Appendix to:

Should Retail Investors' Leverage Be Limited?

(intended for online publication)

Appendix A.1: Currency Pairs Affected by the CFTC Regulation

Table A.1
The CFTC regulation and Leverage-Constraints across Currency Pairs
Description: This table lists the currency pairs affected by the CFTC trading rule restricting the amount of leverage at 50:1 or 20:1.

50:1 leverage				
USD/JPY	AUD/NZD	NZD/CAD	EUR/GBP	GBP/USD
USD/CHF	USD/SEK	CHF/JPY	EUR/JPY	GBP/JPY
AUD/USD	USD/DKK	CAD/JPY	EUR/AUD	GBP/CHF
USD/CAD	USD/NOK	CAD/CHF	EUR/CAD	GBP/CAD
NZD/USD	AUD/CHF	CHF/SEK	EUR/SEK	GBP/NZD
AUD/CAD	NOK/JPY	CHF/NOK	EUR/NOK	GBP/AUD
AUD/JPY	SEK/JPY	EUR/USD	EUR/NZD	GBP/SEK
NZD/JPY	NZD/CHF	EUR/CHF	EUR/DKK	

20:1 leverage				
USD/MXN	USD/CZK	USD/HKD	USD/RUB	ZAR/JPY
EUR/PLN	USD/ZAR	SGD/JPY	EUR/HUF	
USD/PLN	USD/SGB	USD/TRY	USD/HUF	
EUR/CZK	HKD/JPY	EUR/TRY	TRY/JPY	

Appendix A.2: A Stylized Model of the Retail Forex Market

In this appendix, we present a stylized model that motivates our empirical analysis and enables us to evaluate social welfare. We first describe the environment and characterize the equilibrium together with various variables of interest. We then analyze how the leverage-constraint policy affects the equilibrium variables as well as social welfare.

Environment and equilibrium

Agents, information structure, and beliefs.. Consider an economy with a single consumption good (which will be referred to as a dollar) and a single trading period. There is a risk-free asset with gross return normalized to one. There is also a single risky asset (a currency). In the interbank market, the asset currently trades at an exogenous market price normalized to one, $p_0 = 1$. At the next period, the asset will trade at price, p_1, which is a random variable. The objective distribution of the asset price change (that will be reflected in our empirical analysis) is given by $p_1 - 1 \sim N\left(\mu_s^{true}, \sigma^2\right)$. Here, $s \in S$ denotes an aggregate state realized at the beginning of date 0. We let q_s denote the ex-ante probability of the aggregate state (according to each agent) and assume $\sum_s q_s \mu_s^{true} = 0$ so that the price is a martingale under the objective belief.

There are traders, denoted by $i \in I$, that take optimal positions in the asset that will be described below. For simplicity, we normalize the mass of traders to one so that the aggregate and the per-trader outcomes are the same. Traders have dogmatic beliefs and do not learn from prices (formally, traders know each others' beliefs and agree to disagree). Traders' beliefs can also depend on the aggregate state $s \in S$. Specifically, trader i believes the price change is distributed according to, $p_1 - 1 \sim N\left(\mu_s^i, \sigma^2\right)$. Since the objective belief is unique, the heterogeneity in traders' beliefs can be thought of as capturing various behavioral distortions (which we leave unmodeled for simplicity). On the other hand, the dependence of traders' beliefs on the aggregate state allows traders also to be somewhat informed. In particular, to the extent that a trader's belief and the objective belief (μ_s^i and μ_s^{true}) are positively correlated, the trader's positions will tend to generate positive expected return before transaction costs, which we refer to as "information."

Trader i also starts with initial initial wealth given by, n_0^i, and has CARA preferences with coefficient of absolute risk aversion, γ^i. The type of trader i is given by the parameters, $\left(\gamma^i, n_0^i, \mu_s^i\right)$. We let $dF_s\left(\gamma^i, n_0^i, \mu_s^i\right)$ denote the joint distribution function over trader types conditional on the aggregate state. We can be quite general about the shape of this distribution except for a technical condition that we note below. All agents know and agree upon the type distribution, $dF_s\left(\gamma^i, n_0^i, \mu^i\right)$, as well as the probability of aggregate states, $\left(q_s\right)_s$. Their disagreements concern he asset's expected payoff.

There is also a competitive retail brokerage sector that provides intermediation services. Consider a single (representative) broker. For simplicity, the broker is risk neutral and she has the objective belief about the asset payoff. In particular, she believes the price change is distributed according to $p_1 - 1 \sim N\left(\mu_s^{true}, \sigma^2\right)$ conditional on the aggregate state $s \in S$. However, the broker does not observe the aggregate state, and it sets bid and ask prices at the beginning of the period before she can observe endogenous signals about the aggregate state (such as aggregate trading volume). Since traders' beliefs depends on the aggregate state, this might put the broker at an informational disadvantage relative to traders. As in Glosten and Milgrom (1985), the broker will set bid and ask prices that take into account the information content of traders' orders. For simplicity, we assume the broker sets a single bid price and a single ask price, p_0^{bid} and p_0^{ask}, and stands ready to fill sell and buy orders linearly at these prices regardless of the size of the order.[22] We will make assumptions so that, similar to Glosten and Milgrom (1985), the equilibrium bid price will be lower than the ex-ante objective value of the broker (normalized to one) which in turn will be lower than the equilibrium ask price, $p_0^{bid} < 1 < p_0^{ask}$.

Traders' optimal positions.. Trader i takes the bid and ask prices as given and decides to take a long or short position in the risky asset denoted by x_s^i. She invests her residual wealth in the risk-free asset. She can also use leverage on long or short positions without any additional fees but that might be subject to a regulatory limit. Specifically, we require the position (evaluated at the market value) to satisfy, $\left|x_s^i\right| \leq \bar{l} n^i$ where $\bar{l}$ is an exogenous leverage limit

[22] In general, the size of the order can also contain some information about the aggregate state (e.g., larger orders might be associated with better information), and the broker might want to set size-dependent prices that reflect this information. Modeling this feature explicitly could generate additional interesting predictions but it wouldn't change our qualitative conclusions. We therefore restrict attention to linear prices and simplify the analysis.

set by regulation. We also allow for the case $\bar{l} = \infty$, which corresponds to the equilibrium without leverage restriction.[23] The trader's portfolio problem can be written as,

$$\max_{x_s^i \in \left[-\bar{l} n_0^i, \bar{l} n_0^i\right]} E_s^i \left[-\exp\left(n_1^i\right)\right] \text{ where } n_1^i = \begin{cases} n_0^i + x_s^i \left(p_1 - p_0^{ask}\right) & \text{if } x_s^i > 0 \\ n_0^i + x_s^i \left(p_1 - p_0^{bid}\right) & \text{if } x_s^i < 0 \end{cases} \tag{8}$$

In view of the CARA-Normal setup, the trader's optimal position (conditional on the aggregate state realization) is given by,

$$x_s^i = \begin{cases} \min\left(\frac{\mu_s^i - p_0^{ask}}{\gamma^i \sigma^2}, \bar{l} n_0^i\right), & \text{if } \mu_s^i > p_0^{ask} \\ 0, & \text{if } \mu_s^i \in \left(p_0^{bid}, p_0^{ask}\right) \\ \max\left(-\bar{l} n_0^i, \frac{\mu_s^i - p_0^{bid}}{\gamma^i \sigma^2}\right), & \text{if } \mu_s^i < p_0^{bid} \end{cases}. \tag{9}$$

The broker's problem and bid-ask spreads.. The broker is subject to two types of costs. First, as we already mentioned, the broker can be subject to informational costs since traders might on average have some information. The broker takes the opposite side of traders' (possibly informed) positions, and keeps the positions on its balance sheet, which exposes it to potential losses.[24] Second, the broker also incurs technological costs that capture the infrastructure and the employees utilized to facilitate intermediation. For simplicity, we assume these costs grow linearly in the size of traders' positions: specifically, intermediating each unit of long or short position costs the broker $c > 0$ additional dollars. Using these assumptions, the broker's expected certainty equivalent wealth (under its objective belief) conditional on the aggregate state is given by,

$$CE_s^b = \begin{array}{l} \int_{i, x_s^i > 0} -x_s^i \left(E_s^{true}\left[p_1\right] - \left(p_0^{ask} - c\right)\right) dF_s\left(\gamma^i, n_0^i, \mu_s^i\right) \\ + \int_{i, x_s^i < 0} -x_s^i \left(E_s^{true}\left[p_1\right] - \left(p_0^{bid} + c\right)\right) dF_s\left(\gamma^i, n_0^i, \mu_s^i\right) \end{array}. \tag{10}$$

We assume there are a large number of identical brokers that compete a la Bertrand to set bid and ask prices, p_0^{bid} and p_0^{ask}. Competition drives down the broker's expected profit from both buy and sell orders to zero, that is,

$$\sum_{s \in S} q_s \int_{i, x_s^i < 0} -x_s^i \left(E_s^{true}\left[p_1\right] - \left(p_0^{ask} - c\right)\right) dF_s\left(\gamma^i, n_0^i, \mu_s^i\right) = 0,$$

$$\text{and } \sum_{s \in S} q_s \int_{i, x_s^i < 0} -x_s^i \left(E_s^{true}\left[p_1\right] - \left(p_0^{bid} + c\right)\right) dF_s\left(\gamma^i, n_0^i, \mu_s^i\right) = 0.$$

This also implies that the broker's total expected profit is zero, $\sum_{s \in S} q_s CE_s^b = 0$. After rearranging these expressions and using $E_s^{true}\left[p_1\right] = 1 + \mu_s^{true}$, we obtain,

$$p_0^{ask} = 1 + m^{long} + c \text{ and } p_0^{bid} = 1 - m^{short} - c \tag{11}$$

where

$$m^{long} = \frac{E\left[x_s^i \mu_s^{true} | x^i > 0\right]}{E\left[x_s^i | x_s^i > 0\right]} \text{ and } m^{short} = \frac{E\left[x_s^i \mu_s^{true} | x_s^i < 0\right]}{E\left[-x_s^i | x_s^i < 0\right]}. \tag{12}$$

Here, the expectation operator $E[\cdot]$ is taken with respect to the distributions dF_s and q_s (on which there is no disagreement). The terms m^{long} and m^{short} reflect traders' average information: their expected profit per unit position on respectively long and short trades. In particular, m^{long} is positive if the traders on average purchase the asset when it has a positive expected return. Likewise, m^{short} is positive if the traders' on average sell the asset when it has a negative expected return.

[23] In practice, there might also be endogenous restrictions on the leverage ratio as in Geanakoplos (2009) or Simsek (2013). We abstract away from these endogenous leverage limits since they do not affect our qualitative results.

[24] One could wonder whether the broker could avoid this outcome by outlaying the position immediately to the interbank market. Our empirical analysis shows that the bid-ask spreads in the interbank market are on average very similar to the bid-ask spreads in the retail market. This means that outlaying the position to the interbank market is on average not profitable, arguably because similar intermediation costs also apply in the interbank market.

Eq. (11) says that, similar to Glosten and Milgrom (1985), the broker takes into account the information content in buy and sell orders. If m^{long} is positive, then the broker that receives a buy order faces adverse selection. In equilibrium, it increases the ask price so as to break even (otherwise, it would consistently lose money). Symmetrically, if m^{long} is negative, then the broker faces an advantageous selection and lowers its ask price (due to competitive pressure) while still breaking even. Similar considerations explain the relationship between traders' market-timing profit on the short trades, m^{short}, and the broker's bid price.

Definition of equilibrium.. The equilibrium in this model is a collection, $\left(\left(p_0^{ask}, p_0^{bid}\right), \left(\left(x_s^i\right)_{i\in I}\right)_{s\in S}\right)$, such that the positions satisfy (9) given the bid-ask prices, and the bid-ask prices satisfy (11) given the positions and the cumulative distribution function $F_s\left(\gamma^i, n_0^i, \mu_s^i\right)$. We assume there exists a unique equilibrium that also satisfies the inequality, $p_0^{bid} < 1 < p_0^{ask}$ (the bid price is lower than the ex-ante expected payoff which is lower than the ask price). This would be the case under a mild technical assumption on the distribution F_s.[25]

Trading volume.. We next characterize traders' expected profit as well as their expected utility and the social welfare. As we will see, trading volume plays a central role in these characterizations. Therefore, we define the long, the short, and the total trading volume as respectively,

$$V^{long} = E\left[x_s^i | x_s^i > 0\right] \left\{\sum_{s\in S} q_s \int_{i.x_s^i>0} dF_s\left(\gamma^i, n_0^i, \mu_s^i\right)\right\} \tag{13}$$

$$V^{short} = E\left[-x_s^i | x_s^i < 0\right] \left\{\sum_{s\in S} q_s \int_{i.x_s^i<0} dF_s\left(\gamma^i, n_0^i, \mu_s^i\right)\right\}$$

$$\text{and} \qquad V = V^{long} + V^{short}.$$

Here, the terms in set brackets capture the fraction of traders that take respectively long or short positions.[26] The expressions illustrate that the trading volume reflects the fraction of long or short trades as well as the expected size of each trade.

Traders' expected profit.. Under the objective distribution, trader i's expected profit is given by, $\sum_{s\in S} q_s x_s^i \left(E_s^{true}\left[p_1\right] - p_0^{ask}\right)$, if he takes a long position, and a similar expression if he takes a short position. Aggregating these positions, traders' overall expected profit is given by,

$$\Pi = \begin{array}{c} \sum_{s\in S} q_s \int_{i,x_s^i>0} x_s^i \left(E_s^{true}\left[p_1\right] - p_0^{ask}\right) dF_s\left(\gamma^i, n_0^i, \mu_s^i\right) \\ +\sum_{s\in S} q_s \int_{i,x_s^i<0} x_s^i \left(E_s^{true}\left[p_1\right] - p_0^{bid}\right) dF_s\left(\gamma^i, n_0^i, \mu_s^i\right) \end{array}.$$

After substituting $E_s^{true}\left[p_1\right] = 1 + \mu_s^{true}$, together with the definitions of the volume and market-timing profit in Eqs. (12) and (13), we can further rewrite this as,

$$\Pi = V^{long}\left(1 + m^{long} - p_0^{ask}\right) + V^{short}\left(p_0^{bid} - \left(1 - m^{short}\right)\right) \tag{14}$$

The intuition behind this expression is that the typical long position pays $1 + m^{long}$ and costs the ask price, p_0^{ask}. Likewise, the typical short position pays the bid price, p_0^{bid}, and it costs $1 - m^{short}$. The expression illustrates that the traders' expected profit is increasing in their average information and decreasing in the bid-ask spreads.

In equilibrium, the bid and ask prices are given by Eq. (11). Substituting this into Eq. (14), traders' expected profit in equilibrium is given by,

$$\Pi = -cV^{long} - cV^{short} = -cV. \tag{15}$$

[25] To illustrate this, suppose $F_s\left(\gamma^i, n^i, \mu_s^i\right)$ is independent of s. In this case, traders' positions contain no information about the aggregate state, the information terms drop out of (11). Then, there exists a unique equilibrium which also satisfies the inequality, $p_0^{bid} < 1 < p_0^{ask}$ (since $c > 0$). By a continuity argument, there exists a unique equilibrium that satisfies the same inequality as long as the dependence of the distribution $F_s\left(\gamma^i, n^i, \mu_s^i\right)$ on the aggregate state s is sufficiently small. We could parameterize this dependence and formalize the assumption but this is not necessary for our purposes.

[26] These two fractions do not necessarily sum to one since there are also traders that take a zero position (see Eq. (9)).

That is, the equilibrium profit depends negatively on the trading volume. Intuitively, since the competitive broker breaks even, the technological intermediation costs are ultimately passed through to traders via bid-ask spreads. The more traders trade, the more they incur these costs. Perhaps more surprisingly, traders' average information does not affect their equilibrium profit. The intuition is that the market maker sets bid and ask prices to neutralize information. For instance, if the traders' average information improves, then the market maker widens the bid-ask spreads (otherwise, it would consistently make losses and go out of business). Once the broker adjusts, the improved information does not affect traders' profits but it is reflected in bid and ask prices.

The broker's expected revenue and size.. Recall that the broker breaks even in equilibrium. In particular, its expected intermediation revenues are equal to the technological intermediation costs, cV. Recall that we view these costs as capturing the infrastructure and the labor the brokerage employees. Hence, the brokerage's intermediation revenues and size depend positively on the trading volume.

Belief-neutral social welfare.. We next characterize the social welfare in equilibrium. Since there are heterogeneous beliefs about the asset payoff, the social welfare will generally depend on the belief used to calculate agents' utilities. The standard Pareto welfare criterion would correspond to maximizing each agent's utility under her own belief. However, it is unclear whether perceived gains from speculation should be counted towards social welfare since they capture a collective form of irrationality: while all agents believe they have the correct belief, at most one of them could be right.

An alternative is to evaluate investors' beliefs under the objective belief distribution (which in this model corresponds to the broker's belief distribution). While appropriate, this approach faces a challenge in practice: The planner might not know who has the correct belief. Following Brunnermeier et al. (2014), we instead assume the planner evaluates the welfare under a fixed belief h, but she also makes the welfare comparisons robust to the choice of the belief. Specifically, we allow h to be an arbitrary convex combination of the traders' beliefs or the broker's (objective) belief.

We also focus on a utilitarian social planner that maximizes the sum of agents' certainty-equivalent wealth,

$$W^h = \sum_{s \in S} q_s \left(CE_s^{b,h} + \int_i CE_s^{i,h} dF_s\left(\gamma^i, n_0^i, \mu_s^i\right)\right). \tag{16}$$

Here, $CE_s^{b,h}$ denotes the broker's certainty-equivalent payoff and $CE_s^{i,h}$ denotes trader i's certainty-equivalent payoff conditional on the aggregate state. In view of the CARA-Normal setting, restricting attention to traders' certainty-equivalent payoffs is without loss of generality. Assigning all traders as well as the broker the same Pareto weight is slightly more restrictive but it provides a natural benchmark.[27]

Combining Eqs. (9) and (8), trader i's certainty-equivalent payoff under belief h can be calculated as,

$$CE_s^{i,h} = n_0^i + \begin{cases} x_s^i \left(E_s^h\left[p_1\right] - p_0^{ask}\right) - \frac{1}{2}\gamma^i \left(x_s^i\right)^2 \sigma^2, & \text{if } \mu_s^i > p_0^{ask} \\ 0 & \text{if } \mu_s^i \in \left(p_0^{bid}, p_0^{ask}\right) \\ x_s^i \left(E^h\left[p_1\right] - p_0^{bid}\right) - \frac{1}{2}\gamma^i \left(x_s^i\right)^2 \sigma^2, & \text{if } \mu_s^i < p_0^{bid} \end{cases}. \tag{17}$$

Thus, traders' certainty-equivalent payoff reflects their expected profits under belief h as well as their risk aversion and portfolio variance. Likewise, the broker's certainty-equivalent payoff under belief h can be calculated as,

$$CE_s^{b,h} = \begin{array}{l} \int_{i, x_s^i > 0} -x_s^i \left(E_s^h\left[p_1\right] - \left(p_0^{ask} - c\right)\right) dF_s\left(\gamma^i, n_0^i, \mu_s^i\right) \\ + \int_{i, x_s^i < 0} -x_s^i \left(E_s^h\left[p_1\right] - \left(p_0^{ask} - c\right)\right) dF_s\left(\gamma^i, n_0^i, \mu_s^i\right) \end{array}. \tag{18}$$

This is similar to Eq. (10) with the difference that the expected asset payoff is calculated according to a general belief h (which is not necessarily the true belief).

[27] In fact, this assumption is also without loss of generality as long as we allow the planner to do one-time ex-ante transfers among the agents. In this case, an allocation x that leads to greater utilitarian welfare, W^h, than another allocation y can be also made to Pareto dominate the allocation y (under belief h) after combining it with appropriate ex-ante transfers.

Combining Eqs. (16), (17), and (18), we can calculate the social welfare as,

$$W^h = -cV + \sum_{s\in S} q_s \int_i \left(n_0^i - \frac{1}{2}\gamma^i \left(x_s^i\right)^2 \sigma^2\right) dF_s\left(\gamma^i, n_0^i, \mu_s^i\right)$$
$$= E\left[n_0^i\right] - cV - \frac{1}{2}E\left[\gamma^i \left(x_s^i\right)^2 \sigma^2\right]. \quad (19)$$

Here, the expectation operators in the second line are taken with respect to the distributions dF_s and q_s (on which there is no disagreement). Hence, Eq. (19) illustrates that the welfare does not depend on the belief h used for the calculation (i.e., the expected price, $E_s^h[p_1]$, drops out of the welfare calculations). This is because, under any fixed belief h, the expected gain of an agent is the expected loss of another agent. This captures the idea that speculation transfers wealth among agents without creating social value. Once properly accounted for, these transfers do not affect social welfare.[28] As in Brunnermeier et al. (2014), the planner can evaluate the effect of speculation on social welfare without taking a stand on whose belief is correct. We refer to $W \equiv W^h$ as the belief-neutral welfare.

Eq. (19) also illustrates that the belief-neutral welfare is decreasing in the expected intermediation costs, $-Vc$, as well as their expected (risk-aversion weighted) portfolio variance, $E\left[\gamma^i \left(x^i\right)^2 \sigma^2\right]$. Intuitively, every intermediated position requires technological costs, which reduces social welfare as the resources or people used for intermediation could also be used elsewhere. These costs are naturally increasing in trading volume. In addition, to the extent that speculation induces investors to take riskier positions, the resulting portfolio risks also reduce social welfare.

Comparative statics of the leverage-constraint policy

We next characterize the effect of the leverage restriction policy on the equilibrium variables. It is useful to break this exercise into two steps: a partial equilibrium exercise in which brokers' bid-ask spread remain at their pre-policy levels, and a general equilibrium exercise in which the spreads also adjust. In practice, brokers are unlikely to change their spreads in the very short run (which we view as a month or so) due to inertia or optimization frictions.[29] Hence, we view our short-run empirical results as testing the partial equilibrium predictions. In the longer run (which we view as several months), brokers would arguably adjust their bid-ask spreads to their new equilibrium levels. Thus, we view our longer-run empirical results as testing the general equilibrium predictions. We denote the partial equilibrium with hatted variables, and the general equilibrium (after the policy change) with starred variables.

Partial equilibrium effects on trading volume.. Before the leverage-constraint policy, traders' positions are given by Eq. (9) with $\bar{l} = \infty$ and the volume is given by Eq. (13). Now suppose the leverage-constraint policy is imposed. In partial equilibrium, traders' positions are still given by Eq. (9) but with a finite $\bar{l}$ (but still evaluated at the same bid and ask prices). That is, we have,

$$\begin{cases} \hat{x}_s^i = \bar{l} n_0^i < x_s^i, & \text{if } x_s^i = \frac{\mu_s^i - p_0^{ask}}{\gamma^i \sigma^2} > \bar{l} n_0^i \\ \hat{x}_s^i = -\bar{l} n_0^i > -x_s^i, & \text{if } x_s^i = \frac{\mu_s^i - p_0^{bid}}{\gamma^i \sigma^2} < \bar{l} n_0^i \end{cases}.$$

In particular, the long and short positions that violate the leverage-constraint are downscaled to satisfy the leverage-constraint. Thus, the expected size of the long and short positions both decline. By Eq. (13), the trading volumes decline, that is,

$$\hat{V}^{long} \le V^{long}, \hat{V}^{short} \le V^{short} \text{ and } \hat{V} \le V.$$

Partial equilibrium effects on portfolio risks.. In partial equilibrium, the average portfolio risks decline, $E\left[\gamma^i \left(\hat{x}_s^i\right)^2 \sigma^2\right] \le E\left[\gamma^i \left(x_s^i\right)^2 \sigma^2\right]$, since the risky positions that violate the leverage-constraint are reduced, $\hat{x}_s^i \le x_s^i$.

[28] Likewise, the bid and ask prices, p_0^{ask} and p_0^{bid}, do not affect social welfare since they represent transfers between the traders and the brokers. In particular, Eq. (19) would apply not only with the equilibrium bid and ask prices given by Eq. (11)—which ensure that brokers break even, but also with other bid and ask prices that might generate net profits or net losses for the brokers.

[29] Since the leverage-constraint changes the set of trades the broker intermediates, it might take a while for the broker to figure out its overall profits and losses in this new market, and to adjust its bid-ask spreads appropriately.

Partial equilibrium effects on traders' expected profit.. In partial equilibrium, Eq. (14) still applies and implies that traders' expected profit becomes,

$$\hat{\Pi} = \hat{V}^{long}\left(1 + \hat{m}^{long} - p_0^{ask}\right) + \hat{V}^{short}\left(p_0^{bid} - \left(1 - \hat{m}^{short}\right)\right).$$

However, the bid-ask spreads are still at their old equilibrium levels,

$$p_0^{ask} = 1 + m^{long} + c \text{ and } p_0^{bid} = 1 - m^{short} - c.$$

Combining these expressions with Eq. (15), which applies before the policy, we obtain,

$$\hat{\Pi} - \Pi = -c\left(\hat{V} - V\right) + \hat{V}^{long}\left(\hat{m}^{long} - m^{long}\right) + \hat{V}^{short}\left(\hat{m}^{short} - m^{short}\right).$$

Here, the first term captures the effect of the constraint via trading volume. Since $\hat{V} \leq V$, the leverage-constraint tends to improve traders' profits through its effect on volume. The second and the third terms capture the effect via changes in average information. To the extent that the leveraged positions are associated with a different level of information than other positions, then the policy would also affect traders' (partial equilibrium) profit by improving or worsening their average information.

General equilibrium effects on bid-ask spreads.. In general equilibrium, bid and ask prices adjust to neutralize the changes in traders' average information. More specifically, Eq. (11) implies,

$$p_0^{ask,*} - p_0^{ask} = m^{long,*} - m^{long} \text{ and } p_0^{bid} - p_0^{bid,*} = m^{short,*} - m^{short}.$$

Recall also that the equilibrium level of the bid and ask prices are determined as a fixed point. Under the regularity assumptions we made (that ensure unique equilibrium), the signs of the price changes are determined by the sign of the partial equilibrium change in average information, respectively, $\hat{m}^{long} - m^{long}$ and $\hat{m}^{short} - m^{short}$. Hence, the model predicts that the bid-ask spreads should eventually increase (resp. decrease) if the policy increases (resp. decreases) the average information in traders' orders.

Note also that, once the bid-ask spreads adjust, traders' payoffs are given by Eq. (15), and the effect of the policy on these payoffs is given by, $M^* - M = -c(V^* - V)$. In general equilibrium, the leverage-constraint policy affects traders' payoffs only by its effect on trading volume.

General equilibrium effects on the broker... Recall that, in general equilibrium, the broker's revenue and its size are determined by the technological intermediation costs, cV. Hence, by lowering the trading volume, the leverage-constraint policy lowersthe revenue as well as the size of the brokerage sector.

Effects on social welfare.. Before the leverage-constraint policy, the belief-neutral social welfare is given by Eq. (19). After the policy, the social welfare is given by the same expression but evaluated with the partial equilibrium (hatted) variables or the general equilibrium (starred) variables. In particular, the welfare effect of the policy is characterized by its effect on trading volume and average portfolio risks. Recall that in partial equilibrium, the model predicts that the policy reduces the trading volume as well as average portfolio risks. Hence, the model also predicts that the policy improves belief-neutral social welfare in partial equilibrium. This prediction also applies in general equilibrium as long as the endogenous price response is not strong to overturn the sign of the partial equilibrium effects on trading volume and portfolio risks.

In our empirical analysis, we find that the policy substantially reduces the trading volume in the short run as well as in the longer run (see Section 4.1). In unreported results, we also analyze the effect on the volatility of traders' portfolio returns, and find mixed evidence that seems to point toward the policy lowering portfolio volatility (see Footnote 15 in Section 4.2). Hence, from the lens of this model, our empirical evidence suggests that the leverage-constraint policy improves social welfare.

Obviously, our model is too stylized to capture all potential reasons for trade in the forex market. For instance, some traders could be trading to hedge their background risks, as in Simsek (2013a). Others might be enjoying the sensation from trading. If we were to model these other motives for trade, they would show up as additional terms in Eq. (19). Moreover, restricting leverage (and trade) would typically tend to lower social welfare through these terms. In an empirical analysis, it is impossible to capture all possible reasons for trade. We view our analysis as capturing a key driving force for trade (monetary pursuits from speculation). Our empirical analysis suggests that through this channel the leverage restriction policy had a large positive impact on social welfare. This can also be viewed as setting a (very high) threshold that other rationales for trade would have to exceed to overturn our qualitative conclusion that the leverage restriction policy improved social welfare.

Finally, recall that according to our welfare criterion the bid and ask prices do not matter for the social welfare. This is because they represent transfers among investors, which is ignored by a utilitarian planner that uses a single belief and puts equal weight on all agents (see also Footnote 28).[30] While this provides a reasonable benchmark, one could imagine reasons for why the social planner might also care about bid-ask spreads. For instance, suppose some traders are trading for non-speculative reasons, e.g., to hedge their background risks. Higher bid-ask spreads would reduce these traders' welfare. To the extent that the planner overweights such traders' welfare (and underweights the welfare of the "speculative" traders), then higher bid-ask spreads could also lower social welfare. More generally, bid-ask spreads reflect market quality, which the planner might care about in addition to social welfare.

As we noted, our theoretical analysis suggests the policy increases the bid-ask spreads over the longer run if and only if it improves traders' average information. In our empirical analysis, we find no significant effect on bid-ask spreads. We also find (in back-of-the-envelope calculations) that the policy does not substantially change traders' gross returns, which provides a measure of their average information. Hence, the zero result on bid-ask spreads can also be reconciled with our model, and it suggests that the leverage-constraint policy does not have an adverse effect on social welfare through bid-ask spreads.

[30] This is also why the social welfare is characterized by the same equation (19) in partial as well as general equilibrium, even though the corresponding equilibrium allocations differ in terms of the bid and ask prices.

Appendix A.3: Evidence on the Representativeness of the myForexBook Data

This section provides evidence that the trade level data from myForexBook provides a good representation of the population of retail forex traders. The myForexBook web platform provides a social networking environment for traders who have accounts with at least one of around fifty partnering brokerages. Because traders choose to use the myForexBook platform, these traders could be unrepresentative of the overall population.

We first compare the myForexBook traders' performance to that of the population of traders on the brokerage eToro, one of the market's largest off-exchange brokerages. The eToro data includes all transactions between June 2, 2013, and July 14, 2014.[31] The data include over 11 million transactions from retail traders located in nearly 200 countries and independent territories.

Our analysis is specifically interested in how the availability of leverage affects traders' wealth. Though our eToro sample comes from a time period different than the myForexBook data, eToro traders also have worse returns on positions that use more leverage (A.2, columns 3 and 4). On average, they lose between 65 and 75% ROI per trade for every additional 100 units of leverage. myForexBook traders during our main sample window lose around 28% per trade for every 100 unit increase in leverage. These results are consistent with the myForexBook data being plausibly representative of how traders in the market respond to having less leverage.

Table A.2
Correlation Between Leverage and Trader Returns Across Different Data Sets
Description: This table reports OLS estimates of the following regression:

$$\text{ROI}_{jit} = \gamma_i + \gamma_t + \beta_1 \text{leverage}_{jit} + \varepsilon_{jit}$$

where i is a trader, j is a trade, and t is a day (trades are recorded by the second). The dependent variable is *ROI,* which is per-trade return on investment. Columns (1) and (2) use the myForexBook sample that is used throughout the paper. Columns (3) and (4) use the entire population of trades on eToro between June 2, 2013 and July 14, 2014. Standard errors are double-clustered by day and trader, and *, **, and *** denote significance levels $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

data set:	myForexBook		eToro (June 2013 - July 2014)	
dep var: ROI	(1)	(2)	(3)	(4)
leverage / 100	-0.281**	-0.278**	-0.660***	-0.744***
	(0.12)	(0.12)	(0.16)	(0.13)
trader FE	x	x	x	x
day FE		x		x
Number of trades	270,051	270,051	11,580,789	11,580,789
R^2	0.038	0.040	0.079	0.082

We also find that the myForexBook data is similar to the CFTC data (the brokerages in the CFTC reports account for about 95% of the U.S. market for retail forex). We calculate the total number of trades per month in the myForexBook data and take the log first difference. We also take the log first difference of aggregate retail foreign exchange obligations in the CFTC reports. We multiply the forex obligations time series by negative one, because we would expect the brokerage's obligations to decrease when there is more trading; on average, traders lose money when they trade, which would reduce the value of the traders' accounts (lower the brokerages' obligations). These series overlap from November 2011 to April 2012. The Pearson's correlation coefficient between these series is 0.41, which suggests a reasonably strong correlation between the myForexBook and CFTC data sets. Figure A.1 plots these times series.

[31] The data come from the brokerage, eToro. Per our NDA, eToro maintains the right to approve the use of the company's name in the description of the data prior to any publication.

Fig. A.1. Correlation Between the myForexBook Data and the CFTC Brokerage Reports

Description: This figure plots time series of the number of trades in the myForexBook data set and retail foreign exchange obligations for brokerages in the CFTC reports. The time series are transformed to the logarithm of monthly first differences. Retail forex obligations are multiplied by negative one.

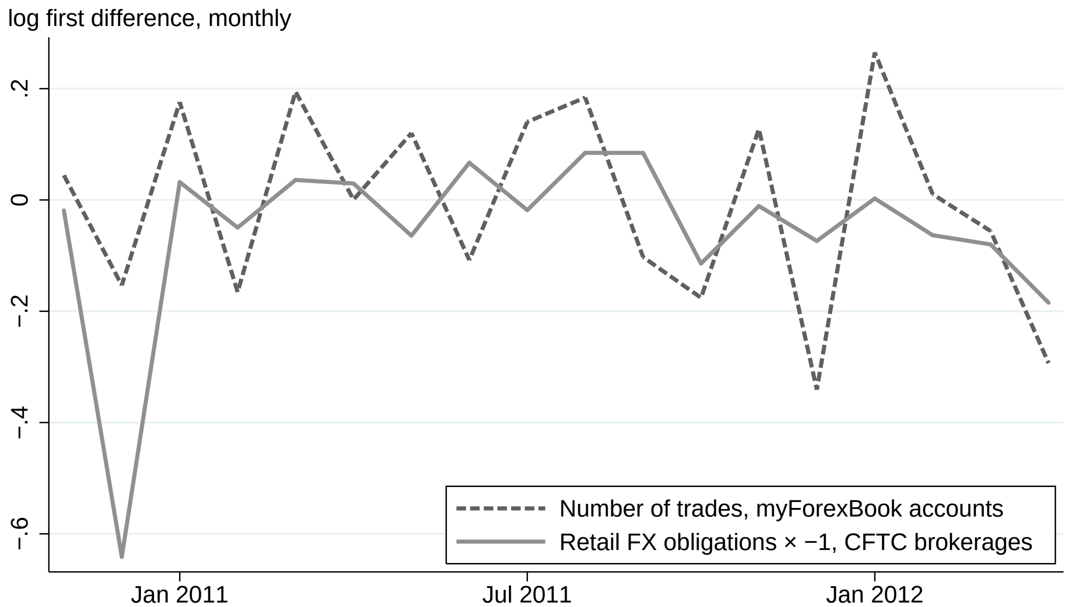

Appendix A.4: Awareness of the CFTC Regulation Announcement

On January 13, 2010, the CFTC announced in the Federal Registrar their intent to restrict foreign exchange dealers' provision of leverage at 10:1. Our analysis shows that this announcement did not affect trader returns, brokerage capital, or the spreads charged by forex brokerages. One plausible explanation for these results is that the announcement could have gone unnoticed by traders, and therefore did not significantly affect trader behavior.

Figure A.2 plots the time series of Google search volume index (SVI) for the search term "forex leverage." Google SVI is often used by the literature to measure attention. There is a substantial increase in attention on forex leverage that occurs as a result of the CFTC's announcement, which is consistent with traders being aware that they were going to have less available leverage.

Fig. A.2. Attention on the CFTC Announcement of Leverage Regulation

Description: This figure plots a time series of U.S. Google search volume index (SVI) for the search term "forex leverage." Google SVI is the ratio of searches for a particular term to the total number of Google searches, normalized on a scale from 0 to 100. The data is at a weekly frequency.

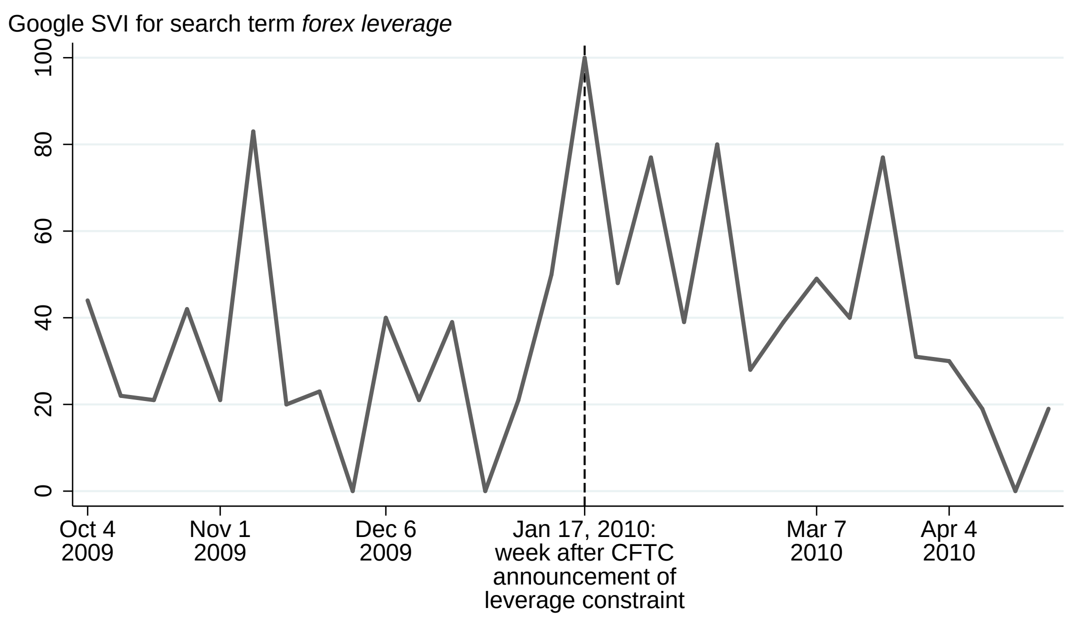

Appendix A.5: Trading Costs and Traders' Gross Returns

Section presents traders' net portfolio returns. Table A.3 presents estimates of traders' gross portfolio returns. To make these back-of-the-envelope estimates, we have to make assumptions about the transaction costs paid by retail forex traders. Trading costs come from traders paying the bid-ask spread on each transaction (to our knowledge, no brokerages charge fixed per-fee costs presently, or during the period of our study). Unfortunately, our transaction-level data set does not tell us the spreads paid by traders, and the amount of trading we observe is too thin to estimate spreads by matching buy to sell orders (for example, many studies use trade and quote database (TAQ) quotes to estimate spreads).

Therefore, our approach is to make assumptions about average spreads and then apply these assumptions to traders' net portfolio returns. Specifically, we believe that average spreads paid are between 3 to 4 pips, where a pip is one one-hundredth of one percent (for example, it would cost three to four dollars to execute the modal trade in our data ($1,000)). We come to this conclusion by noting that most brokerages advertise spreads that are as low as 1 to 2 pips. This headline number is presumably in reference to the most liquid currency pair, the EUR/USD, but other currency pairs cost more to trade. Spreads can also change depending on market conditions: spreads increase by as much as 10 times during episodes of high volatility (see for example, the live spreads presented by the brokerage Oanda: www.oanda.com/forex-trading/markets/recent). Additionally, price slippage would increase the spreads traders actually pay, and the National Futures Association (NFA) found that, during the period we study, at least a few brokerages had computer systems designed to take advantage of slippage (reference). Finally, in support of our assumption that transactions cost an average of 3 to 4 pips, MarketWatch's May 2011 review comparing retail forex brokerages writes that the only brokerage to offer fixed spreads was FX Solutions, which offered 3 pips per EUR/USD transaction (reference).

Under these assumptions, when traders are charged 2 pips per trade, transaction costs explain about 60% of high-leverage traders' about transaction costs (e.g., prior to the leverage-constraint high-leverage traders lose 44% on net and 18% gross). The high-leverage traders still perform worse than low-leverage traders, but not by nearly as much as the difference in their net returns. If we assume that traders are charged 3 to 4 pips per trade, there would be no difference between high- and low-leverage traders gross returns.

Table A.3

Back-of-the-envelope Calculation of Trading Costs' Effect on Gross Returns

Description: This table extends the results on traders' portfolio returns presented in Table 4. Monthly returns are calculated using the account's balances at the beginning and end of the month, excluding deposits. The columns in this table present gross returns calculated after adjusting net returns by the assumed amount of transaction costs paid by traders. Transaction costs in retail forex are the spreads paid by traders. We assume that average spreads during the period we study fell within a range of 2 to 5 pips, where a pip is one one-hundredth of one percent. Stars *, **, and *** denote significance levels $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

U.S. Traders' Portfolio Returns

			gross returns							
assumed per-trade spreads:	net returns (from Table 4)		2 pips		3 pips		4 pips		5 pips	
pre- or post-constraint:	pre-	post-	pre-	post-	pre-	post-	pre-	post-	pre-	post-
sample average	-0.174	-0.095	-0.063	-0.026	-0.007	0.008	0.043	0.040	0.102	0.074
leverage quintile										
high	-0.444	-0.195	-0.176	-0.061	-0.028	0.002	0.092	0.060	0.246	0.122
low	-0.032	-0.020	-0.002	-0.004	0.012	0.008	0.025	0.020	0.038	0.033
high minus low	**-0.412*** **	**-0.175*** **	**-0.174*** **	**-0.058***	**-0.040**	**-0.006**	**0.067**	**0.040**	**0.207** **	**0.089**
	(5.75)	(4.07)	(3.43)	(1.86)	(0.71)	(0.19)	(0.97)	(0.90)	(2.42)	(1.54)

Appendix A.6: Trade-level Returns

This section shows that the leverage-constraint improves traders' returns using tests at the trade-level. These tests look exclusively at the narrow window around the dates of the leverage-constraint (the sample window September 1, 2010 - December 1, 2010) and the regulation's announcement (December 1, 2009 - March 1, 2010). Table A.4 presents summary statistics for trade-level outcomes in this window. Table A.5 uses difference-in-difference regressions to show that the leverage-constraint reduces per-trade losses by about 20 percentage points. A.3 and A.4 plot the impulse-response of the treatment effect of the leverage-constraint in calendar-time and trade-time, respectively. These tests show that the U.S. treatment group and European control group have common trends prior to the regulation. A.5 presents placebo tests for false dates of the regulation. These tests produce few false positive results, indicating that our tests are unlikely to suffer from Type I error. Table A.6 shows that the CFTC's regulation announcement does not significantly affect trade-level outcomes.

Table A.4
Trade-level summary statistics
Description: This table presents summary statistics from the myForexBook account-level database trimmed according to the criteria described in Section 3. The sample includes trades executed by U.S. and European retail forex traders. Return on investment (*roi*) for long (short) positions equals the difference between the nominal value of the currency pair when the position is closed (opened) and when it is opened (closed), divided by the trader's dollar stake in the trade. *Post constraint* equals one if the trade was opened after October 18, 2010, the date by which brokerages needed to comply with CFTC regulation limiting the leverage available to U.S retail forex traders at 50:1, zero otherwise. *Post announcement* equals one if the trade was opened after the CFTC's announcement in the Federal Registrar on January 13, 2010 of their intent to restrict traders' leverage to 10:1, zero otherwise. *High leverage trader* equals one if trader i uses at least 50:1 leverage on at least one trade prior to the CFTC regulation, zero otherwise. *Holding period* is the length of time in hours between when the position is opened and when it is closed.

Panel A: ***sample window around leverage-constraint (Sep 1 - Dec 1, 2010)***

variable	mean	std dev	median	10^{th} %tile	90^{th} %tile
Dependent variables					
Return on investment (ROI)	-0.26	4.81	0.016	-2.33	1.79
trade uses leverage > 50:1 (= 1)	0.084				
Treatment variables					
US trader (= 1)	0.45				
Post constraint (= 1)	0.48				
High leverage trader (= 1)	0.49				
Additional Controls					
log trade size (USD)	0.57	2.24	0.69	-2.30	3.04
log holding period (hours)	0.16	2.43	0.073	-2.93	3.39
Number of trades	270,595				

Panel B: ***sample window around regulation announcement (Dec 1, 2009 - Mar 1, 2010)***

variable	mean	std dev	median	10^{th} %tile	90^{th} %tile
Dependent variables					
Return on investment (ROI)	-0.22	3.93	0.087	-3.21	2.44
trade uses leverage > 10:1 (= 1)	0.42				
Treatment variables					
US trader (= 1)	0.48				
Post announcement (= 1)	0.59				
High leverage trader (= 1)	0.63				
Additional Controls					
log trade size (USD)	1.41	1.83	1.61	0	3.40
log holding period (hours)	0.041	2.50	-0.083	-2.99	3.18
Number of trades	167,035				

Table A.5
Leverage-Constraints and Trade-Level Outcomes
Description: This table reports OLS estimates of the following regression:

$$Y_{jit} = \gamma_i + \gamma_t + \beta_1 \text{US trader}_i \times \text{post constraint}_t + \beta_2 \text{trade}_{jit} + \varepsilon_{jit}$$

where i is a trader, j is a trade, and t is the day trades are opened (execution of trades are recorded at the second). In **Panel A**, the dependent variable is *trade uses leverage > 50:1*, which equals one if the trade uses at least 50:1 leverage. In **Panels B** and **C**, the dependent variable is *ROI*, which is per-trade return on investment. *US trader* equals one if the trade is executed by a trader located in the U.S. and equal to zero if located in Europe. *Post constraint* equals one if the trade was opened after October 18, 2010, the date by which brokerages needed to comply with CFTC regulation limiting the leverage available to U.S retail forex traders at 50:1, zero otherwise. *High leverage trader* equals one if trader i uses at least 50:1 leverage on at least one trade prior to the CFTC regulation, zero otherwise. The sample period is from September 1 to December 1, 2010. Standard errors are double-clustered by day and trader, and *, **, and *** denote significance levels $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

Panel A: ***The Binding Effect of the October 2010 Leverage-Constraint on Trading***

dep var: trade uses leverage > 50:1 (=1)	(1a)	(2a)	(3a)	(4a)
US trader (=1) × post constraint (=1)	-0.0491**	-0.0523**	-0.0520**	-0.0520**
	(0.022)	(0.021)	(0.021)	(0.021)
log(trade size)			0.0148***	0.0145***
			(0.0043)	(0.0043)
log(holding period)				-0.00138**
				(0.00058)
trader FE	x	x	x	x
day FE	x	x	x	x
broker-pair FE		x	x	x
Number of trades	270,595	270,541	270,541	270,541
R^2	0.51	0.53	0.54	0.54

Panel B: ***Oct 2010 Leverage-Constraint & Performance, Euro Traders Control Group***

dep var: per-trade ROI	(1b)	(2b)	(3b)	(4b)
US trader (=1) × post constraint (=1)	0.191**	0.207**	0.204**	0.204**
	(0.094)	(0.098)	(0.098)	(0.099)
log(trade size)			-0.110***	-0.124***
			(0.024)	(0.025)
log(holding period)				-0.0649***
				(0.017)
trader FE	x	x	x	x
day FE	x	x	x	x
currency risk-free rate differential	x	x	x	x
std dev of trader's weekly returns	x	x	x	x
broker-pair FE		x	x	x
Number of trades	270,595	270,541	270,541	270,541
R^2	0.037	0.041	0.042	0.042

Panel C: ***Constraint and Performance; Alt. Control Group – High- vs. Low-Leverage Traders***

dep var: per-trade ROI	(1c)	(2c)	(3c)	(4c)
high leverage trader (=1) × post constraint (=1)	0.237**	0.252***	0.247***	0.253***
	(0.098)	(0.093)	(0.094)	(0.093)
log(trade size)			-0.110***	-0.124***
			(0.024)	(0.025)
log(holding period)				-0.0652***
				(0.016)
trader FE	x	x	x	x
day FE	x	x	x	x
currency risk-free rate differential	x	x	x	x
std dev of trader's weekly returns	x	x	x	x
broker-pair FE		x	x	x
Number of trades	270,595	270,541	270,541	270,541
R^2	0.037	0.041	0.042	0.042

Fig. A.3. Impulse Response of Treatment Effect on Per-Trade Returns
Description: This table reports OLS estimates of the following regression:

$$\text{ROI}_{jit} = \gamma_i + \gamma_t + \sum_{k=T-l}^{T+l} \beta_{1k} \text{US trader}_i \times I_{T+k=t} + \varepsilon_{jit}$$

where i is a trader, j is a trade, and t is a week (trades are recorded by the second). The dependent variable is *ROI,* which is per-trade return on investment. *US trader* equals one if the trade is executed by a trader located in the U.S. and equal to zero if located in Europe. T is the date of the regulation, i.e. October 18, 2010. $I_{T+j=t}$ is an indicator variable for weeks surrounding the regulation. Therefore, β_j for $j = \{-T, ..., T\}$ is the sequence of treatment effects, and hence maps out the impulse response. Standard errors are double-clustered by day and trader, and the dashed lines are 95% confidence intervals around the point estimate of β_j.

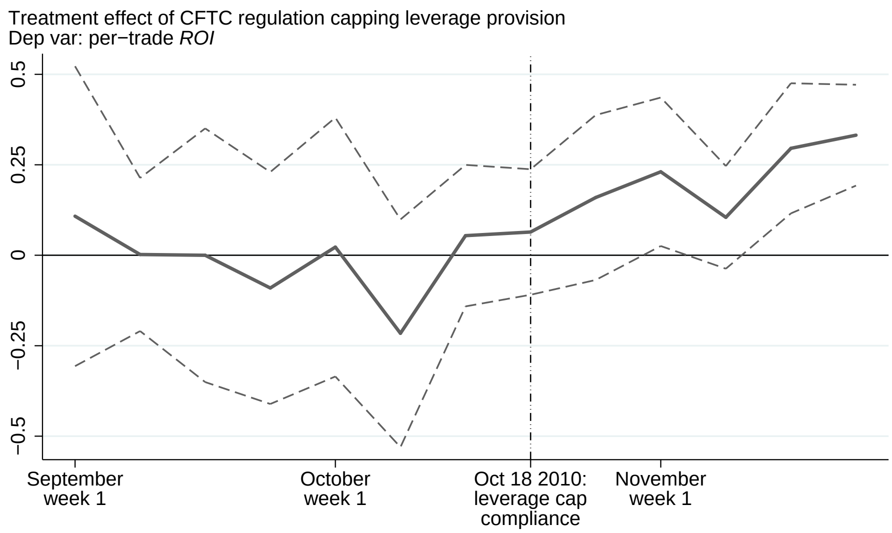

Fig. A.4. Impulse Response of Leverage-Constraint on Per-Trade Returns Using Trade-time
Description: This table reports OLS estimates of the following regression:

$$\text{ROI}_{jit} = \gamma_i + \gamma_t + \sum_{k=T-l}^{T+l} \beta_{1k} \text{US trader}_i \times I_{T+k=t} + \varepsilon_{jit}$$

where i is a trader, j is a trade, and t is a week (trades are recorded by the second). The dependent variable is *ROI,* which is per-trade return on investment. *US trader* equals one if the trade is executed by a trader located in the U.S. and equal to zero if located in Europe. T is the date of the regulation, i.e. October 18, 2010. $I_{T+j=t}$ is an indicator variable for weeks surrounding the regulation. Therefore, β_j for $j = \{-T, ..., T\}$ is the sequence of treatment effects, and hence maps out the impulse response. We sort trades into quartiles, within a trader's account, according to their distance from the leverage-constraint. The omitted coefficient is the interaction between *US trader* and the indicator for the fourth quartile in distance prior to the leverage-constraint. We restrict this sample to traders that use greater than 50:1 leverage on at least one trade prior to the leverage-constraint. Standard errors are double-clustered by day and trader, and the dashed lines are 95% confidence intervals around the point estimate of β_j.

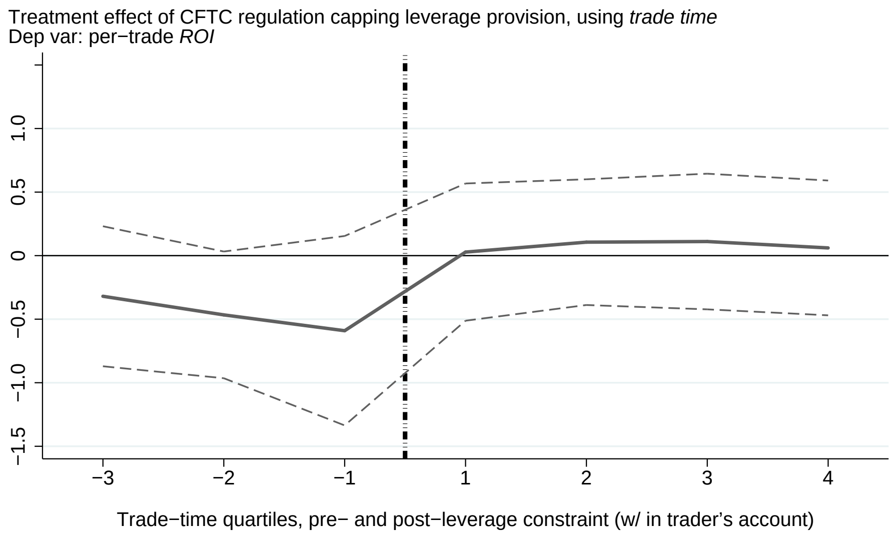

Fig. A.5. Placebo test for the effect of the leverage-constraint

Description, Panel A: This figure illustrates the placebo exercise described in Section **??** and below.

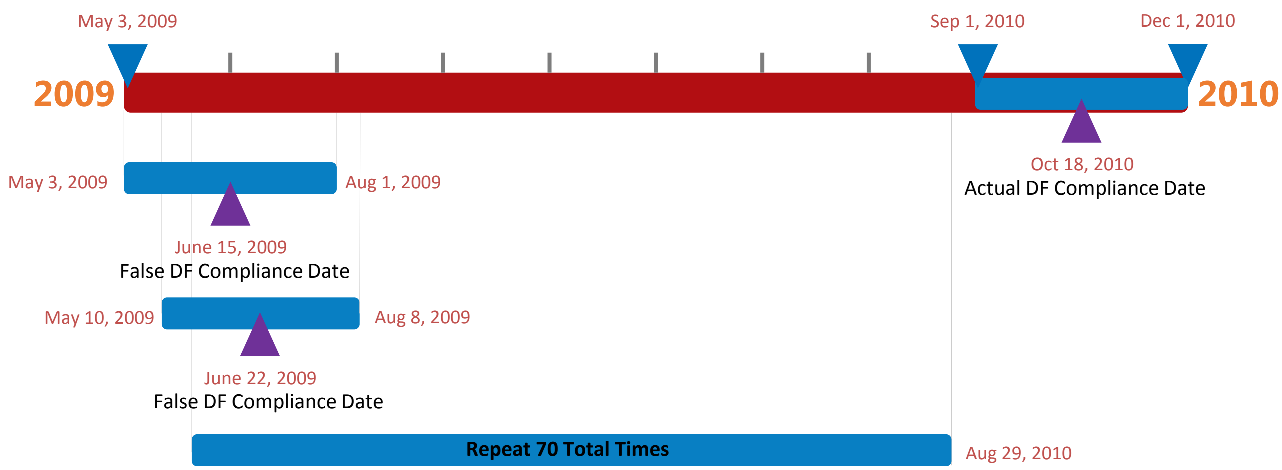

Description, Panel B: This figure plots kernel density estimates using the Epanechnikov kernel function and a histogram of β_1's from a series of placebo tests for the effect of the CFTC regulation on trading outcomes. We run the following regression 70 times

$$Y_{jit} = \gamma_i + \gamma_t + \beta_1 \text{US trader}_i \times \text{post constraint}_t + \beta_2 \text{trade}_{jit} + \varepsilon_{jit}$$

collecting the coefficient, β_1 after each iteration. For each iteration, we change the date of *post constraint*, starting from Sunday, May 3, 2009 rolling forward a week at a time until Aug 29, 2010. Prior to each iteration, we trim the sample using the procedure described in Section 3. This restricts the sample to include only traders that execute trades before and after the false date for *post constraint*.

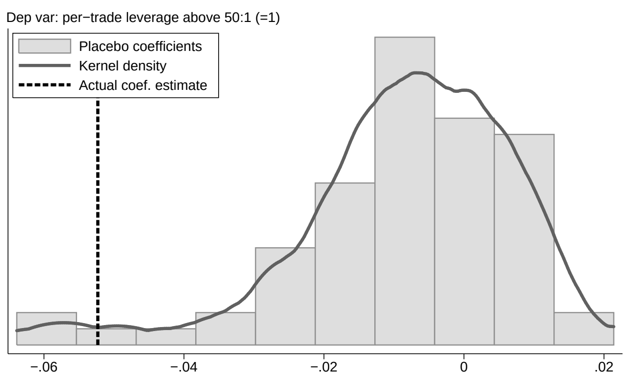

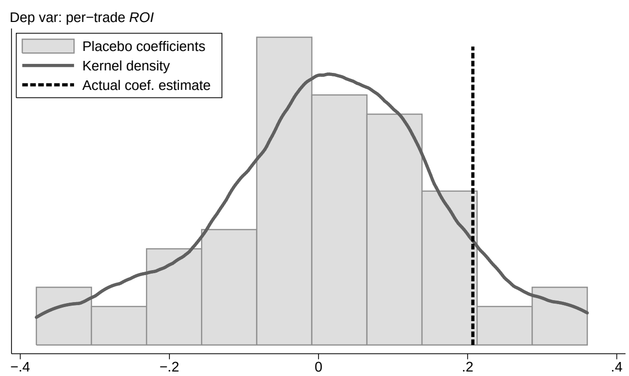

Table A.6
The Announcement of Regulation and Trade-Level Outcomes
Description: This table reports OLS estimates of the following regression:

$$Y_{jit} = \gamma_i + \gamma_t + \beta_1 \text{US trader}_i \times \text{post announcement}_t + \beta_2 \text{trade}_{jit} + \varepsilon_{jit}$$

where i is a trader, j is a trade, and t is a day (trades are recorded by the second). In **Panel A**, the dependent variable is *trade uses leverage > 50:1*, which equals one if the trade uses at least 10:1 leverage. In **Panels B** and **C**, the dependent variable is *ROI,* which is per-trade return on investment. *US trader* equals one if the trade is executed by a trader located in the U.S. and equal to zero if located in Europe. *Post announcement* equals one if the trade was opened after the CFTC's announcement in the Federal Registrar on January 13, 2010 of their intent to restrict traders' leverage to 10:1, zero otherwise. *High leverage trader* equals one if trader i uses at least 50:1 leverage on at least one trade prior to the CFTC regulation, zero otherwise. The sample period is from December 1, 2009 to March 1, 2010. Standard errors are double-clustered by day and trader, and *, **, and *** denote significance levels $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

Panel A: ***The January 2010 Regulation Announcement and High-Leverage Trading***

dep var: trade uses leverage > 10:1 (=1)	(1a)	(2a)	(3a)	(4a)
US trader (=1) × post announcement (=1)	0.0443	0.0169	0.0205	0.0203
	(0.033)	(0.030)	(0.026)	(0.026)
log(trade size)			0.144***	0.143***
			(0.017)	(0.017)
log(holding period)				-0.00299
				(0.0023)
trader FE	x	x	x	x
day FE	x	x	x	x
broker-pair FE		x	x	x
Number of trades	167,035	166,985	166,985	166,985
R^2	0.54	0.56	0.61	0.61

Panel B: ***The Regulation Announcement and Performance, using Euro control group***

dep var: per-trade ROI	(1b)	(2b)	(3b)	(4b)
US trader (=1) × post announcement (=1)	0.00425	-0.0125	-0.0152	-0.0236
	(0.075)	(0.079)	(0.078)	(0.081)
log(trade size)			-0.0649	-0.0946*
			(0.047)	(0.048)
log(holding period)				-0.103***
				(0.017)
trader FE	x	x	x	x
day FE	x	x	x	x
currency risk-free rate differential	x	x	x	x
std dev of trader's weekly returns	x	x	x	x
broker-pair FE		x	x	x
Number of trades	167,035	166,985	166,985	166,985
R^2	0.053	0.057	0.057	0.060

Panel C: ***Performance; Alternative Control Group – High- vs. Low-Leverage Traders***

dep var: per-trade ROI	(1c)	(2c)	(3c)	(4c)
high leverage trader (=1) × post announcement (=1)	-0.0544	-0.0696	-0.0763	-0.0665
	(0.072)	(0.075)	(0.075)	(0.076)
log(trade size)			-0.0655	-0.0950*
			(0.047)	(0.048)
log(holding period)				-0.103***
				(0.017)
trader FE	x	x	x	x
day FE	x	x	x	x
currency risk-free rate differential	x	x	x	x
std dev of trader's weekly returns	x	x	x	x
broker-pair FE		x	x	x
Number of trades	167,035	166,985	166,985	166,985
R^2	0.053	0.057	0.057	0.060

Appendix A.7: Alternative Treatment Groups for Tests of Brokerage Capital

Table 6 shows that the CFTC regulation reducing the provision of leverage to retail traders reduced the amount of capital held by brokerages. The table establishes this finding by comparing CFTC-regulated brokerages that have retail forex obligations to those that do not. However, a plausible concern with this test is that – despite having similar trends prior to the regulation – brokerages without forex obligations are different in unobservable ways, and are therefore not suitable to be a control group. Brokerages with and without forex brokerages could diverge following the regulation because of factors that are unrelated to the leverage restrictions. We address this concern by showing that the regulation has the strongest effect on brokerages that provided more leverage to traders prior to the regulation.

To do so, the following variation on Table 6 sorts brokerages into the amount of leverage they offer traders. We define brokerages as *high leverage* (*low leverage)* if they were providing traders with above (less than) 400:1 leverage around the time of the October 2010 CFTC regulation. We assign brokerages to these classifications by manually searching internet archives, and most of the brokerages were listed on the website: www.100forexbrokers.com.[32] We choose 400:1 leverage as a cutoff, because the website specifies 400:1 as the minimum for a broker's inclusion in their list of "high leverage brokers". Seven brokerages classify as *high leverage* and sixteen as *low leverage.*

Columns (1) and (2) of Table A.7 run difference-in-difference regressions that compare *high leverage* brokerages against brokerages without forex obligations. Columns (3) and (4) use *low leverage* brokerages. The dependent variable is log brokerage excess capital. The point estimate on the difference-in-difference coefficient is between -0.36 and -0.51 for the *high leverage* brokerages and -0.19 to 0.23 for *low leverage* brokerages. These estimates are close to being significant at the 10% level. The lack of statistical significance is presumably due to having few brokerages with forex obligations after conducting the sample splits. Regardless, the effect of the constraint is larger for brokerages that provide more leverage, consistent with the CFTC regulation affecting brokerage excess capital through its effect on retail trader leverage.

Table A.7
Leverage-Constraints and the Excess Capital of High-Leverage-Brokerages
Description: This table reports OLS estimates of the regression

$$\log(\text{excess capital})_{bt} = \gamma_b + \gamma_t + \beta_1 \text{FX broker}_b \times \text{post constraint}_t + \varepsilon_{bt},$$

where b is a broker and t is a month. The data comes from monthly CFTC Futures Commission Merchants Financial Reports. Excess capital is the capital in excess of the regulatory requirement, for each brokerage in the CFTC data set. *FX broker* equals one if the brokerage has any retail forex obligations after they were required to report these obligations starting in November 2010. *Post constraint* equals one in months starting in November 2010, and zero otherwise. Appendix 7 describes how FX brokerages are sorted into high- and low-leverage. Standard errors are double-clustered by broker and month, and *, ** and *** denote significance at the $p < 0.1$, $p < 0.05$ and $p < 0.01$ levels, respectively.

dep var: brokerage excess capital	(1)*	(2)*	(3)†	(4)†
FX broker high leverage (=1) × post constraint (=1)	-0.367	-0.512		
	(0.29)	(0.36)		
FX broker low leverage (=1) × post constraint (=1)			-0.234**	-0.190
			(0.12)	(0.13)
log net capital requirement		-0.274		-0.292
		(0.17)		(0.18)
brokerage FE	x	x	x	x
month FE	x	x	x	x
N (broker-month)	1,332	1,332	1,427	1,427
Number of high (or low) leverage brokers	7	7	16	16
R^2	0.99	0.99	0.99	0.99

* sample includes high-leverage FX brokerages and CFTC regulated brokerages w/ no-FX obligations

† sample includes low-leverage FX brokerages and CFTC regulated brokerages w/ no-FX obligations

[32] An alternative approach to this classification would be to assign brokerages to *high leverage* or *low leverage* using the amount of leverage used by traders in the myForexBook data set. However, there are only seven brokerages that are common to the CFTC's data set and the myForexBook data set.

Appendix A.8: Trader Flows

This section tests for the effect of the leverage-constraint on the entry and exit rates of traders into the retail forex market. Unfortunately, the CFTC brokerage reports do not list the number of trader accounts. So, we use the myForexBook account-level data set to approximate account flows. We define trader entry as the first month a trader is in the data. We define trader exit as the last month that they trade. We then collapse the indicators for trader entry and exit to the brokerage-month-location level, where location is either traders from the U.S. or from Europe.

Table A.8 presents difference-in-differences regressions that compare the number of new (or exiting) U.S. traders to European traders, as a result of the leverage-constraint. The logarithm of new traders is the dependent variable in columns (1) and (2), and the logarithm of exiting traders is the dependent variable in columns (3) and (4). The coefficient of interest is the interaction of *post constraint* and *US traders* – an indicator that equals one if the traders come from the U.S. and zero if they come from Europe. The regressions include month, brokerage, and trader location fixed effects. Columns (2) and (4) have brokerage fixed effects interacted with a time trend, which accounts for the possibility that the growth and exit rates of new traders can vary by brokerage. This also helps control for the unconditional growth rate of the membership of the myForexBook website during this period (the website started in 2009 and its population grew to a peak of around 10,000 traders by the middle of 2011).

The leverage-constraint caused a reduction in trader inflows for new U.S. traders. The constraint reduced trader outflows, but the estimate is not statistically distinguishable from zero. Moreover, the reduction in inflows of U.S. traders is larger than the reduction in outflows. The difference-in-difference coefficient for trader inflows is -0.118, and the sample average of monthly inflows is 1.21, which suggests a 0.118 / 1.21 = 9.8% reduction in inflows. Using the same calculation, the reduction in outflows is 8.5%. Furthermore, columns (2) and (4) use distributed lags around the regulation date to test for pre-trends. The effect of inflows is close to zero before the regulation, but the coefficient falls to around -0.1 persistently thereafter. On the other hand, the effect on outflows is noisy around the regulation date.

Table A.8
The Effect of the Leverage Regulation on Trader Flows
Description: This table uses account level data from the myForexBook data set collapsed to observations at the level of brokerage, month, and the geography of traders within the brokerage. It reports OLS estimates of the regression

$$\log(\text{number of traders})_{bgt} = \gamma_b + \gamma_g + \gamma_t + \beta_1 \text{US flows}_g \times \text{post constraint}_t + \varepsilon_{bgt},$$

where b is a broker, g is trader geography (either U.S. or Europe), and t is a month. *US flows* equals one if the number of traders are from the U.S. and equal to zero if they are from Europe. *Post constraint* equals one in months starting in November 2010, and zero otherwise. The sample period is May 2010 to April 2011 Standard errors are double-clustered by broker and month, and *, ** and *** denote significance at the $p < 0.1$, $p < 0.05$ and $p < 0.01$ levels, respectively.

	trader inflows		**trader outflows**	
dep var:	log(# new traders + 1)		log(# exiting traders + 1)	
	(1)	(2)	(3)	(4)
US flows (=1) × post constraint (=1)	-0.118*		-0.134	
	(0.065)		(0.080)	
US flows (=1) × Sept 2010 (=1)		0.0160		-0.0459
		(0.072)		(0.042)
US flows (=1) × Oct 2010 (=1)		-0.0122		0.111***
		(0.14)		(0.019)
US flows (=1) × Nov 2010 (=1)		-0.101		0.0258
		(0.066)		(0.072)
US flows (=1) × Dec 2010 (=1)		-0.141***		0.136
		(0.044)		(0.090)
US flows (=1) × Jan 2011 (=1)		-0.104		-0.0699***
		(0.061)		(0.0063)
month FE	x	x	x	x
US flows FE	x	x	x	x
brokerage FE	x		x	
time trend × broker FE		x		x
mean of dependent variable	1.21	1.48	1.57	1.81
Broker-month-trading region obs.	393	286	345	260
R^2	0.75	0.87	0.78	0.84