

Raykov, Radoslav

Working Paper

Systemic risk and portfolio diversification: Evidence from the futures market

Bank of Canada Staff Working Paper, No. 2021-50

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Raykov, Radoslav (2021) : Systemic risk and portfolio diversification: Evidence from the futures market, Bank of Canada Staff Working Paper, No. 2021-50, Bank of Canada, Ottawa,
<https://doi.org/10.34989/swp-2021-50>

This Version is available at:

<https://hdl.handle.net/10419/248311>

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.

Systemic Risk and Portfolio Diversification: Evidence from the Futures Market

by Radoslav Raykov

Financial Stability Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

[rraykov@bankofcanada.ca](mailto:r raykov@bankofcanada.ca)



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.

ISSN 1701-9397

©2021 Bank of Canada

Acknowledgements

The author would like to thank, in no particular order, Albert Menkveld, Charlie Kahn, Rod Garrat, Jason Allen, Toni Ahnert, Virginia France, Meryem Duygun, and seminar participants at the Bank of Canada for useful feedback and suggestions. Alexander Chaudhry provided excellent research assistance.

Abstract

This paper explores the extent to which correlated investments in the futures market concentrated systemic risk on large Canadian banks around the 2008 crisis. We find that core banks took positions against the periphery, increasing their systemic risk as a group. On the portfolio level, position similarity was the main systemic risk driver for core banks, while cross-price correlations drove the systemic risk of noncore banks. Core banks were more diversified, but their portfolios also overlapped more. By contrast, non-core banks were less diversified, but also overlapped less. This significantly nuances the debate on concentration versus diversification as systemic risk sources.

Topics: Financial markets; Financial institutions

JEL codes: G10, G20

1 Introduction

With the decline of the traditional banking model, an increasing number of studies have linked the systemic risk of banks to their use of derivatives. For example, Sjostrom (2009) shows how exposures in credit derivatives during the 2008 crisis inflicted large losses on the financial sector and crippled the largest US insurer, AIG, while Brunnermeier, Dong and Palia (2020) and Lepetit, Nys, Rous and Tarazi (2008) find that trading exposures including derivatives increased the systemic risk of banks. The greater use of derivatives has been part of a slow, but pronounced shift in banks' business models away from traditional lending towards the use of trading to generate income (Falato, Iercosan and Zikes, 2019). This new business model has created new interlinkages and stronger comovements among banks.

Derivatives trading creates a common exposure channel across banks because it makes them more vulnerable to contemporaneous losses. One key aspect of contemporaneous losses that has remained outside of the literature's focus is their cross-sectional distribution. All else equal, a trading loss distribution concentrated predominantly on large banks has different implications for systemic stability than a distribution more evenly spread out across the financial system. Specifically, if the systemically important banks all hold similar portfolios against the rest of the system, then they are much more likely to suffer contemporaneous losses as a group, compared to the case where they take opposite positions against each other. This is true regardless of whether these banks' portfolios are concentrated on a single asset or diversified, as long as they remain similar in cross-section. Banks can undertake similar investments for various reasons, including limited liability (Acharya, 2001; Acharya and Yourlmazer, 2005), incentives to be bailed out (Acharya and Yorulmazer, 2006 and 2007), competition levels (Silva Buston, 2019), regulation (Zhou, 2013), or systemic portfolio liquidation costs (Wagner, 2010, 2011).

This paper explores the extent to which correlated investments in the futures market translated into correlated portfolio returns around the Great Financial Crisis, using proprietary positions data on the Canadian futures market. The paper accomplishes four main goals. Firstly, it decomposes observed portfolio comovements into two drivers: portfolio sim-

ilarity and cross-price correlations, exploring the magnitude of each. Secondly, it explores whether the effect of portfolio similarity on portfolio returns is heterogeneous for core versus non-core banks. Thirdly, we decompose portfolio similarity into concentration versus diversification components, thus providing an empirical test of the competing propositions of Wagner (2010, 2011) and Menkveld (2017) on diversification versus concentration as sources of systemic risk. And fourthly, we explore how these parameters evolved around the 2008 crisis. To our knowledge, this is the first paper to explore these questions with the use of daily positions data from banks’ actual trading books. Our study focuses on the portfolio returns of all proprietary trading participants in the Canadian futures market between January 2003 and March 2011, provided by the Canadian Derivatives Clearing Corporation.¹ Futures are the most representative segment of the exchange-traded derivatives market in Canada, comprising about 80% of its value for the sample period.

The financial markets literature has long focused on the cross-sectional comovement of portfolio returns as a source of systemic risk (Menkveld, 2017; Cruz Lopez et al., 2017; Perez-Saiz and Li, 2018). This literature regards large contemporaneous losses across market actors as undesirable, regardless of the identity or importance of those actors; as a consequence, it has focused on aggregate risk measures.² By contrast, our approach focuses on the overlap of individual position characteristics as an underlying source of systemic risk.

Intuitively, if the returns of two bank portfolios are positively correlated, this is either because the portfolios themselves are similar in terms of asset mix and direction, or else because the price movements of assets not held in common are correlated. We use a commonly used position similarity metric (cosine similarity) to explicitly control for similarity in banks’ daily portfolios. Having determined whether position similarity is typical of Canadian core banks, we further test whether their systemic risk is driven by concentrating on the same assets (Menkveld, 2017), or else is attributable to diversification. Diversified portfolios, when sufficiently similar, can also serve as a source of systemic risk (Wagner, 2010, 2011).

¹We thank George Kormas and Olivier Léon from CDCC.

²For example, Duffie and Zhu (2011) propose measuring contemporaneous losses with the market’s mean aggregate loss, while Menkveld (2017) suggests using the tail of the aggregate loss.

The 2008 financial crisis provides a suitable test environment for this study, because monetary response abruptly depressed interest rates, increasing incentives for market participants to bet on the future path of interest rates and the economy through derivatives. For this study, we take advantage of a confidential regulatory dataset containing the proprietary investment positions of all participants in the Canadian futures market, which comprises about 80% of the value of all exchange-traded derivatives traded in Canada. The sample contains 113 distinct futures contracts active between January 2, 2003, and March 31, 2011.

Since in a derivatives market, trading profits and losses always sum up to zero, we first qualify the existing literature, according to which loss correlations are seen as a general source of risk. Since, unlike in a stock market, on any given day, exactly half of the trading counterparties will have a loss offset by the opposite side's profit, from the standpoint of systemic risk, the real issue is not whether multiple simultaneous losses occur on a given day (they always do on the losing side of each contract), but *how they are distributed over the set of market actors*. Our first question therefore is whether that distribution is concentrated on systemically important institutions whose failure could cause economy-wide instability. If contemporaneous losses are dispersed evenly across diverse actors in the financial system, it is likely much more resilient to market turmoil compared to the case where simultaneous losses accrue exclusively to the systemically important banks (DSIBs). Whether this will occur or not depends on whether DSIBs' positions are similar, and whether DSIBs as a group (i.e., the financial system core) collectively take opposite positions against the periphery.³ To find this out, we focus on their positions collectively as a group. Since tracking pairwise comovements across banks is impractical, we follow the standard methodology of the systemic risk literature (e.g., Acharya et al., 2017; Brunnermaier and Adrian, 2017) and use an index of the most important banks (the core banks) against which we measure the comovements of all remaining banks. We therefore measure banks' portfolio return correlations against an index portfolio composed of the core banks. Empirically, core banks are highly correlated and

³For example, if three of Canada's Big Six banks trade in a contract with the remaining three, no more than half of the Big Six can have a simultaneous loss, whereas if all six banks take similar positions against the periphery, they can all have a loss at the same time.

very similar, thus offering a natural benchmark for the rest of the system. Thus, measuring a bank’s correlation against the core captures the possibility that it can be affected by the same shock as the remaining large banks, resulting in simultaneous losses. We use correlations against the core as our main systemic risk metric to cleanly relate to the literature identifying correlated losses as a source of systemic stress (Menkveld, 2017; Perez-Saiz and Li, 2018). This approach also allows us to draw valuable conclusions for the systemically important banks as a group, while preserving the confidentiality of participants’ trading books.

We find that, as a whole, the financial system core took positions against the periphery, thereby increasing the systemic risk of the group of core banks. This is consistent with the theory models of Acharya (2001), Acharya and Yourlmazer (2005, 2006, and 2007), and Wagner (2010, 2011) predicting that large banks have incentives to undertake correlated investments. We find that position similarity significantly increased during the crisis for both bank types, but find no evidence this was due to excessive concentration over the same contracts; instead, diversification played a bigger role during the crisis, while concentration played a larger role in the normal period.

These results are consistent with information contagion among large banks or “behavioral” herding effects (Duygun, Tunaru and Vioto, 2021; Nofsinger and Sias, 1999; Welch, 2000; Hwang and Salmon, 2004). However, increased commonality could also be a rational response to market illiquidity when portfolio liquidation costs are systemic (Wagner, 2011). We explore the time pattern of cross-sectional return correlations for each bank around the crisis, and compare it with the corresponding pattern of position similarity both graphically and statistically. We find position similarity is highly significant in determining cross-sectional return correlations, but its effect is heterogeneous across bank types. Similarity drove systemic risk predominantly for core banks, while for non-core banks, cross-price correlations had a bigger effect. Although core banks were more diversified overall, those among them with more concentrated positions had a higher systemic risk increase; thus, position concentration on specific contracts had a heterogeneous effect on systemic risk. By contrast, we find no evidence that concentration had a heterogeneous effect on the systemic

risk of non-core banks. This nuances our understanding of the debate on diversification versus concentration as systemic risk sources in these markets (Wagner, 2010 versus Menkveld, 2017), showing that certain aspects of both propositions may be correct, but for different banks.

Interestingly, increased portfolio similarity during the crisis did not directly translate into more correlated returns as one would expect. The reason for this is that crisis-induced volatility added high-frequency noise and decorrelated the usually tightly correlated contract returns. The pairwise correlations across contracts dropped by as much as 19 percentage points during the crisis, thus largely offsetting the effect of increased portfolio similarity. As a result, while individual risk increased, *systemic* spillover risk in this market not only did not increase with the crisis, but even scored a modest reduction. On the other hand, return correlations among the group of core banks were already high to begin with. Overall, in the pre-crisis period, the cross-sectional loss distribution was more concentrated on core banks than it was in the crisis, but crisis positions were more similar in cross-section and crisis losses were larger.

This paper contributes to several literatures. On the one hand, a long strain of banking papers have argued that banks have incentives to undertake correlated investments that increase their comovements. For example, Acharya (2001) and Acharya and Yourlmazer (2005) argue that banks do not internalize the costs of joint failure because of limited liability, thus creating incentives to undertake correlated investments. In Acharya and Yorulmazer (2006) and Acharya and Yorulmazer (2007), banks make correlated investments to increase the likelihood of failing simultaneously in order to induce a regulator to bail them out. In Wagner (2010), banks dislike being correlated, but interbank commonality arises as an unwelcome side effect of diversification. By contrast, a different literature focused on herding (Duygun et al., 2021; Nofsinger and Sias, 1999; Welch, 2000; Hwang and Salmon, 2004) perceives common investments as a phenomenon arising from investors believing that the rest of the market has superior information. In line with this literature, we find that similar investments increased during the crisis. Our paper does not take a stance on which of the above

mechanisms causes correlated investments, since their implications are not distinguishable on the portfolio level. However, we do test the two empirically distinguishable propositions of Menkveld (2017), who finds evidence of similarity due to concentration, versus Wagner (2010, 2011), who predicts commonality due to diversification.

The idea of interbank commonality is closely related to that of position crowding. Menkveld (2017) explores the risk from crowded positions, defining crowding as a situation where multiple banks' positions are allocated to the same asset or combination of assets. In practice, the crowding on Nokia stock he discusses refers to a single asset, illustrating systemic risk due to concentration; hereafter, we will use the term “crowding” in that sense to distinguish it from crowding due to diversification as in Wagner (2010, 2011). The fact that Menkveld's (2017) crowding index allows both concentration and diversification to drive the comovement motivates us to measure these two risk sources separately and on the individual bank level.

We see our contribution as disentangling several related issues. Firstly, the paper relates the trading losses of large banks to whether they take similar or opposing positions against each other — a factor with potentially large implications for financial stability. Secondly, we explore the degree to which banks' portfolio return correlations are driven by portfolio similarity, as opposed to cross-price correlations outside of the banks' control. Thirdly, we decompose portfolio similarity into concentration versus diversification components, thus providing an empirical test of the competing propositions of Menkveld (2017) and Wagner (2010, 2011) on correlated investments. Fourthly, we look at the time pattern of return correlations and portfolio similarity for each bank type. While portfolio similarity increased during the 2008 crisis, we also find that the crisis decorrelated asset returns, thereby counteracting the systemic risk effect of increased portfolio similarity. Overall, the paper provides new information about the systemic risk sources in derivatives markets.

The rest of the paper is organized as follows. Section 2 describes the data and provides background about the Canadian futures market. Section 3 describes the empirical strategy. Section 4 focuses on the results, and Section 5 concludes.

2 Market Background and Data

Futures account for about 80% of the value of all exchange-traded derivatives traded in Canada, so this market is highly representative of exchange-traded derivatives. The marketplace for proprietary futures trading in Canada is the Montreal Exchange, owned by the TMX Group. It includes the largest Canadian banks (RBC, TD, Bank of Montreal, CIBC, Scotiabank, and National Bank), smaller domestic institutions, and important international institutions such as J.P. Morgan, Goldman Sachs, and Merrill Lynch. Nonetheless, proprietary trading in this market is dominated by Canadian banks (Raykov, 2021). Since not all Montreal Exchange participants engage in proprietary trading, the active institutions covered in our sample are listed in Table 1; the analysis anonymizes them due to data disclosure requirements.⁴ The market is centrally cleared through the Canadian Derivatives Clearing Corporation (CDCC), a TMX Group subsidiary, which acts as a central counterparty. In 2020, the average notional amount traded daily in Canadian futures was \$113.46 billion in approximately 15 broad contract categories (short-term interest rate futures, bond futures, various share futures, index futures, repo and index swap futures, and sector index futures). Total futures trading accounted for an average volume of 321,386 contracts daily, with an average open interest of 2.47 million contracts monthly. These figures reflect both banks' proprietary trading positions and trades undertaken on behalf of investment clients for which Montreal Exchange member banks serve as a conduit.

For the analysis, we use the end-of-day proprietary positions data on 113 distinct futures contracts traded at the Montreal Exchange between January 2, 2003, and March 31, 2011, belonging to three general types. These are Canadian bankers' acceptance futures (BAX), 10-year Government of Canada bond futures (CGB), and S&P/TSX 60 index standard futures (SXF), together accounting for more than 90% of the open interest in the sample period. BAX and CGB are interest rate futures over short and long rates, respectively, and SXF is an index future over an index of the 60 most liquid Canadian stocks. By entering

⁴An institution is considered active in the proprietary trading market if its firm account has at least one open position on at least half of the trading days in the sample period.

a sequence of contracts expiring on specific dates, a participant is able to bet on the future path of interest rates and the stock market. For example, a participant who is short on a BAX expiring one month from now and long on a BAX contract expiring in three months is effectively betting that the underlying bankers acceptance’s price will increase between one and three months in the future, and the associated short interest rate will drop. Similarly, a participant can bet on specific paths for the long interest rates and the stock market. Since they are linked through the yield curve, the short and long interest futures’ returns (BAX and CGB) are highly correlated, but they are only weakly negatively related to the SXF stock market future. However, crisis-related phenomena, such as yield curve inversions and high volatility, reduced the correlatedness across all three asset types; we comment on this extensively in Section 4.1. Collectively, these three contracts are highly representative of Canada’s futures market around the 2008 crisis (Campbell and Chung, 2003; TMX Montreal Exchange, 2013a, 2013b, 2013c). The summary positions statistics of the contracts in our sample, their returns, and respective correlations are shown in Tables 2 and 3.

3 Empirical Strategy

Channels leading to correlated returns. Our first goal is to disaggregate participants’ portfolio return comovements into those due to position similarity versus price correlations across assets not held in common. To understand this distinction, note that there are only two channels through which two portfolios can comove: the position similarity channel and the cross-price correlation channel. They are illustrated in Figure 1. The position similarity channel refers to portfolios containing the same contracts and in the same direction. For example, a position with 100 long BAX contracts, 50 long CGB contracts, and 50 short SXF contracts (which we can briefly record as $(BAX, CGB, SXF) = (100, 50, -50)$) will display positively correlated returns with the portfolio $(80, 60, -20)$ because it contains the same assets and in the same direction (Figure 1). Note that the effect of bank-specific characteristics influencing portfolio choice is already reflected through this channel.




Figure 1: The two channels leading to correlated returns.

However, sometimes portfolios with no assets in common can still display correlated returns. For example, consider the two “orthogonal” portfolios $(BAX, CGB, SXF) = (100, 0, 0)$ and $(0, 100, 0)$, each holding only BAX or only CGB contracts. If BAX and CGB, which are both interest rate futures, have correlated returns, this will carry over to the portfolio level too.⁵ We refer to this as the cross-price correlation channel (Figure 1, bottom panel). Cross-price correlations can occur mechanically even between unrelated assets and vary significantly over time, as we document; as shown by Figure 1, they are not informative of portfolio composition. Thus, just observing that two portfolios have correlated returns is not necessarily indicative of high interbank commonality. We need a better and more organized way of measuring similarity across bank portfolios in a way that separates it from cross-price correlations.

Measuring portfolio similarity. The empirical literature on portfolio choice (e.g., Sias, Turtle, and Zykaj, 2016; Girardi, Hanley, Nikolova, Pelizzon, and Sherman, 2018; Bech, Bergstrom, Rosvall, and Garratt, 2015) offers a portfolio similarity metric that segregates these two aspects, allowing us to capture portfolio characteristics separately from the price characteristics of the assets involved. The metric is called cosine similarity and is computed as the cosine distance between the two portfolios, represented as vectors in n -space. In this framework, the example portfolios $(100, 50, -50)$ and $(80, 60, -20)$ from Figure 1 are treated

⁵For instance, the 3-month and 10-year interest rates futures BAX and CGB are positively related through the yield curve, but yield curve inversions can weaken this relationship significantly during turmoil periods. See the extended discussion in Section 4.1.




Figure 2: Cosine similarity in 3D-space (not to scale).

simply as vectors forming an angle θ in 3-space, and their cosine similarity is defined as the cosine of that angle (see Figure 2). The benefit of this metric is that it separates the portfolio composition aspect from the price comovements of the assets involved, allowing the researcher to distinguish mechanical portfolio comovements from those due to similar portfolio composition. This metric varies from -1 (for opposite positions) to 1 for identical ones, and is zero when there are no assets in common.

Algebraically, for two arbitrary vectors \mathbf{a} and $\mathbf{b} \in \mathbb{R}^n$, their cosine similarity is defined as

$$\text{sim}(\mathbf{a}, \mathbf{b}) = \cos(\theta) = \frac{\mathbf{a} \cdot \mathbf{b}}{\|\mathbf{a}\| \|\mathbf{b}\|}. \quad (1)$$

In the specific instance of the example portfolios $(100, 50, -50)$ and $(80, 60, -20)$ from Figure 1,

$$\text{sim}(100, 50, -50), (80, 60, -20)) = 12,000 / (\sqrt{15,000} \sqrt{10,400}) = 0.96. \quad (2)$$

Thus, according to this metric, the two example portfolios from the top panel of Figure 1 are indeed similar in composition, independent from how their asset prices comove. If the two portfolios were scalar multiples of each other, their similarity would be 1 .

We can likewise compute the similarity between the two portfolios $(100, 0, 0)$ and $(0, 100, 0)$ from the bottom of Figure 1. In 3-space, these two vectors are orthogonal, forming an angle

of $\pi/2$ whose cosine ought to be 0; and indeed,

$$\text{sim}((100, 0, 0), (0, 100, 0)) = 0/(\sqrt{10,000}\sqrt{10,000}) = 0. \quad (3)$$

Accounting for these two portfolios as orthogonal helps us to know that any return correlation observed between the two is entirely due to cross-price correlations, since they have no assets in common. We will use this feature to explicitly capture cross-price correlations later.

Analogously, if two banks are taking exactly opposite positions in each asset (or scalar multiples thereof), it is easy to show that their cosine similarity equals -1 . We will use this feature to underscore the fact that two banks with positive cosine similarity are clearly *not* counterparties to the same trade (in which case their similarity would be -1).⁶ Since we do not get to observe the counterparty on the other end of each position, this inference is important in determining whether the Big Six Canadian banks on average trade amongst each other as a closed group, or take positions against the rest of the financial system.

The above examples are stylized for a world with just three assets. In reality, within each futures type, there are distinct contracts expiring on different dates and corresponding to different economic bets. For example, a bank that is 100 contracts short on a BAX future expiring in one month and 100 contracts long on a BAX contract expiring in three months is effectively betting that the underlying bankers' acceptance price will go up between one and three months into the future, and the associated three-month interest rate, the CDOR, will drop in that period. Netting out these positions to zero would be incorrect, as a zero net position would equate this bet to the *opposite* bet of its counterparty (that the short interest rate will *increase* between one and three months in the future), or to the situation where no bets are made. For this reason, we compute similarity on the individual contract level, rather than on the broad contract type. There are a total of 113 distinct futures contracts across the three types in our sample.

Correlation and similarity to the core. Whether the loss distribution concentrates

⁶This argument holds exactly in a one-asset world and in on average in a multi-asset world.

on the Big Six systemically important banks depends on whether they as a group take similar positions against the periphery, or instead take positions largely against each other. In the former case, all six big banks can potentially suffer contemporaneous losses; in the latter case, the losses of three of them in each asset would be the remaining three’s gains.

To find which of these two cases occurs, it makes sense to consider core banks’ positions collectively as a group. For similar reasons, it is burdensome to track pairwise return correlations when there are multiple banks; some sort of benchmark is needed. To resolve both problems at once, we adopt the approach already established in the systemic risk literature (e.g., Acharya et al., 2017, or Brunnermaier and Adrian, 2017) and create a value-weighted index of the most important banks, against which we measure the comovements of all remaining banks.⁷ The natural choice of benchmark banks is the financial core (the Big Six), which account for a large fraction of the proprietary market and form exactly half of the sampled institutions. These core banks are both highly correlated and very similar, thus offering a natural benchmark for the rest of the system. Measuring a bank’s correlation against the core captures the possibility that it can be affected by the same shocks as the remaining systemic banks, therefore increasing the systemicity of a crisis. Therefore, we use correlations against the core as our main systemic risk metric. This is motivated by two reasons. Firstly, we want to relate cleanly to the derivatives literature identifying correlated losses as a source of systemic stress (Menkveld, 2017; Perez-Saiz and Li, 2018, etc.); and secondly, because our interest here is more in the possibility of risk spillover across banks rather than the magnitude of this spillover. (Spillover magnitudes in this market are studied by Raykov (2021).)

To implement this strategy, we construct a synthetic Core portfolio in each contract k as a value-weighted index of the core banks’ positions in the same asset:

$$CorePosition_{k,t} = \sum_i w_{i,k,t} Position_{i,k,t}, \quad \text{where } i \in BigSix, \quad (4)$$

⁷When we need to compare a core bank to the remaining big banks, we exclude this individual bank from the core index to prevent a mechanical relation between the bank and the index.

and aggregate contract-specific returns across contracts to compute the core's next-day portfolio return, $R_{c,t+1}$, resulting from time- t positions. Next we proceed to compute the co-movements of each bank's portfolio return $R_{i,t+1}$ against the core portfolio return $R_{c,t+1}$ by computing their pairwise correlation $\rho_{t+1}(R_i, R_c)$. When a bank i already belongs to the core, we compute the core return $R_{c,t+1}$ by excluding i from the core to prevent a mechanical correlation between the bank and the index.

Finally, we construct each bank i 's similarity to the (rest of the) core, $SimCore_{i,t}$, as a value-weighted average of the cosine similarities between i and each remaining core bank.

Cross-price correlations. As illustrated by Figure 1 (bottom panel), portfolios with no assets in common can still display correlated returns if asset prices are correlated. To correctly capture the effect of the cross-price correlation channel, the amounts of such assets held by each bank versus the core need to be controlled for, as they mechanically transfer cross-price correlations to the portfolio level. For example, if bank i 's portfolio is 100% long on BAX and the core holds 100% long CGB, then their cross-portfolio return correlation equals exactly the return correlation between BAX and CGB (in this case, 78%). But if i held 100% short on BAX instead, then the portfolio correlation would be -78% . Clearly, the effect of this needs to be controlled for when we regress portfolio correlations on position characteristics. Therefore, for each ordered pair $j := (X, Y)$ of correlated assets X and Y , we create a cross-price control variable XP_j as the product of the portfolio shares of X and Y held by bank i versus the core:

$$XP_{i,j} := (FracX_{i,t} * FracY_{c,t}), \text{ where } X, Y \in \{BAX, CGB, SXF\}, \rho(R_X, R_Y) \neq 0, \quad (5)$$

where $FracX_i$ and $FracY_c$ are the respective value shares of the two correlated assets in the bank's and the core portfolio (signed positively for long positions and negatively for short positions). Thus, if i 's portfolio is 100% long on BAX and the core holds 100% long CGB, then $XP_{i,(BAX,CGB)} = 1 \cdot 1 = 1$. Likewise, if both i and the core are 100% short on BAX and CGB, respectively, the $XP_{i,(BAX,CGB)} = -1 \cdot (-1) = 1$, allowing the regression coefficient in

front of $XP_{i,j}$ to capture the effect of the cross-price comovement separately from that of position similarity.⁸

Panel data model and estimation. We are interested in the extent to which portfolio similarity drives return comovements across banks, whether position concentration (or diversification) plays a role, and whether that effect is heterogeneous across bank type (core versus non-core). To answer these questions, we regress return correlations on portfolio characteristics in variations of the following panel data model, constructed at daily frequency:

$$\begin{aligned} \rho_{t+1}(R_i, R_c) = & \alpha_i + \beta_1 SimCore_{i,t} + \beta_2 (SimCore_{i,t} * Crisis_t) + \\ & + \beta_3 (SimCore_{i,t} * Concentrated_{i,t}) + \\ & + \beta_4 (SimCore_{i,t} * Concentrated_{i,t} * Crisis_t) + \sum_j \gamma_j XP_{i,j,t} + \varepsilon_{i,t+1}, \end{aligned} \quad (6)$$

where $\rho_{t+1}(R_i, R_c)$ is the return correlation between bank i and the core, $SimCore$ is the cosine similarity between bank i and the core, $Concentrated$ is a measure of average portfolio concentration based on a Herfindahl concentration index (HHI), and $Crisis$ is a dummy equal to 1 during the period of heightened volatility in this market (September 1, 2008 – December 31, 2009). Since the return correlation ρ has to be estimated over a (moving) time window, we convert all position characteristics on the right-hand side to moving averages over a corresponding time window following Falato et al. (2019). This is standard in this literature and has the benefit of filtering out high-frequency noise. The $Concentrated$ and $Crisis$ variables are included only as interactions because neither of them can independently affect the cross-sectional correlation between two portfolios outside of the two channels in Figure 1, for which we already control.⁹ The terms $XP_{i,j,t}$ are the cross-price correlation controls defined above, capturing the degree of return correlation between bank i and the core arising purely from price comovements across assets they do not have in common. We include this control for asset pairs whose correlations are significantly different than zero.

⁸We do not further distinguish between individual contracts, since their returns are above 99% correlated within each type.

⁹However, for robustness, we also ran a specification including the $Crisis$ and $Concentrated$ variables independently, too. Neither of their coefficients was significantly different from zero.

We do not use bank-specific controls for this channel decomposition, since the observed daily portfolios already capture the influence of bank-specific characteristics on portfolio choice; those characteristics cannot influence interbank portfolio correlations outside of portfolio choice. Nonetheless, we still allow for bank fixed effects.

To estimate return correlations with sufficient precision, we use a rolling time window with length of one quarter (90 calendar days) in estimating $\rho_{t+1}(R_i, R_c)$. However, the results are robust to shorter window sizes all the way down to two weeks (14 calendar days). To account for the autocorrelation in the y -variable introduced by overlapping time windows, we estimate the panel in equation (7) with Driscoll-Kraay standard errors (Driscoll and Kraay, 1998), which are heteroscedasticity- and autocorrelation-robust, and work especially well for panels with large T and small N like this one.

Endogeneity and identification. The above estimation is based on the premise that the two channels in Figure 1 operate in the directions shown (position characteristics and cross-price correlations influence portfolio correlations). To ascertain this, we first note that, since portfolio choice always precedes the return realization, realized time $t + 1$ returns cannot influence portfolio characteristics chosen at time t . However, one could still argue that observed correlations in past returns could influence two banks to choose their *current* portfolios in a certain way. A key point here is that banks do not observe each other’s proprietary trading returns and positions; their trading strategies are among the best-guarded industry secrets, breaches of which have resulted in lengthy jail terms.¹⁰ Thus, it is not tenable to assume banks knew the direction and composition of each other’s portfolios (especially not in a market with a central counterparty).¹¹ On the other hand, one could argue that core banks, by perhaps unwittingly taking on similar positions, could still move market prices and change cross-price correlations, creating an unforeseen interaction between the two channels. However, to do that, they would need to have sufficient market share to move prices in the entire market, and not just its proprietary segment. As shown in Raykov (2021), about 80%

¹⁰For example, see Chellel and Hodges (2018).

¹¹A central counterparty (CCP) interposes itself between buyers and sellers, fully anonymizing each trade’s counterparties to each other. In a centrally cleared market, transactors deal exclusively and only with the CCP even under extreme outcomes such as participant default.

of total market positions are held by investment clients of large American and Canadian banks; thus, the Big Six’s proprietary market share, even if they coordinated to move one single asset, would scarcely reach 10% of the overall market. Thus, it appears highly unlikely that even similar positions of large banks can swing the market. Endogeneity should therefore not be a major concern in this setting.

4 Results

We begin by documenting a few stylized facts about return correlations and portfolio similarity across bank types. Then, we explore how much of the return correlation is contributed by portfolio similarity versus cross-price correlations, and whether this effect is heterogeneous across bank type. Thirdly, we decompose the portfolio similarity channel into concentration versus diversification components, thus providing an empirical test of the competing propositions of Wagner (2010, 2011) and Menkveld (2017). And fourthly, we explore whether the 2008 crisis changed any of our previous results.

4.1 Stylized facts

We begin by examining the return correlations. Figure 3 shows the average return correlation against the core, taken cross-sectionally by bank type. The figure shows a significant heterogeneity in the return correlations of core versus non-core banks. Even before the start of the 2008 crisis, core banks were tightly correlated, with correlation coefficients against the core exceeding 0.5 during most of the pre-crisis period. This is consistent with the conventional wisdom that the accumulation of risks leading to the 2008 crisis had to do with correlated investments, as has already been documented in other financial markets (such as the market for credit default swaps and mortgage-backed securities). This also reinforces the notion of considering core banks as a separate group, as they clearly faced much higher risks of contemporaneous losses than the rest of the banking system.¹² By contrast, non-

¹²In certain quarters, some core banks’ correlations against the rest of the core exceed 0.95.




Figure 3: Bank portfolio correlations against the core portfolio (cross-sectional average by bank type). The data is smoothed as a one-quarter moving average. Source: Author's calculations.

core banks' return correlations against the core were much lower throughout (the red line in Figure 3 is centered only slightly above zero). While portfolio correlations drop somewhat before the onset of the 2008 crisis, the period 2008–09 shows a clear return to the previous levels; towards the end of the crisis and going forward, core banks aggressively differentiate their portfolios so that their return correlations drop significantly, as confirmed by the next figure.

Figure 4 shows banks' average portfolio similarity to the core portfolio by bank type, allowing us to visually gauge the degree to which portfolio similarities drove the correlated returns in the previous figure. Figure 4 confirms that, until the crisis, the strong observed portfolio correlations were indeed related to portfolio similarity, and more so for core than for non-core banks. Core banks' portfolios were much more similar to each other and featured higher-correlated returns than non-core banks, and the similarity's overall pattern of variation matched the variation in return correlation until the start of the crisis.

In 2008–09, however, return correlation and similarity metrics diverge. Figure 4 shows a clear increase in portfolio similarity from about 0.2 to 0.55 for core banks, and from about 0.1 to 0.3 for non-core banks, whereas the return correlations during the same period drop




Figure 4: Bank portfolio similarity against the core portfolio (cross-sectional average by bank type, smoothed as a one-quarter moving average). Source: Author’s calculations.

from about 0.5 to 0.1 (for core banks) and from 0.3 to around zero (for non-core banks). Since the two channels (portfolio composition and cross-price correlations) are exhaustive in explaining changes in cross-sectional correlations, the trend towards similarity during the crisis ought to have been offset by the second channel — that of cross-price correlation. Table 3, Panel B, confirms this is indeed so. The onset of high volatility in September 2008 created high-frequency noise in all contracts’ returns (Figure 5), therefore upsetting the hitherto stable cross-price correlation structure. In practice, all three asset types became much less correlated with each other. The correlation between the two interest rate futures (BAX and CGB) dropped from 0.78 down to 0.70, while the SXF’s correlations against BAX and CGB went from -0.14 and -0.17, respectively, all the way to -0.33 and -0.34. All of this ultimately decorrelated banks’ portfolio returns despite increasing portfolio similarity. The high-volatility period lasted from approximately September 2008 to the end of 2009.

Another observation from Figure 4 allows us to draw a rough inference about what groups of banks traded with each other in this sample period. Figure 4 shows that the core banks were very similar to each other — much more so than non-core banks. Based on our definition of cosine similarity, this implies that core banks were largely *not* taking




Figure 5: Contract returns for each of the three types of contracts. Source: Author’s calculations.

opposite positions against each other for most of the sample period. If half of the banks in the core were taking positions against the other half, then their average similarity would be zero. The same argument holds for the non-core banks when their similarity to the core is positive (especially around the financial crisis). Thus, it appears that the core banks were largely *not* taking positions against each other throughout the sample period, and likewise not trading much against non-core Canadian banks during the crisis (since the two similarity lines in Figure 3 are both positive). Since the only remaining type of market participant are the investment clients of large banks, this means that the core banks were trading against a periphery consisting mostly of investment clients¹³ and (partly) of non-core banks, rather than amongst themselves.¹⁴ That made core banks much more susceptible to contemporaneous losses. This result is consistent with previous studies, which have found that investment clients of US firms account for significant portion of this market (Raykov,

¹³Due to the format of the data, these clients’ positions are not individually observable in this data set.

¹⁴We define “periphery” relative to the Canadian core banks whose default would matter for financial stability in Canada. Some foreign participants – e.g., large American banks – may still be core banks in the US, but they are not in Canada. None of the American market participants is a significant depository institution in Canada, and most operate in the Canadian market through subsidiaries such as security trading firms that can provide access to the Canadian market to US retail investors. As the 2008 crisis confirmed, Canada’s banks remained largely immune to the default risk of US core banks. Hence, automatically counting US institutions without major presence in Canada as “core” for Canada would be a mistake.

2021); our finding additionally clarifies that these clients were also on the opposite side of the core banks.

Finally, we note the fact that average cosine similarity across core banks peaked in late 2008 to early 2009 at a value around 0.55 (0.3 for non-core banks). This could be due to a number of different mechanisms proposed in the literature. For example, a strain of papers focused on herding (Duygun et al., 2021; Nofsinger and Sias, 1999; Welch, 2000; Hwang and Salmon, 2004) attributes this to investor beliefs that the rest of the market has superior information. In our setting, this would imply information contagion amongst core banks, and less so amongst non-core banks. However, this theory does not explain the source of information contagion among banks, given the extraordinary precautions they take to protect their proprietary trading strategies.¹⁵ Instead, information contagion appears much more prevalent amongst retail investors (Duygun et al., 2021).¹⁶ Other explanations are offered by the theories of Acharya (2001) and Acharya and Yourlmazer (2005, 2006, 2007), where limited liability or incentives to be bailed out can induce banks to undertake correlated investments – in the former two papers, because banks do not internalize the costs of joint failure, and in the latter two, because banks are more likely to be bailed out if they fail simultaneously. These theories make no specific prediction as to whether interbank commonality arises due to concentration or diversification. By contrast, Wagner (2010, 2011) suggests that systemic portfolio liquidation costs induce banks to diversify, leading to interbank commonality. At the same time, Menkveld (2017) shows empirically that position concentration, rather than diversification, created systemic risk in the European equity market.

Most of these theories are not empirically distinguishable from each other based on their predictions about portfolio choice.¹⁷ For this reason, our paper does not take a stance

¹⁵Core banks are also unlikely to materially affect market prices, as per their market share discussion in Section 3.

¹⁶A recent example is the herding on GameStop stock in early 2021, coordinated by retail investors on public forums.

¹⁷For example, proving that positions were similar due to diversification would not help tell apart the mechanisms of Acharya (2001), Acharya and Yourlmazer (2005), Acharya and Yourlmazer (2006), Acharya and Yourlmazer (2007) and Wagner (2010), although showing similarity due to concentration would be

on which of the above mechanisms underlies correlated investments. Instead, we are more interested in the effect of portfolio similarity on correlated returns, and the extent to which interbank commonality is driven by concentration versus diversification. The latter question does test the two empirically distinguishable propositions above: that of Menkveld (2017), who empirically finds evidence of position concentration, versus Wagner (2010, 2011), who predicts commonality due to diversification.

4.2 Statistical analysis

Figures 3 and 4 are indicative of the overall relationship between portfolio similarity and correlated returns, but they do not answer how much of the observed return correlation is contributed by portfolio similarity versus cross-price correlations, and whether the effect is heterogeneous across bank type. They also do not answer whether it is portfolio concentration or diversification that drives risk through the portfolio similarity channel. To answer these questions, we estimate the panel data model

$$\begin{aligned} \rho_{t+1}(R_i, R_c) = & \alpha_i + \beta_1 SimCore_{i,t} + \beta_2 (SimCore_{i,t} * Crisis_t) + \\ & + \beta_3 (SimCore_{i,t} * Concentrated_{i,t}) + \\ & + \beta_4 (SimCore_{i,t} * Concentrated_{i,t} * Crisis_t) + \sum_j \gamma_j XP_{i,j,t} + \varepsilon_{i,t+1} \end{aligned} \quad (7)$$

separately for core and non-core banks.

First, we construct the cosine similarity between each bank i 's portfolio and the core portfolio, defined as the variable $SimCore_{i,t}$, and proceed to measure i 's position concentration to understand whether similar positions were similar due to concentration or diversification. Depending on the specification, we measure position concentration by a Herfindahl concentration index (HHI), a dummy $Concentrated$ equal to 1 for above-median HHI concentration values, or a dummy $HiConcentrated$ equal to 1 for above-75th-percentile HHI values. We construct these concentration measures first on the asset level (BAX, CGB, SXF) and then

evidence against Wagner (2010) but not the rest.

on the individual contract level, testing for both. To set apart cross-portfolio correlations due purely to price comovements of non-overlapping, but correlated assets in the two portfolios, we also construct the cross-price correlation controls XP_j defined in equation (5) for each (ordered) pair of assets held by bank i and the core that display a significant pairwise correlation.¹⁸ The estimation is performed by panel OLS with Driscoll-Kraay standard errors, clustered at the bank level.

We first explore the influence of position characteristics on the systemic risk measure $\rho_{t+1}(R_i, R_c)$. We decompose banks' return correlations into those coming from portfolio similarity versus cross-price correlations. The effect of the former channel is captured by the coefficient β_1 , while the relevant cross-asset correlations are captured by the respective γ_j coefficients. Table 4 presents the effects of position characteristics on systemic risk separately by bank type.

The results in Table 4 show that portfolio similarity was the main driver of correlated returns for core banks, whereas cross-price correlations were the main driver behind the return correlations of non-core banks. Columns 1 and 2 of Table 4, for example, show that similarity to the core has a coefficient of 1.324 for core banks, versus only 0.603 for non-core banks (both significant at the 1% level). At the same time, the effects γ_j of the cross-price correlations XP_j are insignificant for core banks across all specifications, whereas the BAX-CGB cross-portfolio control is always above 1 and highly significant for non-core banks. This can be seen, for example, in column 2, where this control has a coefficient of 1.037, significant at 1%. The remaining specifications show very similar results: the estimated γ_j coefficients for non-core banks in columns 2, 4, and 6 range between 1.03 and 1.07 and are all significant at 1%. Thus, cross-price correlations are the dominant channel behind return correlations of non-core banks, while portfolio similarity is the dominant channel causing correlated returns in core banks.

These effects are also economically significant. For example, one standard deviation increase in portfolio similarity increases the return correlation of the average core bank to

¹⁸For further details, see the Section 3.

the rest of the core by 0.41 standard deviations, and one standard deviation increase in the joint fraction of correlated assets held, XP_{BAX_i, CGB_e} , increases a non-core bank’s correlation to the core by 0.16 standard deviations in column 2.

We next explore whether the portfolio similarity channel additionally increased systemic risk because of portfolios concentrating on the same asset versus diversifying in a similar way. For this purpose, we interact the *SimCore* and concentration measures used in the different specifications. In Table 4, we measure concentration on the asset level (BAX, CGB, SXF) by an HHI index and two dummies *Concentrated* and *HiConcentrated* based on it. The coefficients of this interaction are insignificant across all specifications, showing no evidence of additional risk due to concentration — something also supported by the summary statistics in Table 2. There is no evidence of heterogeneity across bank type on this dimension.

As discussed in Section 4.1, the 2008 crisis decorrelated cross-asset returns, resulting in a drop in the correlatedness of banks’ portfolio returns despite increased similarity. As a result, portfolio similarity and return correlation metrics in Figures 3 and 4 move in opposite directions during the crisis. Capturing this, all interactions of the *Crisis* and *SimCore* variables produce negative and often significant coefficients in Table 4. (We explore separately whether the crisis-related similarity surge was driven by concentration or diversification.) Thus, although crisis-level P&Ls increased in absolute terms (see Figure 5), their correlatedness, captured by the $\rho_{t+1}(R_i, R_c)$ metric, did not increase and even scored a modest reduction.

There are two angles from which one can view this result. On the one hand, the crisis did not increase the linkage across institutions, because their increasing common exposure through portfolio similarity was offset by a drop in the correlation of asset prices. On the other hand, the pre-crisis return correlations were already high, leaving little room for further increase: Figure 3 shows that the average return correlation of core banks before the crisis often reached 0.75, which means that some core banks must have exceeded that value. At the same time, by the end of 2009, core banks began to aggressively differentiate their portfolios. In 2010, the average similarity amongst core banks was three times lower than at

the peak of the crisis, and fell to near zero in January 2011 (see Figure 4). Unsurprisingly, return correlations followed, and dropped below zero given the negative correlation between the SXF future and the remaining two. Preliminary data from 2020, not included in this analysis, seems to confirm this pattern going forward. Thus, the financial crisis seems to have ended a period of high correlation and high similarity across core banks; this is consistent with the view of the pre-2008 financial system as one typified by correlated investments.

Table 5 repeats the same decompositions, measuring concentration more granularly on the individual contract level. This allows us to perform a finer test of Menkveld’s (2017) empirical proposition about single-asset crowding versus the diversification prediction of Wagner (2010, 2011), and allows us to answer whether the crisis-related surge in similarity was itself due to concentration or diversification.

Table 5 confirms the overall picture from Table 4. As before, the portfolio similarity channel remains the dominant systemic risk driver for core banks, resulting in coefficients with similar positive sign, magnitude, and significance to those in Table 4. The cross-price correlation channel remains the dominant systemic risk driver for non-core banks, with coefficients around 1 and highly significant in columns 2, 4, and 6 of Table 5.

The main difference introduced by contract-level concentration measure is that interactions between *SimCore* and concentration are now positive and highly significant for core banks, providing evidence that portfolio concentration on the same contracts had a heterogeneous effect on their systemic risk. The large and significant coefficients of 1.497 for the interaction *SimCore * Concentrated* and of 4.623 for the interaction *SimCore * HHI*, both significant at 1%, suggest that portfolio similarity had a bigger effect on the systemic risk of core banks with more concentrated positions. By contrast, concentration did not change the effect of portfolio similarity on the systemic risk of non-core banks.

The negative relationship between core similarity and portfolio returns during the crisis somewhat obfuscates whether the crisis-related increase in similarity was caused by increased concentration or diversification. To answer this, Table 6 lists position concentration statistics by period and bank type. The table shows that, if anything, the crisis was accompanied by

a slight *decrease* in position concentration for both bank types. For example, the mean (contract-level) HHI fell from 50% to 46% for core banks, and from 65% to 62% for non-core banks. Both differences are of similar size and statistically significant at the 1% level. Based on this, there is no evidence that the 2009 spike in similarity across banks was triggered by increased concentration; rather, banks seem to have increased both their diversification *and* the degree of overlap between their diversified portfolios, in line with Wagner’s (2010, 2011) predictions. Systemic portfolio liquidation costs play a key role in Wagner’s theory, and it is indicative that banks increased their diversification precisely when portfolio liquidation costs became more systemic.

Table 6 reveals another heterogeneity across bank types when considered in the context of our previous results. In Table 5 (lines 3, 5, and 7), we observed that the interaction between concentration and similarity had a strong positive effect on core banks’ systemic risk during the normal period. At the same time, core banks were the less concentrated of the two groups. This shows that, while core banks showed a lower tendency towards concentration, to the extent that some were more concentrated than others, the more concentrated ones bet on the same set of contracts, thereby increasing their similarity and systemic risk. By contrast, non-core banks were more concentrated overall but less often concentrated on the same contracts, resulting in lower overall similarity and insignificant interactions with the similarity measure.

These findings significantly nuance our understanding of the concentration versus diversification debate on the origins of systemic risk, showing that aspects of both propositions may be true, but for different banks. On the one hand, core banks were both more similar to each other *and* more diversified overall, consistent with Wagner’s (2010, 2011) predictions. On the other hand, core banks were also the group where concentration had a significant heterogeneous effect on risk, i.e., where banks with more concentrated positions suffered a larger risk increase. This heterogeneity is not predicted by Wagner’s model. On the other hand, the fact that the more concentrated, non-core banks *less* often crowded on the same contracts goes against Menkveld’s (2017) example of concentration causing systemic risk.

Thus, it seems that existing models of systemic risk sources in derivatives markets do not account for the full variety of behaviors observed in these markets. Nonetheless, aspects of these theories are testable and provide some support for both. For example, while the heterogeneous effect of concentration for core banks is more consistent with the story in Menkveld (2017), banks' simultaneous increase in similarity and diversification during the crisis is more consistent with Wagner's (2010, 2011) theory of interbank commonality due to systemic portfolio liquidation costs. The timing of this diversification, precisely when liquidation costs became the most systemic, lends further evidence in favor of Wagner's theory.

5 Conclusion

Banks' participation in derivatives markets has been linked consistently to their systemic risk. This paper explores the role derivatives markets play in contributing to systemic risk by causing contemporaneous losses in different agents in the financial system. We take advantage of a data set containing the proprietary open positions of all participants in the Canadian futures market.

We begin by qualifying the existing literature with the observation that exactly half of the counterparties in a derivatives market will sustain simultaneous losses on any given day. Thus, from the standpoint of systemic risk, the important question is not whether same-day losses occur, but how they are distributed over the set of market actors. If contemporaneous losses are dispersed evenly across diverse actors in the financial system, the latter is likely more resilient to market turmoil compared to the case where simultaneous losses accrue exclusively to the important core banks. Whether this will occur or not depends on whether core banks' positions are similar and whether they as a group collectively take opposite positions against the periphery. This leads us to explore the influence of position similarity and other portfolio characteristics on the correlation across bank portfolios.

We perform two decompositions. First, we decompose banks' portfolio comovements

into those coming from portfolio similarity versus those due to cross-price correlations, and explore whether these effects are heterogeneous across bank type. Secondly, we explore whether it is portfolio concentration or diversification that drives systemic risk through the portfolio similarity channel, thus providing an empirical test of the competing propositions of Wagner (2010, 2011) and Menkveld (2017). Finally, we explore whether the 2008 crisis changed any of our previous results.

We find that core banks as a whole took positions against the periphery, thereby increasing their systemic risk. Portfolio similarity was the main driver of correlated returns for core banks, whereas cross-price correlations were the main driver behind the return correlations of non-core banks, revealing an important heterogeneity across bank type. Further decomposing the portfolio similarity channel into concentration and diversification components shows that, although core banks were more diversified overall, those among them with more concentrated positions had a larger systemic risk increase; thus, position concentration on specific contracts had a heterogeneous effect within that group. By contrast, we find no evidence that concentration had a heterogeneous effect on the systemic risk of non-core banks.

This nuances our understanding of the crowding versus diversification debate on the origins of systemic risk (e.g., Menkveld 2017 versus Wagner 2010, 2011). Our findings show that aspects of both theories may be true, but for different types of banks. While core banks were less prone to position concentration, those among them with more concentrated positions flocked to the same contracts, which further increased their systemic risk. By contrast, non-core banks had more concentrated positions overall, but they concentrated less often on the same contracts, so concentration did not add to their systemic risk; both bank types further diversified during the 2008 crisis. The latter outcome is more consistent with Wagner’s (2010, 2011) diversification theory based on portfolio liquidation costs, suggesting that they may play an under-appreciated role in generating systemic risk across banks.

Finally, we note the utility of decomposing portfolio similarity into concentration versus diversification. This has allowed us to perform an empirical test of the competing propo-

sitions of Menkveld (2017) versus Wagner (2010, 2011), which is impossible to do with aggregate crowding measures such as Menkveld’s (2017) *CrowdIx*. The purpose of the latter is to detect linear combinations of assets causing portfolios to comove, allowing both concentration and diversification to drive the comovement. By contrast, measuring cosine similarity and concentration separately and on the bank level allows us to distinguish between these two systemic risk sources and draw heterogeneous conclusions across bank types.

6 Tables

Table 1: Alphabetical List of Financial Institutions

Name	Country of Headquarters
Bank of Montreal*	Canada
Bank of Nova Scotia*	Canada
CIBC*	Canada
Desjardins	Canada
J.P. Morgan	USA
Laurentian Bank	Canada
Merrill Lynch	USA
MF Global	USA
National Bank*	Canada
Newedge Canada Inc.	Canada
Royal Bank of Canada*	Canada
Toronto Dominion (TD)*	Canada

This table presents the financial institutions with active firm accounts in the Canadian futures market during the sample period (January 2, 2003, to March 31, 2011) and their country of headquarters. Core banks (members of the Big Six) are flagged with an asterisk (*). An account is considered active if it has a non-zero open position on at least half of the trading days during the sample period. Any subsidiaries are subsumed under the parent institution and their positions consolidated with those of the parent if participating through more than one entity.

Table 2: Position Summary Statistics

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.	Mean	Std. Dev.		
2003	4.7	6.6	3.1	6.3	0.7	0.9	458.1	635.3	341.1	680.8	294.3	367.4						
2004	5.3	8.6	5.3	9.8	1.1	1.1	518.4	837.4	579.3	1080.9	508.0	538.8						
2005	5.2	8.9	3.8	6.4	1.0	0.9	507.5	859.6	430.0	727.8	567.5	496.7						
2006	5.0	9.0	2.6	3.2	0.7	0.8	479.0	861.5	288.4	357.7	446.6	543.1						
2007	6.7	13.2	1.7	1.9	0.5	0.6	635.6	1256.2	197.0	218.3	368.8	445.5						
2008	5.1	5.8	2.1	2.4	0.6	0.9	492.0	567.2	247.1	288.5	468.1	736.2						
2009	10.6	11.9	2.9	4.4	0.3	0.5	1053.9	1180.6	358.1	533.4	204.7	277.6						
2010	8.7	10.3	2.7	4.2	0.5	0.6	862.4	1014.4	334.0	516.9	370.3	442.2						
2011	11.9	18.3	2.4	2.6	0.4	0.5	1168.9	1796.9	294.9	319.6	331.6	371.0						

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 proprietary accounts. 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 3: Summary Statistics: Futures by Contract Type

Contract Type	Mean	Std. Dev.	Min	Max	N
Panel A. Contract Returns, Percent					
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
Panel B. Cross-Asset Return Correlations					
<i>Non-Crisis</i>					
$\rho(R_{BAX}, R_{CGB})$	0.78				1,731
$\rho(R_{BAX}, R_{SXF})$	-0.14				1,731
$\rho(R_{CGB}, R_{SXF})$	-0.17				1,731
<i>Crisis (Sept. 1, 2008 – Dec. 31, 2009)</i>					
$\rho(R_{BAX}, R_{CGB})$	0.70				334
$\rho(R_{BAX}, R_{SXF})$	-0.33				334
$\rho(R_{CGB}, R_{SXF})$	-0.34				334

This table presents summary statistics for the sampled futures' daily returns (Panel A) and their mutual correlations (Panel B). Panel A reflects the returns of all individual contracts within each type. To compute correlations across contract types in Panel B, individual contracts within each type are lumped together (as within-type returns are over 99% correlated), resulting in a smaller number of observations N . The sample period is from January 2, 2003, to March 31, 2011.

Table 4: Position Characteristics Regressions with Product-level HHI

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable: $\rho_{t+1}(R_i, R_c)$	Core	Non-Core	Core	Non-Core	Core	Non-Core
SimCore	1.324*** (0.365)	0.603*** (0.130)	1.380*** (0.339)	0.638*** (0.112)	1.041* (0.550)	0.619 (0.485)
SimCore*Crisis	-0.537* (0.281)	-0.466 (0.292)	-0.746*** (0.267)	-0.490** (0.229)		
SimCore*Concentrated	0.489 (0.362)	0.0261 (0.139)				
SimCore*Concentrated*Crisis	-1.048** (0.501)	0.00104 (0.232)				
SimCore*HiConcentrated			0.514 (0.519)	-0.227 (0.224)		
SimCore*HiConcentrated*Crisis			-0.450 (0.727)	0.232 (0.546)		
SimCore*HHI					0.725 (0.867)	-0.0224 (0.623)
SimCore*HHI*Crisis					-1.480*** (0.471)	-0.650** (0.291)
XP _{CGB_i,BAX_c}	0.145 (0.605)	0.135 (0.262)	0.251 (0.566)	0.143 (0.296)	0.230 (0.610)	0.129 (0.267)
XP _{BAX_i,CGB_c}	-0.298 (0.677)	1.037*** (0.323)	-0.354 (0.730)	1.068*** (0.318)	-0.322 (0.690)	1.038*** (0.318)
Constant	0.0645 (0.0875)	0.0381 (0.0250)	0.0662 (0.0892)	0.0357 (0.0259)	0.0626 (0.0882)	0.0377 (0.0258)
Observations	11,384	9,426	11,384	9,426	11,384	9,426
R ² (within-group)	.18	.12	.17	.12	.18	.12
Number of groups	6	6	6	6	6	6
Bank-level FE	Yes	Yes	Yes	Yes	Yes	Yes

This table shows the effect of bank similarity to the core and other portfolio characteristics on its return correlation against the core, $\rho_{t+1}(R_i, R_c)$. *SimCore* is bank i 's cosine similarity to the core, defined in Section 2. *Concentrated* is a dummy equal to 1 if the product-level HHI is above its median value. *HiConcentrated* is a dummy equal to 1 if the product-level HHI exceeds its 75th percentile. HHI is a Herfindahl concentration index calculated on the product level. *Crisis* is a dummy equal to 1 from September 1, 2008 to December 31, 2009. XP_j are cross-price correlation controls explained in Section 2, included for assets whose cross-price correlation is significantly different from 0. The estimation is panel OLS with Driscoll-Kraay standard errors. The sample period is from January 2, 2003, to March 31, 2011. The in-sample members are listed in Table 1.

Table 5: Position Characteristics Regressions with Contract-level HHI

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable: $\rho_{t+1}(R_i, R_c)$	Core	Non-Core	Core	Non-Core	Core	Non-Core
SimCore	1.001*** (0.342)	0.626*** (0.134)	1.340*** (0.326)	0.626*** (0.111)	-0.675 (0.463)	0.803*** (0.265)
SimCore*Crisis	-0.451 (0.298)	-0.463* (0.253)	-0.713*** (0.262)	-0.476** (0.228)		
SimCore*Concentrated	1.497*** (0.446)	-0.0221 (0.158)				
SimCore*Crisis*Concentrated	-1.277** (0.650)	-0.0200 (0.233)				
SimCore*HiConcentrated			1.076* (0.623)	-0.0738 (0.166)		
SimCore*HiConcentrated*Crisis			-2.163* (1.253)	0.121 (0.322)		
SimCore*HHI					4.623*** (1.041)	-0.395 (0.432)
SimCore*HHI*Crisis					-1.520*** (0.483)	-0.988** (0.405)
XP _{CGB_i,BAX_c}	0.493 (0.528)	0.137 (0.265)	0.311 (0.556)	0.134 (0.283)	0.544 (0.513)	0.181 (0.280)
XP _{BAX_i,CGB_c}	-0.0854 (0.650)	1.036*** (0.322)	-0.324 (0.663)	1.038*** (0.319)	0.335 (0.631)	0.985*** (0.322)
Constant	0.0469 (0.0803)	0.0377 (0.0248)	0.0554 (0.0883)	0.0372 (0.0255)	0.0163 (0.0773)	0.0354 (0.0252)
Observations	11,384	9,426	11,384	9,426	11,384	9,426
R ² (within-group)	0.23	0.12	0.18	0.12	0.26	0.12
Number of groups	6	6	6	6	6	6
Bank-level FE	Yes	Yes	Yes	Yes	Yes	Yes

This table shows the effect of bank similarity to the core and other portfolio characteristics on its return correlation against the core, $\rho_{t+1}(R_i, R_c)$. *SimCore* is bank i 's cosine similarity to the core, defined in Section 2. *Concentrated* is a dummy equal to 1 if the contract-level HHI is above its median value. *HiConcentrated* is a dummy equal to 1 if the contract-level HHI exceeds its 75th percentile. HHI is a Herfindahl concentration index calculated on the contract level. *Crisis* is a dummy equal to 1 from Sept. 1, 2008 to Dec. 31, 2009. XP_j are cross-price correlation controls explained in Section 2, included for assets whose cross-price correlation is significantly different from 0. The estimation is panel OLS with Driscoll-Kraay standard errors. The sample period is from January 2, 2003, to March 31, 2011. The in-sample members are listed in Table 1.

Table 6: Contract-level HHI Summary Statistics by Bank Type and Period

	Mean	Std. Dev.	Bank-day obs.
<i>Panel A: Core Banks</i>			
Non-Crisis	.495	.002	10,079
Crisis	.457	.005	1,751
Crisis–Non-Crisis Diff.	-.038***	.006	
<i>Panel B: Non-Core Banks</i>			
Non-Crisis	.648	.002	9,672
Crisis	.613	.006	1,923
Crisis–Non-Crisis Diff.	-.034***	.006	

This table presents summary statistics for the product-level HHI index across bank types and time periods. Panel A shows the results for the core banks, and Panel B – for the non-core banks. The crisis period is defined as September 1, 2008 to December 31, 2009. The sample period is from January 2, 2003, to March 31, 2011.

References

- [1] Acharya, V., 2001. A Theory of Systemic Risk and Design of Prudential Bank Regulation. Mimeo, London Business School.
- [2] Acharya, V. V., Pedersen, L. H., Philippon, T., and Richardson, M. 2017. Measuring Systemic Risk. *Review of Financial Studies* 30(1) pp. 2–47.
- [3] Acharya, V., and Yorulmazer, T., 2005. Limited Liability and Bank Herding. Mimeo, London Business School.
- [4] Acharya, V., and Yorulmazer, T., 2006. Cash-in-the-Market Pricing and Optimal Resolution of Bank Failures. *Review of Financial Studies* 21(6), pp. 2705–2742.
- [5] Acharya, V., and Yorulmazer, T., 2007. Too Many to Fail – an Analysis of Time-Inconsistency in Bank Closure Policies. *Journal of Financial Intermediation* 16, 1–31.
- [6] Bech, M.L., Bergstrom, C. T., Rosvall, M., and Garratt, R. J. 2015. Mapping Change in the Overnight Money Market. *Physica A* 424, pp. 44–51.

- [7] Brunnermaier A., and Adrian, T. 2017. CoVaR. *American Economic Review* 106(7), pp. 1705–1741.
- [8] Brunnermeier, M. K., Dong, G., and Palia, D. 2020. Banks’ Non-Interest Income and Systemic Risk. *Review of Corporate Financial Studies* 92, pp. 229–255.
- [9] 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.
- [10] Chellel, K., and Hodges, J. 2018. The Triple Jeopardy of a Chinese Math Prodigy. Bloomberg News, Nov. 19, 2018.
- [11] Cruz Lopez, J., Harris, J., Hurlin, C., and Perignon, C. 2017. CoMargin. *Journal of Financial and Quantitative Analysis* 52(5), pp. 2183–2215.
- [12] Driscoll, J. C., and Kraay, A. C. 1998. Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data. *Review of Economics and Statistics* 80(4), pp. 549–560.
- [13] Duffie, D., and Zhu, H. 2011. Does a Central Clearing Counterparty Reduce Counterparty Risk? *Review of Asset Pricing Studies* 1(1), pp. 74–95.
- [14] Duygun, M., Tunaru, R., and Viotto, D. 2021. Herding by Corporates in the US and the Eurozone through Different Market Conditions. *Journal of International Money and Finance* 110, pp. 1–29.
- [15] Falato, A., D. Iercosan, and Zikes, P. 2019. Banks as Regulated Traders. Finance and Economics Discussion Series 2019-005. Washington: Board of Governors of the Federal Reserve System.
- [16] Girardi, G., Hanley, K. W., Nikolova, S., Pelizzon, L., and Sherman, M. G. 2018. Portfolio Similarity and Asset Liquidation in the Insurance Industry. SAFE Working Paper No. 224, Goethe University, Frankfurt am Main.

- [17] Hwang, S., and Salmon, M. 2004. Market Stress and Herding. *Journal of Empirical Finance* 11(4), pp. 585–616.
- [18] Lepetit, L., Nys, E., Rous, P., and Tarazi, A. 2008. Bank Income Structure and Risk: An Empirical Analysis of European Banks. *Journal of Banking & Finance* 32, pp. 1452–1467.
- [19] Menkveld, A. J. 2017. Crowded Positions: An Overlooked Systemic Risk for Central Clearing Counterparties. *Review of Asset Pricing Studies* 7, pp. 209–242.
- [20] Nofsinger, J. R., and Sias, R. W. 1999. Herding and Feedback Trading by Institutional and Individual Investors. *Journal of Finance* 54 (6), pp. 2263–2295.
- [21] Perez-Saiz, H., and Li, F. 2018. Measuring Systemic Risk across Financial Market Infrastructures. *Journal of Financial Stability* 34, pp. 1–11.
- [22] Raykov, R. 2021. Systemic Risk and Collateral Adequacy: Evidence from the Futures Market. *Journal of Financial and Quantitative Analysis*, forthcoming.
- [23] Sias, R., Turtle, H. J., and Zykaj, B. 2016. Hedge Fund Crowds and Mispricing. *Management Science* 62(3), pp. 764–784.
- [24] Silva-Buston, C. 2019. Systemic Risk and Competition Revisited. *Journal of Banking & Finance* 101, pp. 188–205.
- [25] Sjostrom, W. 2009. The AIG Bailout. *Washington and Lee Law Review*, 66, pp. 943–991.
- [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] Wagner, W. 2010. Diversification at Financial Institutions and Systemic Crises. *Journal of Financial Intermediation* 19, pp. 373–386.
- [30] Wagner, W. 2011. Systemic Liquidation Risk and the Diversity-Diversification Trade-Off. *Journal of Finance* 66(4), pp. 1141–1175.
- [31] Welch, I. 2000. Herding among Security Analysts. *Journal of Financial Economics* 58(3), pp. 369–396.
- [32] Zhou, C. 2013. The Impact of Imposing Capital Requirement on Systemic Risk. *Journal of Financial Stability* 9(3), pp. 320–329.