.

Figure 6 shows the composition of this residual unmitigated risk from the proprietary trades of Canadian vs. US contributors. In contrast to the gross risk analysis, where the biggest risk contributors during the crisis were Canadian banks, here, the contributions are split roughly equally between American and Canadian institutions. In the third quarter of 2008, American institutions make up for about two-thirds of the market shortfall; in the fourth quarter, however, Canadian institutions make up for more than 50% of the risk. These results are consistent with the conclusions of Cruz Lopez et al. (2017), who find that conventional margins over-margin most participants except for a very few but large risk contributors. Starting with 2009:Q1, systemic shortfalls after collateral quickly abate to their near-zero, pre-crisis levels.

The picture changes again when we consider customer trades. As shown by Figure 4(B), customer positions are the source of most of the unmitigated systemic risk in the market (in a ratio of about 5:1 to proprietary trade risk – compare Figures 3(B) and 4(B)). Moreover, Figure 7(B) shows that virtually the entire systemic net shortfall associated with customer trades comes from clients of American institutions. The picture in Total accounts, reflected

by Figures 6 and 8, is similar and implies that almost 100% of the Canadian market’s *unmitigated* systemic exposure came from the US.¹⁶ This is a stark finding, because it shows a drastic composition change before and after the application of collateral. Whereas US and Canadian entities contribute roughly equally to the pre-collateral risk distribution, the unmitigated risk post-collateral comes exclusively from the US.

This clearly demonstrates that conventional margins and default fund deposits do not proportionally buffer risks from banks with different systemic exposures. If that buffering was pro-rata, we would see no change in the relative composition of risk; therefore, one of the key effects of conventionally computed collateral is distributional: it shifts the systemic risk’s composition and buffers that risk unequally, since it ignores risk comovements. This also alters the distribution of risk over time: while gross systemic risk peaks in 2008:Q4 in Figure 5(A), the peak of net systemic risk in Panel B shifts to 2008:Q1. By comparing the two panels of Figure 5, we conclude that this is largely due to the unequal risk mitigation of Bank 7, which experienced a large increase in MES during that quarter, compared to the risk brought on by other banks. By contrast, the relative risk brought on by other banks is adequately reduced – e.g. Bank 3’s shortfall is reduced in half (from about \$100 million to \$50 million) by collateral. Thus, it appears that systemic institutions with quickly evolving MES are not always “captured” by classical margining systems not designed to cope with systemic risk.

5.2.2 Contagion

The risk analysis in the previous section showed that the maximum systemic market shortfall after the application of collateral comes up to \$524.4 million, as reflected in Figure 5(B). Turning to the adequacy of conventional collateral requirements, it is logical to ask to what extent this number is significant – both relative to the size of the market and the assumptions we are making, and whether it contributes to any significant amount of contagion in this market. To that purpose, we look at the costs for market participants in absorbing that shortfall resulting from the specific risk-sharing arrangements present in centrally cleared markets. This cost provides a metric of the contagion expected to occur as a result of the cumulative shortfall, because it occurs on top of any financial strain the institution may already be experiencing. As our baseline calculation indicates, even though \$524.4 million

¹⁶This does not imply that all customers of US institutions were American entities themselves.

may appear like a large number, it does not result in any significant amount of contagion even under the most unfavorable assumptions, so conventionally computed collateral is adequate in containing the crisis systemic risk, despite its distributional effects.

As discussed in Section 2, centrally cleared markets differ from bilateral markets in that losses in excess of collateral are mutualized across market participants; this occurs through consuming the survivors' contributions to the market-wide default fund. This process is controlled by the central counterparty serving the market, which administers the collateral posted for every transaction in the form of margin and default fund deposits. The loss-sharing rules for default losses in excess of the defaulter's collateral are specified in the central counterparty's rulebook. In the Canadian futures market, such excess default losses are shared among banks proportional to their share of initial margin, so the cost for each bank is calculated according to the formula

$$Cost_{i,t} = a_{i,t} \left(\sum_{j \neq i} \sum_k SPS_{j,k,t} - MR_{j,t} - DF_{j,t} \right), \quad (14)$$

where $a_{i,t}$ is bank i 's 60-day share of initial margin.

To size the cost relative to the size of the market and the bank bearing it, we also calculate it as a fraction of the bank's market capitalization. Since most of the market-wide exposure is driven by customer trades, the cost numbers reported below are an upper bound, made on the assumption that customers do not reimburse their clearing bank for their portion of the cost, and that the banks' maximum risk contributions over the quarter all occur on the same day; this is unlikely in practice, so the cost per bank is likely to be lower.

Table 5 displays the *maximum* cost per member resulting from having to share the cost of the peak market-wide shortfall from total accounts of \$524.4 million (in 2008:Q1). The table calculates the loss-sharing costs using two methods. Method 1 divides the shortfall among the sampled banks according to each one's share of initial margin in the total margin required of sampled banks. Method 2 assumes banks with a positive contribution to the net shortfall are unable to make further payments, and so results in higher pro-rata costs for remaining survivors; this is the method used by the CDCC in the event of member default. Both methods show that the maximum cost of risk spillovers above collateral are insignificant relative to the banks' size and financial strength.

The maximum dollar cost per bank calculated according to Method 1 is \$121 million, or

0.19% of the bank’s market capital. Method 2 gives slightly higher amounts of \$159 million, or 0.25% of the capital base. Looking at the cost as a fraction of capital, using Method 1, the highest cost is borne by Bank 7, at 1.01% of its market capital, or Bank 13, using Method 2, at 0.81% of its market capital.

Scaling the costs of risk sharing relative to market capitalization provides a good intuitive way to understand their magnitude. Given a fixed number of shares outstanding, a 1% drop in market capitalization is equivalent to a 1% price drop; thus, none of the costs of systemic risk shortfalls in excess of existing collateral exceed the impact of a 1.01% downward price move. Moreover, this interpretation also allows us to link the costs back to each bank’s MES. Since MES can be equivalently interpreted as the drop in each bank’s market capitalization conditional on a crisis (see Section 4), additional costs of risk sharing totaling maximum 1% of each bank’s market capital cannot further increase the bank’s MES by more than 1 extra percentage point. These calculations are made assuming that each clearing bank bears the entire risk-sharing cost on behalf of its customers, something unlikely to hold in practice; thus, the actual increase in systemic risk is likely to be even lower.

This evidence convincingly shows that even under extreme assumptions aiming to artificially magnify risk spillovers in excess of collateral, systemic risk during the 2008 crisis did not result in significant spillovers that could have further worsened the crisis. Thus, conventional collateral requirements’ inability to address risk comovements across banks appears to be largely offset by their conservative treatment of all losses as potential defaults. However, as shown in Section 5.2.1, this occurs at the cost of unequal buffering of systemic risk across banks, with the highest spillover risk generated precisely by those institutions already holding the most systemic risk: the US institutions hardest hit by the crisis. To make sure that our findings are not driven by scaling or calibration issues, the next subsection repeats this analysis under alternative assumptions.

5.2.3 Spillovers under alternative scenarios

To show that systemic risk spillovers above conventional collateral requirements are small, the above analysis assumed that the resale value of the defaulted contract is zero. Although default auctions do not always fully recover position value, this is still an extreme assumption that amplifies the systemic shortfalls beyond what is ‘extreme but plausible’. Here we consider two alternative scenarios against the event that MES is biased downwards relative

to the true conditional probability of default at its peak. We find that the alternative scenarios result in cost spillovers that are either the same as or smaller than in the baseline scenario.

First we dispense with the zero resale value assumption, and instead assume a more realistic benchmark recovery rate of 75% of the value of each unpaid obligation. As can be seen from the contract haircuts estimated with extreme value theory in Section 6 (Table 8), this assumption is still quite extreme. It corresponds to a price drop occurring more rarely than once in 200 years for the BAX and CGB contracts and once in 50 years for the SXF contract. However, for now we take it for granted, and reserve the analysis with these haircuts and actual default probabilities for Section 6. In the current setting, Alternative scenario 1 further assumes that the observed MES understates the true conditional probability of default by a factor of 2, and Alternative scenario 2 by a factor of 4 times. The latter calibration is based on the fact that the maximum MES observed in-sample is 0.24, which, if understated by a factor of 4, would bring the true default probability up to 0.96; therefore it can be seen as a calibration to match the highest observed MES to a conditional probability of default during a crisis close to 1.

Panels B and C of Table 5 display the results of these calculations; we apply the same cost calculation methods 1 and 2 as in Section 5.2.2. The market-wide shortfalls calculated with these alternative scenarios are either smaller than or the same as in the baseline scenario: \$187 million using Alternative Scenario 1, and \$523.8 million with Scenario 2 (compared to the \$524.4 million in baseline calculation). Consequently, they result in similarly small risk spillovers across banks, regardless of the exact way they are distributed.

Looking at the spillovers as a fraction of the bank's market capitalization, the maximum numbers remain in the same range as before: 0.81% using the (more accurate) Method 2, or 1.01% using the less realistic Method 1. As in the previous section, drops in market capital of this magnitude occur frequently as a result of daily price moves without causing investor concerns, and by the same logic as in the previous section, cannot increase the respective banks' MES by more than approximately 1 percentage point. Recalling that these per-bank cost calculations further assume that clients (who are responsible for the bigger share of market activity) do not share in the cost and leave all payments to their respective clearing institutions, it seems almost certain that the actual costs borne by the banks and their increase in systemic risk will be lower. This confirms the validity of the baseline analysis.

6 Analysis with Conditional Default Probabilities

The previous sections relied on the MES metric to gauge systemic risk. Here, we repeat the collateral adequacy and contagion analysis with the conditional co-crash probability, or ‘tail beta’, commonly used as a proxy for a bank’s conditional probability of default (see Straetmans et al., 2008, De Jonghe, 2010, or Perez-Saiz and Li, 2018). While MES and tail betas capture different aspects of systemic risk, they agree on the paper’s main point: that conventional collateral requirements buffer systemic risk adequately and generate insignificant risk spillovers.

The literature has used the term ‘tail beta’ differently depending on context. Following De Jonghe (2010) and Perez-Saiz and Li (2018), we define tail beta as the probability that a bank’s stock return falls in its lowest 5% quantile, conditional on the S&P/TSX Composite Banks Index return falling within its worst 5% (or equivalently, that the bank’s return $loss\ X$ exceeds its 95th percentile conditional on the index loss Y doing the same). Unlike MES, a tail beta does not capture the extent to which an individual bank stock drops in a crisis, as long as it is below the 5% threshold; thus, tail betas can sometimes attenuate differences between differently affected banks. For example, the steep price drops experienced by Bank 5 in March 2008 and Bank 7 in December 2008, well captured by MES, would not count much towards a higher tail beta, because they are infrequent enough to raise the co-crash probability by much. By construction, tail betas count as most systemic those banks that *repeatedly* fall within the same lowest-return quantile. We therefore use tail beta analysis to complement, rather than replace, our MES analysis.

Relative to MES, tail betas offer the advantage that they are probabilities and can be applied directly to position values to compute expected shortfalls in dollars. Unlike MES, however, they are estimated with extreme value methods requiring a minimum of 5 to 6 years of daily data to produce a single point estimate. Due to the insufficient length of some entities’ price series to produce rolling estimates, instead we estimate a single tail beta for each bank over the entire sample period and then apply the estimates to the crisis years (2007-2010) in calculating systemic shortfalls.¹⁷

We also dispense with the ad hoc position depreciation rates, and use extreme value theory to predict haircuts for each contract type, corresponding to a price drop of severity

¹⁷The results are very similar when using the maximum tail betas for banks permitting rolling estimates.

once in 1 month and once in 10, 50, 100, and 200 years. We then compute the bank-wide and market-wide shortfalls net of collateral, assuming that a defaulted position’s resale value drops by the above haircut, and calculate the cost of mutualizing any shortfalls in excess of collateral as we did in Section 5. To avoid redundancy, we focus only on Total accounts as the most reflective of the collateral adequacy issue.

Table 6 lists the estimated tail betas for the banks in our sample. They provide a different picture of systemic risk than MES, because the two metrics capture different systemic risk aspects. According to tail beta estimates, the three most systemic institutions are Banks 12, 1, and 3 (in that order), whereas MES ranks them as Banks 7, 5, and 14. The reason for this difference is that the latter banks had the deepest price drops conditional on a crisis, while the former ones had the most frequent joint price drops to their bottom 5%. Given the different risk aspects captured by these two metrics, it is of interest whether the collateral adequacy analysis survives when redone with tail betas.

To recalculate net payment shortfalls, we therefore apply the estimated tail betas from Table 6 to each bank’s end-of-day payment obligation $PO_{i,k,t}$ for each contract, prorated by the contract-specific depreciation rate (haircut) d_k estimated with extreme value theory in Panel A of Table 8. The magnitude of d_k is obtained using the inverse of equation (6) according to the formula $\hat{q} = Z_{m-n,n}(m/pn)^{1/\hat{\alpha}}$.

Table 8 estimates a once-in-10-years price drop for the CGB contract at 2.1%, and a once-in-100-years price drop at 4.0%.¹⁸ This is economically significant compared to the CGB mean return of 0.02% and standard deviation of 0.35%. By contrast, the more volatile and less liquid SXF contract has a once-in-10-years price drop estimated at 12.5%, and a once-in-100-years price drop equal to 30.5%, which is again economically significant given its mean return of 0.05% and standard deviation of 1.27%.

The maximum net payment shortfalls for each bank during 2007-2010 are shown in Panel B of Table 8 together with the market total (the net systemic market shortfall). Table 8 shows scenarios corresponding to position value drops with severity once in a month and once in 10, 50, 100, and 200 years, corresponding to probabilities of 5%, 0.04%, 0.008%, 0.004%, and 0.002%, respectively. As can be seen, despite the relatively high tail beta values in Table 6, losses above collateral emerge only beginning with the once-in-50-years scenario,

¹⁸With daily data, the once-in-10-year drop corresponds to a probability of 0.04%, and the once-in-a-century drop to a probability of 0.004%.

where the SXF contract price drops by more than 23%.

Banks 2, 5, and 12 generate losses across all scenarios with severity higher than once-in-50-years, and are joined by Bank 3 only in the once-in-200-years scenario. The largest expected losses above collateral accrue to Bank 5, followed by a combination of Banks 2 and 12, depending on the scenario. Bank 12 is the one with the highest tail beta, so its collateral breaches are unsurprising; the losses of Banks 5 and 2 (which are not among the three most systemic according to Table 6) reflect a combination of higher exposures to the SXF contract and/or lower margin requirement for these banks relative to the estimated position depreciation rate d_k . The total net systemic market shortfalls generated starting with the 50-year scenario are respectively \$30 million, \$67.9 million, and \$131.9 million. The composition of risk across these scenarios is roughly constant and reflects the composite influence of four factors: the conditional probability of default in a crisis, as proxied by each bank's tail beta; the 2-day position depreciation, assumed to follow default on a payment obligation; the bank's position value in each contract; and the value of bank collateral, as calculated by SPAN initial margins and default fund deposits.

Compared to the previous MES analysis, the maximum total market shortfall (\$131.9 million) is four times smaller than the baseline MES scenario from Table 5. Given that tail betas are, on average, higher than crisis-level MES estimates (compare Tables 2 and 6), this is explained by the more realistic position depreciation rates, which, especially for BAX and CGB contracts, do not even approach the assumed 25% or 100% depreciation used in that calculation. Therefore, extreme value theory analysis confirms the ample conservatism of our previous approach and the robustness of our results.

We size the \$131.9 million shortfall relative to the market and repeat the contagion analysis in Table 9, which shows the per-bank cost of sharing losses above collateral. As can be seen from Table 9, once haircuts are computed realistically, the possibility of contagion is even smaller than previous analysis suggests. As shown in Panel C of Table 9, the largest dollar cost to a bank as a result of others' defaults is \$30.5 million, and the largest cost as percentage of a bank's market capital is 0.25% to 0.35% (depending on the cost allocation method). Despite the higher tail betas, both of these costs are lower than the maximum per-bank costs of \$121 million and 1.01% in our MES analysis, once again highlighting its conservatism. Therefore, despite the fact that MES and tail betas measure systemic risk differently, the two metrics fully agree on the central point of this study: that the

inbuilt conservatism of traditional margin and default fund requirements largely offsets their inability to deal with systemic risk. Similarly, both analyses imply minimal if any increase in systemic risk due to the cost of loss-sharing in centralized derivatives markets.

7 Conclusion

Many derivatives markets use collateral requirements calculated with industry standard, but dated methodologies not designed with systemic risk in mind. This has spurred recent interest both in better systemic risk measurement across financial market infrastructures (Perez-Saiz and Li, 2018) and in new margining systems, such as CoMargin (Cruz Lopez et al., 2017). The present paper contributes to this literature by exploring the adequacy of traditional collateral requirements against systemic risk during the financial crisis of 2008, using evidence from the Canadian futures market. On the one hand, conventional collateral requirements conservatively provision for future losses as for potential defaults, aiming to cover them with a high degree of statistical confidence; but on the other, they do not recognize that defaults may be interdependent. We explore whether the inbuilt conservatism of traditional collateral-determining models sufficiently defends them against their inability to account for systemic risk.

To answer this question, we calculate a new systemic risk metric, the expected systemic market shortfall, and analyze its composition across firms both prior to and after the application of conventional collateral requirements calculated with the SPAN method. Our results show that the conservatism of standard methodologies serves well in buffering shortfalls even from interdependent defaults and results in small expected spillovers above collateral, which do not meaningfully add on to banks' pre-existing systemic risk levels. We verify the robustness of this result with two alternative systemic risk measures, MES and tail beta, and explore alternative calibrations of the MES metric relative to the size of the market, some allowing for an implausibly large margin of error. Even under the most extreme scenario, the maximum market-wide shortfall in excess of collateral barely reaches 1% of the banks' market capitalization, and hence does not increase their MES by more than a single percentage point.

At the same time, we also find that traditional collateralization methods result in unequal buffering of systemic risk across banks, with the highest systemic risk contributors

being buffered relatively less than the rest, and therefore being the most likely ones to breach posted collateral. Despite the fact that the amount of breach may be insignificant relative to the market, this strengthens the case for joint margining systems such as CoMargin. At the same time, the paper also argues against the mechanical application of simple P&L correlations as a systemic risk measure for derivatives markets. As argued in the Introduction, correlated losses that are honored in full pose no risk to the market, systemic or otherwise; what really matters is whether such correlated losses *are honored when they occur*. This paper offers a simple and elegant way to reflect this by pro-rating banks' outstanding payment obligations with market-based systemic risk indicators, and calibrating the resulting maximum shortfall. Our results confirm our prior that a mechanical application of P&L correlation-based methodologies overstates the amount of systemic risk, and the minimal collateral breaches we find confirm this view. At the same time, our distributional findings agree with those by more traditional methodologies (Cruz Lopez et al., 2017).

Finally, the paper also contributes to the better understanding of systemic risk sources in derivatives markets. According to our findings, systemic risk brought on by bank customers, rather than direct market members themselves, accounted for the bulk of systemic risk in Canadian futures during the 2008 crisis, and shaped it in important ways. Firstly, systemic risk from customer trades started earlier and tapered off later than that of proprietary bank trades, suggesting it could be a potential lead indicator of future systemic crises. Secondly, we find that systemic risk peaked at different times for banks versus their customers, and that these two participant groups behaved differently during the crisis, with American banks reducing their long positions or flattening out in Canadian futures, and their customers increasing activity and positions instead, thus ultimately bringing the total risk up. Finally, we document large cross-country differences in the behavior of Canadian and US institutions during the crisis, which significantly affected the composition of risk in the Canadian futures market.

References

- [1] Acharya, V. V., Pedersen, L. H., Philippon, T., and Richardson, M. (2017). Measuring Systemic Risk. *Review of Financial Studies*, 30(1) pp. 2–47.
- [2] Adrian, T., and Brunnermeier, M. (2016). CoVaR. *American Economic Review*, 106(7), pp. 1705–1741.
- [3] Artzner, P., Delbaen, F., Eber, J. M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance* 9(3), pp. 203–28.
- [4] Berger, A., and Davies, S. (1998). The Information Content of Bank Examinations. *Journal of Financial Services Research*, 14(2), pp. 117–144.
- [5] Brunnermeier, M., and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies* 22, pp. 2201–2238.
- [6] Campbell, B. and Chung, C. (2003). CGB: Poised for Takeoff. An Analysis of the Ten-Year Government of Canada Bond Future Based on Intraday Trading Data. Working Paper, CIRANO.
- [7] Chicago Mercantile Exchange (2019). CME SPAN: Standard Portfolio Analysis of Risk.
- [8] Cole, R., and Gunther, J. (1998). Predicting Bank Failures: A Comparison of On- and Off-Site Monitoring Systems. *Journal of Financial Services Research* 13, pp. 103–117.
- [9] CPMI-IOSCO. (2012). Principles for Financial Market Infrastructures. Bank of International Settlements: Basel.
- [10] Cruz Lopez, J., Harris, J., Hurlin, C., and Perignon, C. (2017). CoMargin. *Journal of Financial and Quantitative Analysis*, 52(5), pp. 2183–2215.
- [11] De Haan, L., and Ferreira, A. (2006). *Extreme Value Theory: an Introduction*. Springer Verlag.
- [12] De Haan, L., Jansen, D., Koedijk, K., De Vries, C. (1994). Safety First Portfolio Selection, Extreme Value Theory and Long Run Asset Risks. In: Galambos, J., Lechner, J., Simiu, E. (eds.), *Extreme Value Theory and Applications*. Kluwer Academic Publishers, Dordrecht, pp. 471–487.
- [13] De Jonghe, O. (2010). Back to the Basics in Banking? A Micro-analysis of banking System Stability. *Journal of Financial Intermediation* 19, pp. 387–417.
- [14] Duffie, D., and Zhou, H. (2011). Does a Central Clearing Counterparty Reduce Counterparty Risk? *Review of Asset Pricing Studies* 1(1), pp. 74–95.
- [15] Flannery, M. (1998). Using Market Information in Prudential Bank Supervision: A Review of the U.S. Empirical Evidence. *Journal of Money, Credit and Banking* 30(3), pp. 273–305.

- [16] Glasserman, P., Moallemi, C., Yuan, K. (2015). Hidden Illiquidity with Multiple Central Counterparties. Office of Financial Research Working Paper 15-07.
- [17] Heath, A., Kelly, G., Manning, M., Markose, S., Rais, A., and Shaghaghi, A. (2016). CCPs and Network Stability in OTC Derivatives Markets. *Journal of Financial Stability* 27, pp. 217–233.
- [18] Hill, B. (1975). A Simple General Approach to Inference about the Tail of a Distribution. *Annals of Statistics*, 3, pp. 1163–1173.
- [19] Huang, W., and Menkveld, A. (2016). Systemic Risk in Real Time: A Risk Dashboard for Central Clearing Parties (CCPs). SSRN Working Paper.
- [20] Menkveld, A. (2016). Systemic Risk in Central Clearing: Should Crowded Trades Be Avoided? SSRN Working Paper.
- [21] Paddrick, M., Rajan, S., and Young, P. (2016). Contagion in the CDS Market. Office of Financial Research Working Paper 16-12.
- [22] Paddrick, M. and Young, P. (2017). How Safe are Central Counterparties in Derivatives Markets? Office of Financial Research Working Paper 17-06.
- [23] Perez-Saiz, H., and Li, F. (2018). Measuring Systemic Risk across Financial Market Infrastructures. *Journal of Financial Stability*, 34, pp. 1–11.
- [24] Raykov, R. (2018). Risk Mutualization and Financial Stability: Recovering and Resolving a Central Counterparty. *Journal of Financial Market Infrastructures*, 6(2/3), pp. 163–185.
- [25] Straetmans, S., Verschoor, W., and Wolff, C. (2008). Extreme US Stock Market Fluctuations in the Wake of 9/11. *Journal of Applied Econometrics*, 23, pp. 17–42.
- [26] TMX Montreal Exchange. (2013a). BAX Three-Month Canadian Bankers’ Acceptance Futures Descriptive Brochure. TMX Group.
- [27] TMX Montreal Exchange. (2013b). *Government of Canada Bond Futures and Options on Futures Reference Manual*. TMX Group: Montreal.
- [28] TMX Montreal Exchange. (2013c). *Index Derivatives Reference Manual*. TMX Group: Montreal.
- [29] TMX Montreal Exchange. (2017). *Canadian Futures and Options*. TMX Group: Montreal.

8 Tables

Table 1: List of Financial Institutions

Name	Country of Headquarters
Bank of Montreal	Canada
Bank of Nova Scotia	Canada
CIBC	Canada
CreditSuisse	Switzerland
Goldman Sachs	USA
HSBC	UK
J.P. Morgan	USA
Laurentian Bank	Canada
Merrill Lynch	USA
MF Global	USA
National Bank	Canada
Royal Bank of Canada	Canada
Toronto Dominion (TD)	Canada
UBS	Switzerland

This table presents the publicly traded financial institutions participating in the Canadian futures market during the sample period (January 2, 2003, to March 31, 2011) and their country of headquarters. Subsidiaries are subsumed under the parent institution and their Canadian futures positions consolidated with those of the parent if participating through more than one entity.

Table 2: Summary Statistics: MES

	Mean	Std. Dev.	Min	Max
Bank 1	0.0217	0.0171	0.0015	0.1035
Bank 2	0.0208	0.0167	0.0012	0.0978
Bank 3	0.0211	0.0158	1.49×10^{-6}	0.0796
Bank 4	0.0174	0.0216	0	0.1149
Bank 5	0.0254	0.0398	0	0.2337
Bank 6	0.0111	0.012	0	0.0576
Bank 7	0.0423	0.0481	0	0.2395
Bank 8	0.0213	0.0259	0	0.1785
Bank 9	0.0152	0.02	0	0.1199
Bank 10	0.0217	0.017	0	0.0925
Bank 11	0.0104	0.0146	0	0.1000
Bank 12	0.0207	0.0165	0	0.0921
Bank 13	0.0181	0.0152	0	0.0949
Bank 14	0.0234	0.0307	0	0.1930

This table presents summary statistics for the marginal expected shortfall (MES) for publicly traded financial institutions participating in the Canadian futures market during the sample period (January 2, 2003, to March 31, 2011). All statistics are based on publicly available price data using the methodology of Acharya et al. (2017).

Table 3: Summary Statistics: Total Accounts

	Position Size						Position Value					
	BAX		CGB		SXF		BAX		CGB		SXF	
Year	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Total Accounts												
2003	31.7	18.5	6.0	6.2	0.9	0.9	3070.0	1790.2	647.7	672.2	380.2	397.3
2004	41.0	23.2	9.1	8.5	2.4	1.6	3995.2	2264.7	1002.9	943.3	1156.6	779.3
2005	45.4	29.2	9.2	9.3	2.4	1.6	4391.5	2826.0	1051.4	1057.0	1354.1	892.1
2006	69.9	58.8	15.2	14.3	1.6	1.5	6690.6	5621.0	1715.7	1620.6	1107.4	991.3
2007	69.0	50.7	11.8	11.2	1.6	1.5	6590.3	4845.6	1334.1	1255.8	1235.6	1178.8
2008	54.5	45.2	10.3	9.8	1.3	1.2	5292.7	4380.9	1223.7	1164.9	949.5	947.5
2009	47.3	29.2	6.3	6.2	1.3	1.3	4689.7	2904.1	772.9	763.5	788.0	767.3
2010	53.4	35.7	8.5	8.7	1.5	1.4	5262.2	3510.8	1043.5	1060.2	1039.0	996.7
2011	76.0	51.8	9.6	8.2	1.3	1.4	7465.2	5093.5	1159.2	989.1	1040.4	1076.8

Summary open position statistics for in-sample members across sample years and three contract types. The table displays the means and standard deviations of open position size (millions of contracts) and value (millions of Canadian dollars) per contract family and year for each member's total account. Position size and value are shown in absolute terms regardless of direction. The sample period is from January 2, 2003, to March 31, 2011. The in-sample members are listed in Table 1.

Table 4: Summary Statistics: Firm and Customer Accounts

Year	Position Size						Position Value						
	BAX		CGB		SXF		BAX		CGB		SXF		
Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Panel A: Firm Accounts													
2003	5.4	6.8	3.4	6.4	0.8	0.9	523.0	656.5	367.7	700.7	330.4	374.6	
2004	5.7	8.8	5.6	10.1	1.2	1.1	556.5	858.5	621.1	1108.1	575.1	543.6	
2005	5.9	9.2	4.2	6.7	1.1	0.9	569.4	895.9	474.4	753.4	638.5	489.8	
2006	5.8	9.6	2.8	3.3	0.8	0.8	554.7	922.0	316.0	373.1	511.8	554.3	
2007	7.7	14.3	1.6	1.8	0.5	0.6	732.9	1361.4	184.6	205.5	430.1	462.3	
2008	5.5	6.3	2.0	2.5	0.7	0.9	533.1	613.6	233.9	298.1	573.0	786.1	
2009	11.2	12.6	3.0	4.6	0.4	0.5	1110.5	1255.2	359.8	562.2	209.5	280.6	
2010	9.3	10.8	2.8	4.5	0.6	0.6	914.3	1066.2	341.3	550.6	395.6	450.6	
2011	13.4	20.2	2.1	2.4	0.5	0.5	1316.1	1987.7	258.4	286.2	382.0	373.4	
Panel B: Customer Accounts													
2003	31.9	19.0	4.2	4.4	1.0	1.3	3093.4	1836.7	453.1	471.8	424.8	561.4	
2004	40.6	23.0	7.0	6.0	2.0	1.6	3949.0	2236.3	767.2	665.0	972.9	778.5	
2005	46.2	29.7	7.8	8.1	2.0	1.6	4476.1	2881.4	886.5	920.5	1146.7	943.7	
2006	67.9	58.0	16.8	15.1	1.5	1.4	6499.4	5546.7	1893.6	1707.7	1046.0	919.5	
2007	66.6	50.3	13.8	11.3	1.6	1.4	6362.9	4814.4	1553.2	1273.8	1268.8	1098.3	
2008	55.1	45.5	10.3	10.3	1.4	1.2	5352.7	4408.2	1222.4	1219.8	1067.9	992.5	
2009	43.8	27.7	5.7	5.8	1.3	1.2	4343.5	2747.4	700.0	712.0	795.3	743.8	
2010	53.0	35.9	8.7	8.9	1.6	1.6	5220.1	3532.6	1056.2	1082.4	1153.7	1112.3	
2011	69.6	55.7	9.7	8.5	1.5	1.4	6836.4	5470.3	1164.5	1020.9	1160.1	1118.3	

Summary open position statistics for in-sample members across sample years and three contract types. The table displays the means and standard deviations of open position size (millions of contracts) and value (millions of Canadian dollars) for firm accounts in Panel A, and for customer accounts in Panel B. Position size and value are shown in absolute terms regardless of direction. Firm and customer account averages may not add up to the respective total account average since offsetting firm and customer positions are netted. The sample period is from January 2, 2003, to March 31, 2011. The in-sample members are listed in Table 1.

Table 5: Maximum Cost of Shared Losses Above Collateral Requirements, Estimated using MES

Bank	1	2	3	4	5	6	7	8	9	10	11	12	13	14
<i>Panel A. Baseline Scenario: Market Shortfall = 524.4 Million</i>														
Cost per bank (method 1)	121.1	36.3	56.2	0	34.8	0.2	34.4	61.0	0	39.6	0	12.8	50.4	77.7
Cost per bank (method 2)	159.1	47.7	0.0	0	0	0.3	0	80.2	0	52.1	0	16.8	66.2	102.1
Cost/Mkt. cap., % (method 1)	0.19	0.08	0.21	0	0.07	0.02	1.01	0.09	0	0.16	0	0.03	0.62	0.05
Cost/Mkt. cap., % (method 2)	0.25	0.10	0	0	0	0.03	0	0.11	0	0.21	0	0.04	0.81	0.07
<i>Panel B. Alt. Scenario 1: Market Shortfall = 187 Million</i>														
Cost per bank (method 1)	43.2	12.9	20.0	0.0	12.4	0.1	12.3	21.8	0.0	14.1	0.0	4.6	18.0	27.7
Cost per bank (method 2)	49.7	14.9	23.1	0.0	0.0	0.1	0.0	25.1	0.0	16.3	0.0	5.3	20.7	31.9
Cost/Mkt. cap., % (method 1)	0.07	0.03	0.07	0	0.03	0.01	0.36	0.03	0	0.06	0	0.01	0.22	0.02
Cost/Mkt. cap., % (method 2)	0.08	0.03	0.09	0	0	0.01	0	0.04	0	0.06	0	0.01	0.25	0.02
<i>Panel C. Alt. Scenario 2: Market Shortfall = 523.8 Million</i>														
Cost per bank (method 1)	120.9	36.2	56.1	0.0	34.7	0.2	34.3	61.0	0.0	39.6	0.0	12.8	50.3	77.6
Cost per bank (method 2)	158.9	47.6	0	0	0	0.3	0.0	80.1	0.0	52.0	0.0	16.8	66.1	101.9
Cost/Mkt. cap., % (method 1)	0.19	0.08	0.21	0	0.07	0.02	1.01	0.09	0	0.16	0	0.03	0.62	0.05
Cost/Mkt. cap., % (method 2)	0.25	0.10	0	0	0	0.03	0	0.11	0	0.21	0	0.04	0.81	0.07

The table displays the per-bank cost of mutualizing the maximum net systemic market shortfall realized in 2008:Q1 across three scenarios and two calculation methods. Panel A displays the baseline scenario, assuming the resale value of the defaulted contract is zero. Panels B and C display two alternative scenarios assuming the resale value of the defaulted contract drops by 25% and the true MES is respectively 2 or 4 times higher than measured. Method 1 divides this maximum shortfall among banks proportional to each one's 60-day share of initial margin as of the midpoint of 2008:Q1, February 15, 2008, regardless of whether the bank contributes to the market shortfall. Method 2 assumes banks already contributing to the net shortfall are unable to pay further, increasing the cost for remaining survivors, and is the one used by the CDCC in the event of default. The cost per bank is reported in millions of CAD; costs as a fraction of market capitalization are reported in percent. Zero costs for some participants may indicate they are not active in the futures market in 2008:Q1.

Table 6: Estimated Conditional Crash Probabilities (Tail β 's)

	$\Pr(X_i > x_i Y > y)$	Tail Index ($\hat{\alpha}$)	Number of Obs.
Bank 1	0.754	1.076	2,075
Bank 2	0.669	0.982	2,075
Bank 3	0.694	1.131	2,075
Bank 4	0.460	1.347	1,984
Bank 5	0.622	1.232	1,465
Bank 6	0.339	1.657	898
Bank 7	0.461	1.098	1,991
Bank 8	0.394	1.272	2,075
Bank 9	0.407	1.242	1,984
Bank 10	0.674	1.108	2,075
Bank 11	0.361	1.125	2,011
Bank 12	0.868	1.194	2,075
Bank 13	0.535	1.098	2,075
Bank 14	0.553	1.192	1,991

This table presents extreme value theory estimates of the probability of each participating bank's stock return loss (X_i) exceeding the 95th percentile of the loss distribution, conditional on the TSX/S&P Composite Banks Index return loss (Y) exceeding its 95th percentile. x_i and y denote the respective 95% percentiles of the two distributions. Stock price returns and index returns are calculated as the daily percentage changes over the sample period (January 2, 2003, to March 31, 2011). The extreme value theory estimation methodology is described in Section 4.1.2.

Table 7: Futures Contracts' Daily Returns

Contract Type	Mean	Std. Dev.	Min	Max	N
BAX	0.003%	0.07%	-0.51%	0.55%	21,798
CGB	0.02%	0.35%	-1.99%	1.67%	2,947
SXF	0.05%	1.27%	-10.20%	9.53%	2,579

This table presents summary statistics for the daily returns of the futures contracts covered in the sample from January 2, 2003, to March 31, 2011.

Table 8: 2007-2010 Estimated Losses Above Collateral Using EVT

Crash Frequency	1/1 month	1/10 years	1/50 years	1/100 years	1/200 years
<i>Panel A: Estimated Position Depreciation Rates (Haircuts)</i>					
BAX	0.1%	0.4%	0.7%	0.8%	1.0%
CGB	0.6%	2.1%	3.3%	4.0%	4.9%
SXF	1.9%	12.5%	23.3%	30.5%	39.9%
<i>Panel B: Net Systemic Payment Shortfall (millions)</i>					
Bank 1	0	0	0	0	0
Bank 2	0	0	5.6	21.0	39.6
Bank 3	0	0	0	0	17.8
Bank 5	0	0	16.9	30.0	46.2
Bank 6	0	0	0	0	0
Bank 7	0	0	0	0	0
Bank 8	0	0	0	0	0
Bank 9	0	0	0	0	0
Bank 10	0	0	0	0	0
Bank 12	0	0	7.5	16.9	28.3
Bank 13	0	0	0	0	0
Bank 14	0	0	0	0	0
Market Total	0	0	30.0	67.9	131.9

This table presents the maximum net systemic payment shortfall by firm over 2007-2010 using constant conditional default probabilities for the financial institutions from Table 6. Panel A displays univariate haircuts (values at risk) for each contract type, estimated for price drops of frequency ranging from once a month to once in 200 years. Contract returns used are from January 2, 2003 to March 31, 2011. Panel B displays the calculated maximum loss above collateral assuming defaulted position depreciation rates given by the haircuts in Panel A and the constant conditional default probabilities from Table 6. All losses are in millions of Canadian dollars. Institutions inactive in futures during 2007-2010 are omitted.

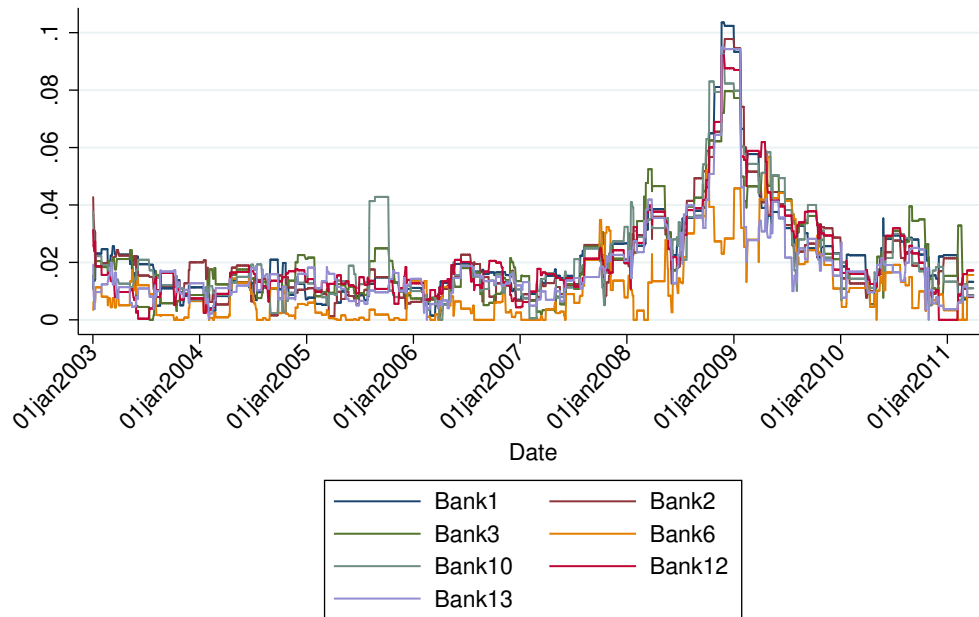
Table 9: Maximum Cost of Shared Losses Above Collateral Requirements, Estimated with Tail β 's

Bank	1	2	3	4	5	6	7	8	9	10	11	12	13	14
<i>Panel A. Market Shortfall = 30.0 Million</i>														
Cost per bank (method 1)	6.9	2.1	3.2	0	2.0	0	2.0	3.5	0	2.3	0	0.7	2.9	4.4
Cost per bank (method 2)	8.2	0	3.8	0	0	0	2.3	4.2	0	2.7	0	0	3.4	5.3
Cost/Mkt. cap., % (method 1)	0.01	0.004	0.01	0	0.004	0	0.06	0.005	0	0.01	0	0.002	0.04	0.003
Cost/Mkt. cap., % (method 2)	0.01	0	0.01	0	0	0	0.07	0.01	0	0.01	0	0	0.04	0.004
<i>Panel B. Market Shortfall = 67.9 Million</i>														
Cost per bank (method 1)	15.7	4.7	7.3	0	4.5	0	4.4	7.9	0	5.1	0	1.7	6.5	10.1
Cost per bank (method 2)	18.7	0	8.7	0	0	0	5.3	9.4	0	6.1	0	0	7.8	12.0
Cost/Mkt. cap., % (method 1)	0.02	0.01	0.03	0	0.01	0	0.13	0.01	0	0.02	0	0.004	0.08	0.01
Cost/Mkt. cap., % (method 2)	0.03	0	0.03	0	0	0	0.16	0.01	0	0.02	0	0	0.10	0.01
<i>Panel C. Market Shortfall = 131.9 Million</i>														
Cost per bank (method 1)	30.5	9.1	14.1	0	8.7	0.1	8.6	15.4	0	10.0	0	3.2	12.7	19.5
Cost per bank (method 2)	41.6	0	0	0	0	0.1	11.8	20.9	0	13.6	0	0	17.3	26.7
Cost/Mkt. cap., % (method 1)	0.05	0.02	0.05	0	0.02	0.01	0.25	0.02	0	0.04	0	0	0.16	0.01
Cost/Mkt. cap., % (method 2)	0.07	0	0	0	0	0.01	0.35	0.03	0	0.05	0	0	0.21	0.02

The table displays the per-bank cost of mutualizing the maximum net systemic market shortfall over 2007-2010, estimated with the constant conditional crash probabilities from Table 6 over three position depreciation scenarios, corresponding to crises with severity/frequency of once in 50, 100, and 200 years. Panel A displays the scenario with depreciation rates corresponding to a once-in-50-years contract price drop. Panel B repeats this calculation with haircuts corresponding to once-in-100-years price drop, and Panel C, with a once-in-200-years price drop. The haircut rates are listed in Table 7 (Panel A) and estimated with extreme value theory. Costs are allocated using 2 methods. Method 1 divides the maximum shortfall among banks proportional to each one's 60-day share of initial margin as of the midpoint of 2008:Q1. Method 2 additionally assumes banks already contributing to the net shortfall are unable to pay further, increasing the cost for remaining survivors, and is the one used by the CDCC in the event of default. The cost per bank is reported in millions of CAD; costs as a fraction of market capitalization are reported in percent. Zero costs for some participants may indicate they are not active in the futures market in 2007-2010.

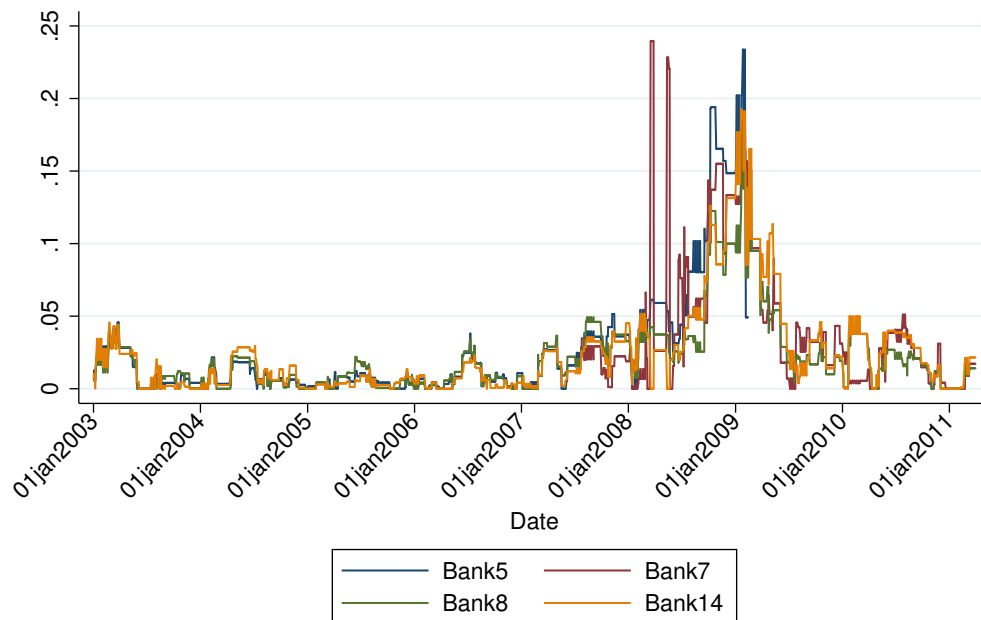
9 Figures

Figure 1. Marginal Expected Shortfall - Canadian Institutions



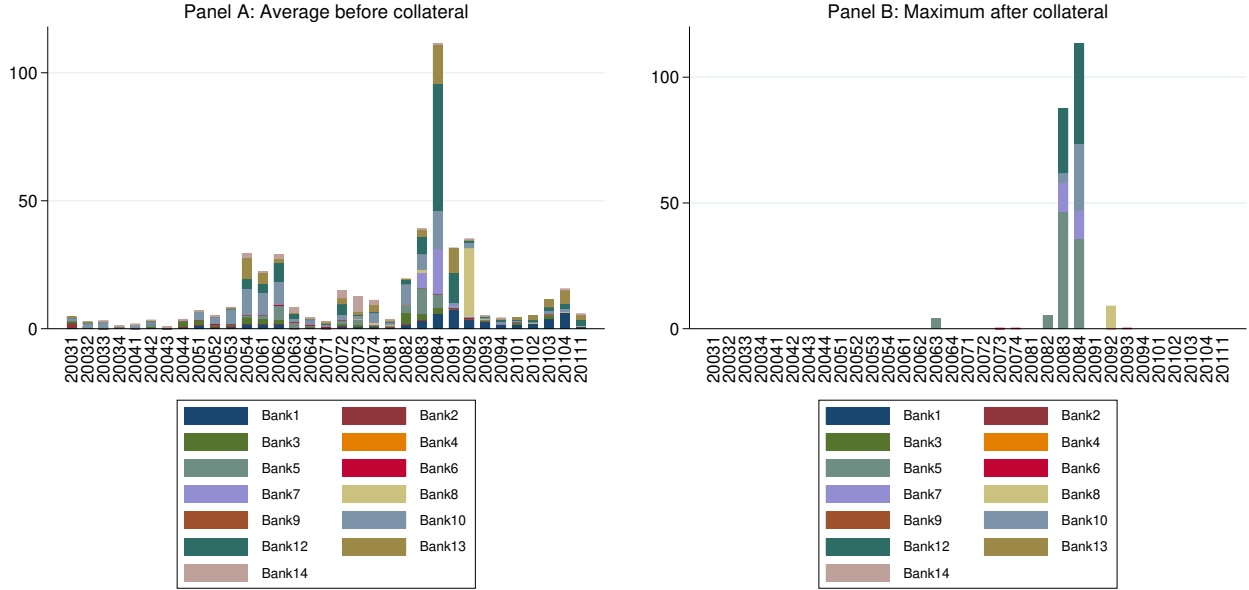
Marginal expected shortfall as calculated in Section 3 for the Canadian sample members listed in Table 1.

Figure 2. Marginal Expected Shortfall - US Institutions



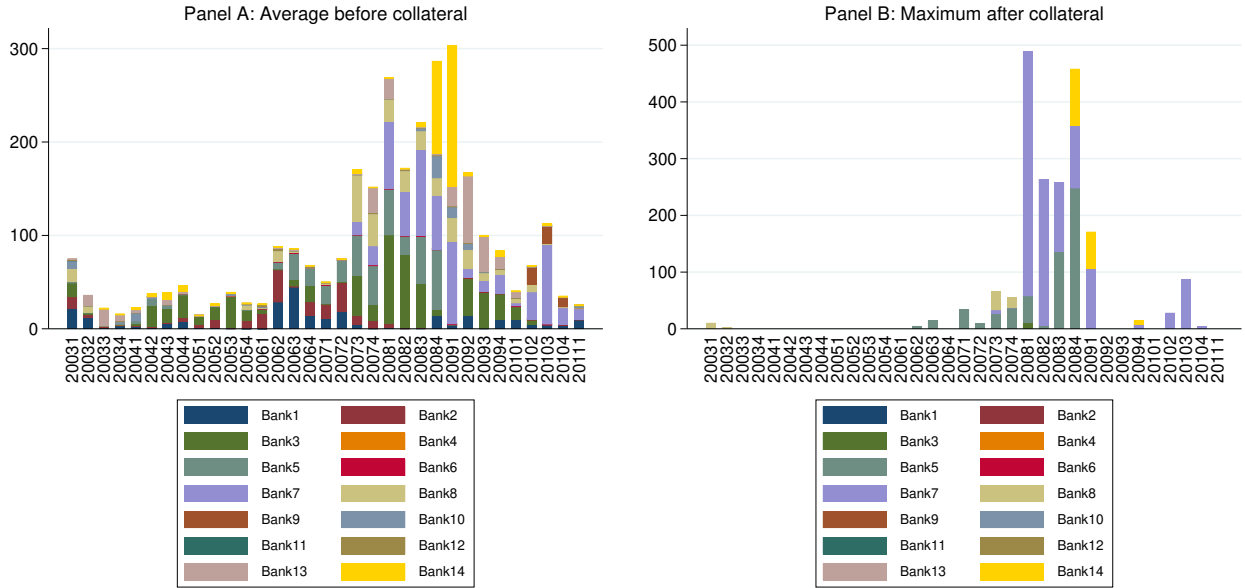
Marginal expected shortfall as calculated in Section 3 for the US sample members listed in Table 1.

Figure 3. Systemic Risk from Proprietary Trades, by Firm



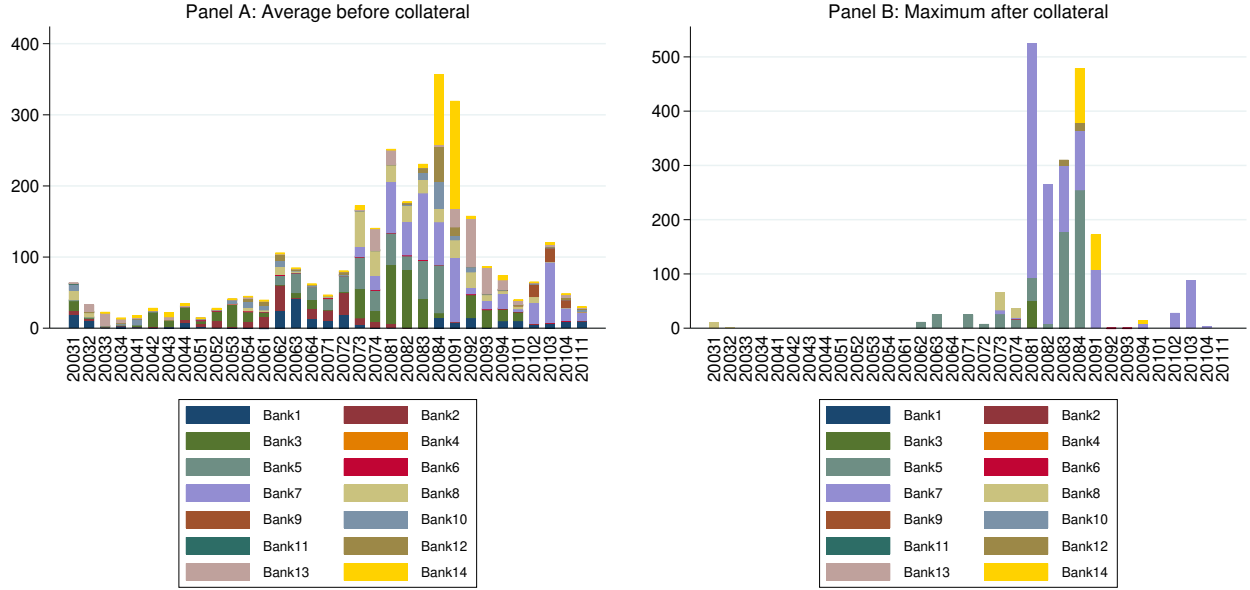
Systemic market shortfall from Firm accounts. Panel A shows the quarterly average systemic market shortfall over days with a positive shortfall. Panel B displays the quarterly maximum net systemic market shortfall obtained after applying collateral.

Figure 4. Systemic Risk from Customer Trades, by Firm



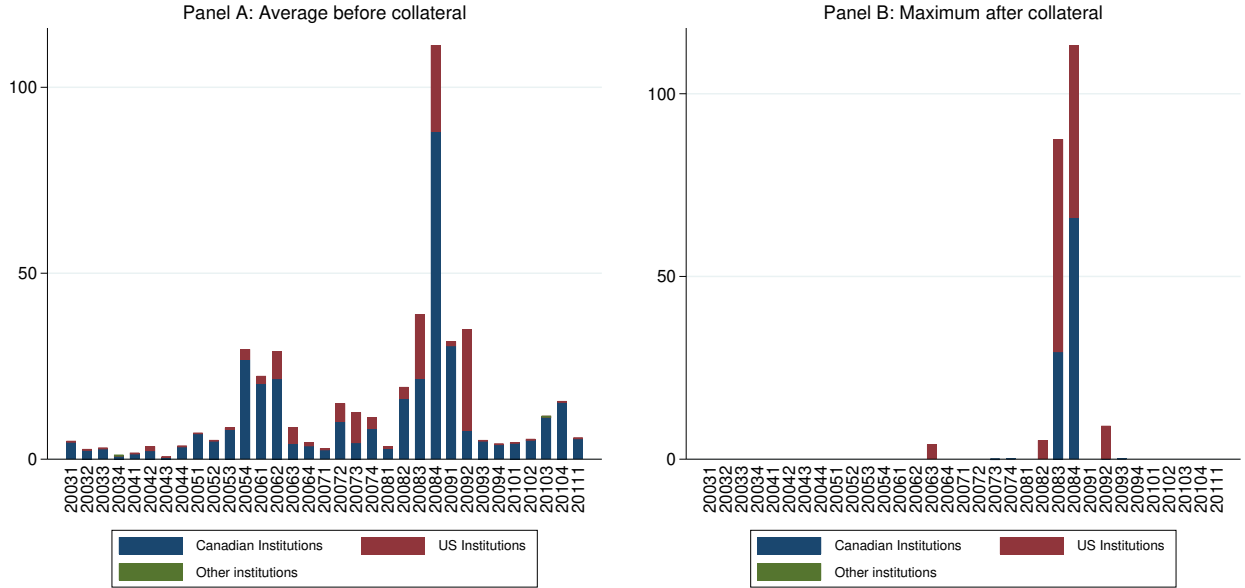
Systemic market shortfall from Customer accounts. Panel A shows the quarterly average systemic market shortfall over days with a positive shortfall. Panel B displays the quarterly maximum net systemic market shortfall obtained after applying collateral.

Figure 5. Systemic Risk from All Trades, by Firm



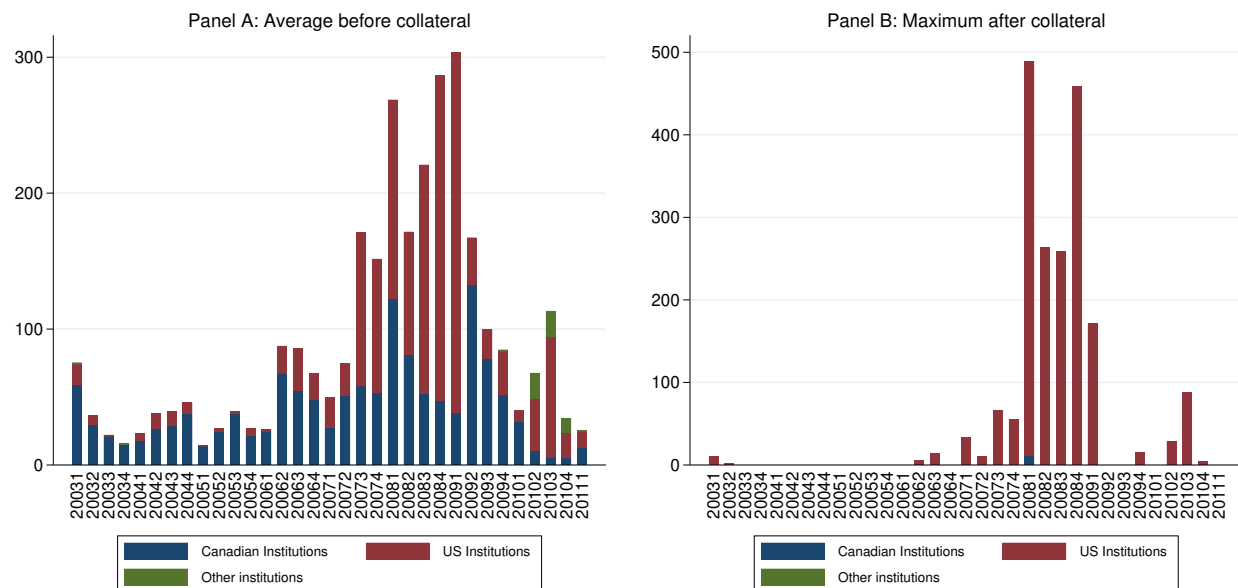
Systemic market shortfall from Total accounts. Panel A shows the quarterly average systemic market shortfall over days with a positive shortfall. Panel B displays the quarterly maximum net systemic market shortfall obtained after applying collateral.

Figure 6. Systemic Risk from Proprietary Trades, by Country



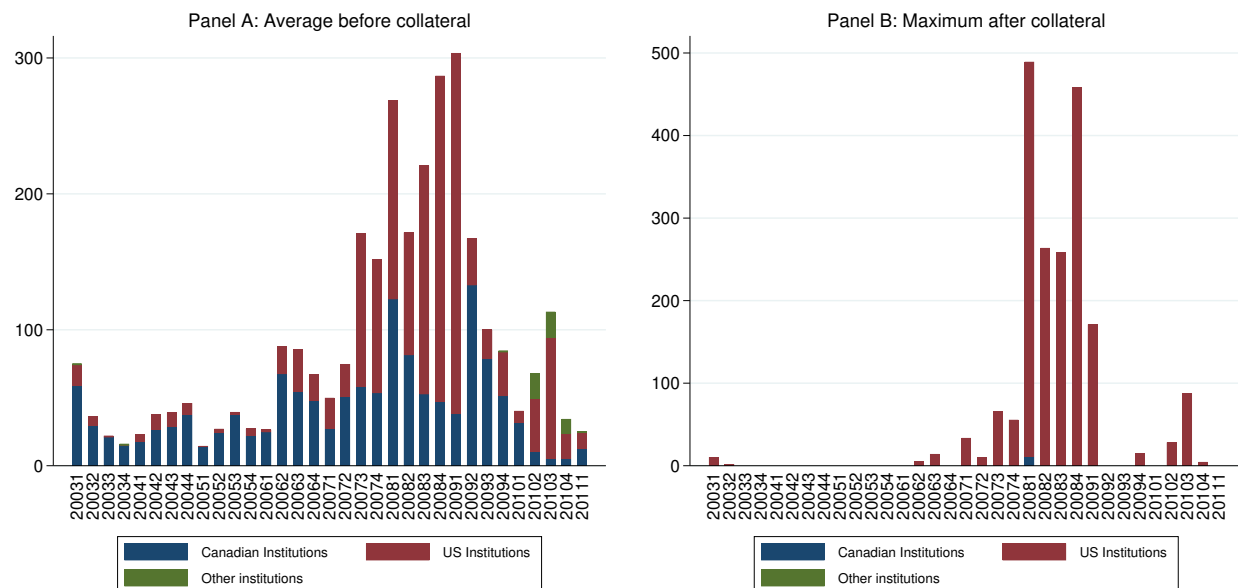
Systemic market shortfall from Firm accounts by country. Panel A shows the quarterly average Systemic Market Shortfall over days with a positive shortfall. Panel B displays the quarterly maximum net systemic market shortfall obtained after applying collateral.

Figure 7. Systemic Risk from Customer Trades, by Country



Systemic market shortfall from Client accounts by country. Panel A shows the quarterly average systemic market shortfall over days with a positive shortfall. Panel B displays the quarterly maximum net systemic market shortfall obtained after applying collateral.

Figure 8. Systemic Risk from All Trades, by Country



Systemic market shortfall from Total accounts by country. Panel A shows the quarterly average systemic market shortfall over days with a positive shortfall. Panel B displays the quarterly maximum net systemic market shortfall obtained after applying collateral.

10 Appendix: The SPAN risk system

The SPAN risk system, introduced by the Chicago Mercantile Exchange in 1988, is an industry-standard method for determining collateral requirements in derivatives trades. It is used by key institutions such as the Chicago Mercantile Exchange, Eurex, LCH.Clearnet, Nymex, Options Clearing Corporation, and CDCC, among others.

SPAN assesses portfolio risk by calculating the maximum likely loss that the portfolio might sustain based on a set of pre-assigned scenarios and risk parameters. SPAN divides the portfolio into contract families having the same underlying asset, called combined commodity groups, and estimates a charge for each combined commodity group independently. Using up to 16 scenarios that vary the price and the volatility of the underlying asset and its time to expiry, SPAN simulates 1-day-ahead changes in the value of each combined commodity group to calculate initial margin requirements.¹⁹ For this purpose, it uses the so-called *margin interval* – an estimate of how difficult it is to liquidate a defaulted position written over an asset with a given volatility within a high degree of statistical confidence. For each combined commodity group, the margin interval (MI) is calculated using the formula

$$MI = 3\sqrt{n} \max \{ \sigma_{20days}, \sigma_{90days}, \sigma_{260days} \}, \quad (15)$$

where n is the number of liquidation days (for futures, $n = 2$), σ is the standard deviation of the underlying asset’s daily price variation over 20, 90, and 260 days, and 3 is equivalent to 99.87% for a one-tail confidence interval under the normal distribution assumption. The margin interval is easiest to interpret as a “three-sigma” value-at-risk-type measure prorated for the number of liquidation days.

The SPAN risk engine uses the margin interval to calculate the maximum price movement reasonably likely to occur for each derivative instrument, using a set of up to 16 risk scenarios, and computes the simulated P&L for each scenario. The base initial margin for futures is then set to cover the maximum P&L over the range of scenarios.²⁰ Initial margin charges computed this way are then aggregated across combined commodity groups to form the Base

¹⁹Depending on the implementation chosen, not every price or volatility scenario may pertain to each contract family. Here we follow the SPAN scenario implementation as used by CDCC during the sample years (2003-2011).

²⁰The detailed risk scenarios and associated methodology are described in detail in the publicly available CDCC Risk Manual.

initial margin. Base initial margin is charged on the financial institution level, based on the institution's combined activity from all accounts, reflected in its Total account. CDCC has provided its monthly margin intervals for 2003-2011 for each contract type, which we utilize to simulate the initial margin for the futures portfolio.

The next step in the SPAN margin computation is to calculate two margin adjustments, known as *intra-commodity and inter-commodity spread charges* for each combined commodity group. Intra-commodity spread charges represent an adjustment to the worst-case scenario loss, meant to account for the fact that contracts of the same type but with different expiry months are treated as equivalent in the scenario simulation. By contrast, inter-commodity spread charges account for mutually offsetting long and short positions in contracts with the same expiry month, written over correlated underlying assets. Intra-commodity and inter-commodity margin adjustments account for a small fraction of margin, involve some discretion by CCP managers, and are rarely implemented consistently across markets, market conditions, or commodities (Cruz Lopez et al., 2017). Due to this, similar to Cruz Lopez et al. (2017), we simulate the SPAN base initial margin without add-ons; this only strengthens our results about the resilience of conventional collateral levels. Additional information about the SPAN methodology can be found in Chicago Mercantile Exchange (2019).