NBER WORKING PAPER SERIES

ENTRY AND EXIT IN OTC DERIVATIVES MARKETS

Andrew G. Atkeson
Andrea L. Eisfeldt
Pierre-Olivier Weill

Working Paper 20416
http://www.nber.org/papers/w20416

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2014

We'd like to thank, for fruitful comments and suggestions, the Editor, Lars Hansen, three referees, as well as Gara Afonso, Saki Bigio, Briana Chang, Jean-Edouard Colliard, Darrell Duffie, Peter Kondor, Ben Lester, Gustavo Manso, Tyler Muir, Martin Oehmke, Ionid Rosu, Tano Santos, Daniel Sanches, Martin Schneider, Shouyong Shi, Randy Wright, Pierre Yared, and numerous seminar audiences. Patrick Kiefer, Pierre Mabille, Omair Syed, Alfredo Reyes, and Semih Uslu provided expert research assistance. We thank the Bank of France, the Fink Center for Finance and Investments, and the Ziman Center for Real Estate for financial support. We thank Liz Galvin for editorial assistance. Part of this project was completed when Pierre-Olivier Weill was a visiting Professor at the Paris School of Economics, whose hospitality is gratefully acknowledged. All errors are ours. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w20416.ack

Entry and Exit in OTC Derivatives Markets
Andrew G. Atkeson, Andrea L. Eisfeldt, and Pierre-Olivier Weill
NBER Working Paper No. 20416
August 2014
JEL No. G0,G1,G2

ABSTRACT

We develop a parsimonious model to study the equilibrium and socially optimal decisions of banks to enter, trade in, and possibly exit, an OTC market. Although we endow all banks with the same trading technology, banks' optimal entry and trading decisions endogenously lead to a realistic market structure comprised of dealers and customers with distinct trading patterns. We decompose banks' entry incentives into incentives to hedge risk and incentives to make intermediation profits. We show that dealer banks enter more than is socially optimal. In the face of large negative shocks, they may also exit more than is socially optimal when markets are not perfectly resilient.

Andrew G. Atkeson
Bunche Hall 9381
Department of Economics
UCLA
Box 951477
Los Angeles, CA 90095-1477
and NBER
andy@atkeson.net

Andrea L. Eisfeldt
UCLA Anderson School of Management
110 Westwood Plaza
Suite C4.10
Los Angeles, CA 90095
and NBER
andrea.eisfeldt@anderson.ucla.edu

Pierre-Olivier Weill
Department of Economics
University of California, Los Angeles
Bunche Hall 8283
Los Angeles, CA 90095
and NBER
poweill@econ.ucla.edu

A data appendix is available at http://www.nber.org/data-appendix/w20416

1 Introduction

We develop a parsimonious model to study the equilibrium and socially optimal decisions of banks to enter and trade in an over-the-counter (OTC) market. Banks differ in terms of their exposure to an aggregate risk factor, their size, and their entry cost, but otherwise are endowed with the same OTC trading technology. In an entry equilibrium, banks' optimal participation decisions determine the structure of the OTC market endogenously: the characteristics of participants, the heterogeneity in their portfolios and trading patterns, the dispersion in their marginal valuations and transaction prices. We argue that equilibrium outcomes reproduce stylized facts about the structure of OTC markets. In particular, large-sized banks endogenously emerge as "dealers" who profit from price dispersion, and in doing so provide intermediation services to middle-sized "customer" banks. We then formalize and explicitly characterize banks' entry incentives, in equilibrium vs. in the corresponding constrained planning problem. This allows us to address policy questions about market size and composition. We show that a bank entering the market as a dealer adds social value, and we formalize the manner in which it mitigates OTC market frictions by facilitating trade among customer banks with dispersed marginal valuations. However, we also show that, in an entry equilibrium, the trading profits of the marginal dealer exceed its marginal social contribution. As a result, dealer banks tend to provide too many intermediation services relative to the social optimum.

We extend our equilibrium concept to study exit. In an exit equilibrium, banks face shocks to their cost of ongoing participation in the OTC market, and they make optimal decisions to stay or exit the market. Crucially, we assume that the market is imperfectly resilient: a bank who has lost some of its trading counterparties due to exit may or may not be able to trade with new ones. In this context, we show that dealers are the most vulnerable to negative shocks: they have the strongest private incentives to exit. However, we find that, relative to the social optimum, dealers exit too much only if the shock is large enough and if the market is not perfectly resilient.

Our model is populated by a continuum of financial institutions, called banks, who contemplate entering an OTC market. Each bank is a coalition of many risk-averse agents, called traders. Banks have heterogeneous sizes and heterogeneous endowments of a non-tradable risky loan portfolio, creating heterogeneous exposures to an aggregate default risk factor. Since traders in banks are risk averse, they will attempt to equalize these exposures by trading derivatives "swap" contracts in the OTC market. First, conditional on their size and initial exposure to aggregate default risk, banks choose whether to pay a fixed cost in order to enter into the market. The fixed cost payment represents the acquisition of infrastructure and specialized expertise required to trade in OTC markets. Second, after entry, all banks are granted access to the same technology to trade swaps. Their traders are paired uniformly, and each pair negotiates over the terms of the contract subject to a uniform trade size limit. This limit represents banks' risk-management constraints on individual trading desk positions in practice. Third, each bank consolidates the swaps signed by its traders and all contracts and loans are paid off.

Banks have two distinct private incentives to enter the OTC market. The first incentive is to hedge their underlying risk exposure. The second incentive arises because, in equilibrium, hedging is imperfect. Imperfect hedging creates dispersion in banks' marginal valuations and price dispersion. This dispersion gives banks incentives to enter in order to earn trading profits, and in doing so provide intermediation services. We show that both incentives are U-shaped functions of a bank's initial risk exposure: they are larger for banks with extreme initial risk exposures, either small or large, and smaller for banks with intermediate exposures. Combined with fixed entry costs, these U-shaped incentives result in entry and trading patterns which are corroborated by empirical evidence. First, small-sized banks cannot spread the fixed entry cost over enough traders, and choose not to enter. Second, medium-sized banks only find it optimal to enter the market if their incentives are large enough, which occurs if their initial risk exposure is sufficiently small or large. They use the OTC market to take a large net position, either long or short, and in this sense act as customers. Third, banks with intermediate exposures have the smallest incentive to enter and so they only pay the fixed cost if they are large enough. Since they start with an intermediate exposure, close to the market-wide average, they do not trade to hedge. Instead, they take many offsetting long and short positions, have large gross exposures, small net exposures, and enter mostly to make intermediation profits. Hence, these large banks endogenously emerge as dealers.

Next, we study the problem of a planner who chooses banks' entry and trading patterns in the OTC market, but is otherwise subject to the same frictions as in the equilibrium. We find that trading patterns are socially optimal conditional on entry, but that entry patterns are not. We show that dealers are socially useful because they facilitate trade between banks with dispersed marginal valuations. However, we find that their profits are larger than their marginal social contribution. Hence, dealers have too great an incentive to enter in equilibrium. In a parametric example, we show that, starting from the equilibrium, a social planner finds it optimal to decrease the entry of dealers and increase the entry of customers. The resulting socially optimal OTC market structure has fewer participants, generates less trading volume, and creates smaller ratios of gross-to-net exposures. Therefore, according to the model, there is a policy role for taxing dealer banks in order to reduce some of the trading volume generated by intermediation activity, while subsidizing the participation of customer banks in order to promote direct customer-to-customer transactions.

In the last part of the paper, we extend our framework to study exit. We assume that entry has already occurred and that banks are faced with an unexpected negative shock: they must incur a cost to continue actively trading in the OTC market, or they must exit. We define market resilience to be the likelihood that traders who lose a counterparty due to exit find a new counterparty with whom to resume trading. In the case of perfect resilience, traders immediately re-match, and the model of exit is equivalent to the entry model. In the case of no resilience, banks that lose a counterparty cannot re-match, as is often assumed. Imperfect resilience then captures the state of the market in the short run, when banks scramble to replace their lost counterparties, some successfully and some not. In our analysis, we allow for variation in market resilience, and describe its crucial role in determining exit outcomes and policy prescriptions. Our main findings are as follows: first, as in the entry model, incentives to participate are U-shaped. This implies that banks who have intermediate initial exposures, which we know from our earlier analysis tend to be large and to act as dealers, have the strongest incentives to exit. Second, we find that, depending on the size of the shock and on market resilience, dealers may exit too little or too much. This is because of two effects going in opposite directions. On the one hand, there can be too little exit because dealers appropriate too much surplus in their bilateral trades relative to the social surplus they create. Just as in the entry model, this gives them too much incentive to stay. On the other hand, and in contrast to the entry model, there is a new effect that can lead to excessive exit: a bank does not internalize that when it exits the market, it lowers the chance of trading for the banks who choose to stay. In a parametric example, we show that either effect can dominate. In particular, when the negative shock is large enough and when the market is not too resilient, then dealer banks exit too much. Thus, according to the model, there is a policy role for subsidizing dealer banks in the short run, during severe financial disruptions.

Related Literature

Our main contribution relative to the literature on OTC markets is to develop a model that is sufficiently tractable to analyze endogenous entry and exit, explain empirical patterns of participation across banks of different sizes, and address normative issues regarding the size, composition and resilience of the market.

One of the most commonly employed frictions used to study OTC markets, following Duffie, Gârleanu, and Pedersen (2005), is the search friction.[1] Kiefer (2010) offers an early

[1] Duffie, Gârleanu, and Pedersen (2002) apply the paradigm to markets for borrowing stocks. Vayanos

analysis of CDS pricing within this framework. Our paper is most closely related to the earlier work of Afonso and Lagos (2012), who focus on high-frequency trading dynamics in the Federal Funds Market. They take entry decisions as given and find that some banks emerge endogenously as intermediaries in the process of reallocating reserves balances over the course of the day. Our main contribution relative to this paper and to this literature more generally is to study the positive and normative implications of entry and exit in OTC markets. To do so, we develop a new model, in the spirit of Shi (1997), in which all trades occur statically in one single multilateral trading session. Our approach preserves the key insight of dynamic models while being much more tractable, providing analytical characterizations of banks' equilibrium and socially optimal entry and exit decisions.

The effects of the trading structure on trading outcomes has been studied in the literature on systemic risk.[2] Allen and Gale (2000) develop a theory of contagion in a circular system, which they use to consider systemic risk in interbank lending markets. This framework has been employed by Zawadowski (2013) to consider counterparty risk in OTC markets. Eisenberg and Noe (2001) also study systemic risk, but use lattice theory to consider the fragility of a financial system in which liabilities are taken as given (see also the recent work of Elliott, Golub, and Jackson, 2014). In addition to proposing a new modeling framework, our work on exit differs from these papers by allowing the OTC market to be imperfectly resilient, i.e., banks may be able to resume trading with others even if their orginal counterparties chose to exit the market. According to our model, accounting for imperfect resilience is crucial to assess the social value of policy intervention: subsidizing dealers is warranted only if the market is not sufficiently resilient.

The paper proceeds as follows. Section 2 presents the economic environment. Section 3 solves for equilibrium trading and entry patterns. Section 4 studies the normative implications of our model. Section 5 extends the model to study exit. Finally, Section 6 provides

and Wang (2007), Weill (2008), and Vayanos and Weill (2008) study cross-sectional asset pricing in multi-asset extensions of the original model. Lagos and Rocheteau (2009) and Gârleanu (2009) demonstrate that relaxing the asset holding constraint can have important impact on asset prices and on the entry decisions of intermediaries. Afonso (2011) studies the externalities involved in investors' entry decisions. Weill (2007) and Lagos, Rocheteau, and Weill (2011) study the liquidity provision of dealers in response to liquidity shocks. Biais, Hombert, and Weill (2014) use this paradigm to analyze equilibrium in limit order markets. Duffie and Strulovici (2012) study limited capital mobility between markets and its impact on asset prices. Pagnotta and Philippon (2011) analyze exchanges competing to offer trading speed. He and Milbradt (2014) study the interaction between default risk and OTC market liquidity. Information diffusion is addressed by Duffie and Manso (2007), Golosov, Lorenzoni, and Tsyvinski (2014); Duffie, Malamud, and Manso (2009); and recently by Babus and Kondor (2013) using different techniques. Adverse selection is studied by Guerrieri and Shimer (2012) and Chang (2012). Trejos and Wright (2012) offer a unified analysis asset pricing in finance and monetary models of bilateral trades. Bolton, Santos, and Scheinkman (2012) show how OTC markets may inefficiently "cream-skim" assets from organized exchange. Gavazza (2013) estimates an applied version of this model and calculates the welfare role of dealers at estimated parameters.

[2]See Stulz (2011) for a discussion of the potential for systemic risk in CDS markets.

additional results, including comparative statics on entry, exit, and market structure with respect to changes in trading frictions in the context of a parametric example, and Section 7 concludes. Proofs are gathered in the online appendix.

2 The economic environment

This section presents the economic environment.

2.1 The agents

The economy is populated by a unit continuum of risk-averse agents, called traders. Traders have utility functions with identical constant absolute risk aversion and they are endowed with a technology to make payments by producing a storable consumption good at unit marginal cost.[3] To model the financial system, we take a novel approach in the literature: we assume that traders are organized into a continuum of large coalitions called banks. Banks are heterogeneous in several dimensions: they differ in their sizes, in their fixed costs of entry in the OTC market, and in their risk-sharing need.

Size and fixed entry costs. We identify the size of a bank with the measure of traders in the coalition, which we denote by S. The fixed cost of entry in the OTC market is denoted by c. Taken together, the distribution of bank sizes and the fixed cost induce a distribution of per-trader entry costs, c/S, in the population of traders. We represent this distribution by the right-continuous and increasing function $\Phi(z)$: the cumulative measure of traders in banks with per-trader entry costs less than z. We assume that this distribution has compact support but otherwise place no further restrictions: the distribution can be continuous, discrete, or a mixture of both.

As will become clear shortly, conditional on its hedging need, a bank's entry decisions will only depend on its per-trader entry cost, c/S. Therefore, from a theoretical perspective, the exact nature of the cost does not matter for our results: we could have alternatively assumed that banks have heterogeneous variable costs with the same distribution $\Phi(z)$. From an empirical perspective, however, fixed costs matter: together with the entry incentives generated by the OTC market, they explain empirical evidence about entry and trading patterns in a cross section of banks sorted by size, an easily observable bank characteristic.[4]

[3]Precisely, if an agent consumes C and produces H, his utility is $U\{C-H\} = -\frac{1}{\eta}e^{-\eta(C-H)}$, for some coefficient of absolute risk aversion $\eta > 0$. Given no wealth effect, an equivalent interpretation is that the agent has a large endowment of a storable consumption good that she uses to make payments.

[4]See Atkeson, Eisfeldt, and Weill (2012) for stylized facts about this cross-section.

Namely, aside from having the obvious consequence that small-sized banks do not participate, fixed costs will create an empirically realistic correlation between size and trading patterns amongst those banks who choose to enter in the market.

Risk-sharing needs. We assume that banks receive heterogeneous initial endowments of a risky asset, which gives them a need to share risk. Given our focus on the long-run structure of the OTC market, we interpret this risky asset endowment as the bank's typical portfolio of illiquid loans, arising from lending activity to households and corporations which we do not model explicitly here. Thus, in this model, an insurance company or a hedge fund would have a small endowment, and a commercial bank a large endowment. For each bank, we denote the per-trader endowment by ω, so that the bank-level endowment is $S \times \omega$. We assume that banks' per-trader endowments, ω, are positive, belong to some finite set, Ω, and are distributed independently from the per-trader entry cost, c/S. That is, the measure of traders in banks with per capita-endowment ω and entry cost less than c/S can be written as a product $\pi(\omega)\Phi(c/S)$, for some positive $\{\pi(\tilde{\omega})\}_{\tilde{\omega}\in\Omega}$ such that $\sum_{\tilde{\omega}} \pi(\tilde{\omega}) = 1$.[5]

In line with our loan portfolio interpretation, we denote the payoff of the asset by $1 - D$, where 1 represents the face value of a typical loan extended by the bank, and D represents its typical aggregate default risk. We assume that D is a (non-trivial) random variable with strictly positive mean and twice continuously differentiable moment generating function. Since D represents aggregate default risk, we assume that its realizations are identical for all banks and all assets.

Before proceeding to the analysis of the OTC market, let us briefly describe what would happen in this environment in the absence of frictions, if banks could trade their risky-asset endowments directly in a centralized market. Then, banks would be able to share their risk fully by equalizing their exposures to the aggregate factor D, they would all trade at the same price, and they would have no incentive to enter the market in order to engage in intermediation by executing a gross volume of trade in excess of their net trades. To depart from these counterfactual predictions, we now consider the OTC market with frictions.

[5]The assumption that size and endowments are independent clarifies the economic forces at play and, importantly, allows us to argue that for banks who choose to the participate in the OTC market, the correlation between size and per-capita endowment is purely endogenous. This being said, our model is flexible enough to handle more general joint distributions of size and per capita endowments. For example, in an earlier version of the paper, we provided a characterization of the post-entry equilibrium when larger banks have more neutral pre-trade exposures than smaller banks, for example through greater internal diversification. We also considered a continuous distribution of ω.

2.2 The OTC market

In our model, banks cannot trade their risky-asset endowments directly and frictionlessly. Instead, they can enter an OTC derivatives market to trade swap contracts, resembling CDS. The timing of entry and trade in the OTC market is as follows.

First, conditional on their size, S, and initial endowment, ω, each bank chooses whether to pay the fixed cost, c, to enter in the OTC market. Then, after entry decisions have been made, all banks that have entered the market are granted access to the same trading technology: their traders are paired uniformly to sign a swap contract subject to a uniform trade size limit. The pairing is uniform in the sense that it occurs in proportion to the distribution of traders present in the market across endowments $\omega \in \Omega$. The swap contract exchanges a fixed payment for a promise to make a payment equal to the realization of the aggregate default factor, D.[6] Finally, after trading, banks consolidate the positions of their traders and all payoffs from loan portfolios and swap contracts are realized.

The key friction shaping trading patterns and entry incentives is the trade size limit on bilateral trades,[7] which ultimately prevents participant banks from fully sharing their risk in an OTC market equilibrium. This leads to equilibrium price dispersion and hence creates incentives for banks to enter and actively engage in intermediation activity.

Moreover, as we shall see below, the trade size limit provides a simple and tractable way of parameterizing the extent to which traders from a single bank with endowment ω can, collectively, direct their trading volume to those counterparties from whom they get the best prices. We shall see that, when the trade size limit increases, trading patterns change. They look less random and more directed in the sense that there is less and less intermediation activity, and gross exposures converge to net exposures.

3 Equilibrium definition and existence

We study an equilibrium in two steps. First, we study an OTC market equilibrium conditional on the distribution of traders in the market arising from banks' entry decisions. We

[6] Our analysis applies more generally to OTC trading of derivatives contracts with payouts indexed to aggregate risk factors. This includes, for example, interest rate swaps, CDS of sovereign entities, CDS indices and, to the extent that default risk is correlated across firms, CDS on single firms.

[7] We do not model the microfoundations of this limit; however, we note that, in practice, traders typically do face line limits. For example, Saita (2007) states that the traditional way to prevent excessive risk taking in a bank "has always been (apart from direct supervision...) to set notional limits, i.e., limits to the size of the positions which each desk may take." In addition, measures such as DV01 or CS1% which measure positions' sensitivities to yield and credit spread changes, as well as risk weighted asset charges, are used to gauge and limit the positions of a particular desk's traders. Theoretically, one might motivate such limits as stemming from moral hazard problems, concerns about counterparty risk and allocation of scarce collateral, or capital requirement considerations.

establish existence and uniqueness of this equilibrium, and show that it is socially optimal conditional on banks' entry decisions. Second, we present the fixed point problem that defines an equilibrium in which banks' entry decisions are chosen optimally. We establish the existence of an equilibrium with positive entry by proving that this fixed point-problem has a non-zero solution.

3.1 OTC market equilibrium conditional on entry

Suppose that banks have made their decisions to enter the OTC market, and let $\mu = \{\mu(\omega)\}_{\omega\in\Omega}$ denote the measures of traders aggregated across banks with per capital endowment ω in the OTC market.[8]

3.1.1 Payoffs

If there is positive entry, $\sum_\omega \mu(\omega) > 0$, then each trader present in the OTC market is paired with a trader from another bank to bargain over a CDS contract. The pairwise matching of traders from different banks is uniform. Thus, for any individual trader from a bank with any given per capita endowment ω, the probability of being paired with a trader from a bank whose per-capita endowment is $\omega \in \Omega$ is

$$n(\omega) \equiv \frac{\mu(\omega)}{\sum_{\tilde{\omega}} \mu(\tilde{\omega})}, \tag{1}$$

the fraction of such traders in the OTC market. We denote the associated cumulative distribution by $N(\omega) = \sum_{\tilde{\omega}\leq\omega} n(\tilde{\omega})$, and its support by $\text{supp}(N)$. The successor of $\omega \in \Omega$ in the support of N is $\omega^+ = \min\{\tilde{\omega} \in \text{supp}(N) : \tilde{\omega} > \omega\}$, with the convention that $\omega^+ = \infty$ and $N(\infty) = 1$ if this set is empty. Similarly, ω^- is the predecessor of $\omega \in \Omega$ in the support of N.

Bilateral exposures. When a trader from a bank of type ω is paired with a trader from a bank of type $\tilde{\omega}$, they bargain over the terms of a derivative contract resembling a CDS. The ω-trader sells $\gamma(\omega, \tilde{\omega})$ contracts to the $\tilde{\omega}$-trader, whereby she promises to make the random payment $\gamma(\omega, \tilde{\omega})D$ at the end of the period, in exchange for the fixed payment $\gamma(\omega, \tilde{\omega})R(\omega, \tilde{\omega})$. If $\gamma(\omega, \tilde{\omega}) > 0$ then the ω-trader sells insurance, and if $\gamma(\omega, \tilde{\omega}) < 0$ she buys insurance. As explained before, traders face a trade size limit: in any bilateral meeting, they cannot sign

[8] The distribution μ is the only relevant aggregate state variable conditional on entry because our model has a natural homogeneity property: in equilibrium, banks' trading and entry incentives only depend on ω. See Footnote 11 and Lemma 5.

more than a fixed amount of contracts, k, either long or short. Taken together, the collection of CDS contracts $\gamma = \{\gamma(\omega, \tilde{\omega})\}_{(\omega,\tilde{\omega})\in\Omega^2}$ must therefore satisfy:

$$\gamma(\omega, \tilde{\omega}) + \gamma(\tilde{\omega}, \omega) = 0 \tag{2}$$

$$-k \leq \gamma(\omega, \tilde{\omega}) \leq k, \tag{3}$$

for all $(\omega, \tilde{\omega}) \in \Omega^2$. Equation (2) is a bilateral feasibility constraint, and equation (3) is the constraint imposed by the trade size limit.

Certainty equivalent payoff. We assume that at the end of the period, traders of bank ω get back together to consolidate all of their long and short CDS positions. This captures a realistic feature of banks in practice: within a bank, some traders will go long and some short, depending on whom they trade with. After consolidation, the per capita consumption of traders in an active bank with per capita endowment ω, entry cost c, and size S is:

$$-\frac{c}{S} + \omega(1-D) + \sum_{\tilde{\omega}} \gamma(\omega, \tilde{\omega})\big[R(\omega, \tilde{\omega}) - D\big]n(\tilde{\omega}), \tag{4}$$

by the law of large numbers. The first term is the per capita entry cost. The second term is the per capita payout of the risky asset endowment. The third term is the per capita consolidated amount of fixed payments, $\gamma(\omega, \tilde{\omega})R(\omega, \tilde{\omega})$, and random payments, $\gamma(\omega, \tilde{\omega})D$, on the portfolio of contracts signed by all ω-traders with their counterparties from banks with endowment $\tilde{\omega}$.

Calculating expected utility, we obtain that the certainty equivalent of (4) is:

$$-\frac{c}{S} + \omega + \sum_{\tilde{\omega}} \gamma(\omega, \tilde{\omega})R(\omega, \tilde{\omega})n(\tilde{\omega}) - \Gamma\left[g(\omega)\right]. \tag{5}$$

The first terms of this certainty equivalent add up the non-stochastic components of (4): the per-capita entry cost, c/S, the face value of the bank's endowment of risky loans, ω, and the sum of all CDS fixed payments, $\gamma(\omega, \tilde{\omega})R(\omega, \tilde{\omega})$. The last term of this certainty equivalent, $-\Gamma\left[g(\omega)\right]$, represents the bank's cost of bearing default risk. Precisely, the stochastic component of (4) is $-g(\omega) \times D$, where

$$g(\omega) \equiv \omega + \sum_{\tilde{\omega}} \gamma(\omega, \tilde{\omega})n(\tilde{\omega}), \tag{6}$$

is the banks' *post-trade exposure* to default risk. It is the sum of the initial exposure, ω, and of all the additional exposures acquired in bilateral trades, $\gamma(\omega, \tilde{\omega})$. Thus, the term $-\Gamma\left[g(\omega)\right]$

in (5) transforms this post-trade exposure to default risk into the certainty equivalent cost of bearing it. Under our assumption that traders have CARA utility:

$$\Gamma\left[g(\omega)\right] \equiv \frac{1}{\eta} \log\left(\mathbb{E}\left[e^{\eta g(\omega) D}\right]\right).$$

One easily shows (see Appendix A.1) that the cost of risk bearing function, $g \mapsto \Gamma[g]$, is twice continuously differentiable, strictly increasing for $g \geq 0$, and strictly convex. Finally, we note that, when D is normally distributed, then $\Gamma[g]$ is a familiar quadratic loss function: $\Gamma[g] = g\mathbb{E}[D] + g^2 \frac{\eta \mathbb{V}[D]}{2}$. The first term is the expected loss due to default risk, $g\mathbb{E}[D]$. The second term is an additional cost arising because banks are risk averse and the loss is stochastic.

3.1.2 Bargaining

Having derived banks' payoffs, we are in a position to discuss how terms of trade are determined via bargaining in the OTC market. Our approach follows the literature which allows risk sharing within families, such as in Lucas (1990), Andolfatto (1996), Shi (1997), Shimer (2010), and others, and assumes that a trader's objective is to maximize the marginal impact of her decision on her bank's utility. Namely, we assume that, when a pair of $(\omega, \tilde{\omega})$ traders bargain over the terms of trade, they take the trading surplus to be:

$$\gamma(\omega, \tilde{\omega})\Bigg(\Gamma'\left[g(\tilde{\omega})\right] - \Gamma'\left[g(\omega)\right]\Bigg). \tag{7}$$

The expression is intuitive. Suppose the ω-trader sells $\gamma(\omega, \tilde{\omega})$ contracts to the $\tilde{\omega}$-trader. Since each trader is small relative to her bank, he or she only has a marginal impact on the cost of risk bearing. Precisely, the cost of risk bearing of bank ω, the seller of insurance, increases by $\gamma(\omega, \tilde{\omega})\Gamma'\left[g(\omega)\right]$, while the cost of risk bearing of bank $\tilde{\omega}$, the buyer of insurance, decreases by $\gamma(\omega, \tilde{\omega})\Gamma'\left[g(\tilde{\omega})\right]$. Hence, the trading surplus (7) measures the net change in the two banks' cost of risk bearing: the number of contracts sold multiplied by the difference between the marginal value of the buyer and the marginal cost of the seller.[9]

[9] One can provide more precise microfoundations for this surplus formula at the cost of lengthening the exposition. For instance, in a previous version of this paper, we assumed that each trader maximizes her expected trading profit, appropriately discounted by the marginal utility of other traders in her bank coalition. Another microfoundation is to assume that a trader maximizes her expected utility and can trade frictionlessly with other traders in her bank coalition. As a result, there is full risk sharing within the coalition, and the CDS price is $\Gamma'\left[g(\omega)\right]$ within bank ω, and $\Gamma'\left[g(\tilde{\omega})\right]$ within bank $\tilde{\omega}$. The trading surplus reduces to (7): indeed every CDS sold by trader ω in the OTC market can be bought later from other traders in her bank coalition at price $\Gamma'\left[g(\omega)\right]$ and, vice versa, every CDS purchased by trader $\tilde{\omega}$ can be sold later to other traders in her bank coalition at price $\Gamma'\left[g(\tilde{\omega})\right]$.

We assume that the terms of trade in a bilateral match between an ω-trader and an $\tilde{\omega}$-trader are determined via symmetric Nash bargaining. The first implication of Nash bargaining is that the terms of trade are bilaterally Pareto optimal, i.e, they must maximize the surplus shown above. Since the marginal cost of risk bearing, $\Gamma'[g]$, is strictly increasing, this immediately implies that:

$$\gamma(\omega, \tilde{\omega}) = \begin{cases} k & \text{if } g(\tilde{\omega}) > g(\omega) \\ [-k, k] & \text{if } g(\tilde{\omega}) = g(\omega) \\ -k & \text{if } g(\tilde{\omega}) < g(\omega). \end{cases} \tag{8}$$

This is intuitive: if the $\tilde{\omega}$-trader expects a larger post-trade exposure than the ω-trader, i.e. $g(\tilde{\omega}) > g(\omega)$, then the ω-trader sells insurance to the $\tilde{\omega}$-trader, up to the trade size limit. And vice versa if $g(\tilde{\omega}) < g(\omega)$. When the post-trade exposures are the same, then any trade in $[-k, k]$ is optimal.

The second implication of Nash bargaining is that the unit price of a CDS, $R(\omega, \tilde{\omega})$, is set so that each trader receives exactly one half of the surplus. This implies that:

$$R(\omega, \tilde{\omega}) = \frac{1}{2}\bigg(\Gamma'[g(\omega)] + \Gamma'[g(\tilde{\omega})]\bigg). \tag{9}$$

That is, the price is halfway between the two traders' marginal cost of risk bearing. As is standard in OTC market models, prices depend on traders' "infra-marginal" characteristics in each match, instead of depending on the characteristic of a single "marginal" trader, as would be the case in a Walrasian market.

3.1.3 OTC market equilibrium conditional on entry: definition and existence

Conditional on the distribution of traders, $n = \{n(\omega)\}_{\omega \in \Omega}$, generated by banks' entry decisions $\mu = \{\mu(\omega)\}_{\omega \in \Omega}$, an *OTC market equilibrium* is made up of CDS contracts, $\gamma = \{\gamma(\omega, \tilde{\omega})\}_{(\omega, \tilde{\omega}) \in \Omega^2}$, post-trade exposures, $g = \{g(\omega)\}_{\omega \in \Omega}$, and CDS prices, $R = \{R(\omega, \tilde{\omega})\}_{(\omega, \tilde{\omega}) \in \Omega^2}$, such that

(i) CDS contracts are bilaterally feasible: γ satisfies (2) and (3);

(ii) CDS contracts are optimal: γ and R satisfy (8) and (9) given g;

(iii) post-trade exposures are generated by CDS contracts: g satisfies (6) given γ.

In what follows, we will also seek to study the efficiency properties of the equilibrium. To that end, we consider the *planning problem conditional on entry*:

$$W(\mu) = \max_{\gamma} \sum_{\omega} \Big\{ - \big[\pi(\omega) - \mu(\omega)\big] \Gamma\left[\omega\right] - \mu(\omega) \Gamma\left[g(\omega)\right] \Big\}, \tag{10}$$

subject to (1), (2), (3), (6), and conditional on the entry decisions summarized by μ. In the planner's objective, the term associated with endowment ω is to be interpreted as follows. Given the assumed distribution of banks over the set Ω, there is a measure $\pi(\omega)$ of traders in banks with pre-trade exposures ω. Conditional on banks' entry decisions as summarized by $\mu = \{\mu(\omega)\}_{\omega \in \Omega}$, a measure $\pi(\omega) - \mu(\omega)$ of these traders are not in the OTC market and keep their exposure ω, incurring the cost of risk bearing $\Gamma\left[\omega\right]$. And a measure $\mu(\omega)$ trade in the OTC market and change their exposures to $g(\omega)$, incurring the cost of risk bearing $\Gamma\left[g(\omega)\right]$.

Given that certainty equivalents are quasi-linear, a collection of CDS contracts and post-trade exposures solve this planning problem if and only if it is Pareto optimal, in that it cannot be Pareto improved by choosing another feasible collection of CDS contracts and post-trade exposures and making deterministic transfers. With this in mind, we find:

Theorem 1. *There exists an OTC market equilibrium conditional on entry. All equilibria solve the planning problem conditional on entry. They all share the same post-trade risk exposures, g, and CDS prices, R. They may only differ in terms of bilateral exposures, γ.*

The theorem shows that all equilibrium objects are uniquely determined, except perhaps the bilateral exposures, γ. Indeed, when two traders with the same post-trade exposures are paired, they are indifferent regarding the sign and direction of the CDS contract they sign.

Note as well that post-trade exposures, g, are uniquely determined even for $\omega \notin \text{supp}(N)$. This is an important property to establish for the analysis of equilibrium entry. Indeed, we shall see that it unambiguously determines an individual bank's entry incentives even when no other bank of the same type enters the market.[10,11]

[10]Dealing with $\omega \in \text{supp}(N)$ creates a technical difficulty because we cannot fully characterize an equilibrium by merely comparing the first-order conditions of the planning problem with the equilibrium conditions. Indeed, the post-trade exposures of $\omega \notin \text{supp}(N)$ are not uniquely pinned down by the planning problem, since these traders are given zero weight in the planner's objective. Nevertheless, we can show that the equilibrium optimality condition (8) uniquely pins down $g(\omega)$ for all $\omega \in \Omega$, including $\omega \notin \text{supp}(N)$.

[11]As is clear from its proof, the theorem would also hold if we allowed bilateral exposures, γ, to depend on any arbitrary vector of bank-level characteristics, such as size. That is, in any equilibrium conditional on entry, banks' per-capita post-trade exposures are uniquely determined and only depend on their per-capita pre-trade exposures, ω.

3.1.4 Post-trade exposures, gross exposures, and net exposures

In this section we establish some elementary results about equilibrium trading patterns conditional on entry. As we shall see shortly, these trading patterns are crucial to understand the economic forces shaping entry incentives.

Our first result concerns post-trade exposures:

Proposition 2. *Suppose that $|supp(N)| \geq 2$. Then, post-trade exposures are increasing and closer together than pre-trade exposures:*

$$0 \leq g(\tilde{\omega}) - g(\omega) \leq \tilde{\omega} - \omega, \text{ for all } \omega < \tilde{\omega}. \tag{11}$$

Moreover, if $n(\omega) + n(\tilde{\omega}) > 0$, then $g(\tilde{\omega}) - g(\omega) < \tilde{\omega} - \omega$. Finally, there is a $\bar{k} > 0$ such that $g(\tilde{\omega}) = g(\omega)$ for all $(\omega, \tilde{\omega}) \in supp(N)^2$ if and only if $k \geq \bar{k}$.

The proposition shows that, as long as k is small enough, then there is partial risk sharing: $g(\tilde{\omega}) - g(\omega)$ is smaller than $\tilde{\omega} - \omega$, but in general remains larger than zero. The proposition also shows that full risk sharing obtains as long as k is large enough.

Appendix A.4 provides further results about post-trade exposures. In particular, we show that if $g(\omega)$ is strictly increasing at ω, then it must be equal to the post trade exposure that arises when traders in an ω bank sell insurance up to their trading limit k to all traders in banks with higher ω and buy insurance up to their trading limit k from all traders in banks with lower ω. This result thus implies that $g(\omega)$ is strictly increasing only when the density of traders in the neighborhood of ω is not too large. If the density of traders in the neighborhood of ω is large, then all traders in that neighborhood share risk locally by trading to a common post-trade exposure. This gives $g(\omega)$ a flat spot in that neighborhood.

Gross vs. net exposures in the cross-section. An important empirical observation in OTC credit derivative markets is that banks' gross exposures can dramatically differ from their net exposures. Banks with large ratios of gross to net exposures act as dealers: they simultaneously buy and sell many CDS contracts, but their long and short positions approximately offset each other. Banks with ratios of gross to net exposures close to one act as customers: they mostly trade in one direction, either long or short.

To see how differences in gross and net exposures arise in our environment, let us consider the gross number of contracts sold and purchased by a bank of type ω, per trader capita:

$$G^{+}(\omega) = \sum_{\tilde{\omega}} \max\{\gamma(\omega, \tilde{\omega}), 0\} n(\tilde{\omega}) \text{ and } G^{-}(\omega) = \sum_{\tilde{\omega}} \max\{-\gamma(\omega, \tilde{\omega}), 0\} n(\tilde{\omega}).$$

One measure of the extent to which gross exposures differ from net exposures is:

$$\min\{G^{+}(\omega), G^{-}(\omega)\}. \tag{12}$$

This represents a bank's intermediation volume: the number of contracts, per-trader capita, that fully offset each other within the bank's portfolio. We have the following proposition:

Proposition 3. *When supp$(N) \geq 3$ and k is small enough, intermediation volume, as defined in (12), is a hump-shaped function of $\tilde{\omega} \in$ supp(N), achieving its strictly positive maximum at, or next to, a median of N. That is, if $\omega \in$ supp(N) is maximum of $\min\{G^{+}(\tilde{\omega}), G^{-}(\tilde{\omega})\}$, then a median of N belongs to $\{\omega^{-}, \omega, \omega^{+}\}$.*

Thus, our model predicts that banks with intermediate pre-trade exposure, ω, will tend to assume the role of dealers.[12] This is intuitive: these banks do not need to change their exposure, since they start with one that is already close to the market-wide average. They can use all their trading capacity to provide intermediation services to others. Banks with extreme exposures assume the role of customers: those with low pre-trade exposures use their trading capacity to sell insurance, while those with high pre-trade exposures use it to purchase insurance.

Gross vs. net exposures in the aggregate. Proposition 3 focuses on small k because, in this case, bilateral exposures, γ, are uniquely determined in equilibrium. Indeed, post-trade exposures are strictly increasing and so the bilateral optimality condition, (8), implies that traders are never indifferent about the size and direction of their trade. For larger k, there may be some indeterminacy in bilateral exposures. As a result, the gross exposure of a bank with endowment ω, $G^{+}(\omega) + G^{-}(\omega)$, may be indeterminate as well. To resolve this indeterminacy, and obtain necessary conditions for gross exposures to exceed net exposures, we consider bilateral exposures that minimize average gross exposures:[13]

$$\mathcal{G}(k) = \min \sum_{\omega} \left[G^{+}(\omega) + G^{-}(\omega)\right] n(\omega), \tag{13}$$

with respect to bilateral exposures, γ, solving the planning problem conditional on entry and given k.[14]

[12] The Proposition can be shown to hold for larger k under additional parametric assumptions, see Atkeson, Eisfeldt, and Weill (2012).

[13] The same logic is used routinely to uniquely pin down trading volume in a frictionless market: one would rule out the possibility that agents buy and sell contracts simultaneously by arguing that arbitrarily small transaction costs lead to positions that minimize their trading volume.

[14] By the Theorem of the Maximum (see Stokey and Lucas, 1989, Theorem 3.6), the set of bilateral exposures solving the planning problem is compact, implying that the above minimization problem has a

Next, we compare gross exposures to net exposures. We note that, unlike its gross exposure, the net exposure of a bank with endowment ω is uniquely determined because it is equal to $|G^+(\omega) - G^-(\omega)| = |g(\omega) - \omega|$. The average net exposure in the market is: $\mathcal{N}(k) = \sum_\omega |G^+(\omega) - G^-(\omega)|\, n(\omega)$. A natural measure of the volume created by intermediation activity is the ratio of gross-to-net exposure:

$$\mathcal{R}(k) \equiv \frac{\mathcal{G}(k)}{\mathcal{N}(k)} \geq 1.$$

When $\mathcal{R}(k) = 1$, gross and net exposures are the same, and there is no intermediation activity. When $\mathcal{R}(k) > 1$, some banks are taking simultaneous long and short positions and intermediation activity arises. Note that, since we consider for this calculation the bilateral exposures that minimize gross exposures, this prediction is robust. That is, intermediation activity arises in all sets of bilateral exposures which are consistent with equilibrium. We obtain:

Proposition 4. *Assume that $|supp(N)| \geq 3$. Then there is some $\hat{k}$ such that $\mathcal{R}(k) > 1$ if and only if $k < \hat{k}$. Moreover $\hat{k} > \bar{k}$, where $\bar{k}$ is the trade size limit needed to equalize post trade exposures as defined in Proposition 2.*

The condition $|\text{supp}(N)| \geq 3$ is necessary because we need at least 3 types to create intermediation activity: indeed, with only two types, each bank would only have one type of counterparty, and would always trade in the same direction. Notice also that $\hat{k} > \bar{k}$: at the point when the OTC market can achieve full risk sharing, all equilibria require some strictly positive amount of intermediation activity.[15]

3.2 Equilibrium entry

We now define and establish the existence of an equilibrium with banks' entry decisions chosen optimally.

solution.

[15]The proposition shows that, by varying k, we effectively vary the extent to which banks are able to direct their trade to their best counterparties. Indeed, when $k < \hat{k}$ is small, frictions are large and trading patterns appear more random: there is partial risk sharing, banks may trade in either direction depending on who they meet, and gross exposures differ strictly from net exposures. When $k > \hat{k}$, frictions are small and trading patterns become directed: there is full risk sharing and each bank trades in only one direction. Lemma B11 in the appendix offers further illustration of this point in the context of a parametric model with three types.

3.2.1 The marginal private value of entry

Given the distribution of traders, n, and the post-trade exposures, g, that arise in the corresponding OTC market equilibrium conditional on entry, we can calculate a bank's net per-trader capita utility of entering given its initial endowment ω, which we call its *marginal private value of entry*:

$$\mathrm{MPV}(\omega) \equiv \begin{cases} 0 & \text{if } \sum_{\tilde{\omega}} \mu(\tilde{\omega}) = 0 \\ \Gamma\left[\omega\right] - \Gamma\left[g(\omega)\right] + \sum_{\tilde{\omega}} \gamma(\omega, \tilde{\omega}) R(\omega, \tilde{\omega}) n(\tilde{\omega}) & \text{if } \sum_{\tilde{\omega}} \mu(\tilde{\omega}) > 0. \end{cases} \tag{14}$$

If no other bank enters in the OTC market, $\sum_{\tilde{\omega}} \mu(\tilde{\omega}) = 0$, the marginal private value is evidently zero. Otherwise, if there is positive entry, $\sum_{\tilde{\omega}} \mu(\tilde{\omega}) > 0$, the marginal private value has two terms. The first term, $\Gamma\left[\omega\right] - \Gamma\left[g(\omega)\right]$, is the bank's change in per capita exposure: it is negative if the bank is a net seller of insurance, and positive if it is a net buyer. The second term is the sum of all CDS premia collected and paid by the bank per trader capita. It is positive if the bank is a net seller, and negative if it is a net buyer.

We now show that $\mathrm{MPV}(\omega)$ is defined unambiguously: it only depends on g, which we know from Theorem 1 is the same in all equilibria. In particular, it does not depend on the particular bilateral exposures, γ, established by banks in the OTC market equilibrium.

Lemma 5. *Given n, in any OTC market equilibrium conditional on entry, the sum of all CDS premia collected and paid by a bank with per-capita endowment ω is uniquely pinned down by the equilibrium post-trade exposures, g:*

$$\sum_{\tilde{\omega}} \gamma(\omega, \tilde{\omega}) R(\omega, \tilde{\omega}) n(\tilde{\omega}) = \Gamma'\left[g(\omega)\right]\left[g(\omega) - \omega\right] + \frac{k}{2} \sum_{\tilde{\omega}} \left|\Gamma'\left[g(\tilde{\omega})\right] - \Gamma'\left[g(\omega)\right]\right| n(\tilde{\omega}).$$

We obtain this formula by adding and subtracting the term $\sum_{\tilde{\omega}} \Gamma'\left[g(\omega)\right] \gamma(\omega, \tilde{\omega}) n(\tilde{\omega})$ to $\sum_{\tilde{\omega}} \gamma(\omega, \tilde{\omega}) R(\omega, \tilde{\omega}) n(\tilde{\omega})$ as defined by equation (9) and then use equation (6) as well as the optimality condition (8) to obtain the first term on the right-hand side of the equation above.

3.2.2 Equilibrium entry: definition and existence

A bank of type ω will find it optimal to enter if and only if

$$\mathrm{MPV}(\omega) \geq \frac{c}{S},$$

where c is the bank's fixed cost of entry, and S is the bank's size. Now recall our distributional assumptions. First, the measure of traders in banks with per capita endowment ω is equal to $\pi(\omega)$. Second, conditional on ω, the measure of traders in banks with per-capita entry costs less than c/S is given by the CDF $\Phi(c/S)$. Thus, the measures of traders in the OTC market must satisfy $\mu \in T[\mu]$, where $T[\mu]$ is the set of measures $\hat{\mu}$ such that:

$$\pi(\omega)\Phi\left[\mathrm{MPV}(\omega)^{-}\right] \leq \hat{\mu}(\omega) \leq \pi(\omega)\Phi\left[\mathrm{MPV}(\omega)\right] \tag{15}$$

for all $\omega \in \Omega$. In the formula, $\pi(\omega)\Phi\left[\mathrm{MPV}(\omega)^{-}\right]$ and $\pi(\omega)\Phi\left[\mathrm{MPV}(\omega)\right]$ are, respectively, the minimum and the maximum measures of type-ω traders in banks that choose to enter the OTC market given the marginal private value of doing so, $\mathrm{MPV}(\omega)$.[16]

An *equilibrium with entry* is, then, a fixed point of the operator T. Based on this definition we establish:

Theorem 6. *There always exists an equilibrium with no entry, $0 \in T[0]$. Moreover, there exists some $b(\eta, k) > 0$, a function of traders' absolute risk aversion, η, and risk limits, k, such that, for any CDF of costs satisfying $\Phi\left[b(\eta,k)^{-}\right] > 0$, there exists an equilibrium with strictly positive entry, i.e. some $\mu \in T[\mu]$ such that $\sum_{\tilde{\omega}} \mu(\tilde{\omega}) > 0$.*

It is obvious that no entry is always an equilibrium: if no other bank enters, then $\mathrm{MPV}(\omega) = 0$ for all ω, and so no bank finds it strictly optimal to enter. The non-trivial part of the theorem is to establish that there exists an equilibrium with *strictly positive* entry. To do so, we note that, for any distribution n, the marginal private value must be strictly positive for at least some type, i.e., $\max_{\tilde{\omega} \in \Omega} \mathrm{MPV}(\tilde{\omega}) > 0$. Intuitively, if $n(\omega) > 0$ for some ω, then any $\tilde{\omega} \neq \omega$ gains from sharing risk with ω. After showing that the marginal private values are continuous functions of the distribution of traders, n, we can take the infimum over all possible n and obtain a strictly positive bound on the marginal private value of at least one type. As long as there are banks with a sufficiently low entry cost, this translates into a strictly positive lower bound $\underline{\mu}$ on the measure of traders in the market. This allows us to apply Kakutani's fixed point theorem on the set of measures μ such that the total measure of traders in the market exceeds this lower bound, $\sum_{\tilde{\omega}} \mu(\tilde{\omega}) > \underline{\mu}$ and, in doing so, find a fixed point with strictly positive entry.[17]

[16]If the distribution of costs is continuous at $\mathrm{MPV}(\omega)$, the equilibrium condition reduces to an equality $\hat{\mu}(\omega) = \pi(\omega)\Phi\left[\mathrm{MPV}(\omega)\right]$. If there is an atom, the equilibrium condition has to be stated with two inequalities, as in (15) above.

[17]Note that the theorem holds when banks' per-trader entry costs are all bounded away from zero: it does not require that there is an atom of banks with infinite size and/or zero per-trader costs, nor that there are banks with arbitrarily large size and per trader costs arbitrarily close to zero.

4 Private vs. social entry incentives

In this section, we study banks' private and social incentives to enter. We first show that private entry incentives are U-shaped functions of banks' initial exposures. We argue that this implies that, in equilibrium, only large enough banks find it optimal to enter the market to act as dealers. Next, we compare private and social entry incentives. We establish that, for banks who act as dealers, the marginal private value of entry is greater than the marginal social value. Thus, these banks have too large an incentive to enter in equilibrium. For banks that assume the role of customers, we obtain the opposite result. Their marginal private value of entry is lower than their marginal social value. Thus, these banks have too small an incentive to enter in equilibrium.

4.1 Properties of the marginal private value of entry

Banks' decisions to enter the OTC market are driven by two motivations. The first motivation is to hedge underlying risk exposure. The second motivation arises because, as long as k is small enough, hedging is imperfect in equilibrium. As a result, the marginal costs of risk bearing are not equalized across banks and prices are dispersed. This gives banks an incentive to enter in order to earn additional trading profits. To isolate the hedging from the trading profit motives in the marginal private value, $\mathrm{MPV}(\omega)$, we extend the decomposition of a bank's trading revenues presented in Lemma 5. We let the hedging motive correspond to the entry incentives of a hypothetical bank whose traders have no bargaining power, and let the trading profit motive correspond to the residual.

For a type-ω bank, the net utility of entry when traders have no bargaining power is

$$\begin{aligned} K(\omega) &\equiv \Gamma[\omega] - \Gamma[g(\omega)] + \sum_{\tilde{\omega}} \Gamma'[g(\omega)]\,\gamma(\omega,\tilde{\omega})n(\tilde{\omega}) \\ &= \Gamma[\omega] - \Gamma[g(\omega)] + \Gamma'[g(\omega)]\big[g(\omega) - \omega\big], \end{aligned}$$

since, in that case, traders would buy and sell CDS at the same price, equal to their bank's marginal cost, $\Gamma'[g(\omega)]$. By the convexity of $\Gamma[g]$, the function $K(\omega)$ is positive. As illustrated in Figure 1, it can be viewed as producer surplus if the bank is a net seller of CDS, and as consumer surplus if it is a net buyer. In what follows, we will refer to $K(\omega)$ as the per-capita *competitive surplus* of bank ω. The competitive surplus measures a bank's fundamental gains from trade, because it is equal to the producer or consumer surplus if the bank would conduct all its trades at the same price, equal to its marginal valuation.

The trading profit motive corresponds to the residual, $\mathrm{MPV}(\omega) - K(\omega)$. It corresponds

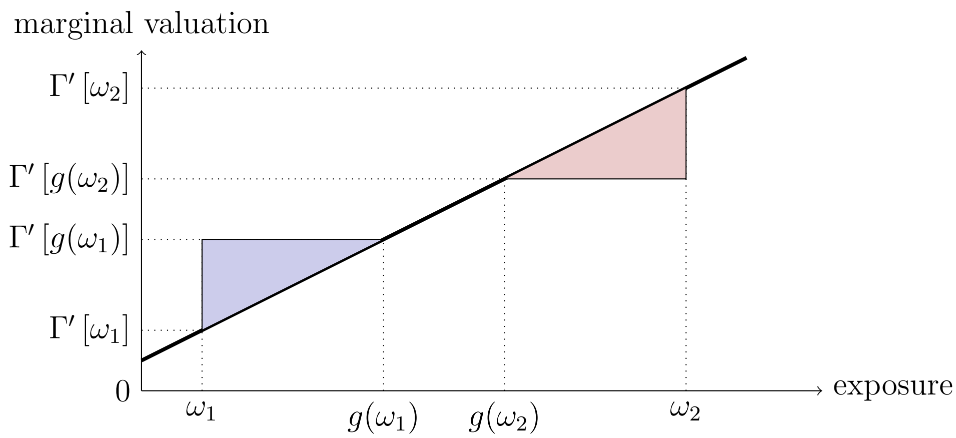

Figure 1: The figure shows the marginal cost of risk bearing for two types of banks. On the left-hand side, a bank of type ω_1 who is a net provider of insurance to other banks in the OTC market, $g(\omega_1) > \omega_1$. For this bank the competitive surplus, $K(\omega)$, corresponds to the shaded area above the marginal cost curve. On the right-hand side, a bank of type ω_2 who is a net demander of insurance. For this bank the competitive surplus corresponds to the shaded area below the marginal cost curve.

to the profits traders can make when risk-sharing is imperfect and the marginal cost of risk bearing is not equalized across banks. This allows traders to negotiate prices above marginal cost when they sell, and below marginal value when they buy. This trading profit motive is given by

$$\begin{aligned}\frac{1}{2}F(\omega) &\equiv \sum_{\tilde{\omega}} \gamma(\omega,\tilde{\omega})R(\omega,\tilde{\omega})n(\tilde{\omega}) - \sum_{\tilde{\omega}} \Gamma'\left[g(\omega)\right]\gamma(\omega,\tilde{\omega})n(\tilde{\omega}) \\ &= \frac{1}{2}\sum_{\tilde{\omega}} \gamma(\omega,\tilde{\omega})\big(\Gamma'\left[g(\tilde{\omega})\right] - \Gamma'\left[g(\omega)\right]\big)n(\tilde{\omega}) = \frac{k}{2}\sum_{\tilde{\omega}} \big|\Gamma'\left[g(\tilde{\omega})\right] - \Gamma'\left[g(\omega)\right]\big|n(\tilde{\omega}),\end{aligned}$$

where the third equality follows from the optimality condition (8). We call $F(\omega)$ the *frictional surplus*, because it represents gains from entering the market over and above the competitive surplus that are purely due to the frictions: the matching and trade size limit that result in imperfect risk sharing. Note that $F(\omega)$ is strictly positive if and only if there is strictly positive price dispersion. Thus, the frictional surplus can be interpreted as the trading profits gained from price dispersion.

By construction, we must have:

$$\text{MPV}(\omega) = K(\omega) + \frac{1}{2}F(\omega).$$

The next proposition shows that the competitive and the frictional surplus are larger at the extremes of the set Ω:

Proposition 7. *The competitive and frictional surplus have the following properties:*

- *The competitive surplus, $K(\omega)$, is equal to zero if and only if $g(\omega) = \omega$. Moreover, when $\Gamma[g]$ is quadratic, it is a U-shaped function of ω.*

- *The frictional surplus, $F(\omega)$, is strictly positive if and only if price dispersion is strictly positive. It is a U-shaped function of ω, and achieves its minimum at any median of the distribution of traders in the market, N.*

The proposition implies that banks with intermediate ω, which we know from Proposition 3 engage in intermediation activity, have the smallest incentives to enter: they tend to have lower competitive and frictional surpluses.

From Figure 1, it is clear that the competitive surplus is smallest if the bank has zero net exposure due to trading, $g(\omega) = \omega$. One can obtain a sharper characterization of $K(\omega)$ in the special case of a quadratic cost of risk bearing. Then, direct calculations reveal that the competitive surplus is quadratic as well, equal to $K(\omega) = \frac{\Gamma''}{2}\left[g(\omega) - \omega\right]^2$, where Γ'' is the (constant) second derivative of $\Gamma[g]$. One sees that the competitive surplus only depends on the distance between post- and pre-trade exposures in this case, and not on their level. It reveals that $K(\omega)$ is larger when $g(\omega)$ and ω are further apart, which in equilibrium occurs when ω is either small or large.

Mathematically, the frictional surplus, $F(\omega)$, measures the average absolute distance, in terms of marginal valuation, between bank ω and other banks $\tilde{\omega} \neq \omega$. As is well known, this average absolute distance is minimized by any median of the distribution of marginal valuation. Since the marginal valuation is increasing in ω, it is also minimized by any median of N.

One insight arising from our analysis is that large banks are more likely to engage in intermediation activity. Intuitively, we have shown that banks with intermediate initial exposures assume the role of dealers (Proposition 3) and, at the same time, have the smallest incentives to enter (Proposition 7). Combined with fixed entry costs, this implies that banks with intermediate exposures who assume the role of dealers must be larger, on average, than banks who assume the role of customers. Formally, we have:

Corollary 8. *Assume that k is small, that the cost of risk-bearing is quadratic, that Ω and $\{\pi(\omega)\}_{\omega\in\Omega}$ are symmetric around some endowment $\omega^\star$, and that $\Phi(z) > 0$ for all $z > 0$. Then, there exists an equilibrium such that:*

- *only large-sized banks enter to assume the role of intermediaries;*

- *middle-sized banks only enter to assume the role of dealers;*
- *small sized banks do not enter.*

The assumptions stated in the corollary ensure that there exists an equilibrium in which the median and the mean of N are both equal to the center of symmetry of the distribution, $\omega^\star$. As a result, $K(\omega)$, $F(\omega)$, and $\mathrm{MPV}(\omega)$ all achieve their minimum at $\omega = \omega^\star$. It should be clear, however, that the economic intuition for the result applies more broadly: it simply relies on the observation that $\mathrm{MPV}(\omega)$ is smaller for intermediate ω and larger for extreme ω.

Before proceeding, let us note that our model delivers additional empirical implications in the cross-section of banks sorted by size. Specifically, in Atkeson, Eisfeldt, and Weill (2012), we showed that large dealer banks tend to have larger gross-exposure per capita, even if they have the same trade size limit k as other banks, as is the case in the data. We also showed that large dealer banks tend to trade amongst each other at less dispersed prices. In the earlier work of Duffie, Gârleanu, and Pedersen (2005), it is exogenously assumed that market makers or dealers trade in a frictionless market at common prices. In our model, such an interdealer market arises endogenously among large banks that are central to the market.

4.2 Properties of the marginal social value of entry

Next, we show that equilibrium entry incentives are not aligned with social interest. To do so, we compare the marginal private value of entry for a bank to its marginal social value from entry as obtained from a social planning problem for banks' entry decisions.

The marginal social value. Consider any bank entry pattern, as represented by measures μ of traders in the OTC market. The social value generated in the OTC market equilibrium conditional on this entry pattern is given by the solution $W(\mu)$ of the planning problem studied in (10). The *marginal social value* of a ω bank is defined as:

$$\mathrm{MSV}(\omega) \equiv \frac{\partial W(\mu)}{\partial \mu(\omega)}, \tag{16}$$

In words, the marginal social value of a type-ω trader is the partial derivative of social welfare with respect to the measure of such traders, given optimal trading behavior in the OTC market. To show that this marginal social value is well defined and provide an explicit expression for it, we can apply an envelope theorem of Milgrom and Segal (2002).

Lemma 9. *For any $\mu \neq 0$, the marginal social value of a bank of type ω is:*

$$MSV(\omega) = K(\omega) + F(\omega) - \frac{1}{2}\bar{F}, \tag{17}$$

where $K(\omega)$ is the competitive surplus, $F(\omega)$ is the frictional surplus, and $\bar{F}$ is the average frictional surplus $\bar{F} \equiv \sum_{\tilde{\omega}} F(\tilde{\omega})n(\tilde{\omega})$.

To understand the formula, it is useful to interpret entry in the OTC market as a process of match creation and match destruction between pairs of traders. Match creation arises because, when a small measure ε of new traders enter, they create ε new matches with incumbent traders. Match destruction arises because, in the absence of entry, these ε incumbents would have been matched together instead of being matched with the ε entrants. Therefore, entry effectively destroys matches amongst these ε incumbents. Given that there are two traders per match, the measure of matches thus destroyed is equal to $\varepsilon/2$. The marginal social value is obtained, using the envelope theorem, by calculating the value of match creation net of the cost of match destruction, holding all bilateral exposures, γ, constant.

The first two terms of (17), $K(\omega) + F(\omega)$, represent the social value of match creation. By entering in the OTC market, ω-traders establish new CDS contracts. In doing so, they create competitive surplus by changing their own exposure from ω to $g(\omega)$, and they also generate frictional surplus with their counterparties. Precisely, if one assumes that, in all matches with ω traders, marginal values and costs are equal to $\Gamma'[g(\omega)]$, then the social value of match creation is equal to the competitive surplus, $K(\omega)$. However, because of imperfect risk sharing, the marginal cost of increasing exposure is always lower than the marginal value, and strictly so if exposures are not equalized. This increases the social value of match creation above and beyond the competitive surplus by the term $F(\omega)$, the average distance between marginal value and marginal cost in all matches involving a type-ω trader.

The last term of (17), $-\bar{F}/2$, represents the social cost of match destruction. When a new trader of type ω enters, $1/2$ matches amongst incumbents are destroyed. The identity of incumbents in these matches that are destroyed is independent of ω and hence the value destroyed per match is the average frictional surplus $\bar{F}$, the average difference between marginal value and marginal cost in all OTC market matches.

The marginal social value is always positive. Next, we establish two results. First, we show that the marginal social value of a bank is always positive, i.e., the value of match creation always exceeds that of match destruction. Thus, absent any entry cost, a planner would always like to make the market as large as possible. Second, we show that, when risk sharing is imperfect, intermediation has strictly positive social value. This means that, if trade size limits are tight enough, a social planner has strict incentives to let some banks with intermediate exposure assume the role of intermediaries.

Lemma 10. *Suppose that supp$(N) \geq 2$. Then, the marginal social value is positive, MSV$(\omega) \geq 0$, and strictly so for at least one ω. Moreover, suppose there is imperfect risk sharing in the OTC market and consider a pure intermediary, i.e., a bank ω such that $g(\omega) = \omega$. Then, for this intermediary, MSV$(\omega) > 0$.*

To see why $\text{MSV}(\omega) \geq 0$, we first note that by convexity, the first term in the marginal social value, $K(\omega)$, is always positive. The second and third term term, $F(\omega) - \bar{F}/2$, which represents the net effect on frictional surplus of match creation and destruction, is also always positive. Indeed, when an additional ω-trader participates in the market, she destroys some existing matches, and simultaneously replaces all these destroyed matches by two matches with herself. Consider, as illustrated in Figure 2, the destruction of a pair of (x, y), with $x \leq y$, and the simultaneous creation of two pairs (ω, x) and (ω, y). The net surplus created is:

$$\Big|\Gamma'\left[g(\omega)\right] - \Gamma'\left[g(x)\right]\Big| + \Big|\Gamma'\left[g(\omega)\right] - \Gamma'\left[g(y)\right]\Big| - \Big|\Gamma'\left[g(x)\right] - \Gamma'\left[g(y)\right]\Big| \geq 0, \tag{18}$$

which is positive by the triangle inequality. The intuition is that, since any direct trade between x and y can be replicated by two indirect trades through the ω-trader, the process of match creation and match destruction cannot destroy any value. To the contrary, the optimal indirect trade through the ω trade can create strictly more value than the optimal direct trade between x and y. This occurs whenever the post-trade exposure of the ω-trader is located either strictly to the left or to the right of the post-trade exposures of the x- and the y-trader. For example if $g(\omega) < g(x) \leq g(y)$, then ω can provide insurance to y at lower cost than x. Even better, ω can also provide insurance to x.

The second part of the Lemma asserts that intermediation is socially beneficial, i.e., it creates strictly positive social value. To see this, consider a pure intermediary, i.e., some bank ω such that $g(\omega) = \omega$. Since risk sharing is imperfect, and since the ω bank does not change its exposure, there must be banks of type x and y such that $n(x) > 0$, $n(y) > 0$ and, $g(x) \leq g(\omega) \leq g(y)$, with at least one strict inequality. In equilibrium, the traders of a bank of type ω provide intermediation, in particular they buy insurance from x and sell insurance to y. To see why $\text{MSV}(\omega) > 0$, note that the entry of a type-ω trader destroys matches of type (x, x) and (y, y) and replaces them by matches of type (x, ω) and (ω, y). In doing so, it creates social value by allowing an indirect trade between x and y, when no direct trade existed before.

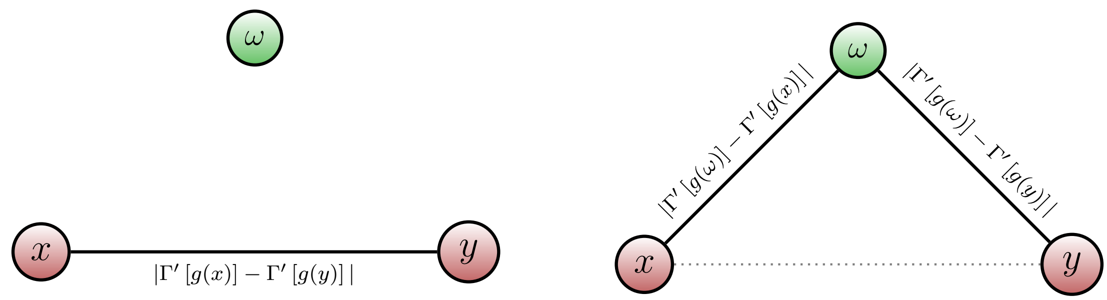

(a) before entry

(b) after entry

Figure 2: An illustration of match creation and match destruction.

Marginal social vs. private value. Finally, we compare $\text{MSV}(\omega)$ and $\text{MPV}(\omega)$. When a ω-trader enters and destroys a match between $x \leq y$-traders, she appropriates half of the surplus:

$$\frac{1}{2}\Big(\big|\Gamma'\left[g(\omega)\right] - \Gamma'\left[g(x)\right]\big| + \big|\Gamma'\left[g(\omega)\right] - \Gamma'\left[g(y)\right]\big|\Big). \tag{19}$$

Subtracting this surplus appropriated from the surplus created, in equation (18), we obtain:

$$\frac{1}{2}\Big(\big|\Gamma'\left[g(\omega)\right] - \Gamma'\left[g(x)\right]\big| + \big|\Gamma'\left[g(\omega)\right] - \Gamma'\left[g(y)\right]\big|\Big) - \big|\Gamma'\left[g(x)\right] - \Gamma'\left[g(y)\right]\big|.$$

One sees from the equation that the difference between the surplus created and appropriated can be either negative or positive, depending on the position of $g(\omega)$ relative to $g(x) \leq g(y)$. One sees in particular that the difference between surplus appropriated and created can be strictly positive even for a pure intermediary. This arises for example if $x = y$ but $g(\omega) \neq g(x)$. Averaging over the $1/2$ incumbent matches destroyed, we obtain that:

$$\text{MSV}(\omega) - \text{MPV}(\omega) = \frac{1}{2}\left[F(\omega) - \bar{F}\right]. \tag{20}$$

The economic intuition for this formula is as follows. The marginal social and private incentives have one term in common: the competitive surplus, which measures the value attached to changes in net positions. Hence, banks' competitive surplus cancels out in computing the difference between the marginal social and private values of entry. In contrast, the consideration of frictional surplus differs in the marginal social and private incentives to enter. A social planner attributes to each marginal entering bank the full frictional surplus that its traders induce by creating new trading opportunities. A marginal entering bank, on the other hand, only considers the portion of that surplus that it can capture through bargaining with its new counterparties. At the same time, however, a social planner also attributes to each marginal entering bank a social cost equal to the loss of frictional surplus in the trading opportunities that are displaced by entry. By direct computations, we have shown that the gap between the marginal social and marginal private incentives, in equation (20), is equal to the portion of the surplus in new trading opportunities created by bank entry that is not captured through bargaining by the entering bank, $F(\omega)/2$, less the expected surplus in trading opportunities among incumbents displaced due to the entry of a new bank, $\bar{F}/2$.

Equation (20) also reveals that banks who are far enough from others, in terms of their marginal cost, have a marginal social value that is greater than their marginal private value. These banks create, on average, more surplus than they appropriate. Now recall that, by Proposition 7, the frictional surplus is U-shaped. Therefore, banks with extreme endowments, who assume the role of customers, have a marginal social value that is greater than their marginal private value. In equilibrium, they have too small an incentive to enter. By contrast, banks with intermediate endowments, who assume the role of intermediaries, have a marginal social value that is lower than their marginal private value. In equilibrium, they have too large an incentive to enter.

As equation (20) makes clear, the finding that $\mathrm{MSV}(\omega) - \mathrm{MPV}(\omega)$ is positive for extreme ω banks and negative for intermediate ω banks follows directly from our assumption of symmetric bargaining weights (of 1/2) for all traders. This result can be altered if we make bargaining weights *asymmetric* across banks. To illustrate this point, in Appendix B.6 we show in the context of a parametric example that if traders from intermediate endowment banks have sufficiently low bargaining power, then it is possible for the gap to be positive instead of negative for middle-ω banks. In this case, there is too little entry of dealer banks in equilibrium.

4.3 A social planning problem for entry

We provide further illustrations of our normative results by studying the socially optimal pattern of entry into the market given the trade size limits and fixed entry costs. Social welfare conditional on the entry patterns μ is equal to $W(\mu)$, as defined in (10). To calculate the total entry cost associated with the entry patterns μ, we first rank traders in terms of their entry costs using the quantile function $\psi(q) \equiv \inf\{z \in [\underline{z}, \bar{z}] : \Phi(z) \geq q\}$, where $[\underline{z}, \bar{z}]$ is the support of $\Phi(z)$. The function $\psi(q)$ represents the cost of the trader positioned at the quantile q of the distribution. It is clear that, all else equal, a social planner always lets traders with the lowest cost enter first. Therefore, when $\mu(\omega) \in [0, \pi(\omega)]$ traders enter the market, they collectively incur the cost:

$$C(\mu) = \sum_{\omega \in \Omega} \pi(\omega) \int_0^{\frac{\mu(\omega)}{\pi(\omega)}} \psi(r)\, dr,$$

which is an increasing, convex, and continuous function. The *social planning problem* is, then, $\max_{\mu} W(\mu) - C(\mu)$, subject to $\mu(\omega) \in [0, \pi(\omega)]$ for all $\omega \in \Omega$. Our main result is:

Theorem 11. *The social planning problem has a solution and has the following properties:*

- *if there is an equilibrium with positive entry, then the planner's problem has a solution with positive entry;*
- *if the planning problem prescribes positive entry, then the first-order necessary conditions for optimality can be written*

  $$\pi(\omega)\Phi\left[MSV(\omega)^-\right] \leq \mu(\omega) \leq \pi(\omega)\Phi\left[MSV(\omega)\right],$$

  where the marginal social value is defined in (16)*;*
- *if the planning problem prescribes positive entry, then it can be implemented as an entry equilibrium provided type-ω banks receive a per-trader subsidy equal to $MSV(\omega) - MPV(\omega) = \frac{1}{2}\left[F(\omega) - \bar{F}\right]$.*

Existence follows because the planner's objective can be shown to be continuous and because the constraint set is compact.

The first bullet point follows because, in any entry equilibrium, for any entrant, individual rationality implies that the utility of entering must be at least as large as the utility of staying out. Adding up all these utilities, and using the fact that all CDS payments add up to zero, we obtain that social welfare in any equilibrium must be at least as large as social welfare with no entry.

The second bullet point shows a close correspondence between the planner's first-order condition and the equilibrium fixed-point equation: they are formally identical, except for the fact that MPV(ω) is replaced by MSV(ω). The intuition is that, if the planner lets type-ω traders enter up until their marginal social value is equal to MSV(ω), then all the traders with cost strictly less than MSV(ω) must have entered, and all traders with cost strictly above than MSV(ω) must have stayed out.

The third bullet point follows directly from the second one: in order to induce optimal entry, a policy maker needs to give a subsidy equal to the difference between marginal social and marginal private value. In particular, banks with extreme exposure, who assume the role of customers, should be subsidized. Banks with intermediate exposures, who assume the role of intermediaries, should be taxed.

One sees clearly from the theorem that inefficiencies arise because of imperfect risk sharing in the OTC market, which creates dispersion in banks' marginal cost of risk bearing. If the planner prescribes full risk sharing in the OTC market, which can be shown to occur regardless of entry costs as long as k is large enough, then $F(\omega) = \bar{F} = 0$, and $\mathrm{MSV}(\omega) = \mathrm{MPV}(\omega)$. Comparing the equilibrium fixed-point problem with the planner's first-order conditions, one then immediately sees that there exists an entry equilibrium that implements the planner's allocation.[18]

In Section 6.1.2, we explicitly compare the planning solution and the equilibrium in a parametric example with three endowment types. We derive conditions such that, starting from the equilibrium, the social planner finds it optimal to decrease the entry of dealers and increase the entry of customers. The resulting socially optimal OTC market structure has fewer participants, generates less trading volume, and creates smaller ratios of gross-to-net positions. Therefore, according to the model, there is a policy role for taxing dealer banks in order to reduce some of the trading volume generated by intermediation activity, while subsidizing the participation of customer banks in order to promote direct customer-to-customer transactions.

Finally, in Appendix B.6, we show in the context of a parametric example that a social optimum can be implemented in an equilibrium with appropriately chosen *asymmetric* bargaining weights. We show that dealers' bargaining weight should be strictly positive, less than $\frac{1}{2}$, and that it should be a decreasing function of the socially optimal fraction of dealers in the market. Given any bargaining weight in between $\frac{1}{2}$ and some strictly positive lower bound, there can be either over- or under-entry of dealers in equilibrium, depending on other model parameters.

5 Exit

The entry problem studied so far focuses on the long run: it provides insights into the structure of OTC markets subject to sunk costs of entry for banks that expect to maintain an average exposure of ω over time. An important question of interest, prompted by the

[18] Whenever the equilibrium implements the planning solution, there are no intermediation profits to be made in the OTC market. Yet, as we saw in Proposition 4, intermediation activity may be necessary to achieve the planner's allocation. Formally, in Appendix A.11 we show that, for some parameter, in the planner's allocation, there is full risk sharing and, at the same time, some banks engage in intermediation activity, in that the ratio of gross exposures to net exposure must be strictly greater than 1.

financial crisis, pertains to the short run: are OTC markets excessively vulnerable to shocks causing the exit of financial institutions?[19] In this section, we study this question in a variation of our framework based on two additional assumptions. First, banks receive shocks to the costs they must incur to continue actively trading in the OTC market. Second, the OTC market is imperfectly resilient, in the sense that traders who lose a counterparty due to exit face difficulties finding a new counterparty with whom to resume trading.

5.1 The model

To model the short-run response of the market to negative shocks, we modify our framework as follows.

Negative shocks. We assume that entry has already taken place, and that all traders in the OTC market have already established bilateral relationships. Then, each bank currently in the market draws a negative shock, denoted by z, raising its costs of continued participation in the market on a per trader basis, and chooses whether or not to exit.

To highlight the similarities between this model of exit and our model of entry, we keep the same notation. We denote the CDF of exit costs by $\Phi(z)$. Of course, this CDF does not need to be the same as the CDF of entry costs we considered before. Likewise, we denote by $\pi(\omega)$ the measure of traders of type ω in the market before exit decisions have been made, with the normalization $\sum_{\tilde{\omega}} \pi(\tilde{\omega}) = 1$. Finally, we denote by $\mu(\omega) \in [0, \pi(\omega)]$ the measure of type-ω traders who remain in the OTC market after exit decisions have been made.

Imperfectly resilient market. The key difference with the model of entry is that, after exit, the OTC market is *imperfectly resilient*: we assume that those traders who lose a counterparty due to exit may not be able to find a new one. More precisely, the traders who lose a counterparty due to exit can re-match amongst each other with probability $\rho \in (0, 1]$. The parameter ρ indexes the resilience of the OTC market: if $\rho \simeq 1$, then the market is very resilient, in that a trader who has lost a counterparty is very likely to establish a new trading relationship with another trader in the market, and vice versa if $\rho = 0$. While we do not model dynamics explicitly, it is also natural to think of $\rho \simeq 0$ as representing the state of the market in the short run, when banks scramble to replace their lost counterparties. At

[19] Dodd (2012) notes that during the financial crisis "Dealers (...) withdrew from the markets. (...) Without the dealers, there was no trading, especially in securities such as collateralized debt obligations, certain municipal securities, and credit derivatives. With no buyers, investors could not reduce losses by trading out of losing positions and they could not sell those positions to meet calls for more margin or collateral to pledge against loans they had taken out to buy those instruments. This illiquidity in OTC markets contributed to the depth and breadth of the financial crisis."

the other extreme, $\rho \simeq 1$ represents the state of the market in the long run, when banks have had the time to replace the counterparties they may have lost. Indeed, we shall see shortly that, when $\rho = 1$, the model of exit becomes identical to the model of entry.

To see how imperfect resilience changes the matching process relative to the entry model, we calculate the probability a given trader in the market has of having counterparty of type ω after exit decisions have been made:

$$\pi(\omega)\frac{\mu(\omega)}{\pi(\omega)} + \sum_{\tilde{\omega}} \pi(\tilde{\omega}) \left[1 - \frac{\mu(\tilde{\omega})}{\pi(\tilde{\omega})}\right] \rho \frac{\mu(\omega)}{\sum_{\tilde{\omega}} \mu(\tilde{\omega})}.$$

The first term is the probability $\pi(\omega)$ that the original counterparty is of type ω, multiplied by the probability that this original counterparty does not exit, $\mu(\omega)/\pi(\omega)$. The second term is the probability of losing a counterparty due to exit, $\sum_{\tilde{\omega}} \pi(\tilde{\omega}) [1 - \mu(\tilde{\omega})/\pi(\tilde{\omega})]$, multiplied by the probability of being re-matched, ρ, multiplied by the probability of re-matching with a counterparty of type ω, $\mu(\omega)/ [\sum_{\tilde{\omega}} \mu(\tilde{\omega})]$. Simplifying the expression, we obtain that the probability of matching with a counterparty of type ω after exit decisions have been made is equal to:

$$\alpha(\mu) \times n(\omega) \text{ where } \alpha(\mu) \equiv \rho + (1-\rho) \sum_{\tilde{\omega}} \mu(\tilde{\omega}), \text{ and } n(\omega) = \frac{\mu(\omega)}{\sum_{\tilde{\omega}} \mu(\tilde{\omega})}.$$

Imperfect resilience impacts the matching probability, $\alpha(\mu)$, in two ways. First, it is in general less than one. As we will see, this changes the equilibrium OTC market conditional on entry. Second, it is now endogenous, and depends on the measure of banks who stay in the market. As we will see, this induces strategic complementarities in exit decisions, and could potentially generate further multiplicity of equilibria. It also creates a new externality that changes the normative analysis in a fundamental way.

5.2 Equilibrium exit

As before, we first define an equilibrium conditional on exit. Relative to our model of entry, the only difference is that the probability of matching, $\alpha(\mu)$, can be less than one. As a result, the post-trade exposure becomes:

$$g(\omega) = \omega + \alpha(\mu) \sum_{\tilde{\omega}} \gamma(\omega, \tilde{\omega}) n(\tilde{\omega}). \tag{21}$$

One sees that, because CDS contracts are consolidated within the bank, all bilateral exposures are scaled down by the matching probability $\alpha(\mu)$. In fact:

Proposition 12. *The tuple $\{\gamma, g, R\}$ is an equilibrium conditional on exit, given matching probability $\alpha(\mu)$ and trade size limit k, if and only if it is an equilibrium conditional on entry with trade size limit $\alpha(\mu)k$.*

In an exit equilibrium, the marginal private value of remaining in the market for a bank with endowment ω becomes:

$$\text{MPV}(\omega) = \begin{cases} 0 & \text{if } \sum_{\tilde{\omega}} \mu(\tilde{\omega}) = 0 \\ K(\omega) + \frac{\alpha(\mu)}{2} F(\omega) & \text{if } \sum_{\tilde{\omega}} \mu(\tilde{\omega}) > 0, \end{cases}$$

which is the same as for the entry equilibrium, except for the fact the frictional surplus is scaled down by the endogenous matching probability, $\alpha(\mu)$. Up to this adjustment, the fixed-point problem for the exit equilibrium is formally the same as the one of Section 3.2.2. Using the same arguments as before, we obtain:

Proposition 13. *There is always an equilibrium in which all banks exit, $0 \in T[0]$. Moreover, there exists some $b(\eta, k) > 0$, a function of traders' absolute risk aversion and risk limits, such that, for any CDF of costs satisfying $\Phi\left[b(\eta, k)^{-}\right] > 0$, there exists an equilibrium in which a strictly positive measure of banks stay, i.e. some $\mu \in T[\mu]$ such that $\sum_{\tilde{\omega}} \mu(\tilde{\omega}) > 0$.*

As before, the marginal private value tends to be smaller for banks with intermediate initial exposures, who in equilibrium provide intermediation services. Therefore, our model implies that, all else equal, intermediaries are the most vulnerable to the negative shocks.

The endogenous probability of matching in the OTC market, $\alpha(\mu)$, creates strategic complementarities in exit decisions. When more banks exit, the measures μ decrease, and so the probability of matching, $\alpha(\mu)$, decreases as well. This reduces the value of remaining in the market to share risk, and so may foster more exit. We shall see in Section 6.2.1, in the context of a parametric example, that these strategic complementarities can create further multiplicity of equilibria.

5.3 Normative implications

We have seen so far that the analysis of equilibrium exit is formally very similar to the one of equilibrium entry. In this section we study the normative implications of exit and show that they can differ significantly from the normative implications of entry. Our main result is:

Lemma 14. *In the model with exit, the marginal social value of a type-ω bank is:*

$$MSV(\omega) = K(\omega) + \alpha(\mu)F(\omega) - \frac{\rho}{2}\bar{F},$$

where $K(\omega)$ is the competitive surplus, $F(\omega)$ is the frictional surplus, and $\bar{F}$ is the average frictional surplus.

There are two differences with the corresponding formula in the entry model, in equation (17). First, the frictional surplus is multiplied by $\alpha(\mu)$, since the matching probability after exit is less than one. Second, the average frictional surplus is multiplied by $\rho/2$. By contrast, if the matching probability after exit were exogenous, then the coefficient would be $\alpha(\mu)/2 > \rho/2$. Put differently, relative to a model in which the matching probability is exogenous, the endogeneity of the matching probability increases the marginal social value of keeping a bank in the market. This is because, if the bank were to exit, some of its counterparties would not be able to re-match. As a result, the difference between marginal social and private value,

$$\text{MSV}(\omega) - \text{MPV}(\omega) = \frac{1}{2}\left[\alpha(\mu)F(\omega) - \rho\bar{F}\right], \tag{22}$$

can now be positive even for pure intermediaries, i.e. banks such that $g(\omega) = \omega$, as long as the market is not resilient enough.

A commonly held view is that there is a policy role for subsidizing intermediaries during financial market disruptions. Indeed, intermediaries may exit too promptly in the face of negative shocks, which has adverse consequences because they are central counterparties to many other players in OTC markets. Accordingly, we have found that intermediaries have the lowest $\text{MPV}(\omega)$, and so the strongest private incentives to exit. Yet, our analysis reveals that these strong private incentives can be either smaller or larger than the corresponding social incentives. It all depends on market resilience, or traders' ability to quickly resume trading with alternate counterparties. If the market is very resilient, then private incentives are smaller than social incentives. To implement a social optimal amount of exit, a policy maker would need to tax intermediaries, not subsidize them. Subsidizing intermediaries is only warranted if the market is not too resilient. We shall see in Section 6.2.2, in the context of a parametric example, that the case for subsidizing intermediaries also depends on the size of the shock. Given any $\rho < 1$, we will obtain that $\text{MSV}(\omega) > \text{MPV}(\omega)$ for large enough shocks.

Finally, we note that, aside from the endogenous probability of trading, $\alpha(\mu)$, the planning problem in the case of exit is formally the same as in the case of entry. Therefore, the results of Theorem 11 continue to hold, after appropriate changes in the formula for marginal private and social values.

6 An illustrative example with three types

In this section we study a version of our model that can be solved analytically. Here we will focus on a few main findings, and we provide a complete analysis in Appendix B.

The model we consider from now on has three symmetric types: $\Omega = \{0, \frac{1}{2}, 1\}$. The distribution of banks in the economy at large is uniform, i.e. $\pi(\omega) = 1/3$ for all $\omega \in \Omega$, and the cost of risk bearing, $\Gamma[g]$, is quadratic. As we have shown before, this arises when the loss upon default, D, is normally distributed. This model is nested in the general framework we studied so far, so all our results apply. Moreover, in Appendix B.2 we show that, in this model, all equilibria and social optima are symmetric, i.e. the fraction of traders in the OTC market with extreme ω are equal so that $n(0) = n(1)$. This makes the model very tractable because the distribution of traders in the OTC market can now be parameterized by a single number, the fraction of traders in $\omega = \frac{1}{2}$ banks, $n(\frac{1}{2})$.

In this symmetric model, the equilibrium in the OTC market conditional on entry is very simple. For instance, in the case of partial risk sharing, the network of trades is shown in Figure 3 (for the other cases, see Appendix B.1). There are direct trades of size k between traders of type $\omega = 0$ and traders of type $\omega = 1$, and indirect trades of size k intermediated by traders of type $\omega = \frac{1}{2}$. Middle-ω banks are pure intermediaries: they meet $\omega = 0$ and $\omega = 1$ banks with equal probability, trade equal amounts, and so they do not change their exposures, $g(\frac{1}{2}) = \frac{1}{2}$.

6.1 Entry with three types

To study entry, we assume that the distribution of bank sizes is Pareto with parameter $1 + \theta > 0$, for some $\theta > 0$, over the support $[\underline{S}, \infty)$.[20] This implies that the measure of traders in banks with size greater than some $S \geq \underline{S}$ is $(S/\underline{S})^{-\theta}$. The associated distribution of per-capita entry costs has support $[0, \bar{z}]$, $\bar{z} \equiv c/\underline{S}$, and CDF $\Phi(z) = (z/\bar{z})^{\theta}$.

6.1.1 Positive results

The first proposition discusses entry patterns:

Proposition 15. *There exists a unique equilibrium with positive entry, μ_E. This equilibrium has the following properties:*

- *entry patterns of extreme-ω banks are symmetric: $\mu_E(0) = \mu_E(1)$;*

[20] Precisely, the fraction of banks with size larger than $S \geq \underline{S}$ is equal to $(S/\underline{S})^{-(1+\theta)}$. To ensure that there is a measure one of traders, we also need to assume that the total measure of bank establishments in the economy at large is $\frac{\theta}{1+\theta}\frac{1}{\underline{S}}$.

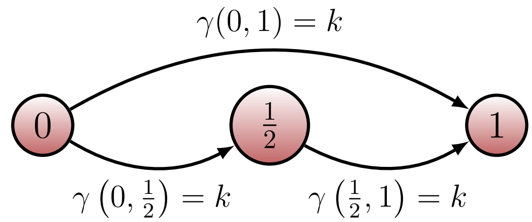

Figure 3: Trading patterns in the three-type model, under partial risk sharing.

- *if $k < 1$, there is partial risk sharing and $n_E(\frac{1}{2}) > 0$;*
- *if $k \geq 1$, there is full risk sharing and $n_E(\frac{1}{2}) = 0$.*

Since the lower bound of the support of the cost distribution is zero, we have that $\Phi(z) > 0$ for all $z > 0$, and so Theorem 6 implies that there always exists an equilibrium with strictly positive entry. Relative to our earlier general existence results, we learn that, in this parametric example, the equilibrium is unique. Given that entry is symmetric, the economic force underlying uniqueness is that the entry decisions of middle-ω banks are strategic substitutes: when more middle-ω banks enter, risk-sharing improves, intermediation profits are eroded, and so middle-ω banks have less incentives to enter.

The proposition also shows that, when trade size limits are small enough, $k < 1$, there is partial risk sharing and extreme-ω banks cannot fully equalize their exposures in the OTC market. As a result, there are intermediation profits to be made and some middle-ω banks operate in the market, $n_E(\frac{1}{2}) > 0$. When $k \geq 1$, there is full risk sharing and $n_E(\frac{1}{2}) = 0$.

Our positive findings below offer insights into the growth in OTC markets by contrasting two possible scenarios: a decline in the entry cost, c, representing more sophisticated trading technologies, or a relaxation of trade size limits, k, representing improvements in operational risk management practices. We first consider a comparative static with respect to the fixed cost of entry, c:

Corollary 16. *In equilibrium, when the fixed entry cost, c, decreases*

- *for all $\omega \in \Omega$, the measure of ω-traders, $\mu_E(\omega)$, increases continuously;*
- *the fraction of middle-ω trader, $n_E(\frac{1}{2})$, increases continuously.*

The first bullet point shows that an improvement in trading technologies, as proxied by a decline in the entry cost, leads to greater entry of all types of banks. The second bullet point shows that this leads to relatively more entry of middle-ω banks, who assume the role of intermediaries, than of extreme-ω banks, who assume the role of customers. This arises because of our assumed distribution of costs together with the observation that middle-ω banks have less incentives to enter than extreme-ω banks. Next, we consider a comparative static with respect to the trade size limit, k:

Corollary 17. *In equilibrium, when the trade size limit, k, increases*

- *the measure of middle-ω traders, $\mu_E(\frac{1}{2})$, varies non-monotonically. It increases with k when $k \simeq 0$, and it goes to zero as $k \to 1$.*
- *the fraction of middle-ω traders, $n_E(\frac{1}{2})$, varies non-monotonically. It is positive when $k = 0$ and equal to zero when $k = 1$. It decreases with k when $k \simeq 0$ and $\simeq 1$, but can increase with k otherwise.*

An increase in k has two opposite effects on middle-ω banks' entry incentives. On the one hand, there is a positive partial equilibrium effect: when k is larger, each trader in a given bank can increase the size of its position and thus earn larger profits. But, on the other hand, there is a general equilibrium effect: risk sharing improves, which reduces intermediation profits. The first effect dominates when $k \simeq 0$, increasing the measure of middle-ω traders. But the second effect dominates when $k \simeq 1$, decreasing the measure of middle-ω traders. To understand the effects on the *fraction* of middle-ω traders, note that, when $k \simeq 0$, an increase in k causes both middle-ω and extreme-ω traders to enter. But extreme-ω traders enter more, resulting in a decrease in $n_E(\frac{1}{2})$. When $k \simeq 1$, risk-sharing is almost perfect, the trading profits of middle-ω banks are close to zero, and so $n_E(\frac{1}{2})$ decreases towards zero.[21]

These comparative statics translate into predictions about the evolution of gross exposures and net exposures as frictions decrease.

Corollary 18 (Exposures)**.** *A reduction in frictions has the following effects on exposures:*

- *When c decreases, both the average gross exposure, $\mathcal{G}$, and the ratio of gross-to-net exposures, $\mathcal{R}$, increase.*
- *When k increases, $k \simeq 0$ or $k \simeq 1$, the average gross exposure increases, but the ratio of gross-to-net exposure, $\mathcal{R}$, decreases.*

Hence, in all cases, reducing frictions causes the market to grow, in the sense of increasing gross notional outstanding per capita. But predictions differ markedly in other dimensions. For example, when trading technologies improve, as proxied by a decline in c, the market grows through an increase in intermediation activity, and the gross-to-net exposure ratio increases. In contrast, when risk-management technologies improve, as proxied by an increase in k, the market can grow through an increase in customer-to-customer trades. As a result,

[21]When k goes to zero, the *measures* of all traders' types, μ_E, go to zero. However, they all do so at the same speed, so that in the $k \to 0$ limit the *fraction* of middle-ω traders, $n_E(\frac{1}{2})$, remains bounded away from zero.

the gross-to-net exposures ratio can decrease. Thus, according to the model, the evolution of the gross-to-net exposures ratio can help distinguish a decrease in frictions due to an improvement in trading versus risk-management technologies. These comparative statics can also help to guide policies aimed at reducing the importance of central dealers. Our analysis indicates that improvements to risk management technologies, rather than reductions in entry costs, are more likely to decrease intermediated trade and excess volume.

6.1.2 Normative results

One benefit of the three-types model is that the planner's solution can be characterized analytically as well. Thus, we can explicitly compare the planner's solution to the equilibrium.

Proposition 19. *In the three-types model, the planner's problem has a unique solution, μ_P, which is symmetric, $\mu_P(0) = \mu_P(1)$. In this solution, relative to the equilibrium:*

- *the fraction middle-ω traders is smaller: $n_P(\frac{1}{2}) \leq n_E(\frac{1}{2})$;*
- *the ratio of gross-to-net exposure, per capita, is smaller: $\mathcal{R}_P \leq \mathcal{R}_E$;*
- *the measure of extreme-ω traders is greater: $\mu_P(0) + \mu_P(1) \geq \mu_E(0) + \mu_E(1)$;*
- *if $\theta \leq 1$, the market is smaller: $\sum_\omega \mu_P(\omega) \leq \sum_\omega \mu_E(\omega)$.*

The main message of the proposition is that the planner chooses a market structure which generates a larger fraction of direct trades between extreme-ω banks, who assume the role of customers, and a smaller fraction of indirect trades intermediated by middle-ω banks, who assume the role of dealers. To do so, the planner lets more extreme-ω banks enter the OTC market and, when $\theta \leq 1$, less middle-ω banks. The net effect is, under the sufficient condition that $\theta \leq 1$, that market size and trading volume are smaller.[22]

The proposition shows that, in equilibrium, the market is too large due to excessive intermediation activity. While intermediation activity should not be fully eliminated, it should be reduced. Taxing banks who assume the role of dealers and subsidizing banks who assume the role of customers would shrink the market and reduce trading volume. Yet, welfare would increase because some intermediated trades would be replaced by direct trades with customers.

[22]Numerical computations, not reported here, suggests that this result can also hold for $\theta \geq 1$.

6.2 Exit with three types

Next, we turn to the model with exit. We assume that the distribution of traders' types in the market is arises from the entry model described above. We show in Appendix B.5 that this implies the restriction $\pi(\frac{1}{2}) \leq \min\{\frac{1}{3}, \frac{1}{k} - 1\}$. We assume that there is a strictly positive measure of dealers to start with, $\pi(\frac{1}{2}) > 0$, which implies that $k\left[1 + \pi(\frac{1}{2})\right] < 1$. We consider a simple discrete distribution of exit costs: all banks have to pay the same cost, z, per trader, in order to resume trading in the OTC market.

6.2.1 Positive results

Our main existence result is:

Proposition 20. *All equilibria are symmetric. If $\rho < 1$, for some z, there exist multiple equilibria with positive participation. Otherwise, if $\rho = 1$, for all z small enough, there exists a unique equilibrium with positive participation. Finally, for all z, the equilibrium with highest participation maximizes utilitarian welfare amongst all equilibria.*

The proposition shows that there can be multiple equilibria: this is due, as noted earlier, to strategic complementarities in exit decisions. When other banks exit, traders are more likely to end up without a counterparty, which increases an individual bank's incentives to exit. Multiple equilibria only arise when $\rho < 1$: thus, according to the model, coordination failures in OTC markets are more likely to arise in the short run, when these markets are not too resilient. The next proposition abstracts from these coordination failures and focuses on the equilibrium with highest utilitarian welfare:

Proposition 21. *There are three cost thresholds, $z_{1E} < z_{2E} < z_{3E}$, such that, in the highest welfare equilibrium:*

- *if $z \in [0, z_{1E}]$, all banks stay:*

$$\mu_E(0) = \pi(0),\ \mu_E(1) = \pi(1),\ \text{and } \mu_E(\tfrac{1}{2}) = \pi(\tfrac{1}{2});$$

- *if $z \in (z_{1E}, z_{2E})$, extreme-ω banks stay and middle-ω banks exit partially:*

$$\mu_E(0) = \pi(0),\ \mu_E(1) = \pi(1),\ \text{and } 0 < \mu_E(\tfrac{1}{2}) < \pi(\tfrac{1}{2});$$

- *if $z \in [z_{2E}, z_{3E}]$, extreme-ω banks stay and middle-ω exit fully:*

$$\mu_E(0) = \pi(0),\ \mu_E(1) = \pi(1),\ \text{and } \mu_E(\tfrac{1}{2}) = 0;$$

- *if $z > z_{3E}$: all banks exit fully:*

$$\mu_E(0) = \mu_E(\tfrac{1}{2}) = \mu_E(1) = 0.$$

Moreover, the measures of extreme-ω and middle-ω traders in the market are continuous and decreasing in z except when $\rho < 1$ at the threshold z_{3E}, where the measure of extreme-ω traders has a downward jump.

The thresholds appearing in the proposition have intuitive interpretations. For example, the first threshold, z_{1E}, is the lowest cost that makes a middle-ω bank indifferent between staying or not, when all other banks stay:

$$z_{1E} = \text{MPV}\left(\tfrac{1}{2}\right)\Big|_{\mu_E(0)=\pi(0),\, \mu_E(1)=\pi(1) \text{ and } \mu_E(1/2)=\pi(1/2)}.$$

When $z = z_{1E}$, $\text{MPV}(0) > \text{MPV}\left(\tfrac{1}{2}\right) = z$, and so all extreme-$\omega$ banks find it optimal to stay.

The proposition shows that middle-ω banks are the most vulnerable to shocks: for any z, a middle-ω bank is more likely to exit than an extreme-ω bank. The reason is, as before, that $\text{MPV}(\frac{1}{2}) < \text{MPV}(0) = \text{MPV}(1)$. We also learn from the proposition that the measure of customer banks is discontinuous in z at $z = z_{3E}$ when $\rho < 1$. This is expected given multiple equilibria: selections of the equilibrium map typically have discontinuities. What is perhaps more interesting is that this discontinuity occurs for relatively large shocks, that it goes downwards so it represents a sudden dry up of trading activity and not a boom, and that it is characterized by a sudden withdrawal of customer banks once all dealer banks have exited. Thus, according to the model, the exit of dealers from the market is associated with more volatile levels of trading activity.

6.2.2 Normative results

In the appendix, we solve the planning problem fully. We show in particular that, in several dimensions, the socially optimal exit patterns resemble the ones that arise in equilibrium. In particular, the exit patterns are characterized by three cost thresholds, just as in Proposition 21, and there also exists a cost threshold at which the measure of customers banks falls discretely. However, the cost thresholds for the equilibrium are different from the one arising in the planner problem. This implies that the planner will choose exit patterns that differ from the equilibrium. In particular, we find:

Proposition 22. *Assume that $\rho < 1$ and let $\mu_E(\frac{1}{2})$ and $\mu_P(\frac{1}{2})$ denote the measure of middle-ω traders who stay in the OTC market, in equilibrium and in the planner's problem. Then, there are three thresholds $0 < z_1 \leq z_2 < z_3$ such that:*

- *if $z \leq z_1$, $\mu_E(\frac{1}{2}) = \mu_P(\frac{1}{2}) = \pi(\frac{1}{2})$;*

- *if* $z \in (z_1, z_2)$, $\mu_E(\frac{1}{2}) > \mu_P(\frac{1}{2})$;
- *if* $z \in (z_2, z_3)$, $\mu_E(\frac{1}{2}) < \mu_P(\frac{1}{2})$;
- *if* $z > z_3$, $\mu_E(\frac{1}{2}) = \mu_P(\frac{1}{2}) = 0$.

Moreover, $z_1 = z_2$ *if* $\rho\left[1 + \pi(\frac{1}{2})\right] \leq 1$.

The proposition shows that from the planner's perspective there is too little exit if the shock is small, $z \in (z_1, z_2)$, and too much exit when the shock is large, $z \in (z_2, z_3)$. Therefore, relative to our earlier result, we find that the case for subsidizing intermediaries not only depends on market resilience, but also on the size of the shock. Subsidizing intermediaries is welfare improving only when the shock is large enough. The intuition for this result is that middle-ω banks are socially more useful when only a few of them remain in the market, which happens when the shock is large.

7 Conclusion

Several observations regarding entry and trading patterns OTC markets for derivatives have recently received considerable attention from policy makers and the public alike. First, the large volume of bilateral trades at varied prices creates an intricate liability structure between participating banks. Second, the largest banks in these markets act as "dealers" by engaging in intermediation activity: they trade contracts in both directions, long and short, and they have very large gross positions relative to their net positions. In contrast, medium-sized banks act as "customers": they trade mostly in one direction, either long or short, so their gross and net positions are close to each other. These trading patterns have raised a number of concerns. Why is much of the intermediation activity provided by large banks? Do they provide too little or too many intermediation services? Does the concentration of intermediation activity in these large banks make OTC markets excessively vulnerable to negative shocks?

In this paper, we highlight the role of entry and exit incentives in shaping the structure and ultimate resilience of OTC markets. We develop a parsimonious model of OTC markets with entry costs and trade size limits to illustrate why only large enough banks enter the market as "dealers," while medium-sized banks enter mainly as "customers." Although we endow all banks with the same trading technology, heterogeneity in trading patterns arises endogenously from banks' varied incentives to trade to hedge their risk exposure vs. to make intermediation profits. Imperfect hedging in equilibrium leads to the price dispersion from which dealers derive their profits. We show that dealers play a socially valuable role in mitigating OTC market frictions. However, they tend to provide too many intermediation services relative to the social optimum: they have a "business stealing" motive for entry. Finally, in the short run, we find that banks who act as dealers are in fact the most vulnerable to negative shocks and have the strongest incentives to exit. Whether these large banks exit more than is socially optimal, however, depends crucially on market resilience. That is, given all other parameters, OTC market participants must find it sufficiently difficult to engage with new counterparties for it to be optimal to subsidize the continued participation of large dealers.

References

Gara Afonso. Liquidity and congestion. *Journal of Financial Intermediation*, 20:324–360, 2011. 5

Gara Afonso and Ricardo Lagos. Trade dynamics in the federal funds market. Working paper, Federal Reserve Bank of New York and NYU, February 2012. 5

Franklin Allen and Douglas Gale. Financial contagion. *Journal of Political Economy*, 1: 1–33, 2000. 5

David Andolfatto. Business cycles and labor-market search. *American Economic Review*, 86 (1):112–132, 1996. 11

Andrew Atkeson, Andrea Eisfeldt, and Pierre-Olivier Weill. The market of otc credit derivatives. Working paper, UCLA, 2012. 6, 15, 22

Ana Babus and Péter Kondor. Trading and information diffusion in over-the-counter markets. Working paper, Federal Reserve Bank of Chicago Central European University, 2013. 5

Bruno Biais, Johan Hombert, and Pierre-Olivier Weill. Equilibrium pricing and trading volume under preference uncertainty. *Review of Economic Studies*, Forthcoming, 2014. 5

Patrick Bolton, Tano Santos, and José Scheinkman. Cream skimming in financial markets. Working paper, Columbia University, 2012. 5

Briana Chang. Adverse selection and liquidity distortion in decentralized markets. Working paper, University of Wisconsin, 2012. 5

Randall Dodd. Markets: Exchange or over-the-counter, March 2012. URL `http://www.imf.org/external/pubs/ft/fandd/basics/markets.htm`. 29

Darrell Duffie and Gustavo Manso. Information diffusion in large population. *American Economic Review, Papers and Proceedings*, 97, 2007. 5

Darrell Duffie and Bruno Strulovici. Capital mobility and asset pricing. *Econometrica*, 80: 2469–2509, 2012. 5

Darrell Duffie, Nicolae Gârleanu, and Lasse H. Pedersen. Securities lending, shorting, and pricing. *Journal of Financial Economics*, 66:307–339, 2002. 4

Darrell Duffie, Nicolae Gârleanu, and Lasse H. Pedersen. Over-the-counter markets. *Econometrica*, 73:1815–1847, 2005. 4, 22

Darrell Duffie, Semyon Malamud, and Gustavo Manso. Information percolation with equilibrium search dynamics. *Econometrica*, 77:1513–1574, 2009. 5

Larry Eisenberg and Thomas H. Noe. Systemic risk in financial systems. *Management Science*, February:236–249, 2001. 5

Matthew Elliott, Benjamin Golub, and Matthew O. Jackson. Financial network and contagion. *American Economic Review*, Forthcoming, 2014. 5

Nicolae Gârleanu. Portfolio choice and pricing in illiquid markets. *Journal of Economic Theory*, 144:532–564, 2009. 5

Alessandro Gavazza. An empirical equilibrium model of a decentralized asset market. *Working Paper*, 2013. 5

Mikhail Golosov, Guido Lorenzoni, and Aleh Tsyvinski. Decentralized trading with private information. *Econometrica*, 83:1055–1091, 2014. 5

Veronica Guerrieri and Robert Shimer. Dynamic adverse selection: A theory of illiquidity, fire sales, and flight to quality. Working paper, University of Chicago, 2012. 5

Zhiguo He and Konstantin Milbradt. Endogenous liquidity and defaultable debt. *Econometrica*, 82:1443–1508, 2014. Working paper, Chicago Booth and MIT. 5

Patrick Kiefer. A search-based pricing model of credit default swaps. Working paper, UCLA, 2010. 4

Ricardo Lagos and Guillaume Rocheteau. Liquidity in asset markets with search frictions. *Econometrica*, 77:403–426, 2009. 5

Ricardo Lagos, Guillaume Rocheteau, and Pierre-Olivier Weill. Crises and liquidity in OTC markets. *Journal of Economic Theory*, 146:2169–2205, 2011. 5

Robert E. Jr. Lucas. Liquidity and interest rates. *Journal of Economic Theory*, 50:237–264, 1990. 11

Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, Oxford, 1995. 44, 45

Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70:583–601, 2002. 22, 57, 58

Emiliano Pagnotta and Thomas Philippon. Competing on speed. Working Paper, NYU Stern School of Business, 2011. URL `http://ssrn.com/abstract=1967156`. 5

Francesco Saita. *Value at Risk and Bank Capital Management*. Academic Press, 2007. 8

Shouyong Shi. A divisible search model of fiat money. *Econometrica*, 65:75–102, 1997. 5, 11

Robert Shimer. *Labor Markets and Business Cycles (CREI Lectures in Macroeconomics)*. Princeton University Press, 2010. 11

Nancy L. Stokey and Robert E. Lucas. *Recursive Methods in Economic Dynamics*. Harvard University Press, 1989. 15, 57

Renee Stulz. Credit default swaps and the credit crisis. *Journal of Economic Perspectives*, 24:73–92, 2011. 5

Alberto Trejos and Randall Wright. Shi-Trejos-Wright Meets Duffie-Garleanu-Pedersen. Working paper, University of Wisconsin, 2012. 5

Dimitri Vayanos and Tan Wang. Search and endogenous concentration of liquidity in asset markets, working paper. *Journal of Economic Theory*, 136:66–104, 2007. 4

Dimitri Vayanos and Pierre-Olivier Weill. A search-based theory of the on-the-run phenomenon. *Journal of Finance*, 63:1361–1398, 2008. 5

Pierre-Olivier Weill. Leaning against the wind. *Review of Economic Studies*, 74:1329–1354, 2007. 5

Pierre-Olivier Weill. Liquidity premia in dynamic bargaining markets. *Journal of Economic Theory*, 140:66–96, 2008. 5

Adam Zawadowski. Entangled financial systems. Working paper, Boston University, 2013. 5