Federal Reserve Bank of New York

December 2015
Volume 21 Number 2

Economic Policy Review

Federal Reserve Bank of New York research publications—the *Economic Policy Review, Current Issues in Economics and Finance, Staff Reports,* and the annual publications catalogue—are available at no charge from the Research Group's website.

www.newyorkfed.org/research

Follow us on Twitter: *@NYFedResearch*

Federal Reserve Bank of New York Economic Policy Review

December 2015
Volume 21 Number 2

Contents

A Primer on the GCF Repo® Service:

Repurchase agreements, or repos, are commonly used by financial entities to access money markets. GCF Repo®, a financial service provided by the Fixed Income Clearing Corporation (FICC), is a particular type of repo in which trades are executed anonymously, with FICC acting as a central counterparty and guaranteeing settlement. In this primer, which consists of an introduction and two articles, the authors explore the effects on GCF Repo of ongoing reforms to the settlement procedures for another type of repo, tri-party repo. Key areas of focus are the impact of the reforms on the use of intraday credit to settle GCF Repo transactions and the strategies followed by dealers in trading GCF Repo—strategies that need to be taken into account when gauging the risks of potential changes to clearance and settlement procedures.

39 The Equity Risk Premium: A Review of Models

Fernando Duarte and Carlo Rosa

The authors estimate the equity risk premium (ERP)—the expected return on stocks in excess of the risk-free rate—by combining information from twenty models for the period 1960-2013. They begin their analysis by categorizing the models into five classes: trailing historical mean, dividend discount, cross-sectional estimation, regression analysis using valuation ratios or macroeconomic variables, and surveys. They find that an optimal weighted average of all models places the one-year-ahead ERP in June 2012 at 12.2 percent, close to levels reached in the mid- and late 1970s, when the ERP was highest in the study sample. The authors note, however, that there is considerable uncertainty in ERP point estimates. The interquartile range across models is 11.6 percent on average, although it reached 6.8 percent in 2012, the lowest level in the study sample. By employing differences across models, the authors argue that the ERP in 2012 is elevated mainly because Treasury yields are low, not because the expected future cash flows from stocks are high.

Adam Copeland

A Primer on the GCF Repo® Service: Introduction

1. Background

Repurchase agreements, or repos, are widely used by financial entities to access money markets. Primary dealers, for example, reported financing $1.9 trillion of securities using repo on July 31, 2013.[1] This primer, which consists of this introduction and two articles, focuses on a particular type of repo, the General Collateral Finance Repo (GCF Repo®). The Fixed Income Clearing Corporation (FICC) introduced this financial service in 1998[2] as a cost-effective way for securities dealers to exchange securities and cash (Fleming and Garbade 2003).

The two articles in this primer concentrate on different aspects of the GCF Repo service. The first article, "The Financial Plumbing of the GCF Repo Service," focuses on how GCF Repo trades are cleared and settled and describes how GCF Repo is affected by the current reforms to the settlement of repos. In particular, the authors lay out the various ways that intraday credit was used pre-reform to facilitate the settlement of GCF Repo trades and why this use of credit was problematic. They then describe the reforms that have been, or are scheduled to be, implemented and the effect of these reforms on the use of intraday credit.

The second article, "An Empirical Analysis of the GCF Repo Service," focuses on how dealers use this financial service. After describing the various strategies that dealers may employ, the authors quantify the predominance of these strategies. For example, they describe the types of dealers seeking funding through GCF Repo and the amount of cash typically borrowed. They also explain how dealers use GCF Repo in conjunction with their other repo transactions, in normal times and during periods of stress.

Although the two articles focus on different aspects of GCF Repo, they are complementary, because the strategies that dealers may follow in trading GCF Repos are influenced by the clearance and settlement procedures in place. Furthermore, in order to gauge the risks of potential changes to the clearance and settlement of GCF Repo, it is important to take into account how GCF Repos are traded.

Both articles rely upon a basic understanding of GCF Repo, so we provide an overview of the topic in this introduction (see also Ingber [2003]). We begin by broadly describing repurchase agreements and then focus on the institutional details of GCF Repo. We end by summarizing the main benefits of GCF Repo.

[1] See the August 8, 2013, release of the Primary Dealer Statistical Releases, published by the Federal Reserve Bank of New York, available at http://www.newyorkfed.org/markets/primarydealers.html.

[2] GCF Repo® is a registered service mark of the Fixed Income Clearing Corporation.

Adam Copeland is an officer at the Federal Reserve Bank of New York.
adam.copeland@ny.frb.org

The author declares that he has no relevant or material financial interests that relate to the research described in this introduction.

2. Repurchase Agreements

Repos are essentially a pair of separate but related transactions between two entities: an agreement to buy a security now (which constitutes the opening leg of the repo), joined with an agreement to sell back the same security in the future at a specified price (the closing leg).[3] Apart from their treatment under bankruptcy, repos often resemble collateralized loans, with the difference in the price of the security across the two legs of the repo transaction translating into an interest rate. In addition to their use in sourcing funds, repos are also used to invest temporary cash balances, for arbitrage purposes, and as a tool for implementing a variety of other strategies. Adopting the view of a repo as a collateralized loan, we designate the entity purchasing (and then reselling) the securities as the cash investor. The other entity is labeled the collateral provider.

Two important elements of the repo agreement are negotiating which securities can be posted as collateral and negotiating the total value of securities to be posted as collateral. When repos are used for funding, which is the more usual case, industry practice is for the value of collateral to exceed the amount of cash.[4] This difference is called the margin, and it measures the amount by which a repo is overcollateralized. The margin protects the cash investor in the event that the collateral provider defaults on its repurchase agreement, by providing a buffer against fluctuations in the value of the securities posted as collateral or a loss in value associated with the quick liquidation of securities.[5]

In Exhibit 1, we present a repo between a collateral provider and a cash investor. The trade terms are a loan of \$1 billion, secured by U.S. Treasury securities, of overnight maturity, with an interest rate of 10 basis points and a margin of 2 percent. In settling this repo, the collateral provider delivers \$1.02 billion in U.S. Treasuries to the cash investor at date t in exchange for \$1 billion in cash. Then at date t+1, the investor releases the Treasuries back to the collateral provider in return for \$1,000,002,777.78 in cash (where \$2,777.78 = \$1 billion × 1/360 × 10 basis points).

[3] Technically, under the master agreements for repos, the requirement is to repurchase the same or equivalent securities, where "equivalent" means fungible.

[4] Repos can also be used to access the collateral markets. In these cases, industry practice in setting margins and interest rates is different from that described in this article, reflecting the different motivation driving the transaction (see Garbade [2006]).

[5] Cash investors cannot make gains when selling these securities. Proceeds from any sale above the principal amount of the repo are returned to the trustee managing the estate of the defaulted dealer. Furthermore, if the proceeds from the sale do not cover the principal, the cash lenders get an unsecured claim on the estate of the defaulted dealer.

Exhibit 1

Settlement of a Repo

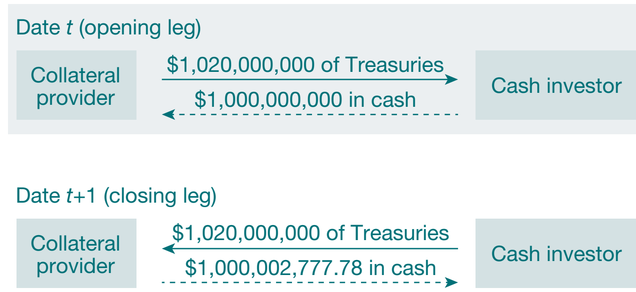

Repo trade details

Term: overnight	Rate: 10 basis points
Collateral type: Treasuries	Margin: 2 percent
Principal: \$1 billion	

Note: \$1,000,002,777.78 = \$1,000,000,000 + (\$1,000,000,000 × 1/360 × 10 basis points).

Repos are flexible enough that cash investors can specify that a particular security be posted as collateral. It is common, however, for a cash investor to specify only that the security underlying the transaction belong to a general asset class (as in the example in Exhibit 1). In these *general collateral* repos, the cash investor agrees to lend cash against a class of securities, such as U.S. Treasuries, agency debentures, or investment-grade corporate bonds.

In the United States, repo trades typically settle in one of two ways. The first is on a bilateral basis. In this case, the collateral provider and cash investor negotiate the repo agreement, which specifies the principal amount, interest rate, margin, term, and class of acceptable collateral. Here, each entity needs to have procedures in place to ensure proper settlement.

In particular, on the opening leg of the repo, the cash investor has to ensure that the collateral provider has sent the securities in the agreed-upon asset class and that the value of the securities satisfies the margin requirement. Tracking and valuing the securities posted as collateral is a major task, especially given the multitude of securities available for use as collateral.[6]

A second way to settle repos, called tri-party, is by using a settlement service provided by a third party—namely, a clearing bank. Tri-party repo services are offered by JPMorgan Chase (JPMC) and Bank of New York Mellon

[6] For example, according to statistics provided to the author by the Federal Reserve Bank of New York, there were over one million mortgage-backed securities outstanding and almost ten thousand agency debt securities outstanding in the Fedwire® Securities Service as of December 2012.

(BNY Mellon), the two institutions that provide clearing and custodial services to the large securities dealers in the U.S. market. Just as in the bilateral case, tri-party repos are negotiated between the collateral provider and the cash investor. Once the terms are agreed upon, the settlement details are transmitted to the clearing bank. The clearing bank then settles the repo agreement on its books, taking care to ensure that the details of the repo agreement are met. In particular, the clearing banks track and value the securities used as collateral and ensure that the proper margining requirements are fulfilled when settling a trade. The clearing banks do not broker transactions or help negotiate the terms; their role is limited to the clearance and settlement of these trades.[7]

Tri-party repos are almost always general collateral repos and they are used by securities dealers to raise funds from cash investors, such as money market mutual funds and investment managers.[8] According to market participants, tri-party repo is one of the main tools through which dealers fund themselves. Indeed, more than $1.6 trillion of collateral was posted for tri-party repo trades on July 10, 2013.[9]

The GCF Repo service exists alongside the tri-party repo market, but is dependent on it. In the next section, we describe the institutional details of this financial service. We then discuss the tight connection between GCF Repo and tri-party repo trades.

3. The GCF Repo® Service

GCF Repo differs from standard repos in a number of ways. Because it is offered by the Fixed Income Clearing Corporation, only institutions deemed eligible by FICC can negotiate GCF Repo trades.[10] While GCF Repo could potentially involve different types of participants, it is used primarily by securities dealers. For expositional clarity, then, we refer to institutions trading GCF Repos as securities dealers, or dealers. GCF Repos are negotiated through interdealer brokers (IDBs) on a blind basis.[11] Dealers tell an IDB the terms under which they are willing to borrow or lend cash. The IDB then tries to broker a trade while maintaining each dealer's anonymity. Once a trade has been brokered, the IDB submits the trade details to FICC, which, acting as a central counterparty, interposes itself and becomes the legal counterparty to each side of the repo transaction for settlement purposes. In this way, GCF Repo provides a way for dealers to anonymously negotiate repos among themselves. Furthermore, dealers do not face counterparty risk from one another, because of FICC's role as a central counterparty.

To protect itself against the risk of a dealer default, FICC, in addition to having eligibility requirements, requires dealers trading GCF Repo to post collateral to a clearing fund.[12] Because dealers post collateral to the clearing fund and because of the guarantee provided by FICC, GCF Repo trades do not include a separate margin requirement. Rather, the value of securities posted as collateral is equal to the amount of cash lent.

To improve liquidity for these trades and to simplify settlement, FICC standardizes GCF Repo trades by defining the acceptable classes of securities used as collateral. Dealers negotiating a GCF Repo transaction are limited to ten general asset classes for collateral (see Table 1). Currently, however, only nine collateral classes are traded, because there are no longer any securities that fall into the FDIC-Guaranteed Corporate Bonds collateral class.[13]

FICC provides netting services for dealers that negotiate GCF Repo contracts. At the end of each trading day, FICC computes for each dealer and each of the general collateral categories the value of securities the dealer has promised to deliver and the value that has been promised to the dealer. The difference between these two values, the net position of a dealer in a collateral category, is settled.

In Exhibit 2, we provide an example of GCF Repo between dealer A and dealer B, to highlight the differences between GCF Repo and the standard repo described in Exhibit 1. In the first step, the dealers negotiate, anonymously, through an IDB. The IDB then sends settlement instructions to FICC,

[7] See Copeland, Martin, and Walker (2014) for more details on tri-party repo.

[8] See Alkan et al. (2013) for more information on cash investors.

[9] The Federal Reserve Bank of New York publishes aggregate statistics on tri-party repo and GCF Repo trades on the Tri-Party Repo Infrastructure Reform webpage, available at https://newyorkfed.org/banking/tpr_infr_reform.html.

[10] Visit http://www.dtcc.com/about/businesses-and-subsidiaries/ficc.aspx for more information about FICC.

[11] In December 2012, 120 entities were eligible to trade GCF Repos. A list of eligible financial entities can be found at http://www.dtcc.com/client-center/ficc-gov-directories. Look in the "FICC GSD Member Directory" for those

Footnote 11 (continued)
members with the "Repo Netting" and "GCF" service designations. The IDBs are also listed; look for those members with "Broker Account" tags next to their name.

[12] The formula used to determine how much collateral a dealer needs to deliver into the clearing fund is laid out in FICC's government securities division rulebook, which is posted on FICC's Rules and Procedures webpage, available at http://www.dtcc.com/legal/rules-and-procedures.aspx.

[13] The Federal Deposit Insurance Corporation's Debt Guarantee Program, developed during the recent financial crisis, generated this special class of corporate bonds. This program is no longer active. For more information, see http://www.fdic.gov/regulations/resources/TLGP/.

Table 1

General Collateral Classes in GCF Repo

Fannie Mae and Freddie Mac adjustable-rate MBS
Fannie Mae and Freddie Mac fixed-rate MBS
FDIC-guaranteed corporate bonds[a]
Ginnie Mae adjustable-rate MBS
Ginnie Mae fixed-rate MBS
Non-MBS U.S. agency securities
U.S. Treasuries with maturities of ten years or less
U.S. Treasuries with maturities of thirty years or less
U.S. Treasury inflation-protected securities
U.S. Treasury STRIPs

Source: Depository Trust and Clearing Corporation (DTCC).

Notes: MBS is mortgage-backed securities; FDIC is Federal Deposit Insurance Corporation; STRIP is separate trading of registered interest and principal. An example of a non-MBS U.S. agency security is agency debentures.

[a] FDIC-guaranteed corporate bonds are no longer a collateral class in GCF Repo.

Exhibit 2

Trade Negotiation and Settlement of a GCF Repo

Notes: IDB is independent broker dealer. FICC is Fixed Income Clearing Corporation. $1,000,002,777.78 = $1,000,000,000 + ($1,000,000,000 × 1/360 × 10 basis points).

which novates the trade (in other words, substitutes the old contract for a new one), becoming the legal counterparty to both dealer A and dealer B for settlement purposes. In this example, at date *t*, dealer A sends $1 billion in Treasury securities to FICC in return for $1 billion in cash. FICC then sends the $1 billion in Treasury securities to dealer B in return for $1 billion in cash. On date *t*+1, these flows are reversed, with the securities being returned to dealer A and the cash plus interest being returned to dealer B, with FICC acting as intermediary. This example is highly stylized; the details of GCF Repo settlement are provided in "The Financial Plumbing of the GCF Repo Service."

GCF Repo is tightly connected to tri-party repo. The opening leg of both types of trades is settled same day on the books of the clearing banks. Further, GCF Repo is settled before tri-party repo, allowing dealers to easily rehypothecate collateral obtained from GCF Repo into tri-party repo—in other words, reuse it for their own borrowing.

4. Benefits of GCF Repo

We conclude with an enumeration of the benefits of the GCF Repo service. A primary benefit of the GCF Repo service is its enhancement of intermediation. As described in more detail in the second article in this series, "An Empirical Analysis of the GCF Repo Service," dealers use GCF Repo to intermediate between cash investors and other dealers. In particular, dealers that can access funding at a low cost may borrow more than they need from cash investors and then lend these extra funds in GCF Repo to dealers with a high cost of funding, taking advantage of the benefit that FICC provides as a central counterparty.

This intermediation is beneficial because it lowers dealers' cost of funding when investors are reluctant to lend to them directly. Without this intermediation, the dealers with a high cost of funding would otherwise need to raise funds through other (more expensive) means, or delever. A lower cost of funding makes dealers more competitive and likely results in lower prices of financial services for households and nonfinancial firms.

An additional benefit of this intermediation occurs in times of stress. In discussions, market participants state that a stressed dealer is likely to rely more upon GCF Repo as a source of funds, taking advantage of FICC acting as a central counterparty. This benefit, however, depends crucially upon FICC managing its risks appropriately.

Two other general benefits associated with GCF Repo are the reduction in transaction costs and the enhancement of liquidity in the interdealer repo market.[14] As described in Fleming and Garbade (2003), relative to standard bilateral repo arrangements, the design of GCF Repo provides these benefits by allowing 1) both legs of the repo to be netted, 2) the repo dealer to decide which collateral to deliver fairly late in the day, and 3) collateral to be substituted easily. The combination of these benefits should spill over and enhance the liquidity of the larger dealer-customer repo market.

[14] Ingber (2003) provides a longer and more detailed list of benefits provided by GCF Repo. In particular, he notes that GCF Repo allows for a longer period of time during the day to trade general collateral repos, lowers the costs of settlement, and allows for "expanded access to the general collateral market to encompass a wider range of financial entities" (p. 48).

References

Alkan, L., V. Chakrian, A. Copeland, I. Davis, and A. Martin. 2013. "Magnifying the Risk of Fire Sales in the Tri-Party Repo Market." Liberty Street Economics, July 13. Available at http://libertystreeteconomics.newyorkfed.org/2013/07/magnifying-the-risk-of-fire-sales-in-the-tri-party-repo-market.html.

Copeland, A., A. Martin, and M. Walker. 2014. "Repo Runs: Evidence from the Tri-Party Repo Market." Journal of Finance 69, no. 6 (December): 2343-80.

Fleming, M., and K. Garbade. 2003. "The Repurchase Agreement Redefined: GCF Repo®." Federal Reserve Bank of New York Current Issues in Economics and Finance 9, no. 6 (June).

Garbade, K. 2006. "The Evolution of Repo Contracting Conventions in the 1980s." Federal Reserve Bank of New York Economic Policy Review 12, no. 1 (May): 27-42.

Ingber, J. 2003. "Gets Confusing Fast: A Review of the GCF Repo® Service." RMA Journal 85, no. 8 (May): 46-51.

Paul Agueci, Leyla Alkan, Adam Copeland, Kate Pingitore,
Caroline Prugar, and Tyisha Rivas

The Financial Plumbing of the GCF Repo® Service

1. Introduction

General Collateral Finance Repo (GCF Repo®) is a popular, well-established service for securities dealers.[1] Its structure provides a way for dealers to exchange government securities for cash among themselves in an anonymous way. Further, the Fixed Income Clearing Corporation, which offers the GCF Repo service, provides netting services and acts as a central counterparty. These benefits have led dealers to enter into a large number of GCF Repo contracts; for example, in the first quarter of 2013, average daily trading was almost $500 billion and average daily net settlement exceeded $250 billion.

GCF Repo trades are cleared and settled on the books of the two large clearing banks, JPMorgan Chase (JPMC) and Bank of New York Mellon (BNY Mellon), with each bank using its own tri-party repo settlement platform. During the 2007-09 financial crisis, weaknesses were revealed in both banks' tri-party repo settlement procedures, and thus in GCF Repo. After the financial crisis, regulators and market participants formed the Tri-Party Repo Reform Task Force, with the aim of producing recommendations to improve the stability of the two banks' tri-party repo settlement platforms (Task Force 2010).[2]

Most of the task force's recommendations focused on reducing the settlement systems' reliance on intraday credit to settle trades. Prior to reform, these systems depended heavily on the clearing banks providing unlimited intraday credit to the institutions entering into tri-party repo and GCF Repo contracts. One of the main goals of the reforms was to develop settlement systems where much smaller amounts of intraday credit are provided and where it is provided in a less discretionary way.

The pre-reform systems were worrisome for two reasons. First, as long as a dealer had securities at the clearing bank to serve as collateral, the clearing bank was willing to extend intraday credit to that dealer to settle tri-party repo trades. Given the size of the larger dealers (with tri-party books of easily more than $100 billion), there was potential for each of the clearing banks to extend an enormous amount of intraday credit relative to its capital base. This situation raised the risk that a clearing bank that could not absorb the impact of a failing dealer would itself be destabilized, leading to an interruption of funding and payment services for all of its other clients. The task force recommended that clearing banks limit intraday credit extensions to no more than 10 percent of the value of a dealer's total tri-party book. With these limits in place, market participants and regulators can be more confident that a clearing bank can handle the default of a large dealer on its tri-party repo obligations.

Second, the discretionary nature of the clearing banks' extension of credit was problematic. In times of stress, a clearing bank

[1] GCF Repo® is a registered service mark of the Fixed Income Clearing Corporation.

[2] For more details on the Tri-Party Repo Reform Task Force and its work, see http://www.newyorkfed.org/banking/tpr_infr_reform.html.

Paul Agueci and Leyla Alkan are senior associates, Adam Copeland an officer, Kate Pingitore a former senior bank examiner, Caroline Prugar a former bank examiner, and Tyisha Rivas a former payments policy analyst at the Federal Reserve Bank of New York. Correspondence: adam.copeland@ny.frb.org

To view the authors' disclosure statements, visit https://www.newyorkfed.org/research/author_disclosure/ad_epr_primer-on-the-gcf-repo.html.

The authors thank Vic Chakrian, Antoine Martin, and Denise Schmedes for their comments. The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

might be unwilling to take on the risk of extending intraday credit to a troubled dealer. Such a move, however, would effectively push the dealer into bankruptcy because it would lose access to planned-for funds. The task force recommended the removal of this discretion. With the reforms, clearing banks' credit extensions are now committed, capped, and collateralized.

Although the clearing banks have made progress in reducing dealers' reliance on intraday credit, most of the improvements have been aimed at the settlement of tri-party repo trades, and not GCF Repo trades. As a result, GCF Repo trades are still settled under systems that rely heavily on the provision of unlimited intraday credit to function.

In this article, we describe in detail the settlement of GCF Repo and the reliance of the settlement process on intraday credit. First, we provide an overview of how GCF Repos are negotiated and cleared. Then we describe how GCF Repo trades were settled up until the first quarter of 2012, the pre-reform state. Since the first quarter of 2012, however, a number of changes have been made to the settlement process as part of the aforementioned reforms to tri-party repo; and so, lastly, we describe the current settlement process. We start with the pre-reform settlement process because an understanding of the former process is important to appreciating how and why the settlement process is changing with the reforms.

The task force also raised concerns about the risk of fire sales. A fire sale is the rapid sale of securities in amounts large enough to cause a temporary decrease in the market prices of those securities. Fire sales are particularly problematic because of the externalities they impose on other dealers. A dealer that is forced to sell its securities in a fire sale faces the difficulty that its actions decrease the prices of the securities, reducing their value. However, other dealers may also be affected if the price declines force those dealers to mark down the same securities on their balance sheets (for example, through mark-to-market accounting practices) or provide more securities as collateral. Such actions may even lead these dealers to sell securities, further depressing prices and reinforcing the fire-sale effect. Little progress has been made on this issue within tri-party repo, however, reflecting both the focus on other objectives and the difficulty in mitigating this risk.[3]

In the latter part of our discussion on the current settlement process, we use the framework presented in Begalle et al. (2013) to discuss the risks of fire sales in GCF Repo. We argue that fire-sale risks in GCF Repo are substantially mitigated by the role of FICC as the central counterparty. An important assumption underlying this argument, however, is the ability of FICC to adequately manage dealer defaults.

2. Overview of GCF Repo

Repos are essentially a pair of related transactions between two entities: an agreement to buy a security now (which constitutes the opening leg of the repo), joined with an agreement to sell back the same security in the future at a specified price (which constitutes the closing leg of the repo). Often, repos effectively serve as collateralized loans, where the difference in the price of the security across the two legs of the transaction translates into an interest rate.

GCF Repo is a service offered by the Fixed Income Clearing Corporation (FICC) and used by dealers that are netting members of FICC's Government Securities Division. The GCF Repo differs from a standard repo in that the trade is completed on a blind-brokered basis, where dealers negotiate their trades through interdealer brokers (IDBs) and thus preserve their anonymity. These repos are general collateral repos, meaning that dealers agree that the securities to be posted as collateral are only required to be in a specific asset class, as opposed to being specific securities. FICC defines ten collateral classes that can be used by dealers, the most popular of which are U.S. Treasuries with maturities of thirty years or less and Fannie Mae and Freddie Mac fixed-rate mortgage-backed securities.[4]

FICC provides two additional types of services for those dealers trading GCF Repos. First, it acts as a central counterparty, absorbing all counterparty risk in these trades. Second, it provides netting services, allowing dealers to offset their repo and reverse repo positions for trades where the securities posted as collateral are of a similar type.[5] These features make the GCF Repo service attractive to dealers, compared with a standard bilateral repo (Fleming and Garbade 2003).

Below, we describe how GCF Repo trades are negotiated and cleared. The details of settlement are then discussed in Sections 3 (the pre-reform state) and 4 (the current state).

[3] See the February 13, 2014, statement "Update on Tri-Party Repo Infrastructure Reform," by the Federal Reserve Bank of New York.

[4] For a list of the collateral classes, see Table 1 of "A Primer on the GCF Repo® Service: Introduction" in this volume.

[5] From the perspective of a dealer, repos are trades in which that dealer has promised to deliver securities against cash, whereas reverse repos are trades in which that dealer has promised to deliver cash against securities.

2.1 How Dealers Trade through IDBs

At a high level, dealers enter into a trade by working through an IDB to negotiate with one another anonymously (see the top panel of Exhibit 1, "Trading"). A dealer states its trading terms to the IDB, which then helps the dealer execute a trade by finding another dealer willing to take the other side, all the while masking the dealers' identities. IDBs offer two basic platforms to help execute trades: electronic and voice.

An electronic platform allows a dealer to see and accept the bid/offer rates that dealers have posted that day according to collateral class and tenor. Further, these platforms have a variety of features that help dealers keep their positions hidden and enable them to manage large orders. Typically, these platforms are used to execute trades quickly on a take-it-or-leave-it basis.

A voice platform involves communicating with a person, namely a broker, at an IDB. Although dealers still may be able to see information about other dealers' bid/offer rates, executing a trade on the voice platform requires going through a broker. An advantage of the voice platform over its electronic counterpart is the ability for a dealer, through a broker, to negotiate the terms of trade. A disadvantage is the slower speed at which trades are executed. Market participants report that electronic platforms are typically used in the morning, when most of the GCF Repo trading occurs and execution speed is highly valued. Voice platforms are typically used in the afternoon, when there is less trading overall and dealers value the ability to negotiate terms.

Exhibit 1

Overview of GCF Repo Clearance and Settlement

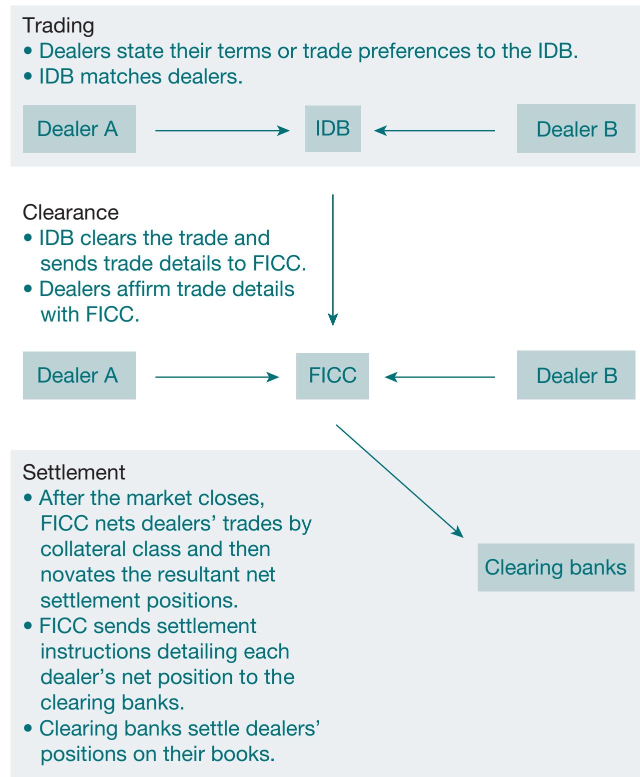

Notes: IDB is independent broker-dealer. FICC is Fixed Income Clearing Corporation.

2.2 The Clearance of Trades

Once two dealers have booked a trade, the IDB becomes the legal counterparty to each dealer. The IDB begins the clearance process by reviewing and confirming the trade details with the dealers (see the middle panel of Exhibit 1, "Clearance"). The IDB, for example, corrects data entry errors that are identified through the confirmation process. The IDB then sends the trade details to FICC and the two dealers. FICC accepts GCF Repo trade details between 7 a.m. and 3 p.m. eastern time.

Once FICC receives the trade details from the IDB, it guarantees the trade, limiting the risk faced by the IDB as the legal counterparty to the trade. As part of the clearance process, dealers are supposed to affirm the details of the trade to FICC. After a trade is affirmed, changes to that trade can only be made if both dealers agree to cancel and rebook the trade. The IDB remains the counterparty to both sides of the trade until the netting process is completed and the resulting net settlement positions are novated, after which FICC becomes the legal counterparty to each dealer for settlement purposes (and the IDB's settlement obligations are eliminated through the netting process).

After two dealers agree to a trade, it takes an IDB only about ten minutes to clear the trade and send the trade details back to the dealers and FICC. In contrast, dealers may take much longer to affirm a trade to FICC. Typically, IDBs will contact dealers if trades are not affirmed within two hours. Dealers can delay only so long; after 3 p.m., FICC automatically affirms all trades it has received from IDBs.[6]

After 3 p.m., when FICC stops accepting trade details from the IDBs, FICC nets down each dealer's trades in a collateral class into a net position. As a consequence of netting, a dealer that promised to deliver and receive securities within the same collateral class over the course of the day only has to settle its net position at the end of the day. FICC then sends settlement

[6] FICC encourages dealers to affirm trades before the 3 p.m. deadline.

instructions to the clearing banks (see the lower panel of Exhibit 1, "Settlement"). Finally, dealers' net positions are settled on the books of the clearing banks at the end of the day.

3. Settlement of GCF Repo Trades—Pre-Reform

In this section, we describe the GCF Repo settlement process as of the first quarter of 2012, or the pre-reform state. We focus on two main processes: the end-of-day settlement and the morning unwind. The end-of-day settlement is the process by which all outstanding GCF Repo positions are settled. The morning unwind is the process whereby the clearing banks return the securities held as collateral for all GCF Repo positions to the repo dealers and return the cash amount to the reverse repo dealers. An advantage of the morning unwind is that it provides dealers with full and unimpeded access to their securities during the business day.

As described above, the clearing banks receive instructions from FICC to settle dealers' net positions, where a net position is the difference between the value of repos and the value of reverse repos that a dealer has traded for a particular collateral class. Dealers have either a zero or nonzero net position for each collateral class. For the nonzero net positions, the dealer has an obligation either to deliver securities that fall within the acceptable class of collateral to FICC and receive cash, or to deliver cash and receive securities.

The clearing banks begin the settlement process by creating "shells," which specify dealers' net repo positions for each of the collateral classes in the GCF Repo service. From the dealer's perspective, a repo shell represents an obligation to deliver securities against cash.

With the creation of these shells, the collateral allocation process begins. In the following section, we describe this process under the simplifying assumption that both dealers involved in the GCF Repo settlement process use the same clearing bank. For this intrabank case, both the securities and cash payments are moving on the books of a single clearing bank. We then detail the extra steps needed to settle GCF Repo allocations that are interbank (settlement between the two clearing banks) in a separate section.

It is important to re-emphasize that the settlement processes described below and illustrated on the accompanying exhibits reflect the pre-reform case (in other words, as of March 2012). With the tri-party reforms, the clearing banks have instituted changes to their settlement processes for GCF Repo trades. These changes are described in Section 4.

3.1 Intrabank GCF Repo Settlement—Pre-Reform

We begin by describing the settlement process for GCF Repo positions when both the repo dealer and the reverse repo dealer use the same clearing bank. We break the settlement process into two parts: First, the end-of-day settlement on day t, when the securities are delivered in exchange for cash. Second, the morning unwind on day $t+1$, when the cash and collateral are returned to the reverse repo and repo dealers, respectively. For overnight trades, end-of-day settlement is the opening leg of the repo and the morning unwind is the closing leg. For trades of longer maturity, the unwind is a mechanism that allows dealers easy and unconstrained access to their securities during the business day. From the perspective of the clearing banks, the term of the GCF Repo trade is irrelevant because all trades are unwound every morning.

End-of-Day Settlement

At the end of the trading day, the clearing banks receive instructions from FICC detailing how to settle dealers' net positions in GCF Repo. For each clearing bank, the settlement process begins with the bank informing dealers of their GCF Repo obligations and creating the appropriate repo shells.[7] The repo dealers then start to allocate collateral from their securities accounts at the clearing bank to the repo shells. A repo shell is said to be "filled" once a dealer has allocated enough securities to fulfill its collateral obligations for that shell. Once all dealers have filled their GCF Repo shells for a specific collateral class—say, Treasuries with a maturity of thirty years or less—the clearing bank moves all of these allocated securities to FICC's securities account at that clearing bank (see Exhibit 2 for a schematic of this process).[8] Simultaneously, the clearing bank credits the relevant dealers' cash accounts and debits FICC's cash account. Because FICC does not typically have cash in its account at the clearing bank, the clearing bank extends intraday credit to FICC to enable this leg of the settlement process, backed by the securities posted as collateral for the GCF Repo positions (see Stage 1 in Exhibit 2).

[7] Copeland et al. (2012) provide details of how dealers allocate collateral to tri-party repo trades. The same methods can be used to allocate collateral to GCF Repo trades because both types of trades are settled on the same tri-party repo settlement platform.

[8] Both clearing banks have the operational capability to move the allocated securities from the dealer to FICC on a shell-by-shell basis. For operational efficiency, however, the clearing banks wait until all the dealers have filled their GCF Repo shells for a specific collateral class, and then move these allocated securities to FICC's account.

Exhibit 2

Intrabank GCF Repo End-of-Day Settlement

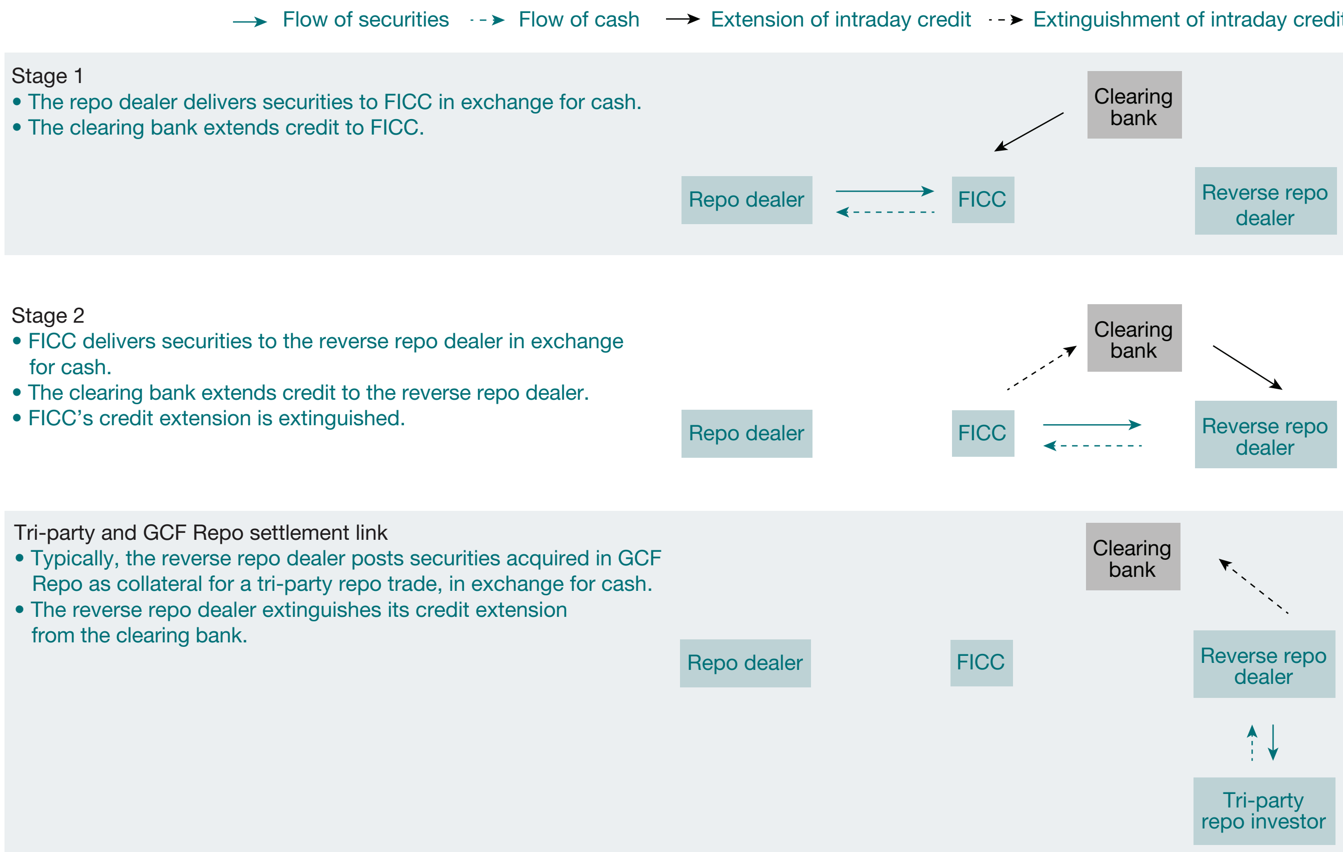

Notes: This exhibit describes the pre-reform settlement process. FICC is Fixed Income Clearing Corporation.

The clearing bank then allocates this set of securities from FICC's securities account into the repo shells characterizing FICC's obligations to deliver collateral to the reverse repo dealers. Note that because of the netting process, the allocation of these securities is not preordained by the day's trading activity. Simultaneously, the clearing bank credits the FICC cash account and debits the reverse repo dealers' cash accounts (see Stage 2 in Exhibit 2). To enable this leg of the settlement process, the clearing bank extends intraday credit to the reverse repo dealers. This credit extension is backed not only by the GCF Repo-related securities posted as collateral, but also by all the unencumbered securities the reverse repo dealers hold at the clearing bank.

The flow of cash from the reverse repo dealers to FICC allows FICC to extinguish its credit extension from the clearing bank. The end result of this process is that securities have moved from the repo dealers' accounts to the reverse repo dealers' accounts, through FICC's account. Simultaneously, there is a corresponding reverse flow of cash. While the movement of securities and cash through FICC's account is a crucial step in the settlement process, typically neither the securities nor cash reside in FICC's account for a significant amount of time.

This settlement process requires the extension of credit by the clearing bank to both FICC and the reverse repo dealers. We label the extension of credit to FICC as *frictional*, because it is extinguished once the end-of-day settlement leg of the GCF Repo position is settled. In comparison, the extension of credit to the reverse repo dealers is extinguished only after the dealers source funds elsewhere—for example, from an investor in the tri-party repo market (see the bottom-right-hand corner of Exhibit 2).[9]

[9] The rehypothecation of GCF Repo-obtained collateral into a tri-party repo trade will not, by itself, generate enough cash to fully pay off the clearing bank's credit extension to the reverse repo dealer, because there are margin requirements for tri-party repo trades. The dealer, then, would need to post more collateral in a tri-party repo trade in order to acquire the necessary amount of cash.

Exhibit 3

Intrabank GCF Repo Morning Unwind

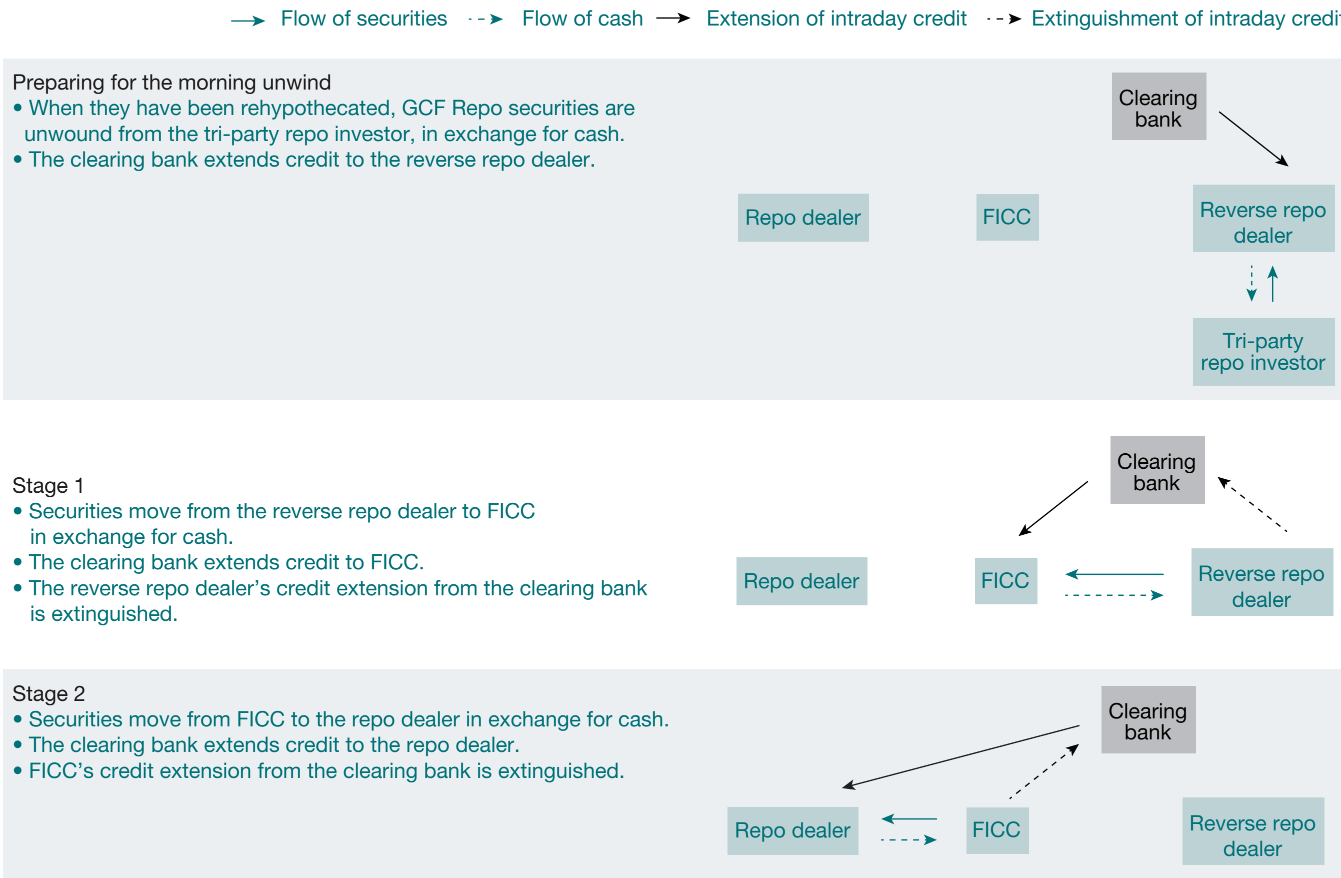

Notes: This exhibit describes the pre-reform settlement process. FICC is Fixed Income Clearing Corporation.

Morning Unwind

Every morning, at approximately 6:30 a.m., the clearing banks begin to unwind all GCF Repo positions, returning collateral to the repo dealers and cash to the reverse repo dealers. Unwinding a GCF Repo position essentially follows the same steps as the end-of-day settlement, but in reverse order. Hence, the first step to the unwind is to ensure that all GCF-related securities are back with the reverse repo dealers (see Exhibit 3 for a schematic of this process). If these securities have been used as collateral in other transactions (for example, rehypothecated to tri-party repo), then the clearing bank extends credit to the reverse repo dealers and recalls the GCF Repo-related securities by substituting cash in place of the desired securities (see "Preparing for the morning unwind" at the top of Exhibit 3). With the securities back in the reverse repo dealers' accounts, the clearing bank begins unwinding the GCF Repo positions. Intraday credit is extended to FICC and the securities are sent to FICC's account in exchange for cash (see Stage 1 of Exhibit 3). With this transfer, the clearing bank's extension of credit to the reverse repo dealers is extinguished (except for the possible differences in margin requirements).

Once the securities are in FICC's account, the clearing bank extends credit to the repo dealers. The securities are then returned to the repo dealers in exchange for cash. The cash is used to extinguish the clearing bank's credit extension to FICC (see Stage 2 of Exhibit 3).

At the end of the unwind, collateral and cash have been returned to the repo and reverse repo dealers, respectively. Dealers now have full access to their portfolios of securities, which they can use for regular trading purposes. In facilitating this unwind, the clearing bank extended intraday credit to both FICC and the repo dealers. As in the end-of-day settlement case, the extension of credit to FICC is frictional. In contrast, the clearing bank extends intraday credit to the repo dealers for the duration of the day. (See Appendix A on net free equity for more details on how the clearing banks manage their credit risk to dealers.) Usually, the dealers wait to extinguish this credit extension until the end of the day, when they are settling their tri-party repo and GCF Repo trades. A straightforward way to extinguish this credit extension at the end of the day is to simply execute an offsetting GCF Repo or tri-party repo trade.

3.2 Interbank GCF Repo Settlement—Pre-Reform

We now extend the above description for the case where the repo and reverse repo dealers use different clearing banks. A key feature of interbank GCF Repo settlement is that the securities posted as collateral by the repo dealer never leave the books of that dealer's clearing bank. This feature forces the clearing banks to coordinate their settlement processes to ensure that all cash flows and credit extensions are properly collateralized. The securities remain on the book of the repo dealer's clearing bank because the system of transferring government securities between institutions, Fedwire® Securities Service, closes at 3 p.m., before GCF Repo settlement begins.[10] Furthermore, it would not be operationally efficient to move securities back and forth across the clearing banks when they unwind each morning.

End-of-Day Settlement

Mirroring the intrabank case, we begin with end-of-day settlement. Suppose that there is a repo dealer at clearing bank 1 (CB1) and a reverse repo dealer at clearing bank 2 (CB2). As in the intrabank case, the repo dealer starts the settlement process by allocating securities to its GCF Repo shell. Once the repo dealer has filled its GCF Repo shell for a specific collateral class, clearing bank 1 moves these securities to FICC's account, extends credit to FICC, and deposits cash into the repo dealer's account (see Exhibit 4 for a schematic of this process). These securities are then moved to a segregated account, which serves as FICC's CB2 account on the books of clearing bank 1. Because the credit extension is secured by the underlying securities, clearing bank 1's credit extension is redirected to this segregated account (see Stage 1 of Exhibit 4).

A message is then sent from clearing bank 1 to clearing bank 2 listing the securities in this segregated account.

[10] Fedwire® is a registered service mark of the Federal Reserve Banks.

Clearing bank 2 then creates copies of these securities, called *securities clones*, in FICC's CB1 account at clearing bank 2, and a cross-clearing bank lien is placed on the securities residing in FICC's CB2 account on the books of clearing bank 1 (ensuring that these securities are not used elsewhere). On clearing bank 2's books, these securities clones are then allocated to FICC's account. To facilitate this transfer, clearing bank 2 extends credit to FICC and deposits cash into FICC's CB1 account at clearing bank 2. The clones are then allocated to the repo shells characterizing FICC's obligations to deliver collateral to the reverse repo dealer. Clearing bank 2 extends credit to the reverse repo dealer and the credit extension to FICC is extinguished (see Stage 2 of Exhibit 4). At this point, FICC has received an intraday credit extension from clearing bank 1 (secured by the securities residing in FICC's CB2 account at clearing bank 1) and has a positive cash balance at clearing bank 2 (residing in FICC's CB1 account at clearing bank 2). To extinguish the credit extension from clearing bank 1, FICC requests that clearing bank 2 wire the cash from FICC's CB1 account at clearing bank 2 to clearing bank 1, using the Fedwire® Funds Service (which is open until 6:30 p.m.). With this cash movement, FICC is once again "flat," in that neither clearing bank is extending intraday credit to it.

At the end of this process, securities (or their clones) have moved from the repo dealer to the reverse repo dealer through FICC, with cash flowing in the opposite direction. Similar to the intrabank case, the clearing banks have extended intraday credit to FICC and the reverse repo dealer to facilitate settlement. The credit extension to FICC is frictional but complicated, owing to its reliance on cross-clearing bank liens. The reverse repo dealer is left with an intraday credit extension from clearing bank 2. As before, this dealer can extinguish this credit extension in a number of ways, including by using the securities it received through GCF Repo to obtain cash in tri-party repo.

While the above example considers one repo dealer at one clearing bank and one reverse repo dealer at the other clearing bank, in reality there are often a number of interbank allocations with repo dealers (in other words, dealers obligated to deliver securities and receive cash) at both clearing banks. This means that, in practice, the clearing banks send information to one another about the securities being delivered by repo dealers. A crucial component of the interbank GCF Repo settlement system is this flow of information. In the pre-reform process, the clearing banks communicate with one another once in the settlement cycle. Specifically, only after repo dealers have filled their GCF Repo shells for all securities classes and these securities have been allocated to the other clearing bank's FICC account does one clearing bank send a message to the other clearing bank with the details necessary to complete settlement of the GCF Repo trades.

Exhibit 4

Interbank GCF Repo End-of-Day Settlement

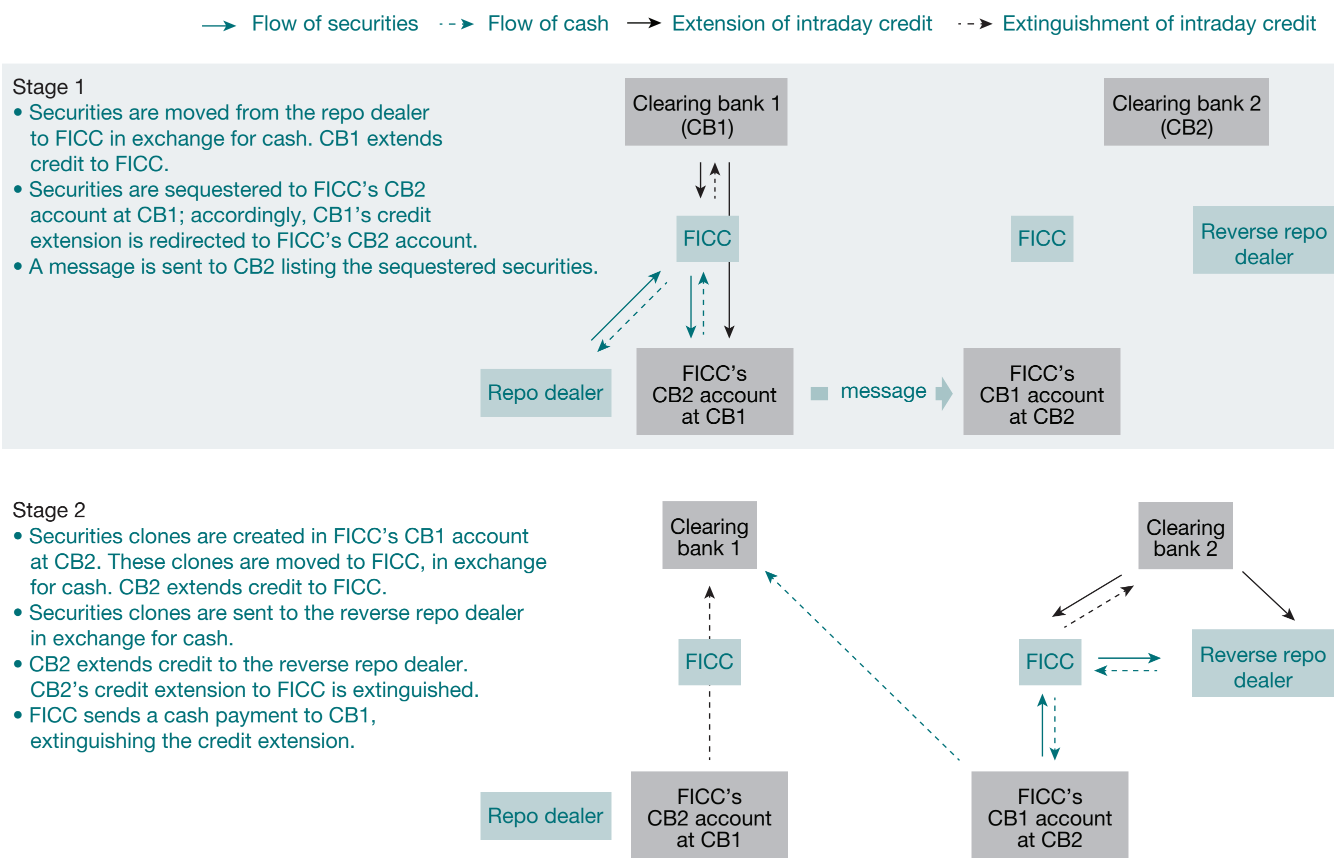

Notes: This exhibit describes the pre-reform settlement process. FICC is Fixed Income Clearing Corporation.

Having repo dealers at both clearing banks obligates FICC to send cash payments from JPMC to BNY Mellon, and vice versa. For operational efficiency, however, FICC sends only one payment between the clearing banks, where this payment is equal to the net flow of cash between the two clearing banks.

Morning Unwind

We now turn to the interbank GCF Repo unwind (see Exhibit 5 for a schematic of this process). Continuing from the example above, suppose that the repo dealer is at clearing bank 1 and the reverse repo dealer is at clearing bank 2. Recall that the actual securities reside on the books of clearing bank 1, in a segregated account (FICC's CB2 account at clearing bank 1) and clearing bank 2 uses clones of these securities on its books.

In most cases, clearing bank 2 begins the unwind by first extending credit to the reverse repo dealer and pulling back all GCF Repo-related securities that have been rehypothecated through tri-party repo using a securities-for-cash substitution mechanism (not shown in Exhibit 5).[11] Clearing bank 2 then extends credit to FICC and moves the securities clones from the reverse repo dealer to FICC's account. The corresponding movement of cash from FICC to the reverse repo dealer enables the dealer to extinguish the credit extension from the clearing bank (ignoring possible differences in margin requirements). The securities clones are then moved to FICC's CB1 account at clearing bank 2. Because clearing bank 2's credit extension to FICC is collateralized by the securities clones, the credit extension to FICC's CB1 account is redirected, as shown in Exhibit 5.

[11] This step is not necessary if the reverse repo dealer has not rehypothecated the securities.

Exhibit 5

Interbank GCF Repo Morning Unwind

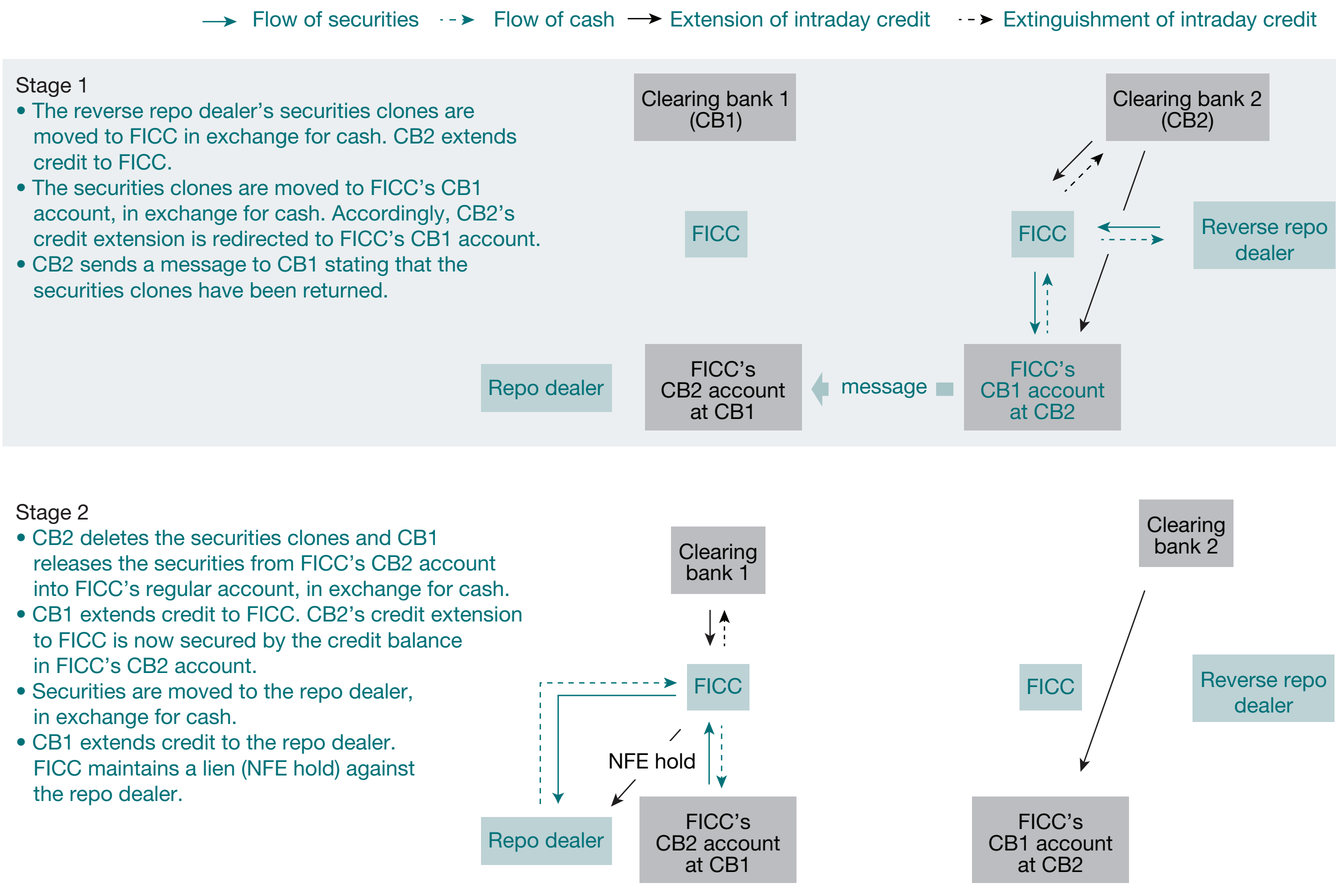

Notes: This exhibit describes the pre-reform settlement process. FICC is Fixed Income Clearing Corporation. NFE is net free equity.

Clearing bank 2 then sends a message to clearing bank 1 stating that all the securities clones have returned to FICC's CB1 account, and these clones are deleted. This enables clearing bank 1 to unwind the securities from FICC's CB2 account on clearing bank 1's books (see Stage 1 in Exhibit 5). Clearing bank 2's credit extension to FICC continues to be collateralized, using a cross-clearing bank lien, by the securities in FICC's CB2 account.

After receiving the message, clearing bank 1 moves the securities from FICC's CB2 account (the special segregated account) to FICC's account. Concurrently, clearing bank 1 debits FICC's account and credits FICC's CB2 account. The securities are then moved from FICC to the repo dealer. To facilitate this movement, clearing bank 1 extends credit to the repo dealer, where this credit extension is secured by a lien that FICC maintains on the repo dealer's unencumbered securities residing at clearing bank 1. This lien, or the net free equity (NFE) hold, is explained further in Appendix A.

So, at the end of the unwind, the securities have been fully unwound to the repo dealers and are available to be used by the dealers for other purposes. The repo dealers at clearing bank 1 have granted FICC a security interest in the unencumbered securities in their accounts, known as the NFE hold. Clearing bank 2 has also extended intraday credit to FICC, which is secured by clearing bank 2's cross-clearing bank lien on the credit balance in FICC's CB2 account at clearing bank 1. FICC is liable for extinguishing the credit extension at clearing bank 1. In the event the repo dealer fails to repay FICC, FICC would liquidate the NFE hold collateral to meet its obligation to clearing bank 1. The intraday credit extensions to FICC and to the repo dealers at clearing bank 1 are not frictional but rather last throughout the day, until the end-of-day settlement process.[12]

[12] With the completion of the unwind, both clearing banks have extended credit to FICC. Clearing bank 2's credit extension is secured by the cash

Exhibit 6

GCF Repo Timeline (Pre-Reform)

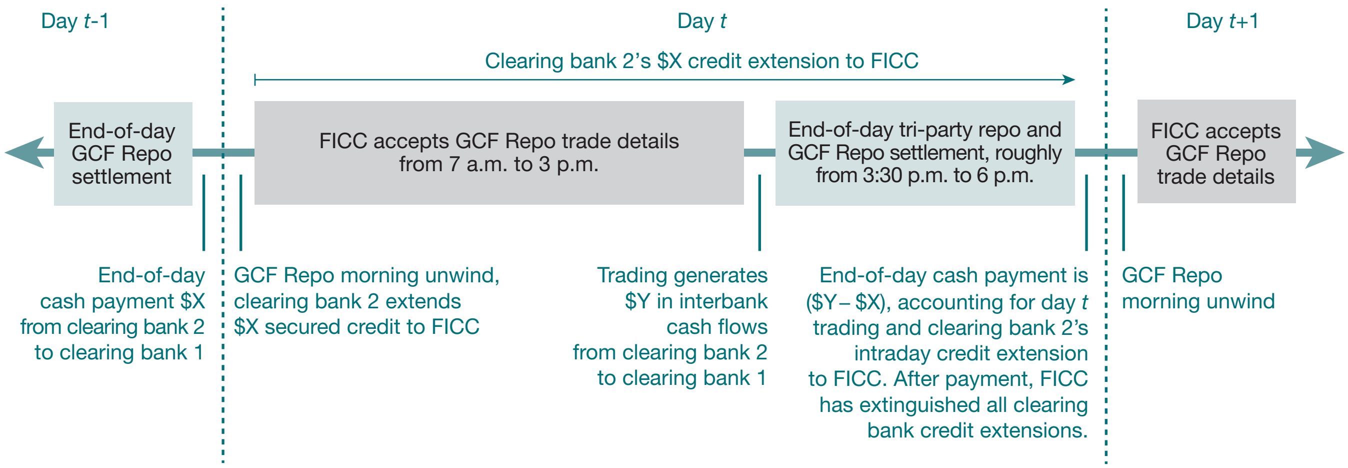

Note: FICC is Fixed Income Clearing Corporation.

The repo dealer typically extinguishes the credit extension from clearing bank 1 (and so lifts FICC's lien on the repo dealer's securities) at the end of the day, when it settles all of its tri-party repo and GCF Repo trades. As mentioned before, it is straightforward for the dealer to raise the necessary cash through another GCF Repo trade. The credit extension by clearing bank 2 to FICC is also extinguished during the end-of-day settlement of interbank GCF Repos.

3.3 Review of the Chronology of GCF Repo Trading, Clearance, and Settlement—Pre-Reform

To facilitate a better understanding of the interactions between trading, clearance, and settlement, in this section we illustrate the chronological flow of activity throughout the day. For GCF Repo, the day starts with the morning unwind, where collateral and cash are returned to the repo and reverse repo dealers, respectively, beginning at around 6:30 a.m. (see Exhibit 6). FICC begins accepting trade details from IDBs at 7:00 a.m.

The majority of trading is completed in the morning, with more than half of trades (in terms of volume) being completed within the first hour of trading. By 10 a.m. on a typical day, three-quarters of trading has been completed. FICC stops accepting trades from IDBs at 3 p.m., and shortly thereafter FICC begins the netting process. Dealers' net positions in GCF Repo are typically settled between 3:30 p.m. and 5:30 p.m.

Exhibit 6 illustrates the clearing banks' credit extensions to FICC that facilitate the unwinding of interbank GCF Repo positions. Suppose that, at the end of day *t*-1, dealers' trading strategies have resulted in dealers at clearing bank 1 sending, on net, $X of securities to dealers at clearing bank 2 in exchange for cash. Consequently, for the morning unwind on day *t*, clearing bank 2 needs to extend $X of intraday credit to FICC (see Stage 1 in Exhibit 5). As illustrated in Exhibit 6, this extension of credit by clearing bank 2 to FICC lasts throughout the day (roughly nine hours).

Now suppose that, on day *t*, trading results in dealers at clearing bank 1 sending, on net, $Y of securities to dealers at clearing bank 2 in exchange for cash. Rather than dealing with the $X and $Y credit extensions separately, FICC and the clearing banks settle the net amount ($Y – $X). To see how this works, consider when $X = $Y. To settle the $Y in net trading for this interbank case at the end of the day, FICC needs to deliver $Y in cash (in exchange for securities) to the group of day *t* repo dealers at clearing bank 1. Because $X = $Y, this cash is supplied entirely by the group of day *t*-1 repo dealers at clearing bank 1, which need to extinguish the credit extension from clearing bank 1 and so release FICC's lien on the *t*-1 repo dealers' securities. (Recall that these dealers received their collateral back in the morning of day *t*.) FICC then delivers the securities from the day *t* repo dealers at

Footnote 12 (continued)
sitting in FICC's CB2 account in clearing bank 1. Clearing bank 1's credit extension is secured by FICC's NFE hold on the relevant repo dealers.

clearing bank 1 to the group of reverse repo dealers at clearing bank 2, in exchange for cash. The $Y in cash that FICC receives is then used to extinguish clearing bank 2's $X credit extension (because $X = $Y) from that morning's unwind process.

For the special case of $X = $Y, no payments need to be made between the two clearing banks. When $Y is not equal to $X, however, FICC will end up with a credit at one clearing bank and an offsetting debit at the other clearing bank. In this case, a payment needs to be sent between the clearing banks to extinguish FICC's credit at the end of the day.

Typically, the payments by FICC to settle the net amount are small relative to the net amount of GCF Repos settled in the interbank case. Nevertheless, it is not uncommon for this net payment to be quite large. For example, when the net flow of cash across the clearing banks changes direction, a payment equal to the absolute value of X plus the absolute value of Y is required to extinguish the intraday credit extension to FICC.

4. Tri-Party Repo Settlement Reforms and GCF Repo

Having described the clearance and settlement of GCF Repo (as of the first quarter of 2012), we now turn to concerns with this financial plumbing. We focus on two potential issues: the heavy reliance on intraday credit to settle GCF Repo positions, including the unwind, and fire-sale risks related to this financial service.

4.1 Use of Intraday Credit to Settle GCF Repos

As reported in Section 1, a main focus of the tri-party repo reforms is to move the clearing banks from a settlement system in which unlimited and discretionary intraday credit is extended, to a settlement system in which intraday credit is capped and committed.

The concerns over clearing banks extending unlimited and uncapped credit continue to exist with the settlement procedures of GCF Repo.[13] During end-of-day settlement, the clearing banks are extending credit to the reverse repo dealers in both the intrabank and interbank cases (see Exhibits 2 and 4). Further, for the intrabank case, the clearing banks extend intraday credit to the repo dealer to facilitate the morning unwind (see Exhibit 3).

[13] See the February 13, 2014, statement "Update on Tri-Party Repo Infrastructure Reform" by the Federal Reserve Bank of New York, available at http://www.newyorkfed.org/newsevents/statements/2014/0213_2014.html.

The clearing banks also extend intraday credit to FICC to settle GCF Repo positions. For the end-of-day settlement in the intrabank and interbank cases, as well as during the morning unwind for the intrabank case, the clearing banks extend frictional credit to FICC. In Exhibits 2, 3, and 4, the frictional aspect of this credit extension is illustrated by the extinguishment of the credit extension to FICC at the end of that particular settlement process. The clearing banks also extend credit to FICC that is nonfrictional—this is done during the morning unwind in the interbank case (see Exhibit 5).

Alongside the tri-party repo reforms, FICC and the clearing banks have implemented (or plan to implement) changes that will reduce the amount of credit extended by the clearing banks to facilitate settlement of GCF Repo positions. In this section, we review these changes in the settlement process and explain the consequent reduction in the amount of credit extended by the clearing banks. We then highlight steps in the settlement process that still require the clearing banks to extend large amounts of intraday credit.

Updates to GCF Repo for the Intrabank Case

For the intrabank case, FICC and both clearing banks are in the process of making changes to GCF Repo that will reduce the amount of credit extended to dealers on typical days. One improvement to the settlement process that has already been implemented is the delay of the unwind from 6:30 a.m. to 3:30 p.m. (mirroring the tri-party repo reforms implemented in August 2011). The advantage of delaying the unwind is that credit extensions to the repo dealer, although still large, are for a much shorter length of time because they are extinguished with the end-of-day settlement process, which begins shortly after the unwind is completed.

Along with the delay of the morning unwind, the clearing banks implemented an intrabank collateral-substitution mechanism that enables dealers to access their securities held as collateral. This mechanism provides access by allowing dealers to replace securities being held as collateral with other securities of equal or greater value that satisfy the terms of the relevant repo contract. Recall that one of the main economic impetuses of the morning unwind is to allow dealers unimpeded access to their securities during the business day. With the collateral-substitution mechanism, dealers can continue to access their securities despite the lack of a morning unwind.[14]

[14] The delay of the morning unwind and the concurrent introduction of a collateral-substitution mechanism mirror what was done for tri-party repo trades as part of the tri-party repo reforms. A description of the delay in the unwind and new collateral-substitution mechanisms can be found in FICC's

A planned improvement to the settlement process is the use of rolling dealers' positions in GCF Repo, or switching to a "Net-of-Net" settlement process. Rolling positions requires the calculation of the net change from one day to the next for each dealer's position in each collateral class. The clearing bank then only settles the daily difference (see Appendix B for a detailed explanation of the rolling position settlement process). If dealers' net GCF Repo positions do not change much from day to day, this process could significantly reduce the amount of securities and cash required to flow among dealers to settle positions. FICC reports that fully implementing the new Net-of-Net process would result in an average reduction in amount settled of 76 percent.[15]

This proposed change in settlement would be both operationally efficient and beneficial in reducing the amount of intraday credit required to settle positions. A potential issue, however, is that if dealers change their trading strategies with the consequence that their net positions fluctuate considerably, the benefits gained through rolling positions, in terms of reducing the amount of credit necessary to settle trades, would be somewhat lessened.

Updates to GCF Repo for the Interbank Case

Less progress has been made on the interbank case than on the intrabank case. The current settlement system for interbank GCF Repo positions still requires the extension of nonfrictional credit to FICC. Reducing or eliminating the extension of intraday credit to settle these positions requires active engagement from the clearing banks, FICC, and the set of dealers that use the GCF Repo service.

For these interbank cases, a planned improvement to settlement is to partially, rather than fully, unwind in the morning.[16] Under the pre-reform system, securities are unwound to the repo dealers and cash is returned to the reverse repo dealers. Under the proposed new arrangement, securities will be unwound to FICC and the repo dealers will access their securities through a collateral-substitution mechanism.

This proposed settlement change impacts the nature of the intraday credit extended by clearing banks, but not the amount. Specifically, the pre-reform, or full, unwind is facilitated by the extension of credit to FICC by the reverse repo dealer's clearing bank and by FICC maintaining a NFE hold on the other clearing bank's repo dealer (see Stage 2 in Exhibit 5 for an illustration of this credit extension). The total amount of credit extended equals the total net position of all interbank GCF Repo trades (see Exhibit 6).

Under the proposed settlement changes, the clearing banks will continue to extend intraday credit to FICC but the credit will be secured by cross-clearing bank liens. These liens will be against specific securities or cash residing in FICC's account at the repo dealer's clearing bank. Importantly, the size of the credit extension will not be changed with these updates to the settlement system.

Where Does That Leave Us?

The GCF Repo settlement process remains overly reliant on intraday credit extensions by the clearing banks. As detailed above, these credit extensions are to dealers and FICC. Below, we analyze the current state of these credit extensions for the intrabank and interbank cases, laying out the difficulties in determining a solution.

For the intrabank case, the proposed process of rolling dealers' positions will require the clearing banks to extend relatively small amounts of credit to dealers under normal circumstances. As previously mentioned, as part of the tri-party repo reforms, the clearing banks plan to establish committed intraday credit lines to dealers. These facilities could be used to provide credit to repo dealers in the GCF Repo intrabank case. A potential problem, however, is that these credit extensions are capped and may be insufficient.[17]

Intrabank settlement also requires frictional credit to FICC. Compared with extending credit to dealers, extending credit to FICC involves different counterparty risks. Specifically, FICC is a financial market utility that has been designated as systemically important. How the clearing banks will handle extending intraday credit to FICC has not yet been determined. But it is important to avoid having a system in which dealers are provided with capped and committed lines of credit to facilitate settlement while FICC has unlimited and uncommitted credit. Such asymmetry in treatment could provide incentives to shift the costs of providing intraday credit from

Footnote 14 (continued)
proposed rule change to the Securities and Exchange Commission, SR-FICC-2011-05, available at http://www.sec.gov/rules/sro/ficc/2011/34-65213.pdf.

[15] See the September 17, 2013, Depository Trust and Clearing Corporation newsletter article "DTCC Improves GCF Repo® End-of-Day Processing to Mitigate Risk and Enhance Efficiencies," by Randy Spencer, available at http://www.dtcc.com/news/2013/september/27/dtcc-improves-gcf-repo-end -of-day-processing.aspx.

[16] The details of this proposed settlement change are given in FICC's proposed rule change to the SEC, SR-FICC-2011-05, cited above. In particular, see section II.B.4, "Substitution on Interbank GCF Repos," on pp. 10-11.

[17] See http://www.jpmorgan.com/pages/jpmorgan/is/products/clearing/bds/resourcecenter/finishline for mention of JPMC's plan to set up a committed and secured credit facility.

the dealers to FICC. For example, for end-of-day settlement, dealers obligated to deliver securities to GCF Repo could perform this action first, and receive cash from FICC, where this cash would be the result of an extension of credit from the clearing banks to FICC. Furthermore, dealers obligated to deliver cash to FICC could delay until the conclusion of the end-of-day settlement process. As a consequence, there would be an infusion of cash into dealers' accounts that could then be used by dealers to facilitate the settlement of their tri-party repo trades.[18] This result, however, effectively shifts the costs of providing intraday credit to settle tri-party repo and GCF Repo trades from dealers to FICC, a result that does little to enhance the stability of the tri-party repo settlement platform in times of stress.

There are many options available to the clearing banks, two of which side-step the issue by eliminating the extension of credit to FICC for the intrabank case. One approach is simply to require the reverse repo dealers to provide the necessary cash up front. A second approach is for the clearing bank to explicitly link the flows of securities and cash between the repo and reverse repo dealers, and so treat the movement of cash and securities through the FICC account (which stands between the repo and reverse repo dealers) as a temporary and intermediary step. With this second approach, the securities would only leave the repo dealer's account when the clearing bank has verified that the reverse repo dealer could provide the necessary amount of cash. With this settlement procedure, credit would not need to be extended to FICC to settle the trade.

For the interbank case, clearing banks extend credit to FICC to unwind all interbank GCF Repo positions. The proposed settlement change outlined earlier does not address this basic issue. There are two unusual aspects to this intraday credit extension to FICC. First, the amount of credit necessary to unwind these transactions is equal to the total net amount of interbank GCF Repo, which can be quite large. In recent years, this amount has been quite variable and has occasionally reached the tens of billions of dollars. Second, the amount of credit extended to FICC is not a result of FICC's actions, but rather of dealers' trading. Consequently, any restrictions on the amount of credit extended to FICC could only be enforced if constraints were placed on dealers' trading behavior. How the clearing banks will handle the intraday credit extensions to FICC to settle interbank GCF Repo trades has not been determined.

[18] This type of strategic behavior with respect to minimizing the costs of intraday credit can be seen with financial institutions using Fedwire Securities. In this security settlement system, the institution sending the security (and receiving cash) initiates the transaction. Given the obligation to deliver a security on a particular day, institutions may send the security early in the day in order to build up their cash balance at the Federal Reserve and so lower the probability of incurring intraday liquidity charges (Mills and Nesmith 2008).

4.2 Fire-Sale Risks

A main objective of tri-party repo settlement reforms is to reduce the risk of fire sales in tri-party repo (Task Force 2010). Little progress has been made on this issue, however, reflecting both the focus on other objectives and the difficulty in mitigating this risk. Borrowing the terminology of Begalle et al. (2013), we highlight two types of fire sales in tri-party repo that concern regulators.

First, there is the pre-default risk of fire sales. Stressed dealers may face difficulties raising funds in tri-party repo because investors may be uncomfortable with the counterparty risk. Losing funding in tri-party repo will cause stressed dealers to delever, selling securities in a bid to raise funds and meet their obligations. The sale of securities will likely cause prices to drop, making it even more difficult for the stressed dealer to raise enough cash to cover its obligations. Further, the fall in prices will impact the entire dealer community through mark-to-market accounting. In particular, the clearing banks use the latest set of prices to value the securities provided as collateral in tri-party repo trades. Falling prices will force all dealers to post more collateral in order to raise the same amount of cash. Enough of a price decline may cause more dealers to become stressed.

Second, there is post-default risk. When a dealer defaults in tri-party repo, its investors receive the securities posted as collateral. Given the large number and wide variety of securities posted, investors are unlikely to coordinate the sales of these securities. Instead, they will likely try to sell them quickly—and this disorderly rush to sell will likely lead to a fire sale.

Fortunately, the role of FICC as a central counterparty in GCF Repo should, in theory, mitigate both types of fire-sale risk. Pre-default risk arises because the entity lending cash is uncomfortable with counterparty risk. But GCF Repo trades are blind-brokered, with FICC standing in as the legal counterparty. With GCF Repo, then, the entities lending cash are not bothered by the possibility of trading anonymously with a stressed dealer.

An important caveat to the above discussion is that dealers must remain confident in FICC and its ability to manage its counterparty risk and absorb the default of a dealer. Conditional on FICC properly managing its counterparty risk (and dealers perceiving that FICC is doing so), there is no pre-default fire-sale risk associated with the GCF Repo service.

Post-default fire-sale risk is also likely to be less of a factor with GCF Repo than with tri-party repo. This is because the structure of the GCF Repo service means that only one entity, FICC, will liquidate the collateral received from a defaulting dealer. Hence, unlike in the tri-party repo market, where cash investors will likely sell the securities held as collateral in an uncoordinated fashion, FICC has the potential to liquidate collateral in an orderly fashion. This control does not completely neutralize the risk of post-default fire sales, however, because FICC faces constraints to quickly sell the securities held as collateral. But the risk is lower relative to the tri-party repo case (all else being equal) because FICC could potentially sell its securities in a coordinated way.

5. Conclusion

Given the popularity and widespread use of GCF Repo among securities dealers, it is important for market participants, regulators, and academics to fully understand the financial infrastructure underpinning this service. This article provides a detailed look at the clearance and settlement of GCF Repo trades, highlights the risks associated with a heavy reliance on intraday credit to settle GCF Repo trades, and discusses how FICC's role as a central counterparty reduces the risk of fire sales associated with this product.

By the end of 2014, both clearing banks had implemented a new settlement system for tri-party repo trades.[19] Some of these planned changes also reduced the reliance on intraday credit for the settlement of intrabank GCF Repo trades. There are, however, open issues regarding how GCF Repo trades will settle for the interbank case. The clearing banks and FICC need to further improve the settlement of interbank GCF Repo trades to minimize the use of intraday credit. Furthermore, whether and how the clearing banks extend intraday credit to FICC to facilitate settlement needs to be decided. These financial plumbing decisions are important because they will likely influence the extent to which dealers rely on intraday credit from the clearing banks and perhaps impact dealers' ability to conduct interbank GCF Repo trades. As such, all parties—the clearing banks, FICC, and the dealers that use GCF Repo—need to remain actively engaged in these issues.

[19] Links to each clearing bank's plans can be found on the Tri-Party Repo Infrastructure Reform webpage, available at http://www.newyorkfed.org/banking/tpr_infr_reform.html.

Appendix A: How the Clearing Banks Manage Intraday Credit Risks (Pre-Reform)

Net free equity (NFE) is a risk management tool used by both clearing banks to ensure that the intraday credit a clearing bank extends to dealers is secured by collateral to which that clearing bank has a right of offset in the case of default. A dealer's NFE is the difference between the value of cash and collateral the dealer holds in various accounts at the clearing bank, taking into account "haircuts," and the amount of intraday credit (overdrafts) the clearing bank is currently extending to it. (The haircut is the value of the collateral in excess of the value of the cash exchanged in a repo.)

Not all of a dealer's accounts are included in its NFE; for example, a dealer's segregated client accounts are excluded. Thus, NFE refers to the total value of cash and collateral to which a dealer has title and which is unencumbered by existing obligations to the clearing bank or others. As part of their risk management processes, the clearing banks continuously monitor each dealer's NFE to ensure that their extension of intraday credit does not push any dealer's NFE below zero.

GCF Repo allocations between two dealers that use different clearing banks create two main risk management challenges for FICC and the clearing banks. The first risk management obstacle occurs with the end-of-day settlement of GCF Repo, and the second occurs with the GCF Repo morning unwind. The solutions to these two risk management problems are different.

Starting first with the end-of-day settlement, the underlying problem is that the securities posted as collateral are not transferred across the clearing banks. This is because the Fedwire Securities system closes at 3 p.m., which is about the time that the GCF Repo settlement process begins. Even if Fedwire Securities were to remain open, it is operationally inefficient to move large numbers of securities across the two clearing banks every day.

As illustrated in Exhibit 4, the difficulty lies in the fact that one clearing bank (clearing bank 2 in the exhibit) has to extend credit to the reverse repo dealer in order to facilitate settlement of the reverse repo leg of the transaction. This securities-for-cash exchange, however, relies on securities residing on the books of the other clearing bank (clearing bank 1 in the exhibit).

The solution is to create a cross-clearing bank lien, whereby the relevant securities are placed in a segregated account on the books of clearing bank 1 (labeled "FICC's CB2 account" on the books of clearing bank 1 in Exhibit 4). Clearing bank 2 then creates copies of its securities (called securities clones) that serve as proxies for the securities sequestered in FICC's CB2 account at clearing bank 1. Under this arrangement, clearing bank 2 can now execute, on its books, the securities-for-cash exchange between FICC and the reverse repo dealer.

The second risk management problem occurs with the GCF Repo morning unwind. Recall that the goal of the unwind is to return securities to the repo dealer so that the dealer can use the securities throughout the trading day for other transactions. To facilitate the unwind, the reverse repo dealer's clearing bank (clearing bank 2 in Exhibit 5) extends intraday credit to FICC so that FICC can deliver cash to the reverse repo dealer in return for the securities collateralizing the GCF Repo position.

Initially, this credit extension is secured by the securities underlying the GCF Repo position (which are sequestered at FICC's CB2 account on the books of clearing bank 1). The goal of the unwind, however, is to transfer these securities back to the repo dealer. To accomplish this while still maintaining a secured intraday credit extension from clearing bank 2 to FICC, FICC and the clearing banks implement a lien called the "NFE hold." The NFE hold is a legal arrangement whereby FICC has a lien on dealers' NFE at each clearing bank, in this case the repo dealer at clearing bank 1. This means that FICC has a lien on all of a dealer's unencumbered securities in various accounts at the clearing bank.[1] The total amount of the lien is equal to the previous day's interbank net funds borrowed. The lien is placed only on those dealers at the net borrowing clearing bank and is allocated proportionately based on those dealers' net repo amounts.

Consider the example laid out in the table below. Here, the clearing bank 1 dealers have borrowed \$6 billion more than they have loaned (see the last row of the upper panel of the table). This cash comes from clearing bank 2 dealers and is the amount of interbank GCF Repo on day *t*. With the unwind on the morning of day *t*+1, FICC will impose a NFE hold on dealers A through E, proportionate to their total repo activity. The NFE hold on dealer A, for example, is equal to (2/28) × \$6,000,000,000 = \$428,571,429.

As explained in Section 4.1, a proposal has been made to replace the NFE hold with a cross-lien legal arrangement as part of a series of reforms aimed at improving the settlement of GCF Repo positions.

[1] FICC's lien is on all of a dealer's unencumbered securities, a set of assets that includes securities that are not eligible for GCF Repo.

Example Calculation of the NFE Hold (Pre-Reform)

Panel A: Clearing Bank 1 Dealers' GCF Repo Positions on Day *t*

Dealer	Total Repo Amount (Billions of Dollars)	Percentage of Total	Total Reverse Repo Amount (Billions of Dollars)	Net Position (Repo – Reverse Repo) (Billions of Dollars)
A	2	7	6	(-4)
B	4	14	3	1
C	9	32	1	8
D	7	25	7	0
E	6	21	0	6
F	0	0	5	(-5)
Total	28	100	22	6

Panel B: Clearing Bank 1 Dealers and the NFE Hold on the Morning of Day *t*+1

Dealer	NFE Hold (Dollars)
A	428,571,429
B	857,142,857
C	1,928,571,429
D	1,500,000,000
E	1,285,714,286
F	0
Total	6,000,000,000

Notes: NFE is net free equity. The "NFE hold" is equal to the dealer's repo amount as a percentage of total repo activity on the clearing bank multiplied by the net amount of interbank funds ($6 billion, in this example).

Appendix B: How the Settlement Process of Rolling Dealers' Positions Works

Consider the case in which a dealer increases from $10 billion to $11 billion its net position in GCF Repo backed by Treasuries with maturities of 30 years or less. With the rolling positions process, the dealer is obligated to deliver only $1 billion in Treasuries to complete the end-of-day settlement process. Recall that under the previous system, the clearing bank would unwind the entire $10 billion GCF Repo position in the morning and then settle $11 billion later that day. To unwind the position in the morning, the clearing bank would need to extend $10 billion in intraday credit.

The use of rolling positions dramatically reduces the reliance on intraday credit, for most cases. Dealers only need credit when they reduce a repo position or increase a reverse repo position. In both instances, the dealer will receive securities from FICC, against which the dealer must deliver cash.

Importantly, for these two instances, the dealer only needs credit for the change in its net position (hence, this type of settlement is referred to as Net-of-Net). Historically, dealers have maintained similar net positions from day to day, a pattern that suggests that rolling positions will dramatically reduce the amount of intraday liquidity needed to settle GCF Repo positions. When there are large changes to a dealer's net position, however, this rolling process does not significantly improve upon a complete unwind and rewind of a dealer's position, from the perspective of the use of intraday credit. In particular, if a dealer switches from being a net lender to a net borrower (or vice versa) for a particular collateral class in GCF Repo, then rolling positions uses the same amount of intraday credit as a full unwind and rewind approach.

References

Begalle, B., A. Martin, J. McAndrews, and S. McLaughlin. 2013. "The Risk of Fire Sales in the Tri-Party Repo Market." Federal Reserve Bank of New York Staff Reports, no. 616, May.

Copeland, A., D. Duffie, A. Martin, and S. McLaughlin. 2012. "Key Mechanics of the U.S. Tri-Party Repo Market." Federal Reserve Bank of New York Economic Policy Review 18, no. 3 (November): 17-28.

Fleming, M., and K. Garbade. 2003. "The Repurchase Agreement Refined: GCF Repo®." Federal Reserve Bank of New York Current Issues in Economics and Finance 9, no. 6 (June).

Mills, D. C., and T. Nesmith. 2008. "Risk and Concentration in Payment and Securities Settlement Systems." Journal of Monetary Economics 55, no. 3 (April): 542-53.

Tri-Party Repo Infrastructure Reform Task Force. 2010. "Task Force on Tri-Party Repo Infrastructure: Report." Payments Risk Committee, sponsored by the Federal Reserve Bank of New York. Available at http://www.newyorkfed.org/banking/tpr_infr_reform.html.

Adam Copeland, Isaac Davis, and Antoine Martin

An Empirical Analysis of the GCF Repo® Service

1. Introduction

General Collateral Finance Repo (GCF Repo®) is a financial service that allows securities dealers to exchange government securities for cash among themselves on an anonymous basis.[1] GCF Repo plays an important role in the tri-party repo market, a market that is essential to the funding of large broker-dealers in the United States. But because of a paucity of available data, knowledge about participants' GCF Repo trading strategies is mostly anecdotal. Market participants report that GCF Repo can play several roles: For some dealers, GCF Repo is a main source of their repo funding. For other dealers, GCF Repo can be used to perform collateral swap trades, allowing them to acquire Treasury securities, the highest quality collateral, in exchange for agency mortgage-backed securities (MBS), which are of lesser quality. GCP Repo is also said to serve as a "buffer" for some dealers, making it possible for them to obtain more funding or more collateral, if they are affected by an unexpected shock. Using newly available data on the universe of GCF Repo activity, we aim to quantify the behavior of dealers that enter into GCF Repo contracts to see if that behavior is consistent with the anecdotal evidence.

Understanding the role of GCF Repo, and its interactions with the tri-party repo market, is important as the repo market evolves. The tri-party repo market has been affected by the reforms of its settlement process (see the first article in this volume, "The Financial Plumbing of the GCF Repo Service"), which are likely to shape the costs of settling GCF Repo transactions. In addition, Basel III reforms, and in particular the supplementary leverage ratio, are having an effect on the costs of repo activity for broker-dealers. By examining dealers' behavior before the reforms were implemented, we provide a benchmark that can used to understand how reforms might influence GCF Repo activity over time.

In this article, we provide three sets of results on the strategies dealers pursue when entering into GCF Repo contracts. First, we describe daily activity by looking at end-of-day settlement and documenting which groups of dealers use GCF Repo for funding. We find considerable variety among dealers, but, on average, those dealers that are not part of a bank holding company (BHC) consistently borrow cash (against securities) in this market. For some of these dealers, GCF Repo appears to be a main source of repo financing.

Second, we examine activity in the GCF Repo market using two different measures of dealers' net and gross activities. We infer that, on average, 1) 23 percent of dealers use GCF Repos to raise funds, 2) 20 percent of dealers use GCF Repo to source collateral or conduct collateral swaps, and 3) the remaining 57 percent of dealers follow a variety of strategies

[1] GCF Repo® is a registered service mark of the Fixed Income Clearing Corporation.

Adam Copeland is a research officer, Isaac Davis a former senior research analyst, and Antoine Martin a vice president at the Federal Reserve Bank of New York.

Correspondence: adam.copeland@ny.frb.org

To view the authors' disclosure statements, visit https://www.newyorkfed.org/research/author_disclosure/ad_epr_primer-on-the-gcf-repo.html.

The authors thank Paul Agueci, Leyla Alkan, Michele Braun, Vic Chakrian, and Kate Pingitore for their comments and, more generally, for sharing their knowledge of the GCF Repo service and the data used in this analysis. The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

when entering into GCF Repos, including acting as liquidity providers to other participants. We also look at whether dealers engage in collateral transformation—for example, swapping agency MBS for U.S. Treasury securities. Although this type of activity has been increasingly talked about in recent years, we find only modest evidence of collateral swaps in the GCF Repo market.

Third, we examine whether dealers' strategies for entering into GCF Repo and tri-party repo contracts are related. We find a negative correlation between daily changes in the amount of cash a dealer raises using tri-party repo and daily changes in the amount raised using GCF Repo. This correlation suggests that dealers view these two financial services as substitutes at the margin. In other words, GCF Repo appears to play the role of a "buffer" when a dealer experiences an unexpected shock.

2. General Description

We begin by providing a general overview of the GCF Repo service, followed by a description of our data.

2.1 Institutional Details

The Fixed Income Clearing Corporation (FICC) introduced the GCF Repo service in 1998. GCF Repo trades are general collateral repurchase agreements between eligible dealers that are executed through interdealer brokers (IDB), where one dealer is putting cash into the deal and the other dealer is providing securities. (These agreements closely resemble collateralized loans.) These trades are called *general collateral* because the institution providing securities is not required to provide a specific security, but rather any security within a fairly large asset class. FICC defines ten collateral classes that can be used by dealers.[2] Only institutions deemed eligible by FICC are able to trade GCF Repo. In December 2012, 120 entities were eligible to trade GCF Repos.[3]

[2] A list of the collateral classes is provided in Table 1.

[3] Almost all of the financial entities in this market can be considered securities dealers, so for expositional clarity we refer to all GCF Repo participants as securities dealers or simply dealers. A list of eligible financial entities can be found at http://www.dtcc.com/client-center/ficc-gov-directories. Look in the "FICC GSD Member Directory" for those members with the "Repo Netting" and "GCF" service designations.

GCF Repos are negotiated on a blind-brokered basis, where a dealer relies on an IDB to match it with an anonymous counterparty. FICC guarantees settlement once the trade has been compared, which occurs upon FICC's receipt of the trade data from the IDB, in the case of GCF Repo trades.

FICC provides netting services for GCF Repo. At the end of each trading day, FICC computes, for each dealer and each collateral category, the amount of securities the dealer has promised to deliver and the amount that has been promised to the dealer. The difference between these two amounts, the net position of a dealer in a collateral category, is then settled.

FICC also novates the net settlement position and becomes the legal counterparty to both sides of a GCF Repo contract for settlement purposes. Thus, each dealer now has FICC as a counterparty, rather than another dealer. Because of the guarantee that FICC provides, GCF Repo trades are not overcollateralized (unlike most repos). Specifically, they do not include a "haircut." Rather, the market price of securities posted as collateral is equal to the amount of cash lent. Nevertheless, to protect itself against financial loss owing to a potential default, FICC, in addition to its eligibility requirements, requires dealers entering into GCF Repo contracts to post collateral and cash in FICC's government securities division clearing fund.

A variety of securities dealers enter into GCF Repos. In the empirical section that follows, we provide statistics describing these dealers. In general, the dealers are both domestic and foreign-based. A majority of dealers are part of bank holding companies, and there are a few instances where different legal entities of the same BHC trade in this market. For example, both the broker-dealer and commercial bank entities of the same BHC may actively enter into GCF Repos. In addition, there are a few legacy entities that have not been consolidated. We assume that there is an economically meaningful reason why a BHC would have more than one entity trading on GCF Repo, so we treat each entity separately. Finally, most, but not all, dealers also enter into tri-party repo contracts.

Dealers use the GCF Repo service to redistribute cash and eligible securities among themselves. In general, dealers negotiate GCF Repo trades for three purposes: raising funds, sourcing collateral, or, generally speaking, leveraging liquidity. We consider each in turn, but note here that the third purpose is a catch-all category, which incorporates a large variety of potentially different trading strategies.

The strategy of raising funds with GCF Repo reflects the reality that dealers, which present different risks as counterparties, face a range of interest rates when seeking to raise funds from the money markets. These differences in rates provide an opportunity for dealers to intermediate funds among themselves. For example, dealers that can borrow cheaply from tri-party repo investors could borrow more than they need and lend the extra cash through GCF Repo. GCF Repo is an effective tool for dealers to intermediate cash among themselves because FICC, acting as a central counterparty, absorbs counterparty risk.

Dealers also enter into GCF Repo contracts to source collateral. Dealers cannot, of course, source specific securities through GCF Repo because of its general collateral design. Rather, dealers can source types of securities. Such transactions can be useful when dealers are seeking securities to fill other general collateral repos, such as tri-party repos.[4] Market participants claim that GCF Repo plays a crucial role in allowing dealers to alter their stock of securities at the end of the day, balancing investor demands that a dealer borrow a consistent amount over time against the dealer's profit-making activity of purchasing and selling securities over the business day. For example, a dealer may need to post U.S. Treasuries as collateral to a tri-party repo investor, but not have enough Treasuries at the end of the day. Using GCF Repo, the dealer can simply obtain the requisite amount of Treasuries with a reverse repo.[5] Alternatively, the dealer could execute a collateral swap if it has other unencumbered securities, such as agency MBS. A collateral swap requires negotiating two GCF Repo trades. The dealer agrees to 1) deliver agency MBS (borrow cash) and 2) accept Treasuries (lend cash). By executing these two trades, the dealer can meet the demands of its investor rather inexpensively—the cost of these two GCF Repo trades is roughly the difference in rates across the two transactions.

Dealers also enter into GCF Repo contracts to pursue a variety of strategies that do not fit neatly into either the raising funds or sourcing collateral categories. We group these remaining strategies into a third category, called "leveraging liquidity." Many of these alternative strategies take advantage of the funding and transactional liquidity in this interdealer market. For example, dealers may experience fails in securities purchases and as a result seek to reverse in similar securities overnight. Or dealers may want to accommodate cash investors that are seeking to lend more than expected on a particular day; dealers may then place the extra amount of cash in GCF Repo. In both these examples, dealers are leveraging the funding liquidity of this interdealer market to accommodate unexpected inflows or outflows of securities or cash from other types of trades. Another strategy that falls into this category is the provision of transactional liquidity. Because a dealer's GCF Repo position for a given collateral class is netted at the end of the day, it is inexpensive for a dealer to buy and sell securities within one of the ten predefined collateral classes throughout the day. Some dealers take advantage of this netting service and act as liquidity providers to other dealers seeking to raise funds or source collateral. Finally, dealers use GCF Repo to implement arbitrage strategies, taking advantage of the transactional liquidity in this market.

[4] Indeed, the ten predefined general asset classes for GCF Repo match up closely with the asset classes generally accepted by cash investors in tri-party repo.

[5] Repos are trades in which the dealer has promised to deliver securities against cash, while reverse repos are trades in which the dealer has promised to deliver cash against securities.

2.2 Data Description

Our analysis relies on confidential data from the FICC collected by the Federal Reserve Bank of New York, which contains the universe of GCF Repo activity from March 1, 2011, to September 30, 2012. The data are collected daily and aggregated by dealer and collateral class. For each collateral class, we observe the gross value of securities the dealer has committed to deliver (the total repo amount) as well as the gross value of securities the dealer will receive (the total reverse repo amount).[6]

Over the sample period, the daily average total value of GCF Repo trades is $493 billion (see Table 1).[7] There are ten collateral classes traded in our sample, but two collateral classes dominate in terms of gross value traded: Fannie Mae and Freddie Mac fixed-rate MBS, and U.S. Treasuries with maturities of thirty years or less (see Table 1 for a list of the collateral classes and note that these classes are not mutually exclusive). In our sample, these two collateral classes account for 83 percent of all GCF Repos. Currently, there are only nine collateral classes traded, because there are no longer any securities that fall into the FDIC-guaranteed corporate bonds collateral class.[8]

In our sample, there are sixty-five securities dealers active in the market. Of those sixty-five dealers, thirty-three entered into GCF Repo contracts every day of our sample, and forty did so on at least 90 percent of the days. In general, the infrequent participants are much smaller in terms of

[6] Although our sample covers a year-and-a-half of activity, it may be the case that trends highlighted in this data are not representative of activity before 2011 or after 2012.

[7] Every trade creates a repo and a reverse repo transaction. When considering aggregate statistics, we add up only repo transactions to avoid double-counting.

[8] The Federal Deposit Insurance Corporation's Debt Guarantee Program, developed during the recent financial crisis, generated this special class of corporate bonds. This program is no longer active. For more information, see http://www.fdic.gov/regulations/resources/TLGP/.

Table 1
Collateral Classes in GCF Repo

Asset Type	Mean Daily Gross Collateral (Billions of U.S. dollars)	Percentage of Total
Fannie Mae and Freddie Mac fixed-rate MBS	209.72	42.55
U.S. Treasuries with maturities of thirty years or less	199.93	40.56
Non-MBS U.S. agency securities	33.66	6.83
Ginnie Mae fixed-rate MBS	27.74	5.63
Fannie Mae and Freddie Mac adjustable-rate MBS	13.97	2.83
U.S. Treasuries with maturities of ten years or less	2.65	0.54
U.S. Treasury inflation-protected securities	2.98	0.60
FDIC-guaranteed corporate bonds[a]	1.30	0.26
U.S. Treasury STRIPs	0.78	0.16
Ginnie Mae adjustable-rate MBS	0.16	0.03
Total	492.89	100.00

Sources: Confidential Fixed Income Clearing Corporation (FICC) data; authors' calculations.

Notes: The mean daily gross collateral is the average value of all repo trades conducted in each day of the sample. Asset types are ranked from largest to smallest as a percent of the total. MBS is mortage-backed securities and FDIC is Federal Deposit Insurance Corporation. STRIP is separate trading of registered interest and principal. An example of a non-MBS U.S. agency security is agency debentures.

[a]FDIC-guaranteed corporate bonds are no longer a collateral class in GCF Repo.

their mean daily repo activity (on days they are active) than the frequent participants. In our sample, frequent participants conduct an average of $24 billion in trades on every day they are active, compared with $2.4 billion for infrequent participants.

While dealers are fairly heterogeneous in their activity, we find it useful to classify them into two groups: those that are part of a bank holding company, and those that are not. This distinction is economically important because independent dealers (those not part of BHCs) rely solely on capital markets for funding. Dealers that are part of BHCs, in contrast, can also obtain funding from their parent company. Turning to the forty-four dealers that are part of a BHC, we find a wide variety in their size. Defining size as the value of U.S. dollar assets held by the associated BHC, the range of asset holdings is $13 billion to $2.3 trillion.[9] We pick a cutoff value of $500 billion to differentiate between dealers that are part of large and small BHCs. Overall then, dealers fall into three groups: those associated with large BHCs, those associated with small BHCs, and those that are not part of a BHC (which we label non-BHC dealers).

[9] Information on the value of U.S. dollar assets at the BHC level comes from the Federal Reserve Y-9C regulatory filings. For detailed information on these filings, see http://www.federalreserve.gov/apps/reportforms/default.aspx and select Form FR Y-9C.

The remainder of this article uses the above data to compute the degree to which dealers seek to raise funds, source collateral, or leverage liquidity with GCF Repo trades as well as to describe which types of dealers are more likely to use each strategy.

3. Daily Net Activity

We begin by looking at dealers' daily net activity across all collateral classes. This measure can give us a sense of whether a dealer uses the GCF Repo market mainly to borrow or to lend cash. For each day, we compute each dealer's net cash position. This position is equal to the sum of the difference between repos and reverse repo amounts across all collateral categories. Formally, dealer *j*'s net cash position at day *t* is given by:

1) $$netcash_{jt} = \sum_{i=1}^{10} \left(repo_{ijt} - reverse_{ijt} \right),$$

Table 2
Distribution of Net Cash Positions
Billions of U.S. Dollars, except Where Noted

	10th percentile	25th percentile	Median	75th percentile	90th percentile	Mean	Number of Observations
All	-13.0	-2.2	0.2	3.7	9.7	0	20,836
Large BHCs	-14.9	-8.2	-0.3	2.8	9.7	-1.8	3,943
Small BHCs	-15.5	-4.2	-0.1	1.6	8.7	-2.3	9,935
Non-BHCs	0	0.13	1.9	6.2	20.8	4.3	6,958

Sources: Confidential Fixed Income Clearing Corporation (FICC) data; authors' calculations.

Notes: Dealers are categorized as belonging to a large BHC, belonging to a small BHC, or not being part of a BHC (non-BHCs). Net cash position is the amount of dollars a dealer is delivering (if negative) or has been promised (if positive) at the end of the day, after accounting for trading activity across all collateral classes.

where $repo_{ijt}$ and $reverse_{ijt}$ are dealer *j*'s repo and reverse repo position, respectively, in collateral group *i* at date *t*. A positive net cash position means that the dealer is receiving cash at the end of the day, after accounting for the dealer's activity across all collateral classes. A negative number means that the dealer is lending cash. In our sample, dealers' net cash positions range from roughly -\$30 billion to \$40 billion. We also look at the net cash position conditional on dealers being associated with large or small bank holding companies (see the second and third rows of Table 2). Although there is still considerable heterogeneity among dealers in each category, we find that the mean net cash position of dealers that are part of BHCs is negative. As a group, then, dealers associated with bank holding companies typically lend cash using GCF Repo. The average dealer that is part of a large BHC typically lends \$1.8 billion every day. The average dealer that is part of a small BHC lends slightly more—\$2.3 billion each day. This flow of cash can also be seen by noting that the average dealer not associated with a bank holding company (labeled non-BHC) typically has a positive net cash position of \$4.3 billion.

Although dealers vary widely in their net cash positions, they are quite consistent in their strategies regarding borrowing or lending cash using GCF Repo. We find that a dealer that borrows cash today will continue to borrow cash tomorrow with 96.0 percent probability (see the first row of Table 3). Similarly, a dealer that lent cash today will continue to lend cash tomorrow with 95.1 percent probability (see the second row of Table 3).

To gain a better sense of the aggregate flow of cash among dealers, we compute the daily net cash position for each group of dealers. In Chart 1, we plot the monthly average net cash position for each of these three groups. As illustrated by the chart, dealers that are not part of BHCs consistently raise cash. The funding raised by non-BHC dealers has doubled over the sample period, from around \$40 billion to \$80 billion each day. As a group, dealers that are part of small BHCs are often the source of the majority of these funds.

With the flow of cash from dealers associated with BHCs to those that are not, there must be a flow of collateral going in the other direction. To understand the movement of collateral among the three groups of dealers, we focus on the two collateral classes that account for the vast majority of GCF Repo contracts: Fannie Mae and Freddie Mac fixed-rate MBS (henceforth, FFFR MBS) and U.S. Treasuries with maturities of thirty years or less. For each of these two collateral types, we compute each dealer group's net position (that is, total repos minus total reverse repos) for each day. We then compute the average position for the month and plot the results in Chart 2, panels A and B.

From this aggregate viewpoint, we see that non-BHC dealers provide both U.S. Treasuries and FFFR MBS as collateral for their repo trades. Toward the end of the sample, however, non-BHC dealers increasingly post FFFR MBS securities as collateral. Strikingly, dealers associated with small BHCs differ markedly from those that are part of large BHCs. Small BHC-affiliated dealers reverse in U.S. Treasuries each day (Chart 2, panel A). These U.S. Treasuries are delivered to them by both non-BHC-affiliated and large BHC-affiliated dealers. For FFFR MBS securities, we observe that small BHC-affiliated dealers switched from delivering these securities in the beginning of the sample to reversing in these securities at the end of the sample (Chart 2, panel B). In contrast, large BHC-affiliated dealers consistently reverse in these securities. Looking at the behavior of large

Table 3
Transition Probabilities
Percent

	t	
t-1	Net Borrower	Net Lender
Net borrower	96.0	4.0
Net lender	4.9	95.1

Sources: Confidential Fixed Income Clearing Corporation (FICC) data; authors' calculations.

Notes: A net borrower is a dealer whose net cash position is positive. A net lender is a dealer whose net cash position is less than or equal to 0. The cell entries show the probability of transitioning from a net borrower or lender in GCF Repo at time *t*-1 to being a net borrower or lender at time *t*. Each row sums to 100 percent.

BHC-affiliated dealers in both panels of Chart 2, we infer that this group of dealers is pursuing, on the whole, a strategy of delivering Treasuries and receiving FFFR MBS securities, a collateral downgrade strategy, throughout the sample.

4. Comparing Gross and Net Activity

We now analyze gross and net activity to further distinguish the degree to which dealers pursue various strategies. We start by considering dealers' net-to-gross ratios. This ratio allows us to differentiate dealers that are mainly employing liquidity-leveraging or collateral-swapping strategies from those that are mainly pursuing funding or securities-acquiring strategies. We then consider another statistic, a swap ratio, which measures how much dealers swap collateral each day.

We begin by constructing a dealer's net-to-gross ratio for each day in the sample. Because we want to account for activity across collateral classes, we construct this measure based on cash activity. The "gross" part of this ratio is the sum of the value of all repos and reverse repos a dealer trades in a day, and is thus a measure of the totality of a dealer's activity. The "net" part is the sum across all collateral classes of the difference between a dealer's total repo and total reverse repo position (see $netcash_{jt}$, defined earlier). Formally, for dealer *j* at date *t*, the net-to-gross ratio is equal to:

$$NtG_{jt} = \frac{\sum_{i=1}^{10}\left(repo_{ijt} - reverse_{ijt}\right)}{\sum_{j=1}^{10}\left(repo_{ijt} + reverse_{ijt}\right)} = \frac{netcash_{ijt}}{\sum_{t=1}^{10}\left(repo_{ijt} + reverse_{ijt}\right)},$$

where *i* indexes the collateral groups traded in GCF Repo.

The net-to-gross ratio is positive if the dealer is receiving cash at the end of the day, and negative if the dealer is delivering cash. By construction, this ratio is always between -1 and 1, and is equal to 0 when the dealer's net position is exactly offsetting (that is, the dealer is receiving equal amounts of cash and collateral). When the ratio is close to 1, the dealer's predominant strategy is to obtain cash in this market. For the ratio to be equal to 1, the dealer must conduct only repo transactions. Similarly, if the ratio is close to -1, the dealer's predominant strategy is to obtain securities for cash. For the ratio to be equal to -1, the dealer must conduct only reverse repo transactions. Finally, when the ratio is closer to zero, substantial netting is occurring. This could mean that dealers are mostly providing transactional liquidity, conducting repo and reverse repo transactions within the same collateral class. Alternatively, dealers could be seeking to manage their inventories by exchanging collateral (for example, collateral upgrade or downgrade). In these cases, dealers would be conducting repo transactions in one collateral class and conducting reverse repo transactions in another collateral class.

We begin by analyzing the net-to-gross distribution for all dealers and then examine each dealer group separately. The histogram of net-to-gross ratios for all dealers over the whole sample (Chart 3, panel A) highlights the diverse set of strategies followed by dealers. The histogram illustrates that on a typical day in the sample, about 23 percent of dealers conduct only repos (the net-to-gross ratio is equal to 1) and thus use GCF Repo to effectively raise funds. Further, almost 10 percent of dealers conduct only reverse repos (the net-to-gross ratio is equal to -1), using this market to acquire securities.[10] Finally, on a typical day, the remaining 57 percent of dealers are executing both repo and reverse repos, with a substantial number of dealers offsetting their repo and reverse repo trades so as to have net positions close to zero (about 8.5 percent of dealers). Dealers with net-to-gross ratios between -1 and 1 most likely pursue a mixed set of strategies, and it is difficult to disentangle dealers' propensity to rely on one strategy more than another without a more formal analysis.

Analyzing the distribution of net-to-gross distributions by dealer group, however, reveals stark differences in strategies pursued by each group. Confirming the results from the previous section, we find that the vast majority of non-BHC dealers have positive net-to-gross ratios and thus receive cash at the end of the day (Chart 3, panel B). Indeed, on a typical day, non-BHC dealers conducted only repo trades almost 35 percent of the time. But not all non-BHC dealers

[10] We also examined histograms for subsets of the sample, and did not find any interesting variation in the distribution of net-to-gross ratios over time.

Chart 1

Daily Net Cash Position by Dealer Group, Monthly Average

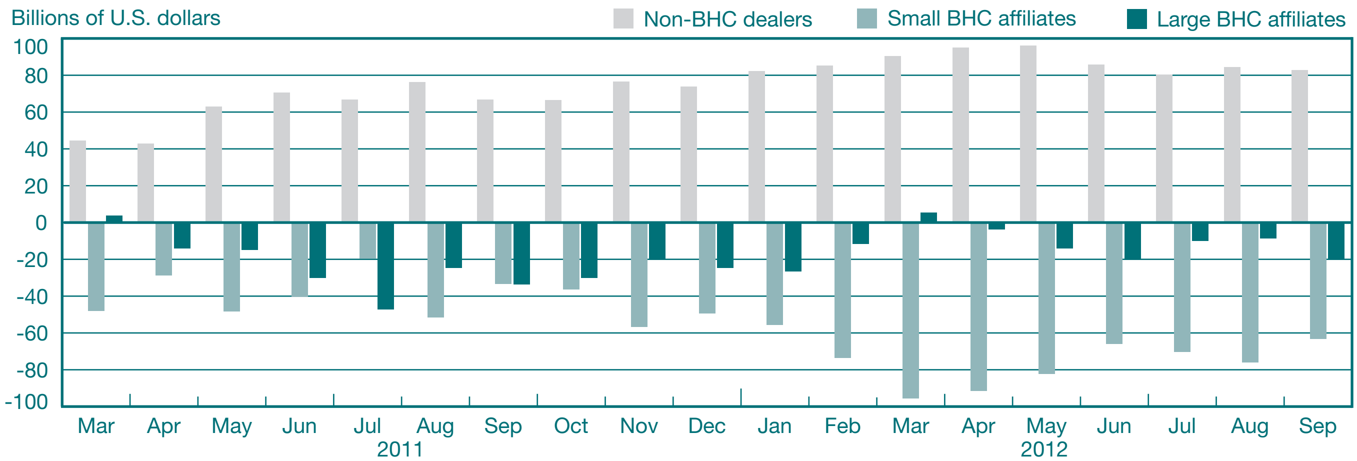

Sources: Confidential Fixed Income Clearing Corporation (FICC) data; authors' calculations.

Notes: Dealers are categorized as belonging to a large BHC, belonging to a small BHC, or not being part of a BHC (non-BHCs). A positive position means that the group of dealers receives cash, on net, each day. A bar represents the daily net cash position for a dealer group averaged over a month.

Chart 2

Daily Net Position of Asset Class by Dealer Group, Monthly Average

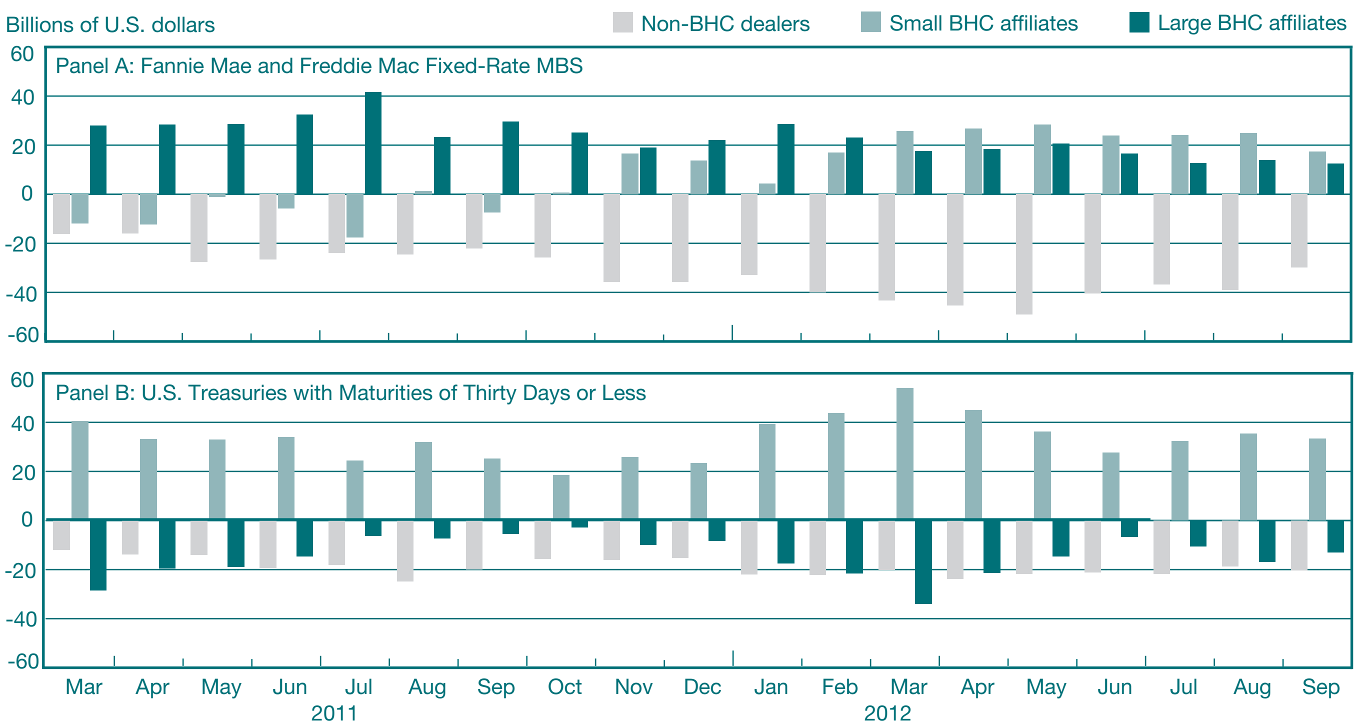

Sources: Confidential Fixed Income Clearing Corporation (FICC) data; authors' calculations.

Notes: Dealers are categorized as belonging to a large BHC, belonging to a small BHC, or not being part of a BHC (non-BHCs). A positive position means that the group of dealers receives cash, on net, each day. A bar represents the daily net cash position for a dealer group averaged over a month.

Chart 3
Distributions of Net-to-Gross Ratios by Dealer and Day

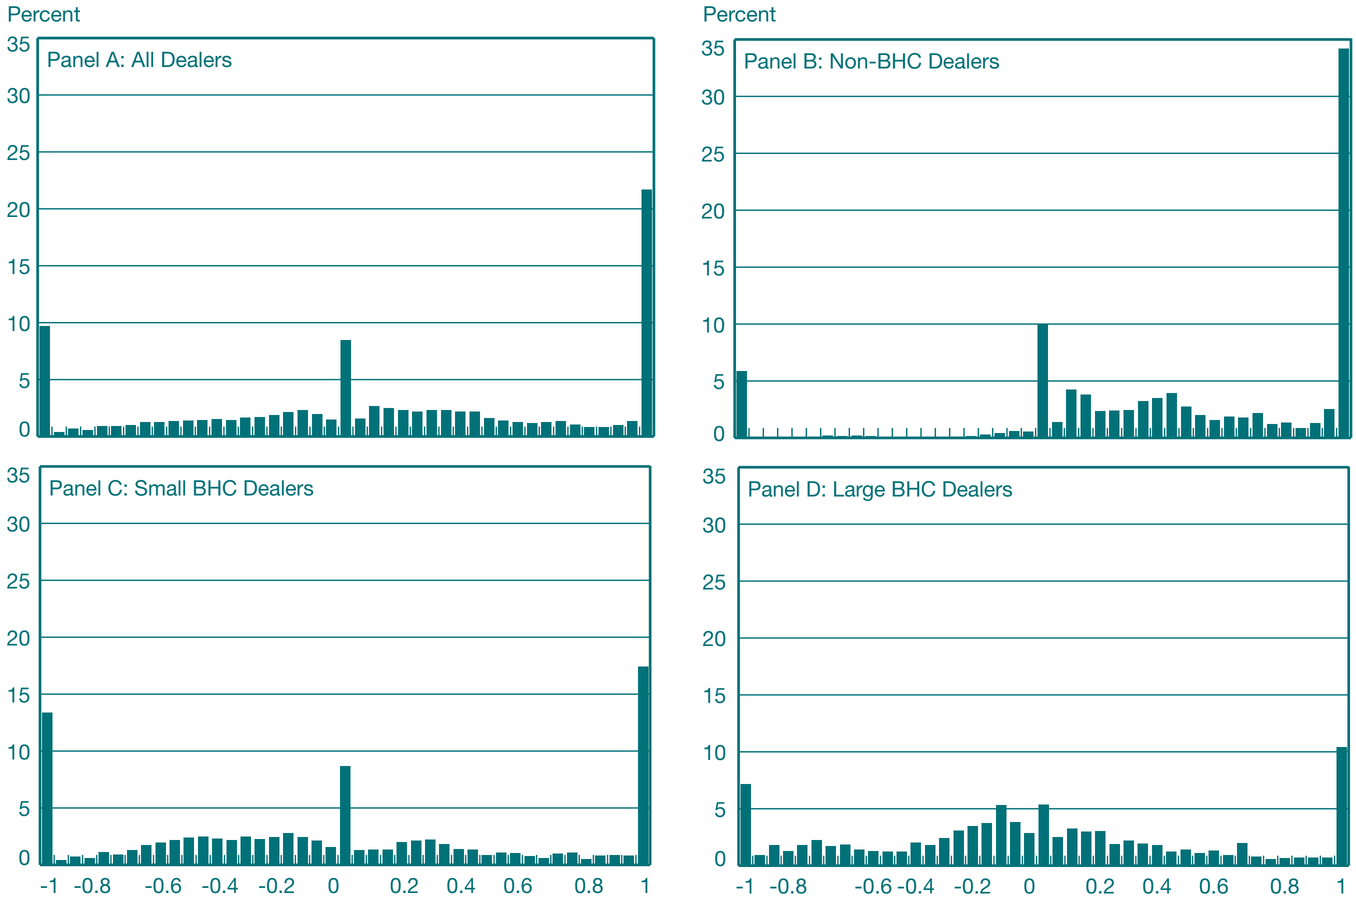

Sources: Confidential Fixed Income Clearing Corporation (FICC) data; authors' calculations.

Notes: An observation is a dealer's net-to-gross ratio for a particular day. The net-to-gross ratio is equal to a dealer's net settlement activity across all collateral groups over a dealer's total trading activity across all collateral groups in a day. A ratio equal to 1 means that the dealer conducts only repos, whereas a ratio equal to -1 means the dealer conducts only reverse repos. There are 20,836 observations for all dealers, and 6,958 for non-BHC dealers, 9,935 for small BHC dealers, and 3,943 for large BHC dealers.

try to raise funds only. On a typical day, a little more than 5 percent of dealers perform only reverse repos, and about 10 percent have net-to-gross ratios near zero, and thus buy and sell roughly equal amounts of repo and reverse repos.

Compared with non-BHC dealers, small BHC dealers have net-to-gross ratios that are more evenly distributed between -1 and 1 (Chart 3, panel C). Small BHC dealers are roughly equally split between borrowing and lending cash at the end of the day (the median value of net-to-gross is -0.01 for small BHC dealers). As with non-BHC dealers, however, there are significant numbers of small BHC dealers that have net-to-gross ratios roughly equal to 1, -1, and 0. On a typical day, small BHC dealers conduct only repos 17 percent of the time, only reverse repos 13 percent of the time, and have net positions close to zero about 9 percent of the time.

Finally, we find that large BHC dealers, relative to all other dealers, are much less likely to have net-to-gross ratios close to 1 or -1 (Chart 3, panel D). Rather, these dealers are much more likely to conduct both repo and reverse repos trades in the same day.

While the net-to-gross ratio reveals whether and to what degree dealers are conducting both repos and reverse

Table 4
Distribution of Swap Ratios

	5th percentile	10th percentile	25th percentile	50th percentile	75th percentile	90th percentile	95th percentile
Swap ratio	0.24	0.42	0.85	1.00	1.00	1.00	1.00

Sources: Confidential Fixed Income Clearing Corporation (FICC) data; authors' calculations.

Notes: A swap ratio is a measure of the amount of collateral swapping that occurs for a dealer in a day. A value of 1 means that there are no collateral swaps, while a value of $0< x<1$ implies that the total value of collateral swapped is equal to $(1-x)$ percent of a dealer's total net position.

repos, this statistic does not allow us to distinguish collateral swapping strategies from other strategies. For example, a net-to-gross ratio near zero can be the result of a dealer executing a repo collateralized by agency MBS alongside a reverse repo collateralized by U.S. Treasuries. These two trades effectively constitute a collateral swap (agency MBS for U.S. Treasuries). To measure what fraction of dealers conduct collateral swaps, we compute a swap ratio for each dealer on each day. This ratio is equal to

$$swap\ ratio_{jt} = \frac{\left|\sum_{i=1}^{10} repo_{ijt} - reverse_{ijt}\right|}{\sum_{i=1}^{10}\left|repo_{ijt} - reverse_{ijt}\right|}$$

This ratio looks at a dealer's net positions across collateral types. When a dealer does not have any offsetting net positions across collateral types, the swap ratio is equal to 1. For example, if a dealer's net position in every collateral position is weakly positive, then the numerator and denominator of the swap ratio are equal. But if a dealer has a positive net position in one collateral class and a negative net position in another collateral class, the denominator will be greater than the numerator. This is because in the numerator the positive net position is summed with the negative net position, while in the denominator the absolute values of both net positions are summed. The closer the swap ratio is to 0, then, the more a dealer is involved in collateral swaps. Because we do not know the true intention of the dealer, we say that the dealer is effectively involved in collateral swapping.

Using our data, we compute the swap ratio for each dealer and for each day and then calculate the distribution of this statistic. We find that the median value of this ratio is equal to 1—in other words, half of the time, dealers are not conducting any collateral swaps (Table 4). This is consistent with the results presented in Chart 3, panel A, where at least 33 percent of dealers conduct only repo or only reverse repo transactions. At the 25th percentile, the swap ratio is equal to 0.85. This value implies that a dealer's effective collateral swaps are equal to 15 percent of the value of a dealer's total net position. It is only at the 10th percentile where collateral swapping becomes a dealer's predominant strategy (a swap ratio of 0.42 implies collateral swaps are equal to 58 percent of a dealer's total net position).

Overall, then, this result indicates that collateral swaps do not occur frequently: a little more than 10 percent of the time, collateral swapping can be said to be the dealer's predominant strategy. The large number of instances where a dealer's net-to-gross ratios are between 1 and -1, then, seem to be primarily driven by netting *within* a collateral group. Such trading behavior could result from dealers providing transactional liquidity to the market by executing repos and reverse repos throughout the day. Dealers might also execute both repos and reverse repos at different times in the day while pursuing different strategies. For example, they begin the day seeking to raise funds. However, settlement fails of other transactions may lead the same dealer to lend cash later the same day.

To further examine the amount of netting that occurs within each asset class, we construct net-to-gross ratios for each dealer, date, and asset class. Hence, the net-to-gross ratio for dealer j, on date t, for asset type i, is equal to

$$NtG_{ijt} = \frac{repo_{ijt} - reverse_{ijt}}{repo_{ijt} + reverse_{ijt}}.$$

Note that this net-to-gross ratio is computed at a lower level of aggregation relative to those displayed in Chart 3.[11] We look at the distribution of this ratio for each asset type. In Table 5, we list the 25th, 50th, and 75th percentiles of these distributions, as well as the percentage of observations equal to -1 and 1. As in Table 1, we list the asset types from largest to smallest, in terms of the dollar value of repos conducted. Strikingly, the vast majority of net-to-gross ratios for the

[11] By construction, there will be a weakly greater share of net-to-gross ratios by asset type equal to -1 or 1 compared with net-to-gross ratios computed across asset types (and so at a higher level of aggregation).

Table 5

Net-to-Gross Ratios by Dealer, Date, and Asset Type

Asset Type	25th percentile	50th percentile	75th percentile	Percentage at -1	Percentage at 1	Number of Observations
Fannie Mae and Freddie Mac fixed-rate MBS	-0.65	0.15	0.75	18	19	15,786
U.S. Treasuries with maturities of thirty years or less	-0.26	0.05	0.64	7	21	17,057
Non-MBS U.S. agency securities	-0.67	0.31	1	20	32	10,965
Ginnie Mae fixed-rate MBS	-1	0.22	1	29	43	7,664
Fannie Mae and Freddie Mac adjustable-rate MBS	-1	0.67	1	40	46	6,798
U.S. Treasuries with maturities of ten years or less	-1	0.33	1	36	48	1,079
U.S. Treasury inflation-protected securities	-1	1	1	43	54	3,980
FDIC-guaranteed corporate bonds[a]	-1	1	1	41	52	2,047
U.S. Treasury STRIPs	-1	1	1	54	45	3,980
Ginnie Mae adjustable-rate MBS	-1	1	1	49	51	491
Total	-0.8	0.17	1	23	31	67,575

Sources: Confidential Fixed Income Clearing Corporation (FICC) data; authors' calculations.

Notes: Percentage at -1 (1) is the percentage of observations equal to -1 (1).The net-to-gross ratio is equal to a dealer's net settlement activity over a dealer's total trading activity for a given collateral group in a day. A ratio equal to 1 means that the dealer conducted only repos, while a ratio equal to -1 means that the dealer conducted only reverse repos. MBS is mortgage-backed securities and STRIPS is separate trading of registered interest and principal.

[a]FDIC guaranteed corporate bonds are no longer a collateral class in GCF Repo.

four smallest asset types are equal to -1 or 1. Dealers, then, are not conducting both repos and reverse repos with these assets types. To a lesser extent, the same result also holds for the asset types of Fannie Mae and Freddie Mac adjustable-rate MBS, and U.S. Treasuries with maturities of ten years or less. Only for the largest three asset types do we see a substantial number of dealers conducting both repos and reverse repos, and enjoying the netting benefits provided by the FICC for GCF Repo.

In summary, in this section we construct two measures describing each dealer's daily activity. From these measures we find that:

- On average, at least 23 percent of dealers use GCF Repo to raise funds (that is, they conduct only repo transactions).
- On average, at least 20 percent of dealers use GCF Repo to manage their inventory of securities. They manage inventories by following two types of strategies:
  - On average, at least 10 percent of dealers focus on purchasing securities (that is, they conduct only reverse repo transactions).
  - On average, at least 10 percent of dealers are predominantly conducting collateral swaps.
- The remaining 57 percent of dealers conduct both repo and reverse repo GCF Repo trades for a variety of reasons, including providing liquidity to other participants.

Our estimate of the fraction of times dealers pursue liquidity-leveraging strategies is, by construction, an upper bound. Dealers who conduct both repos and reverse repos are likely pursuing multiple strategies at the same time. For example, a dealer could seek to raise funds using GCF Repo, while also providing liquidity to the market. More inclusive definitions of funding or inventory management strategies would necessarily lower the 57 percent estimate. Obtaining a more precise estimate of the percentage of time dealers are mainly providing liquidity to the market is something we will explore in future work.

5. Connection between Tri-Party Repo and GCF Repo

Tri-party repo and GCF Repo are intimately connected because they both qualify for same-day settlement on the books of the two clearing banks JPMorgan Chase and Bank of New York Mellon. GCF Repo is settled before tri-party repo, allowing dealers to easily deliver securities to tri-party repo that have been acquired from GCF Repo. Thus, there is ample opportunity for dealers to be strategic when trading both repo products. Of the dealers that we observe actively trading GCF Repo, 85 percent are also actively trading tri-party repo.[12]

Given this close connection, an open question is how dealers choose their trading strategies across tri-party repo and GCF Repo. We start our analysis of dealers' strategic behavior by first looking at the correlation in the change of a dealer's position with each financial product. This correlation informs us whether, at the margin, a dealer views these two types of repos as substitutes or complements.

Drawing on confidential data collected by the Federal Reserve Bank of New York from the two clearing banks that settle tri-party repo contracts, we compute a dealer's change in funding using tri-party repo on consecutive business days. As described in Copeland, Martin, and Walker (2014), these data allow us to measure the value of collateral posted by dealers in this market on a daily basis. Letting tpr_{jt} denote the amount of funding a dealer receives from tri-party repo, the change in a dealer's tri-party repo funding is given by $\Delta tpr_{jt} = tpr_{jt} - tpr_{jt-1}$. Note that tpr_{jt} is always a positive number, because dealers use this product strictly for funding purposes.[13] We measure the change in a dealer's GCF Repo position as the change in the net cash position (see $netcash_{jt}$ defined in equation 1), or $\Delta GCF_{jt} = netcash_{jt} - netcash_{jt-1}$. We then regress the change in GCF Repo position on the change in tri-party repo to measure how dealers jointly alter their positions. Formally, the regression is

2) $$GCF_{jt} = \alpha + \beta \Delta tpr_{jt} + \varepsilon_{jt},$$

where ε_{jt} is an error term. We estimate that β is negative, implying a negative correlation between a dealer's overall position in tri-party repo and GCF Repo (Table 6). This negative (and statistically significant) relationship also holds when looking at all dealers or when focusing on any of the three groups of dealers.

This statistical relationship demonstrates that dealers effectively view these products as substitutes at the margin.[14] The -0.56 coefficient implies that for the average dealer, a decrease of \$100 in tri-party repo is associated with a \$56 increase in that dealer's net cash position in GCF Repo. (Note that increases in the net cash position can mean that a dealer is raising more funds or lending less cash.)

To gain a better sense of the relative magnitude of the changes in dealers' positions across the two products, we look at the absolute value of the change in net cash position in GCF Repo over the sum of the absolute value of the change in net cash in GCF Repo plus the absolute value of the change in tri-party repo funding. Formally, for each dealer and day we compute

$$\frac{|\Delta GCF_{jt}|}{|\Delta GCF_{jt}| + |\Delta tpr_{jt}|}.$$

If this ratio is equal to one-half, then the change in a dealer's net cash position in GCF Repo is equal to the change in the dealer's tri-party repo funding. We compute this ratio for all days and across all dealers and find that the median value of this ratio is 0.69. The average dealer, then, has larger changes in its overall GCF Repo position than in its tri-party repo position. We then look at the 25th and 75th percentiles

Table 6
Relationship across Tri-Party Repo and GCF Repo Trades

	Coefficient	Standard error	Number of Observations
All dealers	-0.56	0.01	15,497
Non-BHC dealers	-0.65	0.05	5,309
Small BHC dealers	-0.55	0.01	7,274
Large BHC dealers	-0.59	0.02	2,912

Sources: Confidential Fixed Income Clearing Corporation (FICC) data; authors' calculations.

Notes: Each row reports the result of a separate regression. The regression estimated the correlation between changes in a dealer's tri-party repo position and changes in the same dealer's GCF Repo position (see equation 2 in the text).

[12] For this analysis, dealers have been aggregated up to the bank holding company level.

[13] For more information on the tri-party repo market, see Copeland, Martin, and Walker (2014) and Krishnamurthy, Nagel, and Orlov (2014).

[14] Mancini, Ranaldo, and Wrampelmeyer (2014) study the European repo market and find some substitutability between the repo market and the unsecured interbank market.

of the distribution of this ratio and find that they are equal to 0.35 and 0.96, respectively. For dealers in the upper quartile, then, the change in the GCF Repo position completely dominates, in terms of size, the change in the tri-party repo position. An interpretation of this high ratio value is that dealers make almost all of their cash and securities adjustments using GCF Repo, as opposed to tri-party repo.

6. Conclusion

This article aims to quantify to what extent dealers pursue various strategies when entering into GCF Repo contracts. We are able to provide some stylized facts and quantify the extent to which different types of behaviors are observed in this market. For the most part, our evidence is consistent with anecdotal evidence provided previously by market participants. That said, we also find that, despite the growing attention collateral transformation has received in recent years, there is only modest evidence of such activity in GCF Repo during our sample. This article also provides a reference point for understanding how both reforms to settlement of GCF Repo contracts and regulatory reforms may influence dealer activity in the future.

We find that, in general, the set of dealers not associated with BHCs raise between $40 billion and $80 billion a day, and so use GCF Repo as a source of funding. Dealers associated with BHCs provide this funding. Further, large BHC-affiliated dealers, as a group, tend to deliver Treasuries to this market and receive Fannie Mae and Freddie Mac fixed-rate MBS securities.

Looking across all dealers and all days, we find that on an average day, at least 23 percent of dealers focus on strategies to raise cash. At least another 20 percent of dealers focus on managing their inventory of securities. This activity involves using GCF Repo to both exclusively source collateral (for example, for re-use in tri-party repo contracts) and perform collateral swaps. Finally, the remaining 57 percent of dealers appear mainly to use other strategies that take advantage of the liquidity in this interdealer market. Our estimates of the percentage of time focused on raising cash and managing inventories are conservative, and so should be viewed as a lower-bound estimate.

We also study dealer behavior across the tri-party repo and GCF Repo services. We find evidence that dealers view these repo services as substitutes at the margin. In particular, changes in a dealer's tri-party repo position are negatively correlated with changes in that dealer's GCF Repo position.

References

Copeland, A., A. Martin, and M. Walker. 2014. "Repo Runs: Evidence from the Tri-Party Repo Market." Journal of Finance 69, no. 6 (December): 2343-80.

Federal Reserve Bank of New York. 2014. "Domestic Open Market Operations during 2013." Available at https://www.newyorkfed.org/medialibrary/media/markets/omo/omo2013-pdf.pdf.

Krishnamurthy, A., S. Nagel, and D. Orlov. 2014. "Sizing Up Repo." Journal of Finance 69, no. 6 (December): 2381-417.

Mancini, L., A. Ranaldo, and J. Wrampelmeyer. 2014. "The Euro Interbank Repo Market." Swiss Finance Institute Research Paper no. 13-71.

Fernando Duarte and Carlo Rosa

The Equity Risk Premium: A Review of Models

- The equity risk premium (ERP) is a key quantity in asset pricing that serves as an indicator of economic activity and a determinant of firms' cost of capital, individuals' savings decisions, and government budgeting plans.

- This study estimates the ERP by combining data from twenty models. It finds that the ERP in 2012 and 2013 reached heightened levels—of about 12 percent—not seen since the 1970s.

- The authors attribute the high ERP to unusually low Treasury yields rather than to expectations that stocks would have high returns.

- One implication of the ERP being driven by bond yields rather than expected stock returns is that traditional indicators of the ERP, such as simple valuation ratios, may not be as good a guide to future excess returns as they have been in the past.

1. Introduction

The equity risk premium—the expected return on stocks in excess of the risk-free rate—is a fundamental quantity in all of asset pricing, both for theoretical and practical reasons. It is a key measure of aggregate risk-aversion and an important determinant of the cost of capital for corporations, the savings decisions of individuals, and budgeting plans for governments. Recently, the equity risk premium (ERP) has also moved to the forefront as a leading indicator of the evolution of the economy, a potential explanation for jobless recoveries, and a gauge of financial stability.[1]

In this article, we estimate the ERP by combining information from twenty prominent models used by practitioners and featured in the academic literature. Our main finding is that the ERP has reached heightened levels. The first principal component of all models—a linear combination that

[1] As an indicator of future activity, a high ERP at short horizons tends to be followed by higher GDP growth, higher inflation, and lower unemployment. See, for example, Piazzesi and Schneider (2007), Stock and Watson (2003), and Damodaran (2012). Bloom (2009) and Duarte, Kogan, and Livdan (2013) study connections between the ERP and real aggregate investment. Offering a potential explanation of the jobless recovery, Hall (2014) and Kuehn, Petrosky-Nadeau, and Zhang (2012) propose that increased risk-aversion has prevented firms from hiring as readily as would be expected in the post-crisis macroeconomic environment. Among many others, Adrian, Covitz, and Liang (2013) analyze the role of equity and other asset prices in monitoring financial stability.

Fernando Duarte is an economist and Carlo Rosa a senior economist at the Federal Reserve Bank of New York.

fernando.duarte@ny.frb.org; carlo.rosa@ny.frb.org

To view the authors' disclosure statements, visit https://www.newyorkfed.org/research/author_disclosure/ad_epr_2015_equity-risk-premium.html.

The authors thank Tobias Adrian and James Egelhof for helpful comments on earlier drafts. All remaining errors are the authors'. This article updates the authors' May 2013 blog post in *Liberty Street Economics,* available at http://libertystreeteconomics.newyorkfed .org/2013/05/are-stocks-cheap -a-review-of-the-evidence.html, and offers a more rigorous and comprehensive treatment than the post. The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

explains as much of the variance of the underlying data as possible—places the one-year-ahead ERP in June 2012 at 12.2 percent, above the 10.5 percent reached during the financial crisis in 2009 and at levels similar to those in the mid- and late 1970s. From June 2012 to the end of our sample in June 2013, the ERP has changed little, despite substantial positive realized returns. It is worth keeping in mind, however, that there is considerable uncertainty around these estimates. In fact, the issue of whether stock returns are predictable is still an active area of research.[2] Nevertheless, we find that the dispersion in estimates across models, while quite large, has been shrinking, potentially signaling increased agreement even when the models differ substantially from one another and use more than one hundred different economic variables.

In addition to estimating the level of the ERP, we investigate the reasons behind its recent behavior. Because the ERP is the difference between expected stock returns and the risk-free rate, a high estimate can be the result of expected stock returns being high or risk-free rates being low. We conclude that the ERP is high because Treasury yields are unusually low. Current and expected future dividend and earnings growth play a smaller role. In fact, expected stock returns are close to their long-run mean. One implication of a bond-yield-driven ERP is that traditional indicators of the ERP like the price-dividend or price-earnings ratios, which do not use data from the term structure of risk-free rates, may not be as good a guide to future excess returns as they have been in the past.

As a second contribution, we present a concise and coherent taxonomy of ERP models. We categorize the twenty models into five groups: predictors that use historical mean returns only, dividend discount models, cross-sectional regressions, time-series regressions, and surveys. We explain the methodological and practical differences among these classes of models, including the diverse assumptions and data sources that they require.

2. The Equity Risk Premium: Definition

Conceptually, the ERP is the compensation that investors require to make them indifferent at the margin between holding the risky market portfolio and a risk-free bond. Because this compensation depends on the future performance of stocks, the ERP incorporates expectations of future stock market returns, which are not directly observable. At the end of the day, any model of the ERP is a model of investor expectations. One challenge in estimating the ERP is that it is not clear what truly constitutes the market return and the risk-free rate in the real world. In practice, the most common measures of total market return are based on broad stock market indexes, such as the S&P 500 or the Dow Jones Industrial Average, which do not include the whole universe of traded stocks and miss several other components of wealth such as housing, private equity, and nontradable human capital. Even if we restricted ourselves to all traded stocks, we would still have several choices to make, such as whether to use value or equal-weighted indexes, and whether to exclude penny or infrequently traded stocks. A similar problem arises with the risk-free rate. While we almost always use Treasury yields as measures of risk-free rates, nominal Treasury securities are not completely riskless since they are exposed to inflation[3] and liquidity risks, even if we were to assume that there is no prospect of outright default. In this article, we focus on how expectations are estimated in different models, and not on measurement issues regarding market returns and the risk-free rate. Thus, we follow common practice and always use the S&P 500 as a measure of stock market prices and either nominal or real Treasury yields as risk-free rates so that our models are comparable with one another and with most of the literature.

While putting the concept of the ERP into practice has its challenges, we can precisely define the ERP mathematically. First, we decompose stock returns[4] into an expected component and a random component:

$$R_{t+k} = E_t[R_{t+k}] + error_{t+k}. \tag{1}$$

In equation (1), R_{t+k} are *realized* returns between t and $t+k$, and $E_t[R_{t+k}]$ are the returns that were expected from t to $t+k$ using information available at time t. The variable $error_{t+k}$ is a random variable that is unknown at time t and realized at $t+k$. Under rational expectations, $error_{t+k}$ has a mean of zero and is orthogonal to $E_t[R_{t+k}]$. We keep the discussion as general as possible and do not assume rational expectations at this stage, although it will be a feature

[2] A few important references among a vast literature are Ang and Bekaert (2007), Goyal and Welch (2008), Campbell and Thompson (2008), Kelly and Pruitt (2013), Chen, Da, and Zhao (2013), and Neely et al. (2014).

[3] Note that inflation risk in an otherwise risk-free nominal asset does not invalidate its usefulness to compute the ERP. If stock returns and the risk-free rate are expressed in nominal terms, their difference has little or no inflation risk. This follows from the following formula, which holds exactly in continuous time and to a first-order approximation in discrete time: real stock returns − real risk-free rate = (nominal stock returns − expected inflation) − (nominal risk-free rate − expected inflation) = nominal stock returns − nominal risk-free rate. Hence, there is no distinction between a nominal and a real ERP.

[4] Throughout this article, all returns are net returns. For example, a 5 percent return corresponds to a net return of 0.05 as opposed to a gross return of 1.05.

Table 1
Data Sources

Fama and French (1992)	Fama-French factors, momentum factor, twenty-five portfolios sorted on size and book-to-market
Shiller (2005)	Inflation and ten-year nominal Treasury yield. Nominal price, real price, earnings, dividends, and cyclically adjusted price-earnings ratio for the S&P 500
Baker and Wurgler (2007)	Debt issuance, equity issuance, sentiment measure
Graham and Harvey (2012)	ERP estimates from the Duke University/CFO Magazine Global Business Outlook Survey
Damodaran (2012)	ERP estimates
Gurkaynak, Sack, and Wright (2007)	Zero-coupon nominal bond yields for all maturities[a]
Gurkaynak, Refet, Sack, and Wright (2010)	Zero-coupon TIPS (Treasury Inflation-Protected Securities) yields for all maturities
Compustat	Book value per share for the S&P 500
Thomson Reuters I/B/E/S	Mean analyst forecast of expected earnings per share
FRED (Federal Reserve Bank of St. Louis)	Corporate bond Baa-Aaa spread and the National Bureau of Economic Research recession indicator

Notes: All variables start in January 1960 (or later, if unavailable for early periods) and end in June 2013 (or until no longer available). CFO surveys are quarterly; book value per share and ERP estimates by Damodaran (2012) are annual; all other variables are monthly. Appendix A provides more details.

[a] Except for the ten-year yield, which is from Shiller (2005). We use the ten-year yield from Shiller (2005) for ease of comparability with the existing literature. Results are virtually unchanged if we use all yields, including the ten-year yield, from Gurkaynak, Sack, and Wright (2007).

of many of the models we consider. The ERP at time t for horizon k is defined as

$$ERP_t(k) = E_t[R_{t+k}] - R^f_{t+k}, \quad (2)$$

where R^f_{t+k} is the risk-free rate for investing from t to $t+k$ (which, being risk-free, is known at time t).

This definition shows three important aspects of the ERP. First, future expected returns and the future ERP are stochastic, since expectations depend on the arrival of new information that has a random component not known in advance.[5] Second, the ERP has an investment horizon k embedded in it, since we can consider expected excess returns over, say, one month, one year, or five years from today. If we fix t, and let k vary, we trace the *term structure* of the equity risk premium. Third, if expectations are rational, because the unexpected component $error_{t+k}$ has mean zero and is orthogonal to expected returns, the ERP is always less volatile than realized excess returns. In this case, we expect ERP estimates to be smoother than realized excess returns.

[5] More precisely, $E_t[R_{t+k}]$ and $ERP_t(k)$ are known at time t but random from the perspective of all earlier periods.

3. Models of the Equity Risk Premium

We describe twenty models of the equity risk premium, comparing their advantages, disadvantages, and ease of implementation. Of course, there are many more models of the ERP than those we consider. We selected the models in our study based on three criteria: the recent academic literature, widespread use of the models by practitioners, and data availability. Table 1 describes the data we use and their sources, all of which are either readily available or standard in the literature.[6] With a few exceptions, all data are monthly from January 1960 to June 2013. Appendix A provides further detail.

We classify the twenty models into five categories based on their underlying assumptions; models in the same category tend to give similar estimates for the ERP. The five categories are: models based on the historical mean of realized returns, dividend discount models, cross-sectional regressions, time-series regressions, and surveys.

All but one of the estimates of the ERP are constructed in real time, so that an investor who lived through the sample would have been able to construct the measures at each point in time using available information only.[7] This helps minimize look-ahead bias and makes any out-of-sample evaluation of

[6] In fact, except for data from I/B/E/S and Compustat, all sources are public.

[7] The one exception is the cross-sectional model of Adrian, Crump, and Moench (2014), which is constructed using full-sample regression estimates.

Table 2
Models Based on the Historical Mean of Excess Returns

Long-run mean	Average of realized S&P 500 returns minus the risk-free rate using all available historical data
Mean of the previous five years	Average of realized S&P 500 returns minus the risk-free rate using only data for the previous five years

the models more meaningful. Clearly, most of the models themselves were designed only recently and were not available to investors in real time, potentially introducing another source of forward-looking and selection biases that are much more difficult to quantify and eliminate.

3.1 Historical Mean of Realized Returns

The easiest approach to estimating the ERP is to use the historical mean of realized market returns in excess of the contemporaneous risk-free rate. This model is very simple and, as shown in Goyal and Welch (2008), quite difficult to improve upon when considering out-of-sample predictability performance measures. The main drawbacks are that it is purely backward-looking and that it assumes the future will behave like the past—in other words, that the mean of excess returns is either constant or very slow-moving over time, giving very little time-variation in the ERP. The main choice is how far back into the past we should go when computing the historical mean. Table 2 shows the two versions of historical mean models that we use.

3.2 Dividend Discount Models

All dividend discount models (DDMs) start with the basic intuition that the value of a stock is determined by no more and no less than the cash flows it produces for its shareholders, as in Gordon (1962). Today's stock price should then be the sum of all expected future cash flows, discounted at an appropriate rate to take into account their riskiness and the time value of money. The formula that reflects this intuition is

$$P_t = \frac{D_t}{\rho_t} + \frac{E_t\,[D_{t+1}]}{\rho_{t+1}} + \frac{E_t\,[D_{t+2}]}{\rho_{t+2}} + \frac{E_t[D_{t+3}]}{\rho_{t+3}} + \dots\,, \tag{3}$$

where P_t is the current price of the stock, D_t are current cash flows, $E_t[D_{t+k}]$ are the cash flows k periods from now expected as of time t, and ρ_{t+k} is the discount rate for time $t+k$ from the perspective of time t. Cash flows to stockholders certainly include dividends, but they can also arise from spinoffs, buyouts, mergers, and buybacks. In general, the literature focuses on dividend distributions because they are readily available data-wise and account for the vast majority of cash flows. The discount rate can be decomposed into

$$\rho_{t+k} = 1 + R^f_{t+k} + ERP_t(k). \tag{4}$$

In this framework, the risk-free rate captures the discounting associated with the time value of money and the ERP captures the discounting associated with the riskiness of dividends. When using a DDM, we refer to $ERP_t(k)$ as the *implied* ERP. The reason for this is that we plug prices, risk-free rates, and estimated expected future dividends into equation (3) and then derive what value of $ERP_t(k)$ makes the right-hand side equal to the left-hand side in the equation—in other words, what ERP value is *implied* by equation (3).

DDMs are forward-looking and are consistent with no arbitrage. In fact, equation (3) must hold in any economy with no arbitrage.[8] Another advantage of DDMs is that they are easy to implement. A drawback of DDMs is that the results are sensitive to how we compute expectations of future dividends. Table 3 displays the DDMs that we consider and a brief description of their different assumptions.

3.3 Cross-Sectional Regressions

This method exploits the variation in returns and exposures to the S&P 500 of different assets to infer the ERP.[9] Intuitively, cross-sectional regressions find the ERP by answering the following question: what is the level of the ERP that makes expected returns on a variety of stocks consistent with their exposure to the S&P 500? Because we need to explain the relationship between returns and exposures for multiple stocks with a single value for the ERP (and perhaps a small number of other variables), this model imposes tight restrictions on estimates of the ERP.

[8] Note that when performing the infinite summation in equation (3), we have not assumed the n^{th} term goes to zero as n tends to infinity, which allows for rational bubbles. In this sense, DDMs do allow for a specific kind of bubble.

[9] See Polk, Thompson, and Vuolteenaho (2006) and Adrian, Crump, and Moench (2014) for a detailed description of this method.

Table 3

Dividend Discount Models

Gordon (1962) with nominal yields	S&P 500 dividend-to-price ratio minus the ten-year nominal Treasury yield
Shiller (2005)	Cyclically adjusted price-earnings ratio (CAPE) minus the ten-year nominal Treasury yield
Gordon (1962) with real yields	S&P 500 dividend-to-price ratio minus the ten-year real Treasury yield (computed as the ten-year nominal Treasury rate minus the ten-year breakeven inflation implied by TIPS [Treasury Inflation-Protected Securities])
Gordon (1962) with earnings forecasts	S&P 500 expected earnings-to-price ratio minus the ten-year nominal Treasury yield
Gordon (1962) with real yields and earnings forecasts	S&P 500 expected earnings-to-price ratio minus the ten-year real Treasury yield (computed as the ten-year nominal Treasury rate minus the ten-year breakeven inflation implied by TIPS)
Panigirtzoglou and Loeys (2005)	Two-stage dividend discount model. The growth rate of earnings over the first five years is estimated by using the fitted values in a regression of average realized earnings growth over the last five years on its lag and lagged earnings-price ratio. The growth rate of earnings from year six and onward is 2.2 percent
Damodaran (2012)	Six-stage dividend discount model. Dividend growth in the first five stages is estimated from analysts' earnings forecasts. Dividend growth in the sixth stage is the ten-year nominal Treasury yield
Damodaran (2012) free cash flow	Same as Damodaran (2012) but uses free-cash-flow-to-equity as a proxy for dividends plus stock buybacks

Sources: See Appendix A and Table 1 for full source details.

The first step is to find the exposures of some assets to the S&P 500 by estimating an equation of the following form:

$$(5) \quad R^i_{t+k} - R^f_{t+k} = \alpha^i \times state\ variables_{t+k} + \beta^i \times risk\ factors_{t+k} + idiosyncratic\ risk^i_{t+k}.$$

In equation (5), R^i_{t+k} is the realized return on a stock or portfolio *i* from time *t* to *t* +*k*. *State variables*$_{t+k}$ are any economic indicators that help identify the state of the economy and its likely future path. *Risk factors*$_{t+k}$ are any measures of systematic contemporaneous covariation in returns across all stocks or portfolios. Of course, some economic indicators can be both state variables and risk factors at the same time. Finally, $idiosyncratic\ risk^i_{t+k}$ is the component of returns that is particular to each individual stock or portfolio that is not explained by *state variables*$_{t+k}$ or *risk factors*$_{t+k}$ (both of which, importantly, are common to all stocks and hence not indexed by *i*). Examples of state variables are inflation, unemployment, the yield spread between Aaa and Baa bonds, the yield spread between short- and long-term Treasury securities, and the S&P 500's dividend-to-price ratio. The most important risk factor is the excess return on the S&P 500, which we must include if we want to infer the ERP consistent with the cross section of stock returns. Other risk factors usually used are the Fama-French (1992) factors and the momentum factor of Carhart (1997). The values in the vector α^i give the strength of asset-specific return predictability and the values in the vector β^i give the asset-specific exposures to risk factors.[10] For the cross section of assets indexed by *i*, we can use the whole universe of traded stocks, a subset of them, or portfolios of stocks grouped, for example, by industry, size, book-to-market, or recent performance. It is important to point out that equation (5) is not a predictive regression; the left- and right-hand-side variables are both associated with time *t*+*k*.

The second step is to find the ERP associated with the S&P 500 by estimating the cross-sectional equations

$$(6) \quad R^i_{t+k} - R^f_{t+k} = \lambda_t(k) \times \hat{\beta}^i,$$

where $\hat{\beta}^i$ are the values found when estimating equation (5). Equation (6) attempts to find, at each point in time, the vector of numbers $\lambda_t(k)$ that makes exposures β^i as consistent as possible with realized excess returns of all stocks or portfolios considered. The element in the vector $\hat{\lambda}_t(k)$ that is multiplied by the element in the $\hat{\beta}^i$ vector corresponding to the S&P 500 is $ERP_t(k)$, the equity risk premium we are seeking.

One advantage of cross-sectional regressions is that they use information from more asset prices than other models. Cross-sectional regressions also have sound theoretical foundations, since they provide one way to implement Merton's (1973) Intertemporal Capital Asset Pricing Model. Finally, this method nests many of the other models considered. The two main drawbacks of this method are that results are dependent on the portfolios, state variables, and risk factors that are used (Harvey, Liu, and Zhu 2014), and that it is not as easy to implement as most of the other options. Table 4 displays the

[10] The vectors α^i and β^i could also be time-varying, reflecting a more dynamic relation between returns and their explanatory variables. In this case, the estimation of equation (5) is more complicated and requires making further assumptions. The model by Adrian, Crump, and Moench (2014) is the only cross-sectional model we examine that uses time-varying α^i and β^i.

Table 4
Models with Cross-Sectional Regressions

Fama and French (1992)	Uses the excess returns on the market portfolio, a size portfolio, and a book-to-market portfolio as risk factors
Carhart (1997)	Identical to Fama and French (1992) but adds the momentum measure of Carhart (1997) as an additional risk factor
Duarte (2013)	Identical to Carhart (1997) but adds an inflation risk factor
Adrian, Crump, and Moench (2014)	Uses the excess returns on the market portfolio as the single risk factor. The state variables are the dividend yield, the default spread, and the risk-free rate

Sources: See Appendix A and Table 1 for full source details.

Table 5
Models with Time-Series Regressions

Fama and French (1988)	Only predictor is the dividend-price ratio of the S&P 500
Goyal and Welch (2008)	Uses, at each point in time, the best out-of-sample predictor out of twelve predictive variables proposed by Goyal and Welch (2008)
Campbell and Thompson (2008)	Same as Goyal and Welch (2008) but imposes two restrictions on the estimation. First, the coefficient b in equation (9) is replaced by zero if it has the "wrong" theoretical sign. Second, the estimate of the ERP is replaced by zero if the estimation otherwise finds a negative ERP
Fama and French (2002)	Uses, at each point in time, the best out-of-sample predictor out of three variables: the price-dividend ratio adjusted by the growth rate of earnings, dividends, or stock prices
Baker and Wurgler (2007)	The predictor is Baker and Wurgler's (2007) sentiment measure. The measure is constructed by finding the most predictive linear combination of six variables: the closed-end fund discount, New York Stock Exchange share turnover, the number of initial public offerings, the average first-day returns on initial public offerings, the equity share in new issues, and the dividend premium

Sources: See Appendix A and Table 1 for full source details.

cross-sectional models in our study, together with the state variables and risk factors they use.

3.4 Time-Series Regressions

Time-series regressions use the relationship between economic variables and stock returns to estimate the ERP. The idea is to run a predictive linear regression of realized excess returns on lagged "fundamentals":

$$R_{t+k} - R^f_{t+k} = a + b \times Fundamental_t + error_t \,. \quad (7)$$

Once estimates $\hat{a}$ and $\hat{b}$ for a and b are obtained, the ERP is obtained by ignoring the error term:

$$ERP_t(k) = \hat{a} + \hat{b} \times Fundamental_t \,. \quad (8)$$

In other words, we estimate only the forecastable or expected component of excess returns. This method attempts to implement equations (1) and (2) as directly as possible in equations (7) and (8), with the assumption that "fundamentals" are the right sources of information to look at when computing expected returns, and that a linear equation is the correct functional specification.

The use of time-series regressions requires a minimal number of assumptions; there is no concept of equilibrium and no absence of arbitrage necessary for the method to be valid.[11] In addition, implementation is quite simple, since it only involves running ordinary least-square regressions. The challenge is to select the variables to include on the right-hand side of equation (7), since results can change substantially depending on the variables that are used to take the role of "fundamentals." Including more than one predictor gives poor out-of-sample performance even if economic theory may suggest a role for many variables to be used simultaneously (Goyal and Welch 2008). Finally, time-series regressions ignore information in the cross section of stock returns. Table 5 shows the time-series regression models that we study.

[11] However, the Arbitrage Pricing Theory of Ross (1976) provides a strong theoretical underpinning for time-series regressions by using no-arbitrage conditions.

Table 6
Surveys

Graham and Harvey (2012)	Since 1996, the Duke University/CFO Magazine Global Business Outlook Survey has asked chief financial officers about the one- and ten-year-ahead ERP. We take the mean of all responses

Sources: See Appendix A and Table 1 for full source details.

3.5 Surveys

The survey approach consists of asking economic agents about the current level of the ERP. Surveys incorporate the views of many people, some of whom are very sophisticated and/or make real investment decisions based on the level of the ERP. Surveys should also be good predictors of excess returns because, in principle, stock prices are determined by the supply and demand of investors such as the ones taking the surveys. However, Greenwood and Shleifer (2014) document that investor expectations of future stock market returns are positively correlated with past stock returns and with the current level of the stock market, but strongly *negatively* correlated with model-based expected returns and future realized stock market returns. Other studies such as Easton and Sommers (2007) also argue that survey measures of the ERP can be systematically biased. In this article, we use the Duke University/CFO Magazine Global Business Outlook Survey of chief financial officers by Graham and Harvey (2012), which, to our knowledge, is the only large-scale ERP survey that has more than just a few years of data (see Table 6).

4. Estimation of the Equity Risk Premium

We now study the behavior of the twenty models under consideration by conducting principal component analysis. Since forecast accuracy can be substantially improved through the combination of multiple forecasts,[12] the optimal strategy to forecast excess stock returns may consist of combining all of these models. The first principal component of the twenty models that we use is the linear combination of ERP estimates that captures as much of the variation in the data as possible. The second, third, and successive principal components are the linear combinations of the twenty models that explain as much of the variation of the data as possible and are also uncorrelated to all of the preceding principal components. If the first few principal components—say one or two—account for most of the variation of the data, then we can use them as a good summary for the variation in all the measures over time, reducing the dimensionality from twenty to one or two. In addition, in the presence of classical measurement error, the first few principal components can achieve a higher signal-to-noise ratio than other summary measures like the cross-sectional mean of all models (Geiger and Kubin 2013).

To compute the first principal component, we proceed in three steps. First, we de-mean all ERP estimates and find their variance-covariance matrix. Second, we find the linear combination that explains as much of the variance of the de-meaned models as possible. The weights in the linear combination are the elements of the eigenvector associated with the largest eigenvalue of the variance-covariance matrix found in the first step. Third, we add to the linear combination just obtained, which has a mean of zero, the average of ERP estimates across all models and all time periods. Under the assumption that each of the models is an unbiased and consistent estimator of the ERP, the average across all models and all time periods is an unbiased and consistent estimator of the unconditional mean of the ERP. The time variation in the first principal component then provides an estimate of the conditional ERP.[13] The share of the variance of the underlying models explained by this principal component is 76 percent, suggesting that little would be gained from examining principal components beyond the first.[14]

We now focus on the one-year-ahead ERP estimates and study other horizons in the next section.

The first two columns in Table 7 show the mean and standard deviation of each model's estimates. The unconditional

[12] See, *inter alia*, Clemen (1989), Diebold and Lopez (1996), and Timmermann (2006).

[13] As is customary in the literature, we perform the analysis using ERP estimates in levels, even though they are quite persistent. Results in first differences do not give economically reasonable estimates since they feature a pro-cyclical ERP and unreasonable magnitudes.

One challenge that arises in computing the principal component is when observations are missing, either because some models can only be obtained at frequencies lower than monthly or because the necessary data are not available for all time periods (Appendix A contains a detailed description of when this happens). To overcome this challenge, we use an iterative linear projection method, which conceptually preserves the idea behind principal components. Let X be the matrix that has observations for different models in its columns and for different time periods in its rows. On the first iteration, we make a guess for the principal component and regress the nonmissing elements of each row of X on the guess and a constant. We then find the first principal component of the variance-covariance matrix of the fitted values of these regressions, and use it as the guess for the next iteration. The process ends when the norm of the difference between consecutive estimates is small enough. We thank Richard Crump for suggesting this method and providing the code for its implementation.

[14] The second and third principal components account for 13 and 8 percent of the variance, respectively.

Table 7
ERP Models

		Mean	Standard Deviation	PC Coefficients $\hat{w}^{(m)}$	Exposure to PC $load_1^{(m)}$
Based on historical mean	Long-run mean	9.3	1.3	0.78	-0.065
	Mean of previous five years	5.7	5.8	0.42	-0.160
Dividend Discount Models (DDM)	Gordon (1962): E/P minus nominal ten-year yield	-0.1	2.1	-0.01	0.001
	Shiller (2005): 1/CAPE minus nominal ten-year yield	-0.4	1.8	-0.10	0.011
	Gordon (1962): E/P minus real ten-year yield	3.5	2.1	0.69	-0.077
	Gordon (1962): Expected E/P minus real ten-year yield	5.3	1.7	-0.78	0.208
	Gordon (1962): Expected E/P minus nominal ten-year yield	0.4	2.3	-0.79	0.077
	Panigirtzoglou and Loeys (2005): Two-stage DDM	-1.0	2.3	0.07	-0.011
	Damodaran (2012): Six-stage DDM	3.4	1.3	-0.26	0.032
	Damodaran (2012): Six-stage free cash flow DDM	4.0	1.1	-0.62	0.053
Cross-sectional regressions	Fama and French (1992)	12.6	0.7	0.80	-0.040
	Carhart (1997): Fama-French and momentum	13.1	0.8	0.81	-0.042
	Duarte (2013): Fama-French, momentum, and inflation	13.1	0.8	0.82	-0.044
	Adrian, Crump, and Moench (2014)	6.5	6.9	-0.05	0.114
Time-series regressions	Fama and French (1988): D/P	2.4	4.0	-0.27	0.069
	Best predictor in Goyal and Welch (2008)	14.5	5.2	-0.07	0.023
	Best predictor in Campbell and Thompson (2008)	3.1	9.8	-0.12	0.081
	Best predictor in Fama and French (2002)	11.9	6.8	-0.72	0.321
	Baker and Wurgler (2007) sentiment measure	3.0	4.7	-0.32	0.184
Surveys	Graham and Harvey (2012) Duke University/ CFO Magazine Global Business Outlook Survey	3.6	1.8	0.72	0.264
	All models	5.7	3.2	0.78	-0.065

Sources: See Appendix A and Table 1 for full source details.

Notes: For each of the twenty models of the equity risk premium, we show four statistics. The first two are the time-series means and standard deviations for monthly observations from January 1960 to June 2013 (except for surveys, which are quarterly). The units are annualized percentages. The third statistic, "PC Coefficients $\hat{w}^{(m)}$", is the weight that the first principal component places on each model (normalized to sum to one). The fourth is the "Exposure to PC $load_1^{(m)}$", the weight on the first principal component when each model is written as a weighted sum of all principal components (also normalized to sum to one). E/P is earnings-to-price. CAPE is cyclically adjusted price-to-earnings. D/P is dividend-to-price.

mean of the ERP across all models is 5.7 percent, with an average standard deviation of 3.2 percent. DDMs give the lowest mean ERP estimates and have moderate standard deviations. In contrast, cross-sectional models tend to have mean ERP estimates on the high end of the distribution and very smooth time series. Mean ERP estimates for time-series regressions are mixed, with high and low values depending on the predictors used, but uniformly large variances. The survey of CFOs has a mean and standard deviation that are both about half as large as in the overall population of models. The picture that emerges from Table 7 is that there is considerable heterogeneity across model types, and even sometimes within model types, thereby underscoring the difficulty inherent in finding precise estimates of the ERP.

Chart 1 shows the time series for all one-year-ahead ERP model estimates, with each class of models in a different panel. The green lines are the ERP estimates from the twenty underlying models. The black line, reproduced in each of the panels, is the principal component of all twenty models. The chart gives a sense of how the time series move together and how much they covary with the first principal component. Table 8 shows the correlations among models. Chart 1 and Table 8 give the same message: despite some outliers, there is a fairly strong correlation within each of the five classes of models. Across classes, however, correlations are small and even negative. Interestingly, the correlation between some DDMs and cross-sectional models is as low as -91 percent. This negative correlation, however, disappears if we look at lower frequencies. When aggregated to quarterly frequency, the smallest correlation between DDM and cross-sectional models is -22 percent, while at the annual frequency it is 12 percent.

Chart 1 also shows that the first principal component covaries negatively with historical mean models but positively with DDMs and cross-sectional regression models. Time-series regression models are also positively correlated with the first principal component, although this is not so

Chart 1

ERP Estimates for All Models

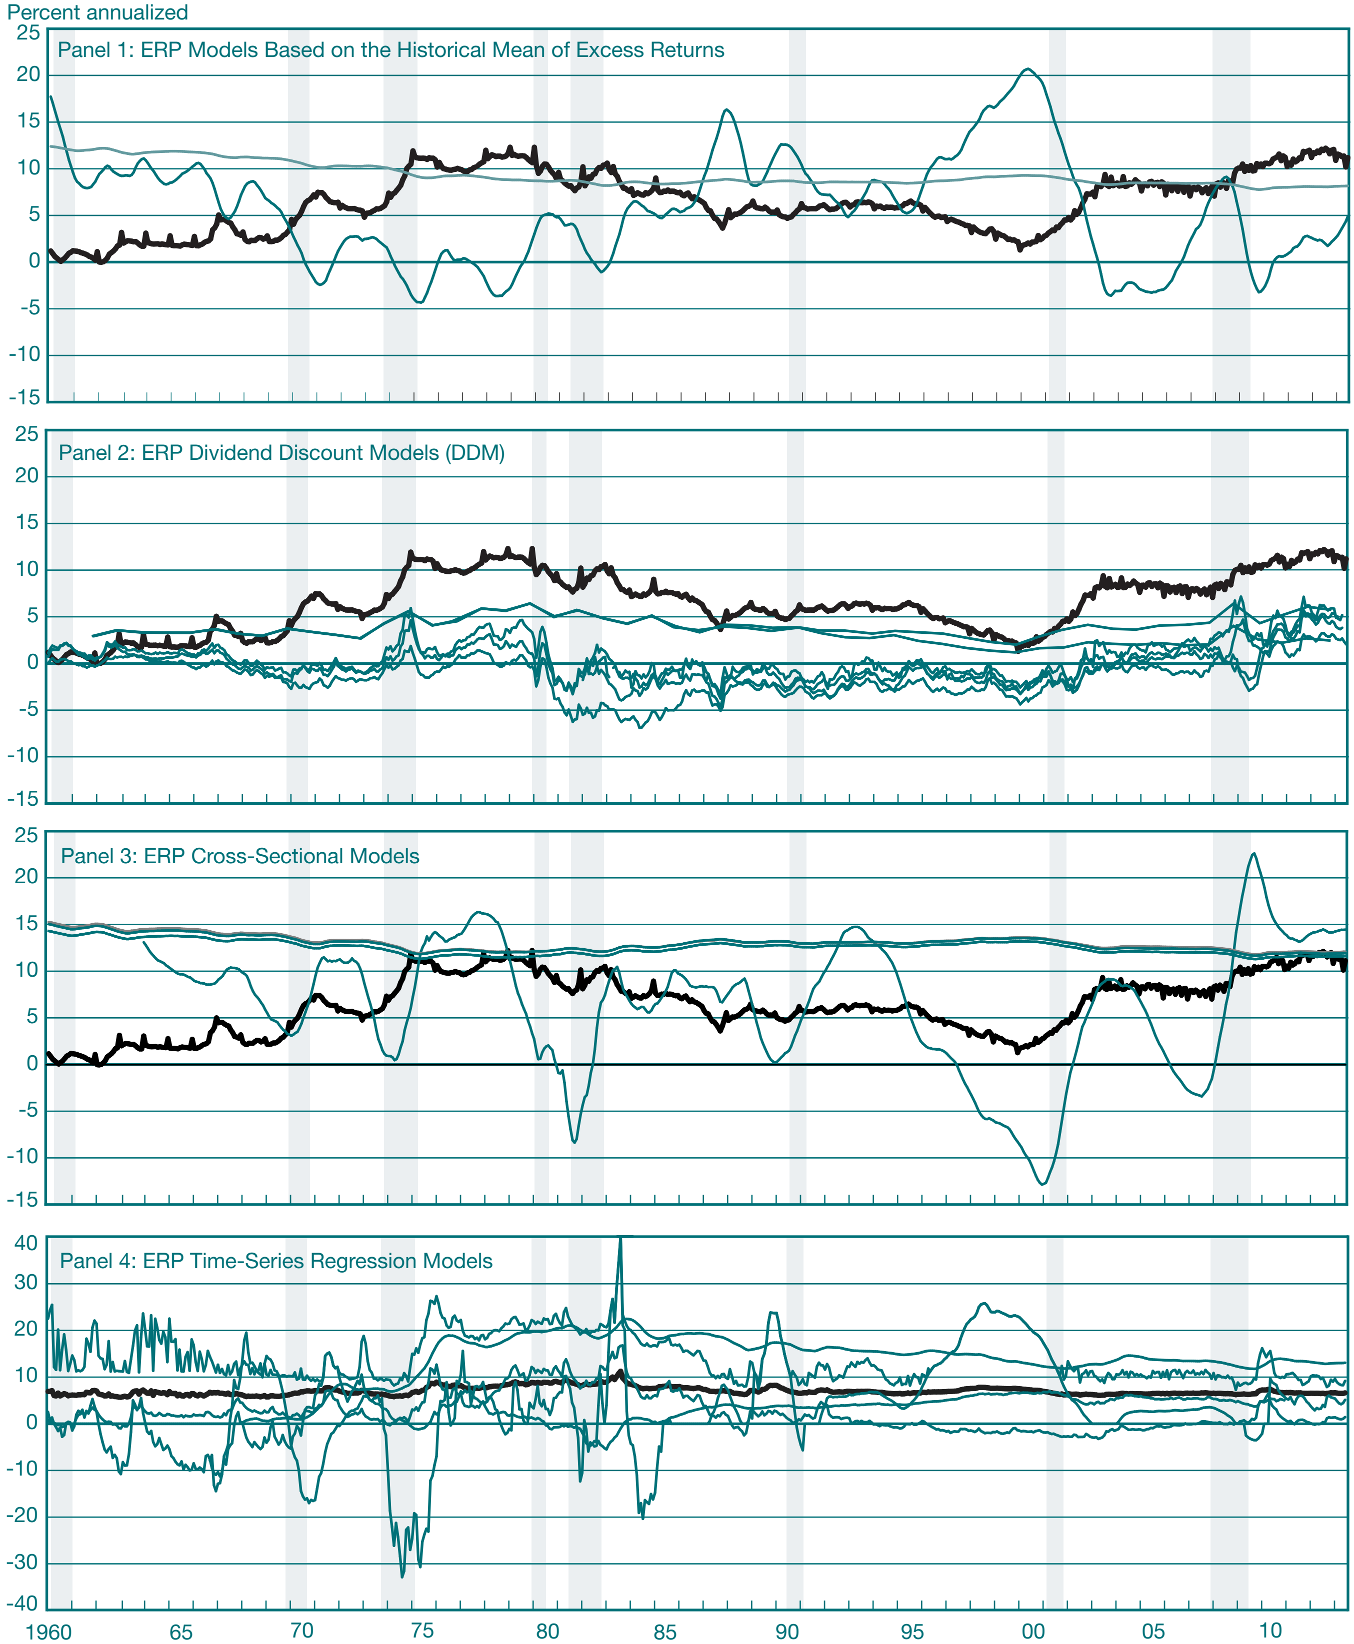

Chart 1 (Continued)

ERP Estimates for All Models

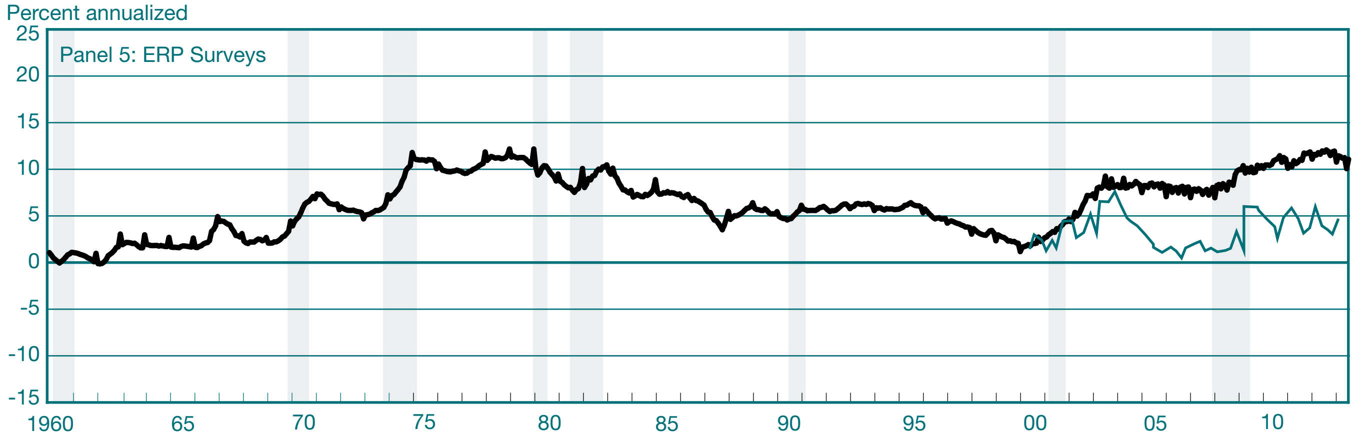

Sources: See Appendix A and Table 1 for detailed source information.

Notes: Each green line gives the one-year-ahead equity risk premium from one of the models listed in Tables 2 through 6. Panel 1 shows the estimates for models based on the historical mean of excess returns; these models are listed in Table 2. Panel 2 shows estimates computed by the dividend discount models listed in Table 3. Panel 3 uses the cross-sectional regression models listed in Table 4. Panel 4 uses the time-series regression models listed in Table 5. Panel 5 shows the estimate obtained from the survey cited in Table 6. In all panels, the black line is the first principal component of all twenty models (it can look different across panels because of different scales used in the y-axis.) The shaded areas indicate periods designated recessions by the National Bureau of Economic Research.

clearly seen in Panel 4 of Chart 1 because of the high volatility of time-series regression ERP estimates. The last panel shows that the survey of CFOs does track the first principal component quite well at low frequencies (for example, annual), although any conclusions about survey estimates should be interpreted with caution given the short length of the sample.

As explained earlier, the first principal component is a linear combination of the twenty underlying ERP models:

(9) $$PC_t^{(1)} = \Sigma_{m=1}^{20} w^{(m)} ERP_t^{(m)}.$$

In the above equation, m indexes the different models, $PC_t^{(1)}$ is the first principal component, $ERP_t^{(m)}$ is the estimate from model m, and $w^{(m)}$ is the weight that the principal component places on model m. The third column in Table 7, labeled "PC Coefficients," shows the weights $w^{(m)}$ normalized to sum up to one to facilitate comparison; in other words, the table reports the weights $\widehat{w}^{(m)}$, where

(10) $$\widehat{w}^{(m)} = \frac{w^{(m)}}{\Sigma_{m=1}^{20} w^{(m)}}.$$

The first principal component puts positive weight on models based on the historical mean, cross-sectional regressions, and the survey of CFOs. It weights DDMs and time-series regressions mostly negatively. The absolute values of the weights are very similar for many of the models, and there is no single model or class of models that dominates. This means that the first principal component uses information from many of the models.

The last column in Table 7, labeled "Exposure to PC," shows the extent to which models load on the first principal component. By construction, each of the twenty ERP models can be written as a linear combination of the twenty principal components:

(11) $$ERP_t^{(m)} = \Sigma_{i=1}^{20} load_i^{(m)} PC_t^{(i)},$$

where m indexes the model and i indexes the principal components. The values in the last column of Table 7 are the loadings on the first principal component ($i = 1$) for each model ($m = 1, 2, \ldots, 20$), again normalized to one for ease of comparability:

(12) $$\widehat{load}_1^{(m)} = \frac{load_1^{(m)}}{\Sigma_{m=1}^{20} load_1^{(m)}}.$$

Most models have a positive loading on the first principal component; whenever the loading is negative, it tends to be relatively small. This means that the first principal component, as expected, is a good explanatory variable for most models. Looking at the third and fourth columns of Table 7 together,

Table 8
Correlation of ERP Models

	Long-run mean	Mean past five years	E/P-ten year	1/CAPE-ten year	E/P-real ten year	Exp E/P-real ten year	Exp E/P-ten year	Two-stage DDM	Six-stage DDM	Free cash flow	Fama and French	Carhart	Duarte	Adrian, Crump, and Moench	D/P	Goyal and Welch	Campbell and Thompson	Fama and French	Sentiment	CFO survey
Long-run mean	100																			
Mean past five years	32	100																		
E/P-ten year	8	15	100																	
1/CAPE-ten year	-9	0	78	100																
E/P-real ten year	-11	25	98	23	100															
Exp E/P-real ten year	-58	42	70	84	60	100														
Exp E/P-ten year	-83	-61	84	95	46	98	100													
Two-stage DDM	17	27	88	54	89	66	79	100												
Six-stage DDM	3	-38	26	39	-30	32	52	-31	100											
Free cash flow	-43	-55	59	70	35	80	94	27	62	100										
Fama and French	69	29	-8	-36	-21	-69	-91	9	-29	-77	100									
Carhart	71	30	-5	-31	-24	-71	-91	10	-25	-75	99	100								
Duarte	71	30	-3	-29	-22	-70	-91	11	-28	-74	99	100	100							
Adrian, Crump, and Moench	-1	-52	36	62	6	54	63	27	23	33	-28	-28	-25	100						
D/P	49	12	27	12	27	42	54	24	74	42	44	54	55	21	100					
Goyal and Welch	25	12	25	21	-7	-36	-60	20	29	-9	7	13	14	-24	61	100				
Campbell and Thompson	27	31	14	-7	81	49	-60	28	-51	-40	60	57	58	-33	54	50	100			
Fama and French	1	-30	-24	-29	37	-27	-37	-18	22	38	36	38	37	-9	40	23	43	100		
Sentiment	-10	33	-4	-20	68	-23	-29	27	-38	-20	18	17	18	-12	-38	-8	21	6	100	
CFO survey	-43	-33	12	30	1	1	13	16	5	-3	-36	-37	-39	60	14	-21	-32	-3	-36	100

Sources: See Appendix A and Table 1 for additional source details.

Notes: This table shows the correlation matrix of the twenty equity risk premium models we consider. Numbers are rounded to the nearest integer. Thick lines group models by their type (see Tables 2-6). Except for the chief financial officer (CFO) survey, the observations used to compute correlations are monthly for January 1960 to June 2013. For the CFO survey, correlations are computed by taking the last observation in the quarter for the monthly series and then computing quarterly correlations. E/P is earnings-to-price. CAPE is cyclically adjusted price-to-earnings. DDM is dividend discount model. D/P is dividend-to-price.

we can obtain additional information. For example, when a model has a very high loading (fourth column) accompanied by a very small PC coefficient (third column), it likely means that the model is almost redundant, in the sense that it is close to being a linear combination of all other models and does not provide much independent information to the principal component. However, if the PC coefficient and loading are both high, the corresponding model is likely providing information not contained in other measures.

Chart 2
One-Year-Ahead ERP

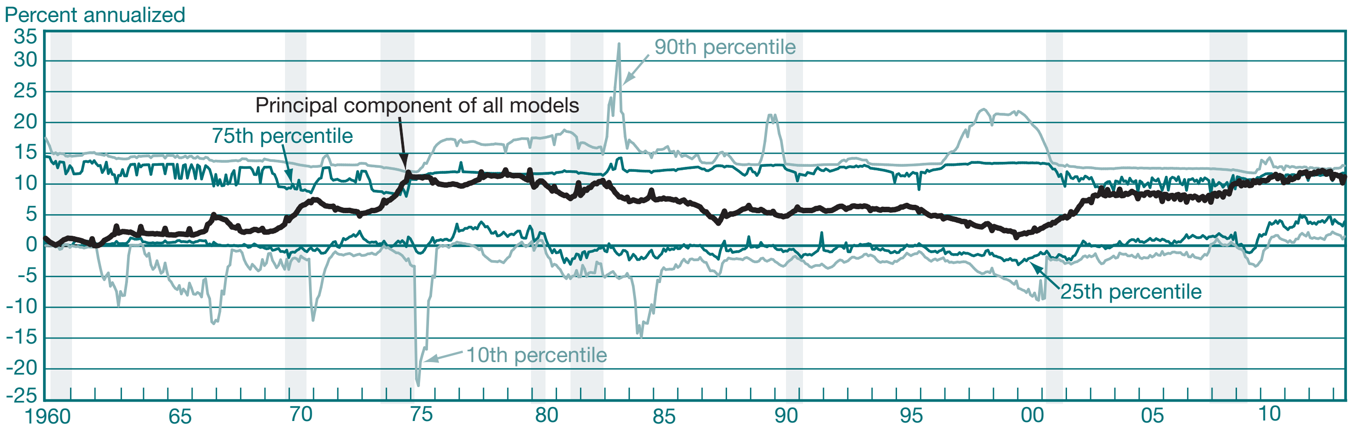

Notes: The black line is the first principal component of twenty models of the one-year-ahead equity risk premium (this is the same principal component shown in black in all panels of Chart 1). The models are listed in Tables 2 through 6. The green lines give the corresponding percentiles of the twenty estimates for each time period. The shaded areas indicate periods designated recessions by the National Bureau of Economic Research.

Chart 2 shows the first principal component of all twenty models in black (the black line is the same principal component shown in black in each of the panels of Chart 1). As expected, the principal component tends to peak during financial turmoil, recessions, and periods of low real GDP growth or high inflation. It tends to bottom out after periods of sustained bullish stock markets and high real GDP growth. Evaluated by the first principal component, the one-year-ahead ERP reaches a local peak in June 2012 at 12.2 percent. The surrounding months have ERP estimates of similar magnitude, with the most recent estimate in June 2013 at 11.2 percent. This behavior is not so clearly seen by simply looking at the collection of individual models in Chart 1, a finding that highlights the usefulness of principal component analysis. Similarly high levels were observed in the mid- and late 1970s, during a period of stagflation, while the recent financial crisis had slightly lower ERP estimates, closer to 10 percent.

Chart 2 also displays the 10th, 25th, 75th, and 90th percentiles of the cross-sectional distribution of models. These bands can be interpreted as confidence intervals since they give the range of the distribution of ERP estimates at each point in time. However, they do not incorporate other relevant sources of uncertainty, such as the errors that occur during the estimation of each individual model, the degree of doubt in the correctness of each model, and the correlation structure between these and all other kinds of errors. Standard error bands that capture all sources of uncertainty are therefore likely to be wider.

The difference in high and low percentiles can also be interpreted as measures of agreement across models. The interquartile range—the difference between the 25th and 75th percentiles—is 11.6 percent on average. It has recently compressed, mostly because the models in the bottom of the distribution have had higher ERP estimates since 2010 while the 75th percentile has remained fairly constant. The lowest value for the interquartile range, 6.8 percent, was reached in 2012. The cross-sectional standard deviation in ERP estimates (not shown in the chart) also decreased from 10.2 percent in January 2000 to 4.3 percent in June 2013, confirming that the disagreement among models has decreased.

Another a priori reasonable summary statistic for the ERP is the cross-sectional mean of estimates across models. In Chart 3, we can see that, by this measure, the ERP has also been increasing since the crisis. However, unlike the principal component, it has not reached elevated levels compared with past values. The cross-sectional mean can be useful, but compared with the first principal component, it has a few undesirable features as an overall measure of the ERP. First, it is procyclical, which contradicts the economic intuition that expected returns are highest in recessions, when risk aversion is high and future prospects look brighter than current ones. Second, it overloads on DDM simply because there is a higher number of DDM models in our sample. And last, it has a smaller correlation with the realized returns it is supposed to predict.

Chart 3
One-Year-Ahead ERP and Cross-Sectional Mean of Models

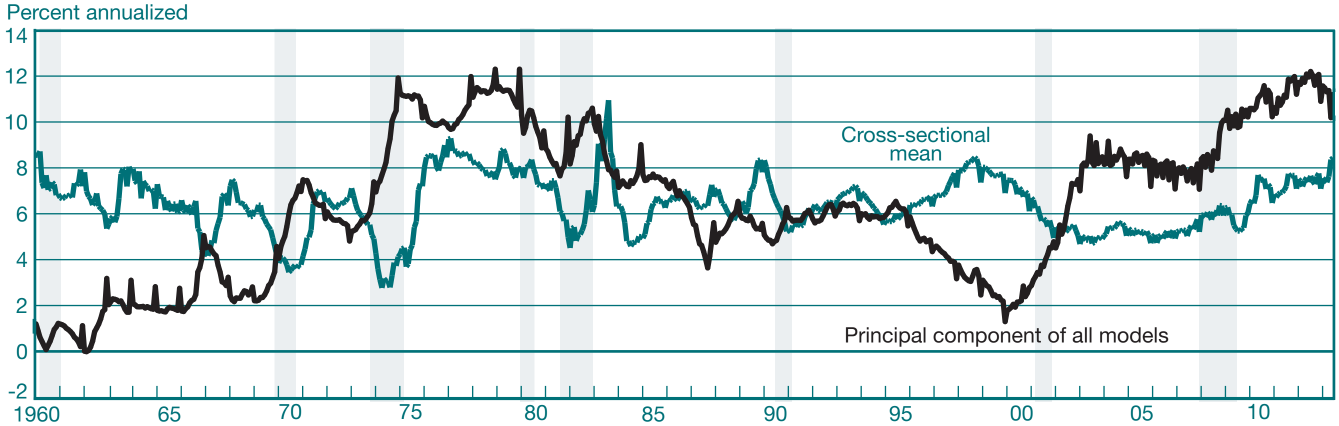

Sources: See Appendix A and Table 1 for detailed source information.

Notes: The black line is the first principal component of twenty models of the one-year-ahead equity risk premium (also shown in Charts 1 and 2). The green line is the cross-sectional average of models for each time period. The shaded areas indicate periods designated recessions by the National Bureau of Economic Research.

5. The Term Structure of Equity Risk Premia

In Section 2, we described the term structure of the ERP—the expected excess returns over different investment horizons. In practical terms, we estimate the ERP at different horizons by using the inputs for all the models at the corresponding horizons.[15] For example, if we want to take the historical mean of returns as our estimate, we can take the mean of returns over a one-month, six-month, or one-year period. In cross-sectional and time-series regressions, we can predict monthly, quarterly, or annual returns using monthly, quarterly, or annual right-hand-side variables. DDMs, on the other hand, have little variation across horizons. In fact, all the DDMs we consider have a constant term structure of expected stock returns, and the only term structure variation in ERP estimates comes from risk-free rates.[16]

Chart 4 plots the first principal components of the ERP as a function of investment horizon for some dates when the ERP

[15] For other ways to estimate the term structure of the ERP using equilibrium models or derivatives, see Äit-Sahalia, Karaman, and Mancini (2014), Ang and Ulrich (2012), van Binsbergen et al. (2014), Boguth et al. (2012), Durham (2013), Croce, Lettau, and Ludvigson (2015), Lemke and Werner (2009), Lettau and Wachter (2011), and Muir (2013), among others.

[16] In equation (3), ρ_{t+k} is assumed to be the same for all k, while risk-free rates are allowed to vary over the investment horizon k in equation (4). Of course, with additional assumptions, it is possible to have DDMs with a nonconstant term structure of expected excess returns.

Chart 4
Term Structure of the ERP

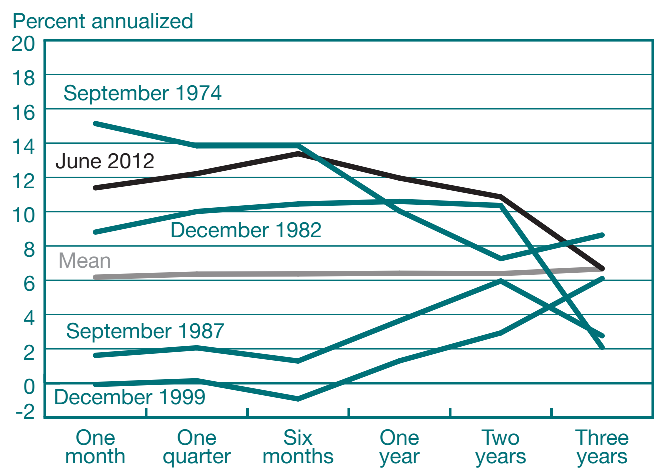

Notes: Each line, except for the gray one, shows equity risk premia as a function of investment horizon for some specific months in our sample. We consider horizons of one month, one quarter, six months, one year, two years, and three years. The gray line (labeled "Mean") shows the average risk premium at different horizons over the full sample, January 1960 to June 2013. September 1987 and December 1999 were low points in one-month-ahead equity premia. In contrast, September 1974, December 1982, and June 2012 were peaks in the one-month-ahead equity premium.

was unusually high or unusually low at the one-month horizon. As was the case for one-year-ahead ERP estimates, we can capture the majority of the variance of the underlying models at all horizons by a single principal component. The shares of the variance explained by the first principal components at horizons of one month to three years range from 68 to 94 percent. The gray line in Chart 4 shows the average of the term structure across all periods. It is slightly upward sloping, with a short-term ERP at just over 6 percent and a three-year ERP at almost 7 percent.

The first observation is that the term structure of the ERP has significant time variation and can be flat, upward sloping, or downward sloping. Chart 4 also shows some examples that hint at lower future expected excess returns when the one-month-ahead ERP is elevated and the term structure is downward sloping, and higher future expected excess returns when the one-month-ahead ERP is low and the term structure is upward sloping. In fact, this is true more generally: there is a strong negative correlation between the level and the slope of the ERP term structure of -71 percent. Chart 5 plots monthly observations of the one-month-ahead ERP against the slope of the ERP term structure (the three-year-ahead minus the one-month-ahead ERP) together with the corresponding ordinary least squares regression line in black. Of course, this is only a statistical pattern and should not be interpreted as a causal relation.

6. Why Is the Equity Risk Premium High?

There are two reasons why the ERP can be high: low discount rates and high current or expected future cash flows.

Chart 6 shows that earnings are unlikely to be the reason the ERP is high. The green line shows the year-on-year change in the mean expectation of one-year-ahead earnings per share for the S&P 500. These expectations are obtained from surveys conducted by the Institutional Brokers' Estimate System (I/B/E/S) and available from Thomson Reuters. Expected earnings per share declined from 2010 to 2013, making earnings growth an unlikely reason for the high ERP in the corresponding period. The black line shows the realized monthly growth rates of real earnings for the S&P 500 expressed in annualized percentages. Since 2010, earnings growth has been declining, hovering around zero for the last few months of the sample. At the end of the sample, it stands at 2.5 percent, which is near its long-run average.

Another way to examine whether a high ERP is caused by discount rates or cash flows is shown in Chart 7. The

Chart 5

Regression of the Slope of the ERP Term Structure on One-Month-Ahead ERP

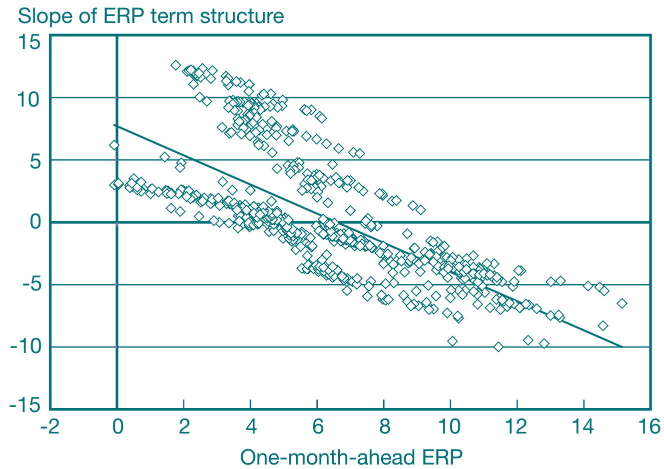

Notes: The chart shows monthly observations and the corresponding OLS regression for the one-month-ahead ERP plotted against the slope of the ERP term structure for the period January 1960 to June 2013. The slope of the ERP term structure is the difference between the three-year-ahead ERP and the one-month-ahead ERP. All units are in annualized percentages. The one-month-ahead and three-year-ahead ERP estimates used are the first principal components of twenty one-month-ahead or three-year-ahead ERP estimates from models described in Tables 2-6. The OLS regression slope is -1.17 (significant at the 99 percent level) and the R^2 is 50.1 percent.

Chart 6

Earnings Behavior

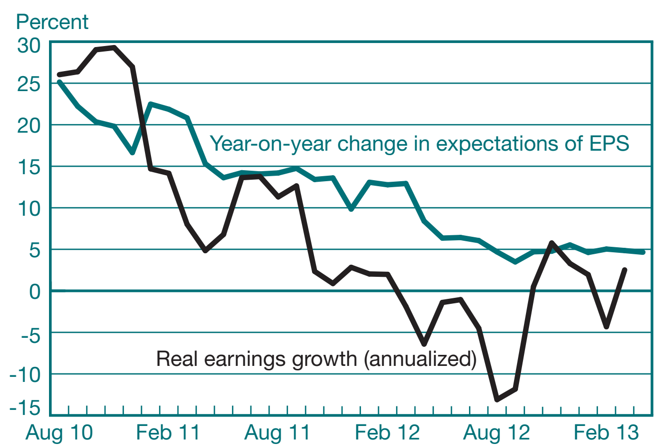

Notes: The black line shows the monthly growth rate of real S&P 500 earnings, annualized and in percentages. The green line shows the year-on-year change in the mean expectation of one-year-ahead earnings per share (EPS) for the S&P 500 from a survey of analysts provided by Thomson Reuters I/B/E/S.

Chart 7
One-Year-Ahead ERP and Expected Returns

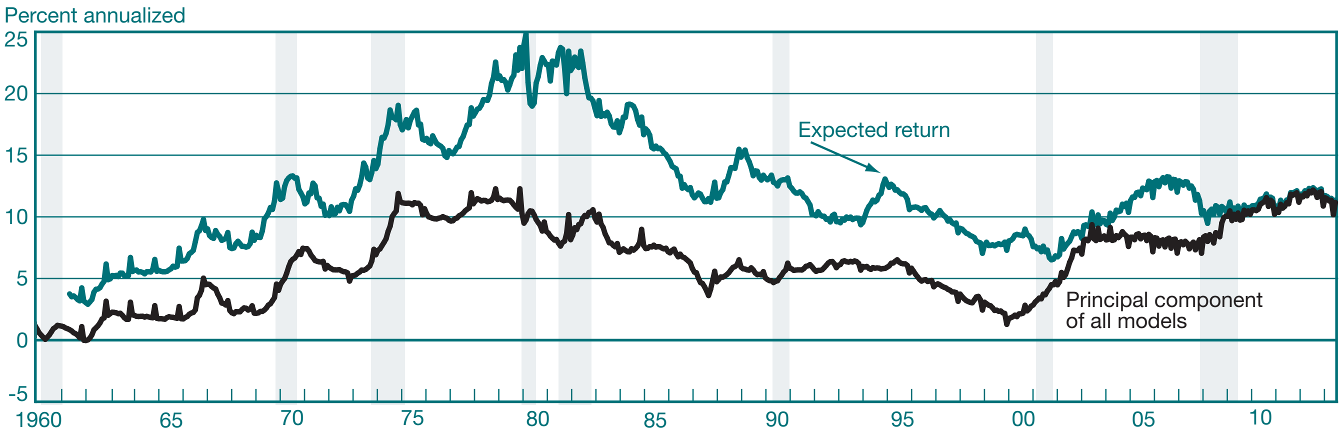

Notes: See Appendix A and Table 1 for detailed source information.

Notes: The black line is the first principal component of twenty models of the one-year-ahead equity risk premium (also shown in Charts 1, 2, and 3). The green line is the one-year-ahead expected return on the S&P 500, obtained by adding the realized one-year maturity Treasury yield from the principal component (the black line). The shaded areas indicate periods designated recessions by the National Bureau of Economic Research.

black line is the same one-year-ahead ERP estimate shown in Chart 2. The green line simply adds the realized one-year Treasury yield to obtain expected stock returns. The chart shows that expected stock returns have increased since 2000, similar to the ERP. However, unlike the ERP, expected stock returns are close to their long-run mean and nowhere near their highest levels, achieved in 1980. The discrepancies between the two lines are the result of exceptionally low bond yields since the end of the financial crisis.

Chart 8 displays the term structure of the ERP under a simple counterfactual scenario, in addition to the mean and current term structures already displayed in Chart 4. In this scenario, we leave expected stock returns unmodified but change the risk-free rates in June 2012 from their actual values to the average nominal bond yields over the period 1960-2013. In other words, we replace R^f_{t+k} in equation (2) by the mean of R^f_{t+k} over t. The result of this counterfactual is shown in Chart 8 in green. Using average levels of bond yields brings the whole term structure of the ERP much closer to its mean level (the gray line), especially at intermediate horizons. This shows that a "normalization" of bond yields, everything else being equal, would bring the ERP close to its historical norm. This exercise shows that the current environment of low bond yields is capable, quantitatively speaking, of significantly contributing to an ERP as high as was observed in 2012-13.

Chart 8
Term Structure of the ERP Using Counterfactual Bond Yield

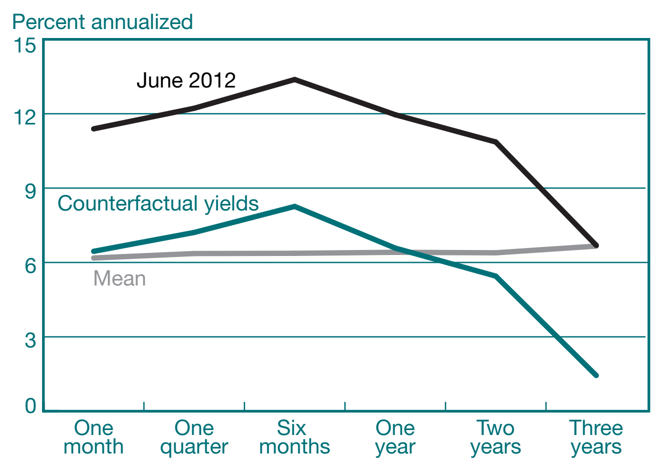

Notes: The gray line (labeled "Mean") shows the mean term structure of the equity risk premium over the full sample, January 1960 to June 2013. The black line (labeled "June 2012") shows the term structure for the most recent peak in the one-month-ahead ERP. These two lines are the same as in Chart 4. The green line (labeled "Counterfactual yields") shows what the term structure of equity risk premia would be in June 2012 if, instead of subtracting June 2012's yield curve from expected returns, we subtracted the average yield curve for January 1960 to June 2013.

7. Conclusion

In this article, we analyze twenty different models of the equity risk premium by considering the assumptions and data required to implement them, and how the models relate to one another. When it comes to the ERP, we find that there is substantial heterogeneity in estimation methodology and final estimates. We then extract the first principal component of the twenty models, which signals that the ERP in 2012 and 2013 is at heightened levels compared with previous periods. Our analysis provides evidence that the current level of the ERP is consistent with a bond-driven ERP: expected excess stock returns are elevated not because stocks are expected to have high returns but because bond yields are exceptionally low. The models we consider suggest that expected stock returns, on their own, are close to average levels.

Appendix A

Data Variables

Fama and French (1992)	http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html Monthly frequency; 1/1/1960 to 6/30/2013. We use twenty-five portfolios sorted on size and book-to-market, ten portfolios sorted on momentum, realized excess market returns, HML, SMB, and the momentum factor.
Shiller (2005)	http://www.econ.yale.edu/~shiller/data.htm Monthly frequency; 1/1/1960 to 6/30/2013. We use the nominal and real price, nominal and real dividends, and nominal and real earnings for the S&P 500, CPI, and ten-year nominal Treasury yield.
Baker and Wurgler (2007)	http://people.stern.nyu.edu/jwurgler/data/Investor_Sentiment_Data_v23_POST.xlsx Monthly frequency; 7/1/1965 to 12/1/2010. We use the "sentiment measure."
Graham and Harvey (2012)	http://www.cfosurvey.org/index.html Quarterly frequency; 6/6/2000 to 6/5/2013. We use the answer to the question "Over the next ten years, I expect the average annual S&P 500 return will be: expected return:" and the analogous question that asks about the next year.
Damodaran (2012)	http://www.stern.nyu.edu/~adamodar/pc/datasets/histimpl.xls Annual frequency; 1/1/1960 to 12/1/2012. We use the ERP estimates from his dividend discount models (one uses free cash flow, the other one does not).
Gurkaynak, Sack, and Wright (2007)	http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html Daily frequency; starting on 6/14/61 for one- to seven-year yields, 8/16/71 for nine- and ten-year yields, 11/15/71 for eleven- to fifteen-year yields, 7/2/81 for sixteen- to twenty-year yields, 11/25/85 for twenty-one- to thirty-year yields. We use all series until 6/30/2013.
Gurkaynak, Refet, Sack, and Wright (2010)	http://www.federalreserve.gov/econresdata/researchdata.htm Monthly frequency; 1/1/1960 to 7/1/2013 for Baa minus Aaa bond yield spread and recession indicator.
Compustat	Book value per share (variable BKVLPS) Annual frequency; 12/31/1977 to 12/31/2012.
Thomson Reuters I/B/E/S	Earnings per share (variables EPS 1 2 3 4 5) Monthly frequency; 1/14/1982 to 4/18/2013 for current and next-year forecasts, 9/20/84 to 4/18/2013 for two-year-ahead forecasts, 9/19/85 to 3/15/2012 for three-year-ahead forecasts, 2/18/88 to 3/15/07 for four-year-ahead forecasts.
FRED (Federal Reserve Bank of St. Louis)	http://research.stlouisfed.org/fred2/graph/?g=D9J and http://research.stlouisfed.org/fred2/graph/?g=KKk Monthly frequency; 1/1/1960 to 7/1/2013 for Baa minus Aaa bond yield spread and recession indicator.

References

Adrian, T., D. M. Covitz, and N. Liang. 2013. "Financial Stability Monitoring." Federal Reserve Bank of New York Staff Reports, no. 601, February.

Adrian, T., R. Crump, and E. Moench. 2014. "Regression-Based Estimation of Dynamic Asset Pricing Models." Federal Reserve Bank of New York Staff Reports, no. 493, June.

Ait-Sahalia, Y., M. Karaman, and L. Mancini. 2014. "The Term Structure of Variance Swaps, Risk Premia, and the Expectation Hypothesis." Available at SSRN: http://ssrn.com/abstract=2136820.

Ang, A., and G. Bekaert. 2007. "Stock Return Predictability: Is It There?" Review of Financial Studies 20, no. 3 (May): 651-707.

Ang, A., and M. Ulrich. 2012. "Nominal Bonds, Real Bonds, and Equity." Unpublished paper, Columbia University.

Baker, M., and J. Wurgler. 2007. "Investor Sentiment in the Stock Market." Journal of Economic Perspectives 21, no. 2 (Spring): 129-51.

Binsbergen, J. V., W. Hueskes, R. Koijen, and E. Vrugt. 2013. "Equity Yields." Journal of Financial Economics 110, no. 3 (December): 503-19.

Bloom, N. 2009. "The Impact of Uncertainty Shocks." Econometrica 77, no. 3 (May): 623-85.

Boguth, O., M. Carlson, A. J. Fisher, and M. Simutin. 2012. "Leverage and the Limits of Arbitrage Pricing: Implications for Dividend Strips and the Term Structure of Equity Risk Premia." Unpublished paper, Arizona State University.

Campbell, J. Y., and S. Thompson. 2008. "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" Review of Financial Studies 21, no. 4 (July): 1509-31.

Carhart, M. M. 1997. "On Persistence in Mutual Fund Performance." Journal of Finance 52, no 1 (March): 57–82.

Chen, L., Z. Da, and X. Zhao. 2013. "What Drives Stock Price Movements?" Review of Financial Studies 26, no. 4 (April): 841-76.

Clemen, R. T. 1989. "Combining Forecasts: A Review and Annotated Bibliography." International Journal of Forecasting 5, no. 4: 559-83.

Croce, M., M. Lettau, and S. Ludvigson. 2015. "Investor Information, Long-Run Risk, and the Term Structure of Equity." Review of Financial Studies 28, no. 3: 706-74.

Damodaran, A. 2012. "Equity Risk Premiums (ERP): Determinants, Estimation, and Implications—The 2012 Edition." Unpublished paper, New York University.

Diebold, F. X., and J. Lopez. 1996. "Forecast Evaluation and Combination." In G. S. Maddala and C. R. Rao, eds., Handbook of Statistics, 241-68. Amsterdam: North-Holland.

Duarte, F. 2013. "Inflation Risk and the Cross Section of Stock Returns." Federal Reserve Bank of New York Staff Reports, no. 621, May.

Duarte, F., L. Kogan, and D. Livdan. 2013. "Aggregate Investment and Stock Returns." Unpublished paper, M.I.T.

Durham, J. B. 2013. "Arbitrage-Free Models of Stocks and Bonds." Federal Reserve Bank of New York Staff Reports, no. 656, December.

Easton, P. D., and G. A. Sommers. 2007. "Effect of Analysts' Optimism on Estimates of the Expected Return Implied by Earnings Forecasts." Journal of Accounting Research 45, no. 5 (December): 983-1016.

Fama, E. F., and K. R. French. 1988. "Dividend Yields and Expected Stock Returns." Journal of Financial Economics 22, no. 1 (October): 3-25.

———. 1992. "The Cross Section of Expected Stock Returns." Journal of Finance 47, no. 2 (June): 427-65.

———. 2002. "The Equity Premium." Journal of Finance 57, no. 2 (April): 637-59.

Geiger, B., and G. Kubin. 2013. "Signal Enhancement as Minimization of Relevant Information Loss." Proceedings of ITG Conference on Systems, Communication, and Coding.

Gordon, M. J. 1962. The Investment, Financing, and Valuation of the Corporation. Homewood, Illinois: Greenwood Press.

Goyal, A., and I. Welch. 2008. "A Comprehensive Look at the Empirical Performance of Equity Premium Predictions." Review of Financial Studies 21, no. 4 (July): 1455-508.

Graham, J., and C. Harvey. 2012. "The Equity Risk Premium in 2012." Unpublished paper, Duke University.

Greenwood, R., and A. Shleifer. 2014. "Expectations of Returns and Expected Returns." Review of Financial Studies 27, no. 3 (March): 714-46.

Gurkaynak, R. S., S. Refet, B. Sack, and J. H. Wright. 2010. "The TIPS Yield Curve and Inflation Compensation." American Economic Journal: Macroeconomics 2, no. 1 (January): 70-92.

Gurkaynak, R. S., B. Sack, and J. H. Wright. 2007. "The U.S. Treasury Yield Curve: 1961 to the Present." Journal of Monetary Economics 54, no. 8 (November): 2291-304.

Hall, R. E., 2014. "The Routes Into and Out of the Zero Lower Bound." Proceedings of the 2013 Jackson Hole Economic Policy Symposium, Federal Reserve Bank of Kansas City.

Harvey, C. R., Y. Liu, and H. Zhu. 2014. ". . . and the Cross Section of Expected Returns." NBER Working Paper no. 20592, October. Available at SSRN: http://ssrn.com/abstract=2249314

Kelly, B., and S. Pruitt. 2013. "Market Expectations in the Cross Section of Present Values." Journal of Finance 68, no. 5 (October): 1721-56.

Kuehn, L. A., N. Petrosky-Nadeau, and L. Zhang. 2012. "An Equilibrium Asset Pricing Model with Labor Market Search." NBER Working Paper no. 17742, January.

Lemke, W., and T. Werner. 2009. "The Term Structure of Equity Premia in an Affine Arbitrage-Free Model of Bond and Stock Market Dynamics." European Central Bank Working Paper no. 1045, April.

Lettau, M., and J. A. Wachter. 2011. "The Term Structures of Equity and Interest Rates." Journal of Financial Economics 101, no. 1 (July): 90-113.

Merton, R. C. 1973. "An Intertemporal Capital Asset Pricing Model." Econometrica 41, no. 5 (September): 867-87.

Muir, T. 2013. "Financial Crises, Risk Premia, and the Term Structure of Risky Assets." Unpublished paper, Yale School of Management.

Neely, C. J., D. E. Rapach, J. Tu, and G. Zhou. 2014. "Forecasting the Equity Risk Premium: The Role of Technical Indicators." Management Science 60, no. 7 (July): 1772-91.

Panigirtzoglou, N., and J. Loeys. 2005. "A Fair-Value Model for U.S. Bonds, Credit and Equities." JP Morgan Investment Strategies Series, no. 11 (January).

Piazzesi, M., and M. Schneider. 2007. "Equilibrium Yield Curves." In NBER Macroeconomics Annual 2006 21: 389-472. Cambridge, Mass.: MIT Press.

Polk, C., S. Thompson, and T. Vuolteenaho. 2006. "Cross-Sectional Forecasts of the Equity Premium." Journal of Financial Economics 81, no. 1 (July): 101-41.

Ross, S. 1976. "The Arbitrage Theory of Capital Asset Pricing." Journal of Economic Theory 13, no. 3 (December): 341–60.

Shiller, R. J. 2005. *Irrational Exuberance.* 2nd ed. Princeton: Princeton University Press.

Stock, J. H., and M. W. Watson. 2003. "Forecasting Output and Inflation: The Role of Asset Prices." Journal of Economic Literature 41, no. 3 (September): 788-829.

Timmermann, A. 2006. "Forecast Combinations." In G. Elliott, C. W. J. Granger, and A. Timmermann, eds., Handbook of Economic Forecasting, vol. 1, 135-96. Amsterdam: Elsevier.

Federal Reserve Bank of New York
Economic Policy Review

Library of Congress Card Number: 77-646559