Large Bets and Stock Market Crashes

Albert S. Kyle* Anna A. Obizhaeva†

This version: November 12, 2020

Previous version: March 3, 2017

First version: March 5, 2012

Abstract

Some market crashes occur because of significant imbalances in demand and supply. Yet, conventional models fail to explain the large magnitudes of price declines. We propose a unified structural framework for explaining crashes, based on the insights of market microstructure invariance. A proper adjustment for differences in business time across markets leads to predictions which are different from conventional wisdom andconsistent with observed price changes during the 1987 market crash and the 2008 sales by Société Générale. Somewhat larger-than-predicted price drops during 1987 and 2010 flash crashes may have been exacerbated by too rapid selling. Somewhat smaller-than-predicted price decline during the 1929 crash may be due to slower selling and perhaps better resiliency of less integrated markets.

JEL Codes: G01, G28, N22.

Keywords: crashes, liquidity, price impact, market depth, systemic risk, market microstructure, invariance.

*Kyle: Robert H. Smith School of Business, University of Maryland, College Park, MD, USA, akyle@rhsmith.umd.edu. Albert Kyle was a staff member of the Task Force on Market Mechanisms ("Brady Commission") in 1987–1988, worked on a research project examining high-frequency trading in the CME S&P 500 E-mini futures market for the Commodity Futures Trading Commission during 2009–2010, and has worked with the Securities and Exchange Commission, the Department of Justice, and the Federal Reserve Bank of Atlanta. He has been a member of the CFTC Technology Committee and is a member of the FINRA Economic Advisory Board. He is a member of the Board of Directors of an asset management company which trades equities on behalf on institutional investors.

†Obizhaeva: New Economic School, Moscow, Russia, aobizhaeva@nes.ru.

After stock market crashes, rattled market participants, frustrated policymakers, and puzzled economists are usually unable to explain what happened. Noticeably heavy selling pressure has been often recorded during crashes, and it is known that large bets move prices in the direction of trades, as discussed by Kraus and Stoll (1972), Grinold and Kahn (1995), and Gabaix (2009). Yet, there is no compelling quantitative explanation for why relatively small quantities sold might have led to such large price dislocations in the highly liquid stock market.

We investigate this issue through the lens of the market microstructure invariance, a conceptual framework recently developed by Kyle and Obizhaeva (2016). By analysing prices and quantities in market-specific business time rather than in calendar time, this approach explains why bets of observed sizes could indeed create market crashes.

We illustrate our approach by studying five crash events, chosen because data on the magnitude of contemporaneous selling pressure became publicly available in their aftermaths:

- After the stock market crash of October 1929, the report by the Senate Committee on Banking and Currency (1934) (the "Pecora Report") attributed the sharp plunge in broker loans to forced margin selling during the crash.

- After the October 1987 stock market crash, the U.S. Presidential Task Force on Market Mechanisms (1988) (the "Brady Report") reported quantities of stock index futures contracts and baskets of stocks sold by portfolio insurers during the crash.

- After the futures market dropped by 20% at the open of trading three days after the 1987 crash, the Commodity Futures Trading Commission (1988) documented large sell orders executed at the open of trading; the press identified the seller as George Soros.

- After the Fed cut interest rates by 75 basis points in response to a worldwide stock market plunge on January 21, 2008, Société Générale revealed that it had been quietly liquidating billions of Euros in stock index futures positions accumulated earlier by rogue trader Jérôme Kerviel.

- After the flash crash of May 6, 2010, the Staffs of the CFTC and SEC (2010b,a) cited as its trigger the large sales of futures contracts by one entity, identified in the press as Waddell & Reed.

We do not study the flash crash events in 1961 and 1989, the LTCM crisis in 1998, the quant crash in August of 2006, or the U.S. Treasury note flash rally in October of 2016 because data on the size of sales which precipitated the events is not available.

Each of the five crashes is associated with a large sell bet, where we think of a "bet" (or "meta-order") as being a statistically independent decision to either speculate on information or to hedge risks by buying or selling significant quantities of risky financial assets, often implemented as sequences of orders executed over time. These bets resulted either from trading by one entity or from correlated trading of multiple entities with the same motivation.

Many practitioners and academics believe that during these five crashes selling pressure was quantitatively too small to induce significant price declines. We call this interpretation "conventional wisdom." Scholes (1972), Harris and Gurel (1986), and Wurgler and Zhuravskaya (2002) illustrate it by saying that the demand for financial assets is elastic in the sense that selling 1% of market capitalization has a price impact of less than 1%. Since annual turnover rates do not vary significantly across stocks, this also implies that sales of 5% of average daily volume is expected to have only modest impact on stock prices. Extrapolated to the market for stock index futures with its extremely high liquidity, these estimates further suggest that selling 5% of daily volume must have even smaller impact there. When applied to market bets during the five crashes, this thinking implies tiny market impact.

Microstructure invariance implies a different way to extrapolate impact estimates from stocks to index futures. This approach implies much bigger price impact of bets in liquid futures markets, because it models trading in market-specific business time, not calendar time.

Indeed, we show below that invariance principles imply business time passes about 225 times faster in the equities market as a whole than in markets for less liquid individual stocks. Thus, one calendar day of trading in stock index futures is equivalent to 225 calendar days of trading in a stock. Intuitively, a bet of 5% of one day's volume for index futures is equivalent to selling of 5% of daily volume each day for 225 consecutive days (not one day) for a stock, or a bet of 1125% ($= 5\% \times 225$) of one day's volume. A bet of 5% of one day's volume for index futures must have much bigger (not smaller) price impact than a bet of 5% of one day's volume for less liquid individual stock.

Invariance principles imply a universal formula for market impact: It is a function of the dollar size of a bet, expected dollar volume, returns volatility, and a couple of invariant parameters. Kyle and Obizhaeva (2016) calibrate these parameters using a database of about 400,000 portfolio transition orders executed during the period 2001–2005 in U.S. stocks. Portfolio transition orders are well suited for calibration of market impact functions because they can be thought of as exogenous shocks to demand and supply. In this paper we extrapolate the estimates from the sample of relatively small individual U.S. stocks to the gigantic U.S. stock market as a whole and from 2001–2005 to other historical periods. The implied estimates are indeed large enough to explain crashes.

Table 1 summarizes our results for each of the five crash events. It shows the actual percentage decline in market prices, the percentage decline predicted by invariance, the percentage decline predicted by conventional wisdom, the dollar amount sold as a fraction of average daily volume, and the dollar amount sold as a fraction of one year's GDP. The estimates implied by conventional wisdom assume that market impact is equal to the percentage of market capitalization sold. Other market impact estimates based on studies by Grinold and Kahn (1995), Torre (1997), Almgren et al. (2005), and Frazzini, Israel and Moskowitz (2018)—discussed later—produce estimates similar to conventional benchmarks.

Table 1: Summary of Five Crash Events: Actual and Predicted Price Declines.

	Actual	Predicted Invariance	Predicted Conventional	%ADV	%GDP
1929 Market Crash	25%	46.43%	1.36%	265.41%	1.136%
1987 Market Crash	32%	16.77%	0.63%	66.84%	0.280%
1987 Soros's Trades	22%	6.27%	0.01%	2.29%	0.007%
2008 SocGén Trades	9.44%	10.79%	0.43%	27.70%	0.401%
2010 Flash Crash	5.12%	0.61%	0.03%	1.49%	0.030%

Table 1 shows the actual price changes, predicted price changes, and bets as percent of average daily volume and GDP.

The 1929 crash, the 1987 crash, and the Société Générale trades of 2008 involve very large sales of more than 25% of average daily volume or more than 0.25% of GDP. By contrast, the sales by Soros in 1987 and the flash crash of 2010 both involve sales of only 2.29% and 1.49% of average daily volume. For all events, the price impact estimates based on conventional wisdom are minuscule in comparison to actual price changes. In contrast, the magnitudes of invariance estimates are more similar to actual price declines.

Yet, there is substantial variation across events in forecasting errors, which may be related to the speed with which events unfold. The actual price decline of 25% during the 1929 crash was smaller than the forecast of 46.43%, perhaps due to efforts made to spread the impact of margin selling out over several weeks rather than several days. The Soros 1987 trades and the 2010 flash crash were both flash-crash events in which prices declined rapidly and then recovered a few minutes later, perhaps due to the unusually rapid rate at which bets were executed.

Except for the time frame of execution, the spirit of invariance suggests that institutional details related to market structure, information asymmetries, or motivation of traders should not affect market impact estimates much. Yet, high variation in the degree of market integration across assets, lack of capital available to take the other sides of large bets, and extreme disruptions to the market mechanism may help explain why some price declines might have deviated from predicted levels. In 1929, smaller-than-predicted price declines may have been reduced by markets being less integrated than today, and potential buyers were keeping capital on the sidelines to profit from price declines widely expected to occur if margin purchases were liquidated. In 1987, larger-than-predicted price declines may have been exacerbated by breakdowns in the market mechanism documented in the Brady Report.

A stock market crash is a tail event in the probability distribution of bet sizes. The invariance-implied distributions of bet sizes calibrated from the size of portfolio transition orders for individual stocks is very close to log-normal with log-variance of 2.53. Extrapolated to market bets, this implies that the two flash crashes are 4.5-standard-deviation events, which are expected to occur several times per year. We conjecture that these bets usually do not cause crashes because they are executed more slowly and under conditions of lower volatility than the two unusual flash crashes. The three other large crashes are approximately 6-standard-deviation events, which are expected to occur only once in hundreds or thousands of years. This suggests either that typical bets have a larger standard deviation than portfolio transition orders or that the tail of the distribution is more consistent with a power law than a lognormal distribution.

The bet-induced crashes differ conceptually from macroeconomic crises including sovereign defaults, banking crises, exchange rate crises, and bouts of high inflation, as catalogued by Reinhart and Rogoff (2009). Bet-induced crashes are likely to be short lived, especially if followed by appropriate government policy. Recovery from economic crises in contrast takes many years, even after significant changes in macroeconomic policies and market regulation. For example, the looser monetary policy implemented by Federal Reserve System immediately after the 1929 crash calmed down the market by the end of 1929. Even though the wealth effect of declining equity prices may have helped trigger a recession by reducing consumption, Friedman and Schwarz (1963) write that the Great depression of the 1930s resulted from a subsequent shift towards a deflationary monetary policy, not from the 1929 crash itself. Similarly, the unwinding of Jérôme Kerviel's gigantic rogue bet in January of 2008 was followed by the collapse of Bear Stearns a few weeks later, but the deep and long lasting recession which unfolded in 2008–2009 was triggered by the bursting of the real estate credit bubble, not from liquidation of his bet.

Unable to find rational quantitative explanations, some researchers believe that market crashes result from irrational behavior. The "animal spirits" hypothesis of market crashes says that price fluctuations occur as a result of random changes in psychology and emotions, which may not be based on economically relevant information or rational calculations. Keynes (1936) said that financial decisions may be taken as the result of "animal spirits—a spontaneous urge to action rather than inaction, and not as the outcome of a weighted average of quantitative benefits multiplied by quantitative probabilities." Akerlof and Shiller (2009) echo Keynes: "To understand how economies work and how we can manage them and prosper, we must pay attention to the thought patterns that animate people's ideas and feelings, their animal spirits." Promptly after the 1987 crash, Shiller (1987) surveyed traders and found that "most investors interpreted the crash as due to the psychology of other investors."

We disagree with the animal spirits theory of crashes. Before the 1929 crash, market participants widely discussed the possibility that forced liquidations of margin accounts would lead to a collapse in prices. Before the 1987 crash, market participants discussed that portfolio insurance sales might lead to a market meltdown. In the absence of quantitative justification, these prescient views were largely dismissed due to a deeply entrenched ideological belief that the demand for equities is elastic. Here we provide quantitative justification for a theory of crashes based on the market impact of large bets, not based on psychology. The margin sales of 1929, portfolio insurance sales of 1987, and liquidation of Kerviel's rogue positions were all large bets resulting from rapid execution of mechanical trading strategies, not from psychology. While the sales of George Soros in 1987 and sales of Waddell and Reed in 2010 may reflect the animal spirits of one person and one entity, the rapid recovery of prices after these events suggest the opposite of market-wide irrationality or psychological contagion.

The remainder of this paper discusses the conventional wisdom in assessing market impact, market microstructure invariance, particulars of each of the crash events, and lessons learned.

1 Market Impact of Large Bets: Previous Literature

Previous studies have used different methodologies to obtain to obtain widely varying estimates of the market impact of large bets. This literature can be divided into two strands. The first strand, which we call "conventional wisdom," examines the price effects of seasoned equity offerings, changes in the composition of the S&P 500 index, and similar events. These studies imply that the price elasticity of demand for stocks is elastic in the sense that selling one percent of the market capitalization of a single stock is associated with a price decline of less than one percent. The second strand examines the price impact of trades by institutional investors. These studies imply that the demand for individual stocks is inelastic in the sense that selling one percent of the market capitalization of a stock is associated with a price decline greater than one percent.

Conventional Wisdom. Based on data on secondary equity distributions, Scholes (1972) claimed that the price impact of large sales of equities is negligible. Harris and Gurel (1986), Wurgler and Zhuravskaya (2002), and others infer from the price response to additions and deletions of stocks to equity indices like the S&P 500 that selling 1% of an individual stock's shares outstanding has a price impact of at most 1%. Wurgler and Zhuravskaya (2002, Table IV, p. 603) provide a summary of demand elasticities inferred from different papers. The elasticities range from 3000, representing an almost infinitely elastic demand schedule from Scholes's study, to 1. These studies all imply an elastic demand for stocks.

When extrapolated from the market for one stock to the stock market as a whole using the same elasticity, these empirical studies support the conventional wisdom that selling pressure does not create stock market crashes. From a theoretical perspective, the conventional wisdom is based on the logic of perfectly competitive capital markets, the capital asset pricing model, and the efficient markets hypothesis. Theoretically, the market risk premium of 5%–7% per year reflects compensation for bearing the risk of the entire stock market for one year. When large bets are executed, market participants taking the other side of these bets are exposed to risks of much smaller magnitude than the entire market and hold these positions over much shorter horizons than one year, usually a few days or minutes. Thus, the compensation required for absorbing these bets should be dramatically smaller than the equity market risk premium.

Indeed, when conventional wisdom was applied to the crash of 1987, prominent financial economists claimed that the price impact of reported sales was up to 100 times too small to generate a crash. Leland and Rubinstein (1988), the academics most closely associated with portfolio insurance in 1987, say, “To place systematic portfolio insurance in perspective, on October 19, portfolio insurance sales represented only 0.2% of total U.S. stock market capitalization. Could sales of 1 in every 500 shares lead to a decline of 20% in the market? This would imply a demand elasticity of 0.01—virtually zero—for a market often claimed to be one of the most liquid in the world.” Miller (1991) makes similar claims about the 1987 crash: “Putting a major share of the blame on portfolio insurance for creating and overinflating a liquidity bubble in 1987 is fashionable, but not easy to square with all relevant facts. ... No study of price-quantity responses of stock prices to date supports the notion that so large a price decrease (about 30%) would be required to absorb so modest (1%–2%) a net addition to the demand for shares.” Indeed, Brennan and Schwartz (1989) calibrated a theoretical model of competitive capital markets and showed that portfolio insurance sales (of 0.63% of market capitalization) would have had an effect on prices about 100 times smaller than the actual size of the 1987 crash (of 32%).

Since price pressure was thought to be too small to explain quantitatively market crashes, some observers of the 1987 stock market crash, including Miller (1988, p. 477) and Roll (1988), sought to explain the large price declines as market reactions to new fundamental information rather than response to trading. Yet, it was difficult to find new information to which market prices would have reacted so drastically.

In our analysis below, we define conventional wisdom using the least conservative conventional estimate and assume a unit demand elasticity for stocks: Selling 1% of capitalization moves prices down by 1%. Mathematicallly, suppose a stock’s price is P, outstanding shares are N, and shares sold are Q. Then, the expected log-percentage market impact $\Delta \ln P$ is Q/N:

$$\Delta \ln P \approx \frac{\Delta P}{P} = \frac{Q}{N}. \tag{1}$$

Throughout this paper, we adopt the convention that Q is unsigned trade size and $\Delta P/P$ is expected unsigned price impact.[1]

[1] The size of market impact $\Delta \ln P$ is either the expectation of the post-trade log-price minus pre-trade log-price for buy bets or the expectation of the pre-trade log-price minus post-trade log-price for sell bets. A similar formula can be written for simple percentage impact $\Delta P/P$, where ΔP is either the difference between post-trade price and pre-trade price for buy bets or the difference between pre-trade price and post-trade price for sell bets.

The conventional market impact function can be also expressed in terms of average daily volume. Let V denote daily volume in shares. Assume for simplicity that an asset's turnover is approximately 100% per year consisting of 250 trading days. Since 1% of market capitalization is approximately equal to 250% of daily volume, the conventional wisdom (1) can be interpreted as

$$\Delta \ln P \approx \frac{\Delta P}{P} = \frac{Q}{250 \text{ days} \cdot V}. \tag{2}$$

This equation implies that in any market, regardless of its liquidity, selling a fixed fraction of daily volume or market capitalization has the same small percentage price impact. For example, selling 25% of average daily volume, which represents 0.10% of shares outstanding, has tiny price impact of about 0.10%, or 10 basis points.

The Brady Report used the same intuition to compare daily volume elasticities in the 1929 crash to the 1987 crash:

> "To account for the contemporaneous 28% decline in price, this implies a price elasticity of 0.9 with respect to trading volume which seems unreasonably high. As a percentage of total shares outstanding, margin-related selling would have been much smaller. Viewed as a shift in the overall demand for stocks, margin-related selling could have accounted realistically for no more than 8% of the value of outstanding stock. On this basis, the implied elasticity of demand is 0.3 which is beyond the bound of reasonable estimates."

Institutional Trades Academic studies of large institutional bets typically find that the demand for stocks in inelastic, not elastic. Kraus and Stoll (1972) study block trades of large NYSE stocks. More recent estimates of market impact from executions of large orders by institutional investors include Chan and Lakonishok (1995, 1997) and Keim and Madhavan (1997). Some studies find nonlinear price impact. The "square root model" or Barra model, described by Grinold and Kahn (1995) and Torre (1997), says that the execution of an order of size Q on average moves price by $\Delta P/P = 1 \cdot \sigma \cdot (Q/V)^{1/2}$. Frazzini, Israel and Moskowitz (2018) estimate a more complicated version of the square root model. Almgren et al. (2005) incorporate information on execution horizon T into a model with concave price impact similar to a square root model. The square root model implies that an order for 25% of one days volume in a stock with 2% daily volatility implies a price impact of 100 basis points. With 100% annual turnover, this implies an elasticity of demand of 0.10, which is similar to the price impact from other studies of large institutional orders.

While the inelastic demand from these models implies more price impact than conventional wisdom, a demand elasticity of only 0.10 is nevertheless not small enough to generate crashes. This is shown in Appendix B, which extrapolates the estimates of these models to crash events. For crashes to result from selling pressure, the elasticity needs to be approximately 0.01, yet another order of magnitude more inelastic than impied by these studies. The square root model makes it especially difficult to explain crashes because its implied concave price impact makes marginal price impact decrease as the size of large bets increases. To explain crashes, we use a linear price impact model, which is popular with finance theorists because is excludes simple forms of arbitrage (e.g., Huberman and Stanzl (2004)).

2 Market Impact of Large Bets: Invariance

Market microstructure invariance can explain stock market crashes because it implies an alternative methodology for extraplating price impact from the less liquid markets for individual stocks to the more liquid market for all stocks. Instead of assuming that the elasticity of demand is constant across assets, invariance implies that the demand for more liquid stocks is more inelastic than for less liquid stocks. Since the equity market as a whole is much more liquid than the market for an individual stock, the demand elasticity becomes small enough—and price impact therefore large enough—to explain stock market crashes. While we disagree with the conventional wisdom that the demand for stocks is elastic, our approach is consistent with inelastic demands estimated for individual stocks based on studies of institutional trades.

Market Microstructure Invariance. Invariance is based on the simple intuition that trading in a speculative market is a game in which fundamental risks are reshuffled among participants in business time. The speed of business time varies significantly across assets and proportionally to the rate at which new bets, or trading ideas, arrive. Trading is fast in liquid markets and slow in illiquid markets.

Invariance consists of two conjectures: (1) The distribution of standard deviations of dollar gains and losses on bets is the same across markets, *when standard deviation is measured in* *units of business time.*[2] (2) The expected dollar costs of executing similar bets are constant across markets, *when similarity of bets is defined in terms of the same dollar risks transferred per unit of business time.* These invariance conjectures have important implications for the rate at which financial markets transfer risks.

We review below the derivation of these scaling laws using a simplified version of Kyle and Obizhaeva (2016).[3] Kyle and Obizhaeva (2020) show how to get the same scaling laws in the context of an equilibrium model of speculative trading with endogenous acquisition of private information and endogenous entry into the market. In the model, scaling laws for bet sizes and transaction costs ultimately follow from the assumption that efforts required to generate private signals do not vary across markets. This is likely to hold, at least approximately, in the equilibrium where trader try to allocate their skills optimally across markets to maximize the value of trading.

Business time. For a given stock, suppose that bets of average size $\bar{Q}$ arrive at rate γ. For a typical stock, we might have $\gamma = 100$ bets per day and $\bar{Q} = 10\,000$ shares. As γ increases, market participants transfer risks more quickly, and business time passes at a faster rate relative to calendar time.

Since individual bets are difficult to observe, it is hard to measure γ and $\bar{Q}$ in practice. Yet, up to some constant which is the same for all stocks, γ and $\bar{Q}$ can be inferred from daily dollar volume $P \cdot V$ and daily returns volatility σ. The proof is based on two simple equations.

First, define trading activity as the product of dollar volume and returns volatility:

$$W := P \cdot V \cdot \sigma. \tag{3}$$

Trading activity W better reflects the rate at which the market transfers risks than dollar volume $P \cdot V$ because it takes into account that trading assets with higher volatility σ transfers more risk per dollar traded. Since bets sum up to volume, $V = \gamma \cdot \bar{Q}$, we can write W in terms of γ and $\bar{Q}$:

$$W = \gamma \cdot \bar{Q} \cdot P \cdot \sigma. \tag{4}$$

[2] This conjecture does not say that dollar returns volatility or returns volatility are constant in business time.

[3] Kyle and Obizhaeva (2018*b*) obtain similar predictions using dimensional analysis and leverage neutrality. Kyle and Obizhaeva (2018*a*) derive them from the meta-model, a system of simple equations inherent to many microstructure models. Kyle, Obizhaeva and Wang (2014) provide illustration using a one-period equilibrium model.

Since dollar volume $P \cdot V$ has units of dollars/day and returns volatility has units per day$^{1/2}$, trading activity W has units dollars/day$^{3/2}$.

Second, the first invariance conjecture says that dollar risk $P \cdot \sigma$ transferred by an average bet of $\bar{Q}$ shares per unit of business time $1/\gamma$ is invariant across markets. Thus, for some dollar constant $\bar{C}$, such as $\bar{C} = \$2000$, we have

$$\bar{Q} \cdot P \cdot \frac{\sigma}{\sqrt{\gamma}} = \bar{C}. \tag{5}$$

Equations (4) and (5) make up a system of two log-linear equations in two unknowns $a := P \cdot \bar{Q} \cdot \sigma$ and γ:

$$a \cdot \gamma = W \qquad \text{and} \qquad a \cdot \gamma^{-1/2} = \bar{C}. \tag{6}$$

The solution for a and γ is

$$a = \bar{C} \cdot \left(\frac{W}{\bar{C}}\right)^{1/3} \qquad \text{and} \qquad \gamma = \left(\frac{W}{\bar{C}}\right)^{2/3}. \tag{7}$$

Now define $H := 1/\gamma$ as the time interval between bets. For example, if $\gamma = 100$ bets per day for some stock, then H is about 4 minutes during trading hours from 9:30 a.m. to 4:00 p.m. Equation 7 implies that average dollar bet size $P\bar{Q}$ and time between bets H are given by

$$P \cdot \bar{Q} = \frac{\bar{C}}{\sigma} \cdot \left(\frac{W}{\bar{C}}\right)^{1/3} \qquad \text{and} \qquad H = \left(\frac{W}{\bar{C}}\right)^{-2/3}, \qquad \text{where} \qquad W := P \cdot V \cdot \sigma. \tag{8}$$

Since W has units of dollars/day$^{3/2}$ and $\bar{C}$ has units of dollars, equations 8 have correct units of dollars for $P \cdot \bar{Q}$ and days for H.

Equations 8 show how to extrapolate the size and number of bets from one stock to another, under the invariance assumption that $\bar{C}$ is constant across markets. Define a benchmark stock as a security with stock price $P^* = \$40$ per share, expected volume $V^* = 10^6$ shares per calendar day, expected percentage returns volatility $\sigma^* = 0.02$ per day$^{1/2}$, and trading activity $W^* = P^* \cdot V^* \cdot \sigma^*$; these parameters would approximately correspond to a stock from the bottom of the S&P 500 index. Suppose the time interval between bet arrives is, say, $H^* = 100$ bets per day.

Equation (8) implies that H is inversely proportional to the 2/3 power of trading activity,

$$\frac{1}{H} = \frac{1}{H^*} \cdot \left(\frac{W}{W^*}\right)^{2/3} = \frac{1}{H^*} \cdot \left(\frac{P \cdot V \cdot \sigma}{P^* \cdot V^* \cdot \sigma^*}\right)^{2/3}. \tag{9}$$

Business time H represents different lengths of calendar time for different assets: 4 minutes for the benchmark stock, an hour for thinly traded stocks, less than one minute for actively traded stocks, and about one second for the market as a whole. The conventional wisdom implicitly makes the mistake of extrapolating from one market to another under the assumption that business time H is constant across markets.

Distribution of Bet Size. The logic of invariance can be applied to the entire distribution of random bet sizes $\tilde{Q}$, not just the means $\bar{Q}$. This logic implies that probability distributions of bet sizes $\tilde{Q}$ must be the same across markets if $\tilde{Q}$ is scaled by trading volume per bet $V \cdot H$, rather than by trading volume per calendar day V as implicitly assumed by conventional wisdom. When bet size $\tilde{Q}$ is scaled by VH, the resulting scaled bet size $\tilde{Z}$ has a mean of one and the same distribution for all stocks:

$$\frac{\tilde{Q}}{V \cdot H} \stackrel{\mathrm{d}}{=} \tilde{Z} \qquad \text{which implies} \qquad \frac{\tilde{Q}}{VH^*} = \left(\frac{W}{W^*}\right)^{-2/3} \cdot \tilde{Z}. \tag{10}$$

Equivalent bets transfer the same dollar risks in business time. In calendar time, they correspond to a smaller fraction of daily volume in markets with trading activity and thus shorter time interval between bet arrivals.

Price Impact. The logic of invariance can be also extended to market impact. Think of bets in two different markets as equivalent if they have the same scaled size $\tilde{Z} = \tilde{Q}/(VH)$. Bet cost invariance conjectures that equivalent bets have the same price impact when scaled by returns volatility in business time. Letitng $\Delta P/P$ denote price impact of a bet of size Q and letting $\sigma\sqrt{H}$ denote volatility in business time, bet cost invariance therefore implies an invariant price impact function $f()$ such that

$$\frac{\Delta P}{P} = \sigma \cdot \sqrt{H} \cdot f(Z), \qquad \text{where} \qquad Z = \frac{Q}{VH}. \tag{11}$$

If the price impact function is modeled as a power function $f(Z) = \alpha \cdot Z^{\beta}$ with proportionality constant α and exponent β, then equation (11) takes the form

$$\frac{\Delta P}{P} = \alpha \cdot \sigma \cdot \sqrt{H} \cdot \left(\frac{Q}{V \cdot H}\right)^{\beta}. \tag{12}$$

Plugging H from equation (9) and assuming linear market impact yields

$$\frac{\Delta P}{P} = \alpha \cdot \left(\frac{W}{\bar{C}}\right)^{1/3} \cdot \sigma \cdot \left(\frac{Q}{V}\right), \qquad \text{where} \qquad W := \frac{PV\sigma}{\bar{C}}, \text{ with } \beta = 1. \tag{13}$$

In comparison with conventional intuition (2) that bets of the same fraction of daily volume Q/V must have the same percentage price impact, holding volatility σ constant, the linear specification (13) has the additional factor $(W/\bar{C})^{1/3}$, which shows up due to the faster pace of business time in markets with higher trading activity. The factor $(W/\bar{C})^{1/3}$ implies that the demand becomes more inelastic as trading activity W increases. Increasing W by a factor of 1000 decreases the elasticity of demand by a factor of 10, reducing demand elasticity from say 0.10 to 0.01. As we shall see, this makes the demand elasticity small enough that observed order imbalances explain market crashes.

Intuition behind Equivalent Bets. We next compare the magnitude of selling pressure during five market crashes with sizes of large institutional orders executed in U.S. equities.

As a yardstick for measuring the size of institutional orders, we rely on the data on 400,000+ portfolio transition orders from Kyle and Obizhaeva (2016). A portfolio transition occurs when assets managed by one institutional asset manager are transferred to another manager. Trades converting the legacy portfolio into the new portfolio are typically handled by a professional third-party transition manager. Portfolio transitions represent some of the largest changes in portfolios held by institutional investors during the year. We estimate the distributions of portfolio transition orders to be symmetric around zero with unsigned order sizes close to log-normal random variables with different log-means and the same log-variance of 2.53:

$$\ln\left(\frac{\tilde{Q}}{V}\right) \sim \mathcal{N}\left(-5.71 - \frac{2}{3} \cdot \ln\left(\frac{W}{W^*}\right), 2.53\right). \tag{14}$$

The empirical distribution of $\ln(\widetilde{Q}/V)$ has a slope of $-2/3$ with respect to change in the log of trading activity $\ln(W)$, as expected given prediction (10).

Figure 1 is the main figure of our paper. It shows two types of extrapolation across markets; one is based on conventional wisdom and another is based on invariance. The vertical axis is the log of order size as a fraction of daily volume $\ln(Q/V)$. The horizontal axis is the log of scaled trading activity $\ln(W/W^*)$. The point $\ln(W/W^*) = 0$ corresponds to the benchmark stock with trading activity W^* from the bottom of the S&P 500 index. Trading activity varies by a factor of about 10^6 (= $\exp(12+2)$) from $\ln(W/W^*) = -12$ for the least actively traded stocks to $\ln(W/W^*) = 2.00$ for the most actively traded stocks such as Apple. Trading activity in the overall stock market is much higher, up to $\ln(W/W^*) = 8.20$; it is about 500 (= $\exp(8.20-2.0)$) times larger than that of the most liquid stocks.

The (black) horizontal lines show the extrapolation direction implied by the conventional wisdom. These iso-quants mark orders of sizes equal to a given percentage of calendar-day volume. For example, the horizontal line $|Q/V| = 5\%$ represents orders equal to 5% of daily volume.

The diagonal (red and green) lines with slopes of $-2/3$ show equivalent bets as implied by invariance. The lowest (red) line identifies log-medians of Q/V for different markets, as implied by invariance (14); this line intersects the vertical axis at -5.71, a point corresponding to a median bet in the benchmark stock equal to $\exp(-5.71)\cdot V$, or approximately $0.33\%\cdot V$. The six (green) parallel lines above the median line mark orders whose log-sizes are one through six standard deviations above the log-median sizes, respectively. Each log standard deviation represents an increase in bet sizes by a factor of $\exp(2.53^{1/2}) \approx 4.90$.

For each of the 60 months from January 2001 through December 2005, the 400,000+ portfolio transition orders are sorted into volume bins based on thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for NYSE-listed common stocks. The 600 blue diamonds in figure 1 represent the largest orders in each of 10 volume bins for each of 60 months. The diamond points form a cloud tilted along invariance-implied iso-lines with slope of $-2/3$. The dots are certainly not on a horizontal line, as would be predicted by the conventional wisdom. Since each bin contains on average about 650 points, invariance and log-normality of order size suggest that these largest portfolio transition orders should lie slightly below the 3-standard-deviation diagonal with predicted slope of $-2/3$. As can be seen visually from the figure, this is approximately the case. The scatter plot of largest portfolio transition orders confirms the predictions of the invariance hypothesis.

Figure 1 also depicts the five crash events by big round red dots. Extrapolating along horizontal lines, the conventional wisdom would say that these five events are not unusual. The two flash crashes are only 2.29% and 1.49% of daily volume. Even the largest crashes—the 1929 crash, the 1987 crash, and liquidation of Kerviel's position—represent "only" 265%, 67% and 28% of daily volume and only 1.36%, 0.63%, and 0.43% of market capitalization, respectively.

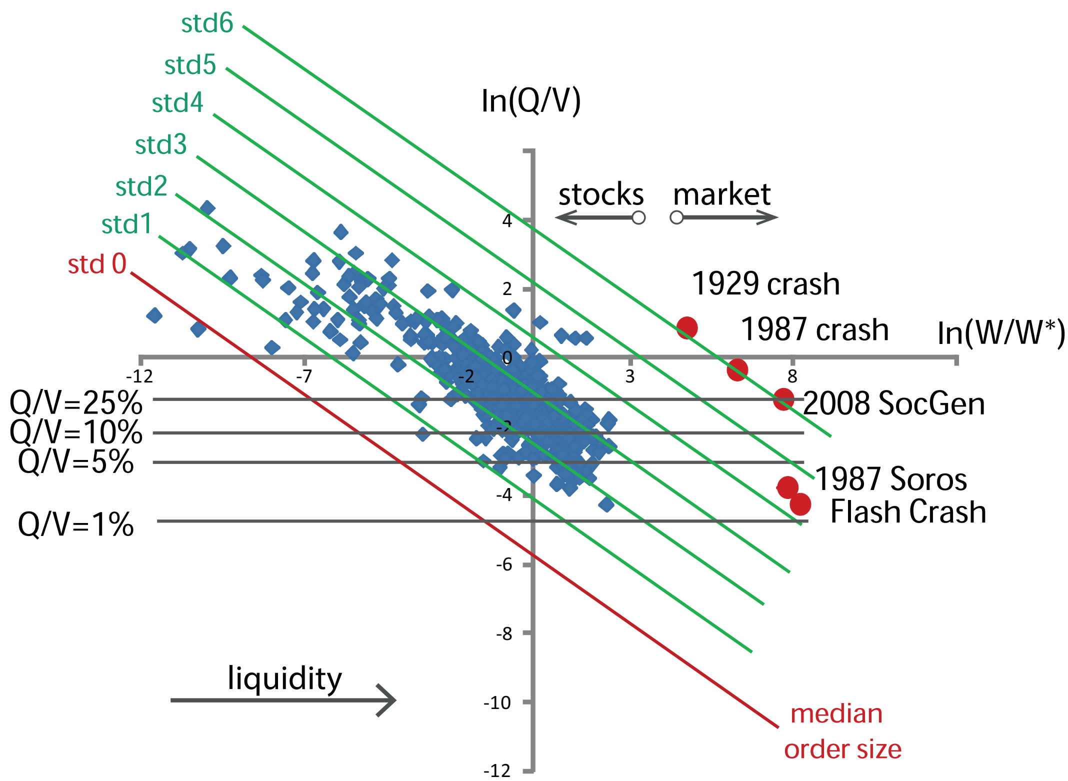

Figure 1: Largest Portfolio Transition Orders and Market Crashes.

This figure shows the largest portfolio transition orders for each month from January 2001 to December 2005 and for each of ten volume groups (blue points) as well as the bets during five market crashes (red points). Volume groups are based on thresholds corresponding to 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of dollar volume for common NYSE-listed stocks. The vertical axis is $|\ln(Q/V)|$. The horizontal axis is $\ln(W/W^*)$, where $W^* = 40 \cdot 10^6 \cdot 0.02$ and $W = P \cdot V \cdot \sigma$. The median order is $-5.71-(2/3)\cdot\ln(W/W^*)$ (red line). The x-standard deviation events are $-5.71-(2/3)\cdot\ln(W/W^*)+x\cdot\sqrt{2.53}$ (green lines).

These percentages of daily volume are not very different from what is seen in the largest portfolio transition orders, which are often about 25% of daily volume in liquid stocks and even larger percentage of daily volume (up to 700%) in illiquid stocks. Therefore, based on horizontal extrapolation, nothing unusual would be expected to happen during crash episodes.

Yet, in the context of invariance, extrapolation along diagonal lines with slope $-2/3$ shows that the crash events are extremely large, even compared to the largest institutional orders. The two flash crashes correspond to about 4.5-standard-deviation events. The 1929 crash, the 1987 crash, and the liquidation of Jérôme Kerviel's positions correspond to about 6-standard-deviation events. This suggests it would not be surprising that they caused significant price dislocations.

Intuition about Turnover and Elasticity. The two invariance principles are consistent with the intuition that all returns volatility results from the price impact of bets. Suppose, for simplicity, that all bets in the benchmark stock, with 100 bets per day, are the same size and equal to 1% of daily volume. Since the daily volatility of 200 basis points results from random impact of 100 independently distributed bets, each bet must have a price impact of 20 basis points ($\sigma/\sqrt{100 \text{ per day}}$).

Now consider what would happen if trading volume increases by a factor of 8, holding turnover constant. For the demand elasticity to remain constant, the price impact of must be a constant proportion of the bet size as a fraction of volume or market capitalization: $\Delta P/P = \alpha \cdot \sigma \cdot \bar{Q}/V$. If the number of bets increases from 100 per day to $\gamma > 100$ per day, then return volatility decreases to the product of the square root of the number of bets $\sqrt{\gamma}$ and the price impact of each bet $\alpha \cdot \sigma \cdot \bar{Q}/V = \alpha \cdot \sigma/\gamma$. We thus obtain daily return volatility equal to $\gamma^{-1/2} \cdot \sigma \cdot \gamma < \sigma$: Volatility is too small unless the number of bets does not change. Indeed, assuming the number of bets does not change is consistent with horizontal extrapolation implied by conventional wisdom in Figure 1.

For more active markets to have more bets with constant turnover, linear price impact requires the price impact formula to have an extra stock-specific impact factor which increases as trading activity increases. The assumptions of invariance accomplish this precisely with the extra factor proportional to $W^{1/3}$ in equation 13. The particular exponent $1/3$ makes the model satisfy the assumptions of invariance.

In our examples, we assume turnover is constant in business time. For the demand elasticity to be constant with linear price impact, it would be hypothetically possible to add an additional invariance principle that turnover is constant in business time, not calendar. This hypothesis is empirically implausible. For example, if dollar volume increases by a factor of 8, resulting in bet arrival γ increasing by a factor of 4, it can be shown that invariance would require the turnover rate to increase by a factor of 4 to keep the demand elasticity constant. This would require market turnover to be about 225 times greater than the turnover of typical individual stocks. Instead of making such a counterfactual empirical assumption, we allow dollar volume, volatility, and turnover to be characteristics which vary arbitrarily across different assets.

Invariance-Implied Market Impact Formulas. We use a log-linear version of the linear impact model (13). The expected percentage price impact from buying or quantity Q of a security with share price P, expected daily volume V and daily expected volatility σ is given by

$$\Delta \ln P = \frac{\bar{\lambda}}{10^4} \cdot \left(\frac{P \cdot V}{P^* \cdot V^*}\right)^{1/3} \cdot \left(\frac{\sigma}{\sigma^*}\right)^{4/3} \cdot \frac{Q}{H^* \cdot V}. \tag{15}$$

This formula assumes benchmark stock values $P^* = \$40$ per share, $V^* = 10^6$ shares per day, $\sigma^* = 0.02$ per day$^{1/2}$, with $H^* = 0.01$ day ≈ 4 minutes. The dimensionless proportionality factor $\bar{\lambda}$ is scaled so that $\bar{\lambda} = 5.00$ implies that the market impact of trading $Q/V = 1\%$ of daily volume in the benchmark stock has a price impact of 5 basis points. Invariance says that this factor $\bar{\lambda}$ is the same for all markets and time periods.

Kyle and Obizhaeva (2016) estimate $\bar{\lambda}$ using data on implementation shortfall of portfolio transition orders.[4] Introduced by Perold (1988), this metric is the difference between execution price and "paper trading" benchmark price recorded before the order was placed. The calibrated value of $\bar{\lambda}$ is equal to 5.00 with standard errors of 0.38. Thus, the predictions about price changes during crashes have percentage errors of about 7%.

We use a log-linear version of the market impact model rather than a simple linear model because our analysis deals with very large orders, sometimes equal in magnitude to trading volume of several trading days. In contrast, Kyle and Obizhaeva (2016) consider relatively smaller portfolio transition orders with an average size of 4.20% and median size of 0.57% of daily volume; for these smaller orders, the distinction between continuous compounding and simple compounding is immaterial.

[4](Kyle and Obizhaeva, 2016, equation (37), p. 1400) estimate an *average* impact cost parameter of $\bar{\kappa}_I = 2.50$ basis points (standard error 0.19 basis points) for transition orders, not price impact coefficient $\bar{\lambda}$ itself. Of course, there is a tight connection between the two concepts. Assuming that orders are broken into pieces and executed at prices which tend to increase along an upward sloping supply schedule, total price impact $\bar{\lambda} = 2 \times 2.50$ must be about twice the average impact cost $\bar{\kappa}_I$. Although invariance also has implications for bid-ask spread costs, these costs are negligible for large bets, and hence we ignore them. The implied standard error of $\bar{\lambda}$ is 2×0.19 basis points, i.e., about 7% of the estimate 2×2.50.

Equation (15) is a universal formula for market impact, which may be applied to different markets and time periods. Having calibrated it on the sample of portfolio transition orders in the individual U.S. stocks for 2001 through 2005, we next extrapolate the same formula to large market bets and to different time periods.

In Appendix A, we discuss several implementation issues that need to be addressed in order to apply invariance to data on the five crash events. The issues include defining boundaries of the market, choosing proxies for expected volume and volatility, and understanding functioning of market institutions. Appendix B reports estimates based on alternative models of market impact based on conventional wisdom and the literature on institutional trades.

Statistical Bias. Invariance explains why statistical estimates of the demand elasticity of stocks ranges from -3000 to -0.01—across more than five orders of magnitude. Scholes (1972) regressed percentage price changes on log-dollar-sizes of secondary stock offerings, obtaining a slope coefficient close to zero. According to microstructure invariance, this regression is misspecified. If secondary offerings are large bets, larger stocks should have offerings of larger dollar size, and such offerings should have lower price impact. This biases the slope coefficient downward. Invariance implies that the misspecified slope coefficient may be zero.[5] Studies of index additions and deletions are also biased to the extent that they do not take into account stock-specific characteristics. Studies of the price impact of institutional trading do tend to divide stocks into volume or capitalization groups. This mitigates the bias and explains why these studies obtain price impacts for individual stocks which are similar, if not identical, to our estimates based on portfolio transition orders.

[5]Consider a simple hypothetical example. There are three stocks with the same volatility and trading volumes which differ by a factor of 8: $5 million, $40 million, and $320 million per day. Consider three bets of size $200 000, $400 000, and $800 000 dollars in the stock with $40 million per day volume. Invariance implies that equivalent bet size doubles when volume increase by a factor of 8. Therefore, consider also three equivalent bets with half the volume in the less active stock ($100 000, $200 000, $400 000) and twice the volume in the more active stock ($400 000, $800 000, $1 600,000). Invariance implies that the price impacts of equivalent bets falls by a factor of 2 when volume increases by a factor of 8. Therefore, assume price impacts of these nine bets are 200, 400, 800 basis points in the least liquid stock; 100, 200, 400 basis points in the middle stock; and 50, 100, 200 basis points in the most liquid stock. Now, consistent with Scholes's methodology, consider a misspecified linear regression of log-price-impact on log-size and a constant term for these 9 data points. The slope coefficient is exactly zero. This is consistent with Scholes's near infinite elasticity estimate. The constant term implies price impact of approximately 200 basis points. (Kraus and Stoll, 1972, footnote 22, p. 577) pointed out that brokerage fees applicable to Scholes's data for secondary offerings were more than 400 basis points, which allowed the underwriter to absorb significant price impact into their fee without the price impact showing up in the immediate market.

Invariance also explains why the square root model may appear empirically reasonable, even if linear impact is the correct. Much of the variation in order size as a fraction of daily volume is crossectional: The fraction is low for more active stocks and high for less active stocks. To illustrate the bias, suppose that all bets for a particular stock are the same size. Invariance implies that increasing volume by a factor of 8 lowers bet size as a fraction of volume by a factor of 4 and lowers price impact by a factor of 2. Thus, a cross-sectional nonlinear regression of price impact on bet size as a fraction of volume yields exactly the square root model of price impact. It is statistically difficult to distinguish linear from square root price impact within stocks.

A Market Crash Scenario. Suppose both illiquid and liquid assets have annual turnover of 100% over 250 trading days. The first is a benchmark stock with $P^* = \$40$ per share, $V^* = 10^6$ shares per day, and $\sigma^* = 0.02$ per day$^{1/2}$. The second is the entire U.S. stock market, which consists of both the stock index futures market and cash stock market. The market has daily dollar volume of about $P \cdot V = \$270$ billion per day, about 6750 ($= 15^3 \cdot 2$) times the dollar volume of the benchmark stock, and daily returns volatility $\sigma = 0.01$ per day$^{1/2}$, one-half of stock volatility. Since business times passes at a rate proportional to $(P \cdot V \cdot \sigma / \bar{C})^{2/3}$, the stock market operates about 225 times faster ($= (6,750 \cdot 1/2)^{2/3}$) than the market for the benchmark stock, implying $H = H^*/225$. If bets in the benchmark stock arrive about once every 4 minutes, bets in the entire market arrive about once per second.

Now let us compare a bet of 25% of daily volume in the benchmark stock with a bet of 25% of daily volume in the market as a whole. For the benchmark stock, the sale would be 250 000 shares worth \$10 million. For the market as a whole, the sale would be slightly less that \$70 billion. Such sales might represent the liquidation of a gigantic institutional position, similar in magnitude to the liquidation of Jérôme Kerviel's rogue trades in 2008, or it might represent many small investors withdrawing equity exposure from index mutual funds or ETFs over a short period of a few days.

The conventional wisdom predicts that the price impact of both bets would be miniscule. Since 100% turnover per year implies daily turnover of 0.40% of market capitalization, a bet of 25% of one day's volume represents 0.10% of market capitalization. Unit demand elasticity therefore imlies a price decline of 10 basis points, which the market would barely notice.

By contrast, the invariance-implied extrapolation (13) leads to very different predictions. For the individual stock, since one percent of daily volume implies a price impact of 5 basis points, linear impact implies that a sale of 25% of daily volume in the benchmark stock has a price impact of 125 basis points. The implied demand elasticity of 0.08 is far smaller than the elasticity of one which represents conventional wisdom and is consistent with the academic literature on the price impact of institutional bets.

The invariance-implied elasticity is much lower for a bet on the entire market, and the price impact is correspondingly greater. Equation 13 implies that since trading activity is higher by a factor of 15^3, price impact would be 15 times higher if volatility were the same. Since market volatility of 1% per day is half of the daily volatility of 2% for the benchmark stock, price impact is reduced by a factor of 2 from 15 to 7.5 times the price impact of 125 basis points for the individual stock. The price impact of a bet about $70 billion in the market as a whole is therefore about 937 basis points, similar to the price declines observed when Kerviel's trades were liquidated. The implied demand elasticity for the market as a whole is only 0.01, about 7.5 times smaller than for an individual stock and 100 times smaller than conventional wisdom.

To summarize, we do not disagree with price impact estimates for individual U.S. stocks based on institutional trades; instead, we suggest an alternative way of extrapolating price impact from individual stocks to the overall stock market.

Our calculations suggest that the overall stock market is much more fragile than most economists believe. Sudden equity index ETF or mutual fund liquidations of $200 billion over a few days would potentially result in a 30% crash in stock prices, matching the crash of 1987.

3 Examples of Five Market Crashes

The actual price changes during crash events reflect not only sales by particular groups of traders placing large bets but also many other events occurring at the same time, including arrival of news, trading by other traders, and functioning of trading infrastructure. We next discuss each of these five episodes.

3.1 The Stock Market Crash of October 1929

The stock market crash of October 1929 is the most infamous crash in the history of the United States. It became associated with even larger stock price declines from 1930 to 1932, bank runs, and the Great Depression.[6]

The Dow Jones average declined by about 25% during the last week of October 1929 (from 305.85 on October 23 to 230.07 on October 29) and 34% during the last three months of 1929 (from 352.57 on September 25 to 234.07 on December 25). These price changes included a 11% drop in the morning on Black Thursday, October 24; a 13% drop on Black Monday, October 28; and another 12% drop on Black Tuesday, October 29.

In the late 1920s, many Americans became heavily invested in a stock market boom. A significant portion of stock investments was made in leveraged margin accounts. Between 1926 and 1929, both the level of margin debt and the level of the Dow Jones average doubled in value. Both the stock market boom and the boom in margin lending came to an abrupt end during the last week of October 1929.

During the week before Black Thursday, October 24, the Dow Jones average fell 9%, including a drop of 6% on Wednesday, October 23, and this led to a self-reinforcing cycle of liquidations of stocks in margin accounts.

To quantify the margin selling which occurred during the last week of October 1929, we follow the previous literature and contemporary market participants by estimating margin selling indirectly from data on broker loans and bank loans collateralized by securities. For the last week of October 1929, we estimate margin selling as $1.181 billion. For the three months from September 30, 1929, to December 31, 1929, we estimate total margin selling as $4.348 billion. Details of the estimations are presented in the Appendix C.

These liquidations exerted downward price pressure on the stock market. To estimate its magnitude, we treat the 1929 stock market as one market, rather than numerous markets for different stocks, and plug estimates of expected dollar volume and volatility for the entire stock market into equation (15).

Historical volatility during the month prior to October 1929 was about 2.00% per day. Historical volume was $342.29 million per day in 1929 dollars. Prior to 1935, the volume reported on the ticker did not include "odd-lot" transactions and "stopped-stock" transactions, which have been estimated to be equal about 30% of "reported" volume (Board of Governors of the Federal Reserve System, 1943, p. 431). We thus multiply reported volume by 13/10, obtaining

[6] Our analysis is based on several documents: Board of Governors of the Federal Reserve System (1929, 1927–1931); Galbraith (1954); Senate Committee on Banking and Currency (1934); Friedman and Schwartz (1963); Smiley and Keehn (1988); Haney (1932).

an estimate of \$444.97 million per day. The margin sales of \$1.181 billion during the last week of October were approximately 265% of average daily volume.

Equation (15) implies that margin-related sales of \$1.181 billion were expected to trigger a price decline of 46.43%:[7]

$$46.43\% = 1 - \exp\left(-\frac{5.00}{10^4}\cdot\left(\frac{444.97\cdot 10^6\cdot 9.42}{40\cdot 10^6}\right)^{1/3}\cdot\left(\frac{0.0200}{0.02}\right)^{4/3}\cdot\frac{1.181\cdot 10^9}{(0.01)(444.97\cdot 10^6)}\right).$$

As a robustness check, Table 2 reports other estimates using historical trading volume and volatility calculated over the preceding m months, with m = 1, 2, 3, 4, 6, 12. Invariance predicts price declines ranging from 26.79% to 46.43%, only slightly larger than the actual price change of 25%.[8]

Table 2: 1929 Stock Market Crash: Implied Price Impact of Margin Sales.

	Months Preceding 24 October 1929					
m:	1	2	3	4	6	12
ADV (in 1929-\$M)	444.97	461.45	436.49	427.20	387.18	390.45
Daily Volatility	0.0200	0.0159	0.0145	0.0128	0.0119	0.0111
10/24–10/30 Sales (%ADV)	265%	256%	271%	276%	305%	302%
Price Impact	46.43%	36.26%	33.75%	29.93%	29.00%	26.79%
9/25–12/25 Sales (%ADV)	977%	942%	996%	1,018%	1,123%	1,114%
Price Impact	89.95%	80.95%	78.04%	73.01%	71.66%	68.28%

Table 2 shows the implied impact of \$1.181 billion of margin sales during the week October 24–30, 1929, and \$4.343 billion of margin sales from September 25 to December 25, along with average daily 1929 dollar volume and daily volatility for m = 1, 2, 3, 4, 6, 12 months preceding October 24, 1929. The conventional wisdom predicts a price decline of 1.36% from October 24–29 and 4.99% from September 25 to December 25. The actual price decline was 25% from October 24–29 and 34% from September 25 to December 25.

[7] To convert 1929 dollars to 2005 dollars, we use the GDP deflator which equates \$1 in 1929 to \$9.42 in 2005. We use the year 2005 as a benchmark because the estimates of Kyle and Obizhaeva (2016) are based on the sample period 2001–2005, with more observations occurring in the latter part of the sample.

[8] Given percentage standard errors of impact estimate of 7%, the 2-standard deviation interval is 46.43%$(1 \pm 2\times$ %7); it is above the actually observed price decline of 25%.

In contrast, since the reduction of broker loans of $1.181 billion was only a very small fraction of the $87.1 billion market capitalization of NYSE issues at the end of September 1929 (Brady Report, p. VIII-13), conventional intuition (1) predicts a price change of only 1.36%, much smaller than actual price decline of 25% and about 40 times smaller than the magnitude predicted by invariance.

We also make price impact calculations for margin sales of $4.348 billion during the last three months of 1929. Conventional wisdom implies a price drop of 4.99%. Invariance implies a much larger price decline ranging from 68.28% to 89.95%, more than the actual price decline of 34% during the last three months of 1929 and the price decline of 44% from high point in late September 1929 to low point in mid November 1929.

3.2 The Market Crash in October 1987

From Wednesday, October 14, 1987, to Tuesday, October 20, 1987, the U.S. equity market suffered the most severe one-week decline in its history. The Dow Jones index dropped 32% from 2,500 to 1,700; as of noon Tuesday, October 20, the S&P 500 futures prices had dropped about 40% from 312 to 185. On Black Monday alone, October 19, 1987, the Dow Jones index fell 23%, and the S&P 500 futures market dropped 29%.

It has long been debated whether the market crash resulted from the sales by institutions implementing portfolio insurance. Portfolio insurance is a trading strategy that replicates put option protection for portfolios by dynamically adjusting stock market exposure in response to market fluctuations. Since portfolio insurers sell stocks when prices fall, the strategy amplifies downward pressure on prices in falling markets. We calculate the price impact of portfolio insurance sales implied by invariance.

We consider the entire stock market to be one market; this is consistent with the Brady Report. Accordingly, we estimate daily volume as the sum of average daily volume in the futures market and the NYSE for the previous month. Some portfolio insurers abandoned their reliance on the futures markets and switched to selling stocks directly because futures contracts became unusually cheap relative to the cash market. We construct a proxy for sales as the sum of portfolio insurance sales in the futures market and the NYSE from tables in the Brady Report, figures 13–16, pp. 197–198, obtaining results similar to Gammill and Marsh (1988).

Over the four days October 15, 16, 19, 20, 1987, portfolio insurers sold S&P 500 futures contracts representing \$10.48 billion in index futures and \$3.27 billion in NYSE stocks. The gross sales amount of \$13.75 billion in futures and stocks are combined for the purpose of analyzing price impact of portfolio insurance sales. Reported values are all 1987 dollars.

In the month prior to the crash, the historical volatility of S&P 500 futures returns was about 1.35% per day, similar to estimates in the Brady Report. The average daily volume in the S&P 500 futures market was equal to \$10.37 billion. The NYSE average daily volume was \$10.20 billion. Portfolio insurance gross sales were equal to about 67% of one day's combined volume.

Plugging portfolio insurance gross sales and market parameters into equation (15) yields a price decline of 16.77%:[9]

$$16.77\% = 1 - \exp\left(-\frac{5.78}{10^4}\cdot\left(\frac{(10.37+10.20)\cdot 10^9\cdot 1.54}{40\cdot 10^6}\right)^{1/3}\cdot\left(\frac{0.0135}{0.02}\right)^{4/3}\cdot\frac{(10.48+3.27)}{(0.01)(10.37+10.20)}\right).$$

Table 3 reports other estimates based on historical trading volume and volatility calculated over the preceding m months, with $m = 1, 2, 3, 4, 6, 12$. These estimates range from 11.87% to 16.77%. For robustness, estimates under several alternative assumptions are presented in Table 3. Details are presented in Appendix C. The similarity between predicted and observed price declines is consistent with our hypothesis that heavy selling by portfolio insurers played a dominant role in the crash of October 1987.

The estimates based on conventional wisdom are much smaller. According to the Brady Report there were 2,257 issues of stocks listed on the NYSE, with a value of \$2.2 trillion on December 31, 1986. Conventional wisdom implies that gross sales of \$10.48 billion in futures and \$3.27 billion in individual stocks, representing 0.63% of shares outstanding in total, would have an impact of only 0.63%. Other alternative models yield estimates not higher than 2%. Citing similar arguments, many experts have rejected the idea that sales of portfolio insurers caused the 1987 market crash.

Invariance-implied estimates are somewhat smaller than the price drops of 32% in the cash equity market and 40% in the S&P 500 futures market observed during the 1987 market crash. The price decline may have been triggered by negative news about anti-takeover legislation as well as trade deficit statistic on October 14. It may be further aggravated by break-downs in the market mechanism which disrupted index arbitrage relationships, as documented in the Brady Report.

[9]The GDP deflator of 1.54 converts 1987 dollars to 2005 dollars.

Table 3: 1987 Stock Market Crash: Effect of Portfolio Insurance Sales.

	Months Preceding 14 October 1987					
m:	1	2	3	4	6	12
S&P 500 Fut ADV (1987-$B)	10.37	11.29	11.13	10.12	10.62	9.85
NYSE ADV (1987–$B)	10.20	10.44	10.48	10.16	10.04	9.70
Daily Volatility	0.0135	0.0121	0.0107	0.0102	0.0112	0.0111
Gross Sells (% ADV)	66.84%	63.28%	63.65%	67.82%	66.53%	70.33%
Price Impact	16.77%	14.18%	12.23%	11.87%	13.20%	13.64%
Price Impact of Net Sales Combined	13.78%	11.62%	10.00%	9.71%	10.81%	11.17%
Price Impact of S&P 500 Sales	14.11%	11.67%	10.10%	10.00%	10.93%	11.45%
Price Impact of NYSE Sales	13.00%	11.18%	9.56%	9.09%	10.32%	10.53%

Table 3 shows the implied impact triggered by portfolio insurers' net sales of S&P 500 futures contracts ($9.51 billion) and NYSE stocks ($1.60 billion), portfolio insurers' gross sales of S&P 500 futures contracts ($10.48 billion) and NYSE stocks ($3.27 billion), portfolio insurers' sales of S&P 500 futures adjusted for purchases of index arbitrageurs ($10.48 billion minus $3.27 billion), and portfolio insurers' sales of NYSE stocks adjusted for sales of index arbitrageurs ($3.27 billion plus $3.27 billion) in 1987 dollars. Average daily dollar volume and daily volatility are based on m months preceding October 14, 1987, with $m = 1, 2, 3, 4, 6, 12$, both for the S&P 500 futures and CRSP stocks. Conventional wisdom predicts price declines of 0.51% for portfolio insurers' net sells and 0.63% for their gross sells. The actual price decline was 32% for the Dow Jones average and 40% for S&P 500 futures.

3.3 Trades of George Soros on October 22, 1987

On Thursday, October 22, 1987, just three days after the 1987 market crash, George Soros lost $60 million in minutes by selling a large number of S&P 500 futures contracts as prices spiked down 22% at the opening of trading. These sales have been attributed to pessimistic predictions that Robert Prechter made based on "Elliott Wave Theory" and similarities between the 1929 crash and the 1987 crash.

The Commodity Futures Trading Commission (1988) issued a report describing the events of October 22, 1987, without mentioning Soros by name. At 8:28 a.m. CT, approximately two minutes before the opening bell at the NYSE, a customer of a clearing member submitted a 1,200-contract sell order at a limit price of 200, more than 20% below the previous day's close of 258. Over the first minutes of trading, the price dropped to 200, at which point the sell order was executed. At 8:34 a.m., a second identical limit order for 1,200 contracts from the same customer was executed by the same floor broker. These transactions liquidated a long position acquired on the previous day at a loss of about 22%, or about $60 million in 1987 dollars. Within minutes, S&P 500 futures prices rebounded and, over the next two hours, the market recovered to the levels of the previous day's close. Within days, Soros's Quantum Fund sued the brokerage firm which handled the order, alleging a conspiracy among traders to keep prices artificially low while his sell orders were executed.

Two other events may have exacerbated the decline in prices in the morning of October 22 by increasing the selling pressure. First, when the broker executed the second order, he mistakenly sold 651 more contracts than the order called for. The oversold contracts were taken into the clearing firm's error account and liquidated at a significant loss to the broker. Second, the Commodity Futures Trading Commission (1988) reports that between 9:34 a.m. and 10:45 a.m. the same clearing firm also entered and filled four large sell orders for another customer—a pension fund—with a total of 2,478 contracts sold at prices ranging from 230 to 241. These additional orders are for almost exactly the same size as Soros's orders. This fact suggests information leakage or coordination regarding the size of these unusually large orders.

We compare the actual price decline of 22% with predictions based on invariance. During the prior month, average daily volatility was 8.63%, and average daily volume in the S&P 500 futures market was $13.52 billion in 1987 dollars. The very high volatility estimate based on crash data is reasonable because market participants expected high volatility to persist. Since Soros's sales started just before the opening of NYSE trading, the arbitrage mechanism which connects stock and futures markets did not have time to work; indeed, futures contracts traded at levels about 20% cheaper than stocks. We thus consider only S&P 500 futures market in this example, not combining it with the market for NYSE stocks.

Each S&P 500 contract had a notional value of 500 times the S&P 500 index. With an S&P 500 level of 258, one contract represented ownership of about $129,000. Soros' sale of 2,400 contracts, or about $309.60 million, was equal to 2.29% of average daily volume. Given the prior month estimates, equation (15) predicts a price decline of 6.27%:[10]

$$6.27\% = 1 - \exp\left(-\frac{5.00}{10^4}\cdot\left(\frac{13.52\cdot 10^9\cdot 1.54}{40\cdot 10^6}\right)^{1/3}\cdot\left(\frac{0.0863}{0.02}\right)^{4/3}\cdot\frac{309.60\cdot 10^6}{(0.01)(13.52\cdot 10^9)}\right).$$

Table 4 presents three sets of estimates based on the historical volume and volatility of S&P 500 futures contracts calculated over the preceding m months, with $m = 1, 2, 3, 4, 6, 12$. Invariance implies (A) price impact of 1.67% to 6.27% based on 2,400 contracts alone; (B) price impact of 2.12% to 7.90% adding 651 error contracts (3,051 contracts in total); and (C) price impact of 3.81% to 13.85% adding 2,478 contracts sold by the pension fund (5,529 contracts in total).

Table 4: October 22, 1987: Effect of Soros's Trades.

	Months Preceding 22 October 1987					
m:	1	2	3	4	6	12
S&P 500 Fut ADV (1987-$B)	13.52	11.72	11.70	10.99	10.75	10.04
Daily Volatility	0.0863	0.0622	0.0502	0.0438	0.0365	0.0271
2,400 contracts as %ADV	2.29%	2.64%	2.65%	2.82%	2.88%	3.08%
Price Impact A	6.27%	4.50%	3.40%	2.96%	2.36%	1.67%
Price Impact B	7.90%	5.68%	4.30%	3.75%	2.99%	2.12%
Price Impact C	13.85%	10.06%	7.66%	6.69%	5.36%	3.81%

Table 4 shows the implied price impact of (A) Soros's sell order of 2,400 contracts; (B) Soros's sell order of 2,400 contracts plus 651 contracts of error trades (3,051 contracts in total); and (C) Soros's sell order of 2,400 contracts, plus 651 contracts of error trades, plus the sell order of 2,478 contracts by the pension fund (5,529 contracts in total). The calculations assume average daily 1987 dollar volume and daily volatility for $m = 1, 2, 3, 4, 6, 12$ months preceding October 22, 1987 for the S&P 500 futures contracts. Conventional wisdom predicts price declines of 0.01%, 0.02%, and 0.03%, respectively. The actual price decline in the S&P 500 futures market was 22%.

The actual price decline of 22% is significantly larger than our estimate. Factors which could have led to large impact include potentially underestimated expected volatility, front-running based on leakage of information about the size of the order, and the peculiarly aggressive execution strategy of placing two limit orders with a limit price of 200, more than 20% below the

[10]The GDP deflator of 1.54 converts 1987 dollars to 2005 dollars.

previous day's close.

Conventional wisdom would imply minuscule price changes. Given the total value of $2.2 trillion of issues listed on the NYSE at the end of 1986, the Soros's order, the erroneous sales, and the sales by the pension fund would be expected to have a combined impact of only 0.03%.

3.4 Liquidation of Kerviel's Rogue Trades in January 2008

On January 24, 2008, Société Générale issued a press release stating that the bank had "uncovered an exceptional fraud." Subsequent reports by Société Générale (2008a,b,c) revealed that rogue trader Jérôme Kerviel had used "unauthorized" trading to place large bets on European stock indices.

Kerviel had established long positions in equity index futures contracts with underlying values of €50 billion: €30 billion on the Euro STOXX 50, €18 billion on DAX, and €2 billion on the FTSE 100. He acquired these naked long positions mostly between January 2 and January 18, then concealed them using fictitious short positions, forged documents, and emails suggesting his positions were hedged. The fall in index values in the first half of January led to losses on these secret directional bets. Internal investigators became strongly suspicious about the nature of the positions on Friday, January 18.

Société Générale informed the heads of the central bank and the Financial Markets Authority (AMF), the French stock market regulator. The AMF allowed the bank to delay public announcement of the fraud for three days, so that Kerviel's positions could be liquidated quietly. The head of the central bank also delayed informing the government. After liquidating the positions between Monday, January 21, and Wednesday, January 23, the bank had sustained losses of €6.4 billion which—after subtracting out €1.5 billion profit as of December 31, 2007—were reported as a net loss of €4.9 billion.

As Société Générale liquidated the positions, prices fell all across Europe. The Stoxx Europe Total Market Index (TMI)—which represents all of Western Europe—fell by 9.44% from the close on January 18 to its lowest level on January 21. On Monday, January 21—a bank holiday with muted U.S. financial markets activity—the Fed held an unscheduled FOMC meeting via conference call at 6:00 p.m. New York time, several days before its scheduled meeting. At 8:30 a.m. the next day, the Fed announced an unprecedented 75-basis point cut in interest rates, which pushed all prices up and helped Société Générale to liquidate the rest of the position on better terms. We do not know whether Fed officials were aware of Société Générale's situation when this decision was made. According to the Fed's Minutes, published five years later, the purpose of the meeting was to "to update the Committee on financial developments over the weekend and to consider whether we want to take a policy action," but there is no mention of Société Générale. In his memoir, Bernanke (2015, pp. 195–196) said the Fed "had no idea the rogue-trading bombshell was coming." Yet, he mentions (p. 195) "a conference call the morning of January 19 Paris time," during which "senior SocGen managers in Paris and New York had told New York Fed supervisors that the bank would report positive earnings for the fourth quarter, even after taking write-downs on its subprime mortgage exposure."

On the one hand, the surprise early announcement of an interest rate cut could have helped the bank obtain more favorable execution prices on some portion of its trades. On the other hand, January 21 was a bank holiday in the United States; in the previous year, the futures markets had only one third of the typical volume on days when U.S. markets were closed. Low volume on the bank holiday could have reduced liquidity, making the unwinding of positions more expensive.

Due to significant correlations among European markets, we perform our analysis under the assumption that all European stock and futures markets are one market. Based on data from the World Federation of Exchanges, the seven largest European exchanges by market capitalization in 2008 (NYSE Euronext, London Stock Exchange, Deutsche Börse, BME Spanish Exchanges, SIX Swiss Exchange, NASDAQ OMX Nordic Exchange, Borsa Italiana) had average daily volume for the month ending January 18, 2008 equal to €69.51 billion.

We also sum average daily volume across the ten most actively traded European equity index futures markets (Euro Stoxx 50, DAX, CAC, IBEX, AEX, Swiss Market Index SMI, FTSE MIB, OMX Stockholm 30, Stoxx 50 Euro) and find average daily futures volume of €110.98 billion. The total daily volume in both European stock and equity futures markets was equal to €180.49.

Our estimate of expected volatility is 1.10%, the previous month's daily standard deviation of returns for the Stoxx Europe Total Market Index (TMI).

According to equation (15), the liquidation of Kerviel's €50 billion position—equal to about 27.70% of the average daily volume in aggregated stock and futures markets—is expected to trigger a price decline of 10.79% across European markets:

$$10.79\% = 1 - \exp\left(-\frac{5.00}{10^4}\cdot\left(\frac{180.49\cdot 1.4690\cdot 0.92\cdot 10^9}{40\cdot 10^6}\right)^{1/3}\left(\frac{0.0011}{0.02}\right)^{4/3}\frac{50}{(0.01)\cdot 180.49}\right).$$

In this equation, we use an exchange rate of \$1.4690 per Euro to convert Euro volume into U.S. dollar volume and convert 2008 dollars into 2005 dollars to be able to use them in our calibrated formulas.[11]

Table 5 shows the estimates of price impact based on historical trading volume and volatility calculated over the preceding m months, with $m = 1, 2, 3, 4, 6, 12$. Invariance predicts price changes ranging from 10.59% to 12.93%. The Stoxx TMI index actually fell by 9.44% from the market close of 316.73 on January 18 to its lowest level of 286.82 on January 21.

Table 5: January 2008: Effect of Liquidating Kerviel's Positions.

	Months Preceding January 18, 2008					
m:	1	2	3	4	6	12
Stk Mkt ADV (2008–€B)	69.51	66.51	67.37	67.01	66.73	66.32
Fut Mkt ADV (2008–€B)	110.98	114.39	118.05	117.46	127.17	121.26
Daily Volatility	0.0110	0.0125	0.0121	0.0117	0.0132	0.0111
Order as %ADV	27.70%	27.64%	26.97%	27.11%	25.79%	26.66%
Price Impact	10.79%	12.66%	11.94%	11.53%	12.93%	10.59%
Total Losses (2008–€B)	2.77	3.27	3.08	2.97	3.34	2.72
Losses: Adj A (2008–€B)	5.08	5.58	5.39	5.28	5.65	5.03
Losses: Adj B (2008–€B)	7.39	7.89	7.70	7.59	7.96	7.34

Table 5 shows the predicted losses of liquidating Kerviel's positions of €50 billion under the assumption that the major European cash and futures markets are integrated, and one Euro is worth \$1.4690. Results are provided based on average daily volume of the major European stock exchanges and index futures as well as daily volatilities of Stoxx Europe TMI, based on m months preceding January 18, 2008, with $m = 1, 2, 3, 4, 6, 12$. Conventional wisdom predicts price decline of 0.43%. The actual price decline in the Stoxx Europe TMI was 9.44%.

In contrast, conventional wisdom predicts that sales of €50 billion would have a much smaller price impact of 0.43%, given that it represents less than one percent of the total capitalization of European markets, which was about €11.752 trillion in December 2007, as reported by Federation of European Securities Exchanges.

We also estimate the dollar costs of liquidating the rogue position to range from €2.72 bil-

[11] The GDP deflator of 0.92 converts 2008 dollars into 2005 dollars.

lion to €3.34 billion under different assumptions about expected volume and volatility. Adding mark-to-market losses sustained prior to liquidation leads to estimated losses ranging from €5.03 billion to €7.96 billion. These estimates are consistent with officially reported losses of €6.30 billion. Appendix C presents the details.

In explaining the costs of liquidating the positions to shareholders already concerned about the bank's losses on subprime mortgages, bank officials blamed "the very unfavorable market conditions" (see the explanatory note about the exceptional fraud released by Société Générale on January 27). Expressing conventional wisdom, the bank announced that its trades accounted for not more than 8% of turnover on any one of the futures exchanges on which they were conducted and thus did not have a serious market impact. When examined through the lens of invariance, the reported losses are of the magnitude expected from the price impact of the trades on European stock markets.

3.5 The Flash Crash of May 6, 2010

During the morning of May 6, the S&P 500 declined by 3%. Rumors of a default by Greece had made markets nervous in a context where there was already uncertainty about elections in the U.K. and an upcoming jobs report in the U.S.

During the five minute interval from 2:40 p.m. to 2:45 p.m. ET, the E-mini S&P 500 futures contract suddenly dropped 5.12% from 1,113 to 1,056. After a pre-programmed circuit breaker built into the CME's Globex electronic trading platform halted trading for five seconds, prices went up 5% over the next ten minutes, recovering losses.

After the Flash crash, the Staffs of the CFTC and SEC (2010a,b) issued a joint report. It said that an automated execution algorithm sold 75,000 S&P 500 E-mini futures contracts between 2:32 p.m. and 2:51 p.m. on the CME's Globex platform, exactly during the V-shaped flash crash. The E-mini contract represents exposure of 50 times the S&P 500 index, one tenth the multiple of 500 for the older but otherwise similar contract sold by portfolio insurers in 1987. Given the S&P 500 index values, the program sold S&P 500 exposure of $4.37 billion. The joint report did not mention the name of the seller, but journalists identified the seller as Waddell & Reed.

Many people did not believe that selling 75,000 contracts could have triggered a price decline of 5%, because they implicitly relied on conventional intuition: The $4.37-billion sale represented only 3.75% of the daily volume of about 2,000,000 contracts per day in the S&P 500

E-mini futures market. Thus, many accused high frequency traders of failing to provide liquidity as prices collapsed.

We examine whether such an order could have resulted in a flash crash. During the preceding month, the volume in E-mini contracts was about $132 billion per day, and the volume in the stock market was about $161 billion per day; the combined daily volume was $292 billion. The historical daily volatility was 1.07%.

Equation (15) implies that the sales of $4.37 billion—equal to about 3.31% of daily volume in S&P 500 E-mini futures market in the previous month or 1.49% for futures and stock market combined—is expected to trigger a price decline of 0.61%:[12]

$$0.61\% = 1 - \exp\left(-\frac{5.00}{10^4}\cdot\left(\frac{(132+161)\cdot 0.90\cdot 10^9}{40\cdot 10^6}\right)^{1/3}\cdot\left(\frac{0.0107}{0.02}\right)^{4/3}\cdot\frac{75,000\cdot 50\cdot 1,164}{0.01\cdot(132+161)\cdot 10^9}\right).$$

Table 6 shows estimates based on historical volume and volatility of S&P 500 E-mini futures contracts calculated over the preceding m months, with $m = 1,2,3,4,6,12$. These estimates range from 0.44% to 0.73%. Appendix C provides estimates under the alternative assumptions of two-percent volatility and less integrated markets. These estimates tend to be higher, ranging from 0.76% to 2.91%.

The predicted price impact of 0.61% is smaller than the actual decline of 5.12%. As discussed later, unusually fast execution may have significantly increased the temporary impact of these trades and led to rapid rebound in prices afterwards. Given that the capitalization of U.S. market was about $15.077 trillion at the end of 2009, conventional wisdom would predict an even smaller price decline of 0.03%.

4 Policy Implications and Lessons Learned

Application of microstructure invariance concepts to intrinsically infrequent historical episodes requires an exercise in judgement to extract appropriate lessons learned. In some cases theory implies values that differ from actual price declines. We discuss next factors that can potentially explain these discrepancies. While speculative in nature, our discussion suggests important lessons for policymakers concerned with measuring and predicting crash events of a systemic nature, for asset managers worried about managing market impact costs associated with exe-

[12] The GDP deflator of 0.90 converts 2010 dollars into 2005 dollars.

Table 6: Flash Crash of May 6, 2010: Effect of 75,000 Contract Futures Sale.

	Months Preceding 6 May 2010					
m:	1	2	3	4	6	12
S&P 500 Fut ADV (2010 $B)	132.00	107.49	109.54	112.67	100.65	95.49
Stk Mkt ADV (2010 $B)	161.41	146.50	142.09	143.03	132.58	129.30
Daily Volatility	0.0107	0.0085	0.0078	0.0090	0.0089	0.0108
Order as %ADV	1.49%	1.72%	1.73%	1.71%	1.87%	1.94%
Price Impact (hist σ)	0.61%	0.49%	0.44%	0.53%	0.55%	0.73%
Price Impact ($\sigma = 2\%$)	1.39%	1.52%	1.53%	1.52%	1.61%	1.65%

Table 6 shows the predicted price impact of sales of 75,000 S&P 500 E-mini futures contracts. Calculations are based on average daily volume and volatility of the S&P 500 E-mini futures for the m months preceding January 18, 2008, with $m = 1, 2, 3, 4, 6, 12$. Conventional wisdom predicts a price decline of 0.03%. The actual price decline in the S&P 500 E-mini futures market was 5.12%.

cution of large trades that might potentially disrupt markets, and for researchers interested in understanding of how financial markets work.

Price Impact is Large in Liquid Markets. Market participants often execute large orders by restricting quantities traded to be not more than five or ten percent of average daily volume over a period of several days. This heuristic strategy is believed to be reasonable for individual stocks and thus certainly for more liquid markets such as markets for stock index futures.

We disagree. While this strategy is reasonable for trading in individual stocks, our analysis shows that it may incur much larger-than-expected costs when implemented in more liquid markets. The price impact of trading a given fraction of daily volume (in volatility units) is proportional to the cube root of dollar volume and returns volatility. For example, if dollar volume $P \cdot V$ increases by a factor of 1000—approximately consistent with difference between a benchmark stock and stock index futures—market impact of a given fraction of volume increases by a factor of $(1000)^{1/3} = 10$ in equation (15). The larger estimates stem from assumptions that invariance hypothesis holds and impact is linear in bet size.

Rapid Execution Magnifies Transitory Price Impact. Our analysis shed some light on how price dynamics in response to large bets depends on the speed of execution. While the long-term impact of bets is likely to depend on their information content, the short-term price dynamics is probably affected by the speed of trading.[13] Speeding up execution exacerbates temporary impact associated with V-shaped price paths, in which prices first sharply plunge and then rapidly recover.

The model of smooth trading of Kyle, Obizhaeva and Wang (2014) provides a theoretical framework for modeling short-term price reactions to unusually rapid execution of large bets.[14] The model implies that markets interpret extremely rapid, heavy selling as an indication that extremely negative information is about to flow into the market. Prices collapse immediately when a heavy rate of selling is detected. When the expected negative information does not materialize, prices rebound, even though much of the heavy selling continues.

The impact formula (15) contains parameters calibrated using the data on portfolio transition orders. Most transitions were executed over a period of a few days, and only the most complex of them were carried out over a period of a few weeks. Their executions at a prudent pace were designed to keep impact costs low. Extrapolating estimates from portfolio transitions to sales during crashes implicitly makes the identifying assumption that crash were also executed at the same "natural" pace.

Yet, during 1987 and 2010 flash crashes, larger-than-predicted price declines followed by rapid price recoveries suggest that transitory price impact may have been exacerbated by the extremely rapid rate at which selling took place. According to the Staffs of the CFTC and SEC (2010b), for example, the May 2010 flash crash order was executed extremely rapidly in just 20 minutes, while earlier two orders of similar magnitudes had been executed over periods of 5 and 6 hours, which is 15 times slower.[15]

[13] Financial crises eventually followed the crash events of 1929 and liquidation of Kerviel's rogue trades in 2008. Whether margin sales in 1929 or Kerviel's trades in 2008 had information content is a difficult question to frame in a meaningful manner. For example, perhaps Kerviel traded against informed traders who correctly foresaw the impending financial crisis, delaying the incorporation of this information into prices until his own trades were liquidated.

[14] The model of smooth trading gives rise in the equilibrium to both endogenous permanent and temporary impacts, λ and κ. In most of traditional models such as Kyle (1985), price impact is permanent and transaction costs of an informed trader do not depend on the speed of trading as long as he trades continuously.

[15] The smooth trading model implies temporary price impact is linear in the speed of trading. Since selling during the flash crash occurred about 15 times faster than normal order execution, the model implies transitory price impact to be 15 times greater than in the case when selling occurs at a "normal" rate, followed by a reversal. Mul-

Policy Responses to Mitigate Crashes. Some policy responses may help to mitigate negative effects of crashes. These policies have to aim at easing flow of credit as well as providing funds that will make up the gap in demand and supply. Our analysis suggests that the amount of funds necessary to counteract the shock must be comparable to the size of the shock itself.

A good example is the 1929 market crash, during which price declines were somewhat smaller than predicted by invariance and the crash was well contained until the end of 1929. First, immediately after the initial stock market break on Black Thursday, a group of prominent New York bankers put together an informal fund of about \$750 million to buy securities—similar in size to the margin sales shock of about \$1.181 billion—in order to support prices. When their decisions were publicized, the sense of panic subsided. Similar actions, for example, were undertaken by J.P. Morgan and other bankers after a crash in 1907.

Second, the New York Fed acted prudently in 1929 as well. In the 1920s, bankers and their regulators were aware that if non-bank lenders suddenly withdrew funds from the broker loan market, there would be pressure on the banking system to make up the difference. By discouraging banks from lending into the broker loan market prior to the 1929 crash, the New York Fed increased the ability of banks to support it after the stock market crashed. During the last week of October 1929, the New York Fed wisely reversed its course and encouraged banks to provide bank loans on securities to their clients as a substitute for broker loans. The unprecedented increase in demand deposits at New York banks gave them plenty of cash to use to finance increased loans on securities. The New York Fed also encouraged easy credit by purchasing government securities and cutting its discount rate twice. Some brokers cut margins from 40% to 20%.

All of these stabilizing policies smoothed the margin selling out and allowed brokers to liquidate the large positions of under-margined stock investors gradually over five weeks, rather than selling collateral off at fire-sale prices over several days. This appears to have helped the market to digest imbalances and reduce temporary price impact, thus avoiding a sudden, brutal bursting of the stock market bubble.

Effect of Large Bets May Propagate Across Integrated Markets. Financial markets are integrated. Our study suggests that heavy selling in one market is likely to affect correlated markets. This connectedness raises the question of how to define the boundaries of the market for the

tiplied by 15, our estimates of 0.61% price decline becomes even larger than the actual decline of 5.12%.

purpose of applying invariance.

For example, during the 1987 crash, not only U.S. markets but also many major world markets experienced severe declines, despite the fact that the portfolio insurance selling was confined to the former. According to Roll (1988), this justifies his opinion that portfolio insurance did not trigger the 1987 crash. The common patterns across markets were also documented during liquidation of Kerviel's positions in January 2008. Even markets where Société Générale did not liquidate positions had very similar performance. The bank thus argued that its own impact on prices was limited, and large price declines in multiple markets had to be attributed to other factors.

We disagree. The commonality of patterns suggests that market impact estimates should take into account how market liquidity is shared across markets in different continents and markets of economically related securities. It supports our preferred strategy for the analysis of the 1987 crash of looking at aggregated stock and futures markets. It also supports our analysis of Société Générale's case where we aggregate data across all European markets rather than focus only on isolated pools of liquidity for countries where the bank liquidated its position.

Efficiency versus Stability. There may be a trade-off between efficiency and stability. In less efficient trading arrangements, more capital is required to sustain orderly trading, but this capital also makes the systems more stable during volatile times. Invariance may help to assess the affect of market integration on liquidity.

The inefficiency of financial markets in 1920s may explain their remarkable resilience. During the 1929 crash, the gigantic amount of selling related to liquidation of margin loans was more than 15 times greater than selling during the 1987 crash, as a percentage of GDP. Yet, the price decline was only half as large.[16]

In the 1920s, speculative capital may have been compartmentalized into numerous separate silos. Speculative trading and intermediation associated with underwriting of new stock issues often took place in "pools." The pools were typically dedicated to trading only one stock, and investors in the pools often had close connections to the company whose stock the pool

[16]The 1987 portfolio insurance trades of $13 billion were equal to only about 0.28% of GDP in that year (1987 GDP was $4.7 trillion); stock prices fell 32%. During the last week of October 1929, the margin related sales of $1.181 billion were equal to about 1% of GDP (1929 GDP was $104 billion), approximately four times the levels of the 1987 crash; yet stock prices fell by only 25%. Inclusion of additional sales equal to about 3% of GDP in subsequent weeks makes margin selling in 1929 to be more than 15 times greater than selling during the 1987 crash, as a percentage of GDP.

traded; there were no prohibitions against insider trading and no SEC requiring firms to disclose material information to the market. These pools traded actively, used leverage, took short positions, and arbitraged stocks against options, particularly when facilitating distribution of newly issued equity. There were no futures markets or ETFs allowing investors to trade large baskets of stocks. When faced with massive liquidations of margin loans, the market may have more speculative capital available to stabilize the situation than in a more "efficiently" leveraged system in which institutional investors can spread their capital across markets by trading hundreds of stocks simultaneously.

Invariance-implied estimates would change significantly if instead of being considered as one large market, the stock market in 1929 were thought of as a set of many small, isolated, and thus less liquid markets for individual stocks. One would expect market impact to be much smaller in these less liquid markets for the same market bet. As a hypothetical illustration, suppose the 1929 stock market consisted of 125 separate markets for 125 different stocks, and assume all of them were of the same size and turnover. Comparing to one large integrated market, 125 small markets would absorb the same shock $125^{2/3} = 25$ times more slowly and its impact would be $125^{1/3} = 5$ times smaller, as implied by equations (9) and (15).

Early Warning Systems May Be Useful and Practical. Some strategies are inherently destabilizing. They have built-in features of negative feedbacks: as prices go down, more selling is required and this pushes prices further down. The more capital is invested into these strategies, the bigger is their potential destabilizing effects on prices. Equipped with quantitative invariance formulas for market impact, one may detect instances when destabilizing strategies become so large that they may put financial markets at risk.

Tuzun (2012) uses invariance to assess the effect of leveraged ETFs on markets. He finds that short ETFs and leveraged long ETFs in financial stocks were close to the tipping point in 2008 and 2009. A price decline of 1% would induce leveraged ETFs to sell about $1 billion; invariance implies that this imbalance would lead to a further price decline of another 1% and thus potentially trigger a downward spiral.

For some of five crash events in our paper, policymakers or stock market participants also had in hand the information required to quantify the price impact and foresee the systemic risks looming from sudden liquidations of large stock market exposures. Yet, they mistakenly trusted in conventional intuition when assessing potential magnitudes of price declines.

Contrary to the beliefs of some, the market crashes in 1929 and 1987 were not completely unexpected. In both cases, data was publicly available before the events. Data on broker loans was published by the Federal Reserve System and the NYSE before the 1929 crash. Estimates of assets under management by portfolio insurers were available before the 1987 crash.

In both cases, the potential price impact of liquidations was a topic of public discussion among policy makers and market participants. In the months prior to the 1929 stock market crash, brokers were raising margin requirements to protect themselves from a widely discussed collapse in prices which might be induced by rapid unwinding of stock investments financed with margin loans. Market participants watched statistics on broker loans carefully, noting the tendency for total lending in the broker loan market to increase as the stock market rose. Markets were aware that margin account investors were buyers with "weak hands," likely to be flushed out of their positions by margin calls if prices fell significantly. Discussions about who would buy stocks if a collapse in stock prices forced margin account investors out of their positions resembled similar discussions in 1987 concerning who would take the opposite side of portfolio insurance trades.

The debate about the extent to which portfolio insurance contributed to the 1987 crash started long before the crash itself. On the day the 1987 crash occurred, academics were holding a conference on a topic of potential "market meltdown" induced by portfolio insurance sales. The term "market meltdown," popularized by then NYSE chairman John Phelan, was used in the year or so before the stock market crash to describe a scenario of cascading portfolio insurers' sell orders resulting in severe price declines and posing systemic risks to the economy. Months before the 1987 crash itself, the SEC's Division of Market Regulation (1987)—responding to worries that portfolio insurance have made the market fragile—published a study describing in some detail a potential meltdown scenario induced by portfolio insurance sales, which closely resembled the subsequent crash in October 1987. Yet, the study dismissed the risk of such an event as a remote possibility, in agreement with conventional wisdom at the time.

Many market participants were firmly convinced that, given the substantial trading volume in the U.S. equity markets—and especially the index futures market—there was enough liquidity available to accommodate sales of portfolio insurers without any major downward adjustment in stock prices. During hearings before the House Committee on Energy and Commerce (1987) on July 23 prior to the 1987 crash, Hayne E. Leland defended portfolio insurance:

> "We indicated that average trading will amount to less than 2% of total stocks and

> derivatives trading. On some days, however, portfolio insurance trades may be a greater fraction. ...In the event of a major one-day fall (e.g., 100 points on the Dow Jones Industrial Average), required portfolio insurance trades could amount to $4 billion. Almost surely this would be spread over 2–3 day period. In such a circumstance, portfolio insurance trades might approximate 9–12% of futures trading, and 3–4% of stock plus derivatives trading."

If regulators had applied simple principles of invariance prior to the crash, they would have been alarmed by Hayne Leland's projection of potential sales of 4% of stock-plus-futures volume over three days in response to a decline in stock prices of about 4% (i.e., 100 points on the Dow Jones average). They would see that the stock market was already close to a tipping point. Historical volume and volatility in July 1987 implied that sales of $4 billion in response to a 4% price decline would lead to another drop in prices, just slightly smaller than 4%. Absent stabilizing trades by investors trading in an opposite direction, potential portfolio insurance sales were already on the verge of triggering precisely the cascade meltdown scenario.

5 Conclusion

Crash-like events continue to occur. The Staffs of the Fed, the CFTC, and SEC (2015) describe the "flash rally" in the U.S. Treasury market on October 15, 2014, during which prices rose rapidly for several minutes and then fell back down. Since the report was not based on audit trail data identifying individual traders, it does not rule out the possibility that the flash rally resulted from rapid buying by one trader. Obizhaeva (2016) describes how the sharp V-shaped devaluation of Russian currency on December 16, 2014, was likely caused by a large multi-billion-dollar bet. The collapse of the Chinese stock market in the summer of 2015 was likely caused by liquidations of margin accounts, as discussed in Bian et al. (2018); this crash was in many ways similar to the crash of 1929 in the U.S. market; in both cases extraordinary steps were taken to stabilize the markets.

Our study of five case studies should not be interpreted as a regression with five data points. Instead, we think that examining these episodes leads to useful insights about why stock market crashes happen, how to prevent them if possible, and how to respond to them if not.

References

Akerlof, George A., and Robert J. Shiller. 2009. *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism.* Princeton University Press.

Almgren, Robert, Chee Thum, Emmanuel Hauptmann, and Hong Li. 2005. "Direct Estimation of Equity Market Impact." *Risk*, 18(7): 58–62.

Bernanke, Ben S. 2015. *The Courage to Act: A Memoir of a Crisis and Its Aftermath.* Norton.

Bian, Jiangze, Zhiguo He, Kelly Shue, and Hao Zhou. 2018. "Leverage-induced Fire Sales and Stock Market Crashes." Working Paper.

Black, Fischer. 1971. "Towards a Fully Automated Exchange, Part I." *Financial Analysts Journal*, 27: 29–34.

Board of Governors of the Federal Reserve System. 1927–1931. *Annual Report of the Board of Governors of the Federal Reserve System.* 13th through 17th Annual Reports Covering Operations for the Years 1926 through 1930.

Board of Governors of the Federal Reserve System. 1929. *Federal Reserve Bulletin.* multiple issues.

Board of Governors of the Federal Reserve System. 1943. *Banking and Monetary Statistics, 1914–1941.*

Board of Governors of the Federal Reserve System. 2008. *Meeting of the Federal Open Market Committee on January 29–30, 2008.*

Bouchaud, Jean-Philippe, Doyne J. Farmer, and Fabrizio Lillo. 2009. "How Markets Slowly Digest Changes in Supply and Demand." In *Handbook of Financial Markets: Dynamics and Evolution.* 47–160. North Holland.

Brennan, Michael J., and Eduardo S. Schwartz. 1989. "Portfolio Insurance and Financial Market Equilibrium." *Journal of Busines*, 62: 455–472.

Chan, Louis K. C., and Joseph Lakonishok. 1995. "The Behavior of Stock Prices Around Institutional Trades,." *Journal of Finance*, 50(4): 1147–1174.

Chan, Louis K. C., and Joseph Lakonishok. 1997. "Institutional Equity Trading Costs: NYSE Versus Nasdaq,." *Journal of Finance*, 52(2): 713–735.

Commodity Futures Trading Commission, Division of Economic Analysis and Division of Trading and Markets. 1988. *Final Report on Stock Index Futures and Cash Market Activity During October 1987 to the U.S. Commodity Futures Trading Commission.*

Division of Market Regulation, Securitie and Exchange Commission. 1987. *The Role of Index-Related Trading in the Market Decline of September 11 and 12, 1986.*

Frazzini, Andrea, Ronen Israel, and Tobias J. Moskowitz. 2018. "Trading Costs." Working Paper.

Friedman, Milton, and Anna Jacobson Schwartz. 1963. *A Monetary History of the United States, 1867-1960. NBER Book Series Studies in Business Cycles,* Princeton University Press.

Friedman, Milton, and Anna Schwarz. 1963. *A Monetary History of the United States.* Princeton University Press.

Gabaix, Xavier. 2009. "Power Laws in Economics and Finance." *Annual Review of Economics,* 121: 255–293.

Galbraith, John Kenneth. 1954. *The Great Crash, 1929.* Boston:Houghton Mifflin Co.

Gammill, James F., and Terry A. Marsh. 1988. "Trading Activity and Price Behavior in the Stock and Stock Index Futures Markets in October1987." *Journal of Economic Perspectives,* 2(3): 25–44.

Grinold, Richard C., and Ronald N. Kahn. 1995. In *Active Portfolio Management.* Probus Publishing.

Haney, H. Lewis. 1932. "Brokers' Loans." *Journal of the American Statistical Association,* 27: 65–72.

Harris, Lawrence, and Eitan Gurel. 1986. "Price and Volume Effects Associated with Changes in the S&P 500 List: New Evidence for the Existence of Price Pressures." *Journal of Finance,* 41(4): 815–829.

House Committee on Energy and Commerce, U.S. Congress. 1987. *Program Trading: Hearing before the Subcommittee on Telecommunications and Finance of the Committee on Energy and Commerce.* Vol. 100th Cong., 1st Sess. Testimony of Hayne E. Leland.

Huberman, Gur, and Werner Stanzl. 2004. "Price Manipulation and Quasi Arbitrage." *Econometrica*, 72(4): 1247–1275.

Keim, Donald B., and Ananth Madhavan. 1997. "Transaction Costs and Investment Style: An Inter-Exchange Analysis of Institutional Equity Trades." *Journal of Financial Economics*, 46(3): 265–292.

Keynes, John M. 1936. *The General Theory of Employment, Interest and Money.* Palgrave Macmillan, United Kingdom.

Kraus, Alan, and Hans R. Stoll. 1972. "Price Impact of Block Trading on the New York Stock Exchange." *Journal of Finance*, 27(3): 569–588.

Kyle, Albert S. 1985. "Continuous Auctions and Insider Trading." *Econometrica*, 53(6): 1315–1335.

Kyle, Albert S., and Anna A. Obizhaeva. 2016. "Market Microstructure Invariance: Empirical Hypotheses." *Econometrica*, 84(4): 1345–1404.

Kyle, Albert S., and Anna A. Obizhaeva. 2018*a*. "Adverse Selection and Liquidity: From Theory to Practice." Available at SSRN: `https://ssrn.com/abstract=3236030` or `http://dx.doi.org/10.2139/ssrn.3236030`.

Kyle, Albert S., and Anna A. Obizhaeva. 2018*b*. "Dimensional Analysis, Leverage Neutrality, and Market Microstructure Invariance." Working Paper, University of Maryland.

Kyle, Albert S., and Anna A. Obizhaeva. 2020. "Market Microstructure Invariance: A Dynamic Equilibrium Model." Working Paper. Available at SSRN: `https://ssrn.com/abstract=3326889` or `http://dx.doi.org/10.2139/ssrn.3326889`.

Kyle, Albert S., Anna A. Obizhaeva, and Yajun Wang. 2014. "Industrial Organization, Order Internalization, and Invariance."

Leland, Hayne E., and Mark Rubinstein. 1988. "Comments on the Market Crash: Six Months After." *Journal of Economic Perspectives*, 2: 45–50.

Miller, Merton H. 1988. "Explaining the Events of October 1987." In *After the Crash: Linkages between Stocks and Futures. AEI Studies*, , ed. Robert J. Mackay. American Enterprise Institute for Public Policy Research, University Press of America.

Miller, Merton H. 1991. *Financial Innovations and Market Volatility.* Cambridge, MA:Blackwell.

Obizhaeva, Anna A. 2016. "The Russian Ruble Crisis of December 2014." *Voprosy Ekonomiki*, 5: 66–86.

Perold, André. 1988. "The Implementation Shortfall: Paper vs. Reality." *Journal of Portfolio Management*, 14(3): 4–9.

Reinhart, Carmen M., and Kenneth Rogoff. 2009. *This Time Is Different: Eight Centuries of Financial Folly.* Princeton University Press.

Roll, Richard. 1988. "The International Crash of October 1987." *Financial Analysts Journal*, 44(5): 19–35.

Scholes, Myron S. 1972. "The Market for Securities: Substitution Versus Price Pressure and the Effects of Information on Share Prices." *Journal of Business*, 45(2): 179–211.

Senate Committee on Banking and Currency, U.S. Congress. 1934. *Stock Exchange Practices: Report Of The Committee On Banking And Currency Pursuant To S.Res. 84 (72nd Congress) and S.Res. 56 and S.Res. 97 (73rd Congress).* Vol. 73rd Congress, Second Session, U.S. Government Printing Office. "The Pecora Commission Report".

Shiller, Robert J. 1987. "Investor Behavior in the October 1987 Stock Market Crash: Survey Evidence." NBER Working Paper 2446.

Smiley, Gene, and Richard H. Keehn. 1988. "Margin Purchases, Brokers' Loans and the Bull Market of the Twenties." *Business and Economic History*, 17: 129–142.

Société Générale. 2008a. "Explanatory note about the exceptional fraud." Paris, January 27, 2008.

Société Générale. 2008b. “Mission Green, Summary Report: Interim conclusions as of February 20, 2008.” General Inspection Department.

Société Générale. 2008c. “Mission Green, Summary Report.” General Inspection Department, May 20, 2008.

Staffs of the CFTC and SEC. 2010a. *Preliminary Findings Regarding the Market Events of May 6, 2010.* Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues. May 18, 2010. “Preliminary Report”.

Staffs of the CFTC and SEC. 2010b. *Findings Regarding the Market Events of May 6, 2010.* Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues. September 30, 2010. “Final Report”.

Staffs of the Fed, the CFTC, and SEC. 2015. *The U.S. Treasury Market on October 15, 2014.* Joint staff report by U.S. Department of the Treasury, Board of Governors of the Federal Reserve System, Federal Reserve Bank of New York, U.S. Securities and Exchange Commission, U.S. Commodity Futures Trading Commission. July 13, 2015.

Torre, Nicolo G. 1997. *BARRA Market Impact Model Handbook.* BARRA Inc., Berkeley.

Torre, Nicolo G., and Mark J. Ferrari. 1999. *The Market Impact Model.* Barra.

Tuzun, Tugkan. 2012. “Are Leveraged and Inverse ETFs the New Portfolio Insurers?” *Working Paper.*

U.S. Presidential Task Force on Market Mechanisms. 1988. *Report of the Presidential Task Force on Market Mechanisms.* Washington:U.S Government Printing Office. The “Brady Report”.

Wurgler, Jeffrey A., and Ekaterina Zhuravskaya. 2002. “Does Arbitrage Flatten Demand Curves for Stocks?” *Journal of Business*, 75(4): 583–608.

Appendix A: Implementation Issues.

In order to apply microstructure invariance to data on the five crash events, several implementation issues need to be addressed.

First, the volume and volatility inputs in our formulas should not be thought of as parameters of narrowly defined markets of a particular security in which a bet is placed but rather as parameters based on the market as a whole. Securities and futures contracts may share the same fundamentals and have a common factor structure. When a large order moves prices in the S&P 500 futures market, index arbitragers usually insure that prices for the underlying basket of stocks move by about the same amount as well. It is difficult to identify the boundaries of the market. Consistent with the spirit of the Brady Report, we take the admittedly simplified approach of adding together cash and futures volume for three of the four crash events in which stock index futures markets existed. In our analysis of the Soros trades, we ignore cash market volume because his trades were executed so quickly that price pressure in the futures market was not transferred to cash markets.

Second, the spirit of the invariance hypothesis is that volume and volatility inputs into the market impact equation (15) are market expectations prevailing before the bet is placed. Expected volume and expected volatility determine the sizes of bets investors are willing to make and the market depth intermediaries are willing to provide. Different price impact estimates are possible, depending on whether volatility estimates are based on implied volatilities before the crash, implied volatilities during the crash, historical volatilities based on the crash period itself, or historical volatilities based on months of data before the crash. For robustness, we present results based on historical data for different windows prior to the crash event.

Third, it is likely that the price impact of an order is related to the speed with which it is executed. The market impact model (15) assumes that orders are executed at a "normal" speed in the relevant units of business time. For example, a very large order in a small stock may be executed over several weeks or even months, while a large order in the stock index futures market may be executed over several hours. The impact model leaves open the possibility that unusually rapid execution of very large orders may increase their temporary price impact, but these effects are hard to quantify properly. We discuss this issue further in Section 4.

Fourth, there have been numerous changes in market mechanisms between 1929 and 2010, including better communications technologies, introduction of electronic handling of orders, a reduction in tick size, and the migration of trading volume from face-to-face trading floors to anonymous electronic platforms. Such changes may have lowered bid-ask spreads, but we believe—in the spirit of Black (1971)—that they have had little effect on market depth, which largely dominates the price impact of large bets. We thus apply estimates of market depth based on portfolio transitions during 2001–2005 to the entire period 1929–2010.

Fifth, Kyle and Obizhaeva (2016) calibrate both linear and square-root impact models consistent with invariance. From an empirical perspective, the square root specification explains price impact somewhat better than the linear model, as consistent with the empirical econophysics literature (Bouchaud, Farmer and Lillo, 2009). Yet, the linear model explains the price impact of the largest one percent of bets in the most active stocks slightly better than the square root model. Crash events are explained by applying invariance to a linear model. To make this point, "invariance" implicitly assumes a linear impact function in the main part of the paper. Due to its concavity, the square root model predicts much smaller price declines during crash events. Appendix C presents these estimates along with estimates based on alternative models.

Appendix B: Estimates for Different Market Impact Models.

We compute estimates of predicted price changes based on several alternative models of market impact. Market impact is expected to depend on market characteristics such as market capitalization N, daily share volume V, returns volatility σ, and the corresponding GDP deflator d_{gdp}; unsigned bet size Q; and perhaps the time horizon T over which the bet is executed.

We consider several specifications when calculating the implied magnitudes of simple (non-logged) market impacts $\Delta P/P$:

- The invariance-implied linear model ("Inv-LIN"), discussed in Kyle and Obizhaeva (2016):

$$\frac{\Delta P}{P} = \frac{2\cdot 2.50}{10^4}\cdot\left(\frac{P\cdot V\cdot d_{\text{gdp}}}{40\cdot 10^6}\right)^{1/3}\cdot\left(\frac{\sigma}{0.02}\right)^{4/3}\cdot\frac{Q}{(0.01)\cdot V}. \tag{16}$$

- The invariance-implied square-root model ("Inv-SQRT"), discussed in Kyle and Obizhaeva (2016):

$$\frac{\Delta P}{P} = \frac{2\cdot 12.08}{10^4}\cdot\left(\frac{\sigma}{0.02}\right)\cdot\left(\frac{Q}{(0.01)\cdot V}\right)^{1/2}. \tag{17}$$

- The conventional model ("Conv-N"), based on market capitalization:

$$\frac{\Delta P}{P} = \frac{Q}{N}. \tag{18}$$

- The conventional model ("Conv-V"), based on daily volume:

$$\frac{\Delta P}{P} = \frac{Q}{250 \cdot V}. \tag{19}$$

- The Barra model, discussed in Torre and Ferrari (1999) and Grinold and Kahn (1995):

$$\frac{\Delta P}{P} = \sigma \cdot \left(\frac{Q}{V}\right)^{1/2}. \tag{20}$$

- Almgren–Chriss model ("AC"), discussed in Almgren et al. (2005):

$$\frac{\Delta P}{P} = 0.314 \cdot \sigma \cdot \frac{Q}{V} \cdot \left(\frac{N}{V}\right)^{1/4} + 2 \cdot 0.142 \cdot \sigma \cdot \left(\frac{Q}{V \cdot T}\right)^{3/5}. \tag{21}$$

- Frazzini–Israel–Moskowitz model ("FIM"), discussed in Frazzini, Israel and Moskowitz (2018) in Table VII, column (9):

$$\frac{\Delta P}{P} = \left(-0.2 \cdot \ln(1 + N \cdot 10^{-9} \cdot d_{\text{gdp}}) + 0.35 \cdot \frac{Q}{0.01 \cdot V} + 9.32 \cdot \left(\frac{Q}{0.01 \cdot V}\right)^{1/2} + 0.13 \cdot \sigma \cdot \sqrt{252} \cdot 100\right) \cdot \frac{2}{10^4}. \tag{22}$$

In the last two models, the estimates are multiplied by a factor of 2 to convert transaction costs estimates to price impact estimates.

The Almgren–Chriss model (21) explicitly depends on the execution horizon T. For the 1929 crash, we assume selling occurred over five days ($T = 5$). For the 1987 crash, we assume selling occurred over four days ($T = 4$). For Soros' trades, we assume selling occurred over 6 minutes from 8:28 a.m. to 8:34 a.m. ($T = 6/420$). For the liquidation of Kerviel's trades, we assume selling occurred over three days ($T = 3$). For the flash crash of 2008, we assume selling occurred over about twenty minutes, or 1/20 of a day ($T = 1/20$).

Panel A of Table 7 presents impact estimates based on six impact models for percentage market impact along with actual price declines for the five crashes. First, all estimates are much

Table 7: Alternative Models.

Panel A: Simple percentage impact $\Delta P/P$.

	Actual	Inv-LIN	Inv-SQRT	Conv-N	Conv-V	Barra	AC	FIM
1929 Market Crash	25.00%	62.56%	3.94%	1.36%	1.06%	3.26%	6.62%	4.95%
1987 Market Crash	32.00%	18.31%	1.33%	0.63%	0.27%	1.10%	1.04%	2.02%
1987 Soros's Trades	22.00%	6.47%	1.58%	0.01%	0.01%	1.31%	3.47%	0.62%
2008 SocGén Trades	9.44%	11.40%	0.70%	0.43%	0.11%	0.58%	0.35%	1.18%
2010 Flash Crash	5.12%	0.61%	0.16%	0.03%	0.01%	0.13%	0.16%	0.24%

Panel B: Log-percentage impact $\Delta \ln P$.

	Actual	Inv-LIN	Inv-SQRT	Conv-N	Conv-V	Barra	AC	FIM
1929 Market Crash	25.00%	46.43%	3.86%	1.36%	1.06%	3.21%	6.41%	4.83%
1987 Market Crash	32.00%	16.77%	1.32%	0.63%	0.27%	1.10%	1.04%	1.99%
1987 Soros's Trades	22.00%	6.27%	1.57%	0.01%	0.01%	1.30%	3.41%	0.62%
2008 SocGén Trades	9.44%	10.79%	0.70%	0.43%	0.11%	0.58%	0.35%	1.17%
2010 Flash Crash	5.12%	0.61%	0.16%	0.03%	0.01%	0.13%	0.16%	0.24%

Table 7 presents actual price declines along with price declines implied by several models for five market crashes. Panel A shows the estimates for models with simple returns $\Delta P/P$, and panel B shows the estimates for models with log returns $\Delta \ln P$.

lower than actual price declines, except for the Inv-LIN estimates. Second, the Conv-N and Conv-V estimates based on the conventional intuition usually generate the smallest estimates among models. Third, calibrated on the sample of institutional transactions, the Barra, AC, and FIM estimates are all similar in magnitude; they are slightly larger than conventional estimates but still much lower than the actual price declines. Fourth, the AC estimate is significantly larger than other alternative estimates for the Soros bet because this estimate explicitly accounts for the very short execution horizon of this bet.

For all five crashes, the Inv-SQRT estimates are quantitatively similar to the Barra estimates. Due to its concavity, the square root models predict much smaller price declines than the linear model. Thus, invariance alone does not explain magnitudes of price declines during crash events; instead, crash events are explained by applying invariance to a linear model.

Panel B of Table 7 presents impact estimates based on six impact models for log-percentage market impact $\Delta \ln P$ along with actual price declines for the five crashes. These estimates are obtained from models (16)–(22), where $\Delta P/P$ on the left-hand side of these equations is replaced with $\Delta \ln P$. The invariance-implied linear model (Inv-LIN) and the conventional model based on market capitalization (Conv-N) for log-impact are the two models discussed on detail in the main part of our paper.

The estimates based on log-returns are smaller than the estimates based on simple returns, but this difference is negligible for most models. The only exception is the Inv-LIN model, for which large estimates based on the simple return are reduced when log-returns are used instead; the biggest difference is observed for the 1929 crash, for which the simple-return model implies price decline of 63% and the log-return model implies price decline of only 46%.

Appendix C: Estimation Details.

Estimation Details for the Crash of 1929

A significant portion of stock investments in the late 1920s was made in leveraged margin accounts. To finance their leveraged purchases of stocks, individuals and non-financial corporations relied either on bank loans collateralized by securities or on margin account loans at brokerage firms. When investors borrowed through margin accounts at brokerage firms, the brokerage firms financed only a modest portion of the loans with credit balances from other customers. To finance the rest, brokerage firms pooled securities pledged as collateral by customers under the name of the brokerage firm (in "street name") and then re-hypothecated these pools by using them as collateral for broker loans. The broker loan market of the late 1920s resembled the shadow banking system of the early 2000s in its lack of regulation, perceived safety, and the large fraction of overnight or very short maturity loans.

The broker loan market was controversial during the 1920s, just as the shadow banking system was controversial during the period surrounding the financial crisis of 2008–2009. Some thought the broker loan market should be tightly controlled to limit speculative trading in the stock market on the grounds that lending to finance stock market speculation diverted capital away from more productive uses in the real economy. Others thought it was impractical to control lending in the market because the shadow bank lenders would find ways around restrictions and lend money anyway. The New York Fed chose to discourage New York banks from lending money against stock market collateral. As a result, loans to brokers by New York banks declined after reaching a peak in 1927.

Attracted by the high interest rates on broker loans—typically 300 basis points or more higher than loans on otherwise similar money market instruments—non-New York banks and non-bank lenders continued to supply capital to the broker loan market. Many of these loans were arranged by the New York banks; sometimes, non-bank lenders bypassed the banking system entirely, making loans directly to brokerage firms.

Investment trusts (similar to closed end mutual funds) placed a large fraction of the newly raised equity into the broker loan market rather than buying expensive common stocks. Corporations, flush with cash from growing earnings and proceeds of securities issuance, invested a large portion of these funds in the broker loan market rather than in new plant and equipment.

To quantify the margin selling which occurred during the last week of October 1929, we follow the previous literature and contemporary market participants by estimating margin selling indirectly from data on broker loans and bank loans collateralized by securities.

In the 1920s, data on broker loans came from two sources. First, the Fed collected weekly broker loan data from reporting member banks in New York City supplying the funds or arranging loans for others. Second, the New York Stock Exchange collected monthly broker loan data based on demand for loans by NYSE member firms. The broker loan data reported by the New York Stock Exchange include broker loans which non-banks made directly to brokerage firms without using banks as intermediaries; such loans bypassed the Fed's reporting system. Since loans unreported to the Fed fluctuated significantly around the 1929 stock market crash, we rely relatively heavily on the NYSE numbers in our analysis below but also pay careful attention to the weekly dynamics of the Fed series for measuring selling pressure during the last week of October 1929.

We calculate weekly proxies for margin sales as follows. (1) We difference the weekly Fed series to construct weekly changes. (2) We interpolate the monthly NYSE series to construct a weekly series by assuming that these loans changed at a constant rate within each month, except for October 1929. For October 1929, the Fed series shows little change, except for the last week, and we therefore assume that the entire monthly change in the NYSE series represents unreported changes in broker loans which occurred during the last week of October 1929. (3) Finally, we add changes in bank loans collateralized by securities to take into account the fact that some changes in broker loans do not represent margin sales because they were converted into bank loans collateralized by securities. The last adjustment also has a significant effect because there was an unprecedented increase in banks loans collateralized by securities during the last week of October 1929, followed by offsetting reductions during November.

Figure 2 shows the weekly levels of the Fed's broker loan series and the monthly levels of the NYSE broker loan series. Two versions of each series are plotted, one with bank loans collateralized by securities added and one without ("Fed Broker Loans," "Fed Broker Loans + Bank Loans," "NYSE Broker Loans," "NYSE Broker Loans + Bank Loans"). The figure also shows the level of the Dow Jones Industrial Average from 1926 to 1930. The time series on both broker loans and stock prices follow similar patterns, rising steadily from 1926 to October 1929 and then suddenly collapsing. According to Fed data, broker loans rose from \$3.141 billion at the beginning of 1926 to \$6.804 billion at the beginning of October 1929. According to NYSE data, the broker loan market rose from \$3.513 billion to \$8.549 billion during the same period. As more and more non-banks were getting involved in the broker loan market, the difference between NYSE broker loans and Fed broker loans steadily increased until the last week of October 1929, when non-bank firms pulled their money out of the broker loan market and the difference suddenly shrank.

During the period 1926–1930, weekly changes in broker loans were typically small and often changed sign, as shown in the tiny bars at the bottom of figure 2. Starting with the last week of October 1929, there were five consecutive weeks of large negative changes, almost twenty times larger than changes during preceding weeks. This de-leveraging erased the increase in broker loans which had occurred during the first nine months of the year.

For the last week of October 1929, we estimate margin selling as \$1.181 billion (the difference between the estimated reduction in broker loans of \$2.440 billion from \$8.549 billion to \$6.109 billion and increase in bank loans on securities of \$1.259 billion from \$7.920 billion to \$9.179 billion). For the three months from September 30, 1929, to December 31, 1929, we estimate margin selling as \$4.348 billion (the difference between the reduction in NYSE broker loans of \$4.559 billion from \$8.549 billion to \$3.990 billion and an increase in bank loans on securities of \$0.211 billion from \$7.720 billion to \$7.931 billion).

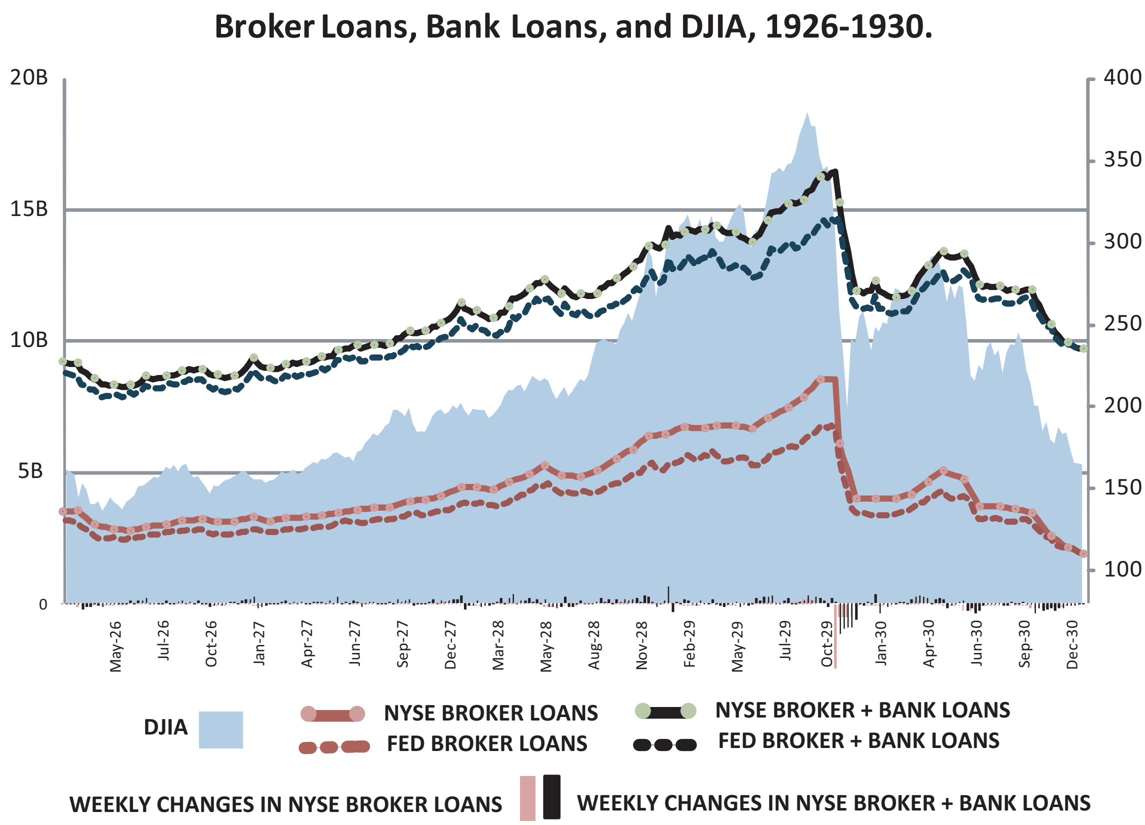

Figure 2: Broker Loans and 1929 Market Crash.

The figure shows weekly dynamics of seven variables from January 1926 to December 1930: NYSE broker loans (red solid line), Fed broker loans (red dashed line), the sum of NYSE broker loans and bank loans (black solid line), the sum of Fed broker loans and bank loans (black dashed line), changes in NYSE broker loans (red bars), changes in the sum of NYSE broker loans and bank loans (black bars), and the Dow Jones average (in blue). Monthly levels of NYSE broker loans are marked with solid dots. Weekly levels of NYSE broker loans are obtained using a linear interpolation from monthly data, except for October 1929, when all changes in NYSE broker loans are assumed to occur during the last week.

Estimation Details for the Crash of 1987

Along with our main estimates in Table 3, we present several other estimates for robustness. First, some of the market participants classified as portfolio insurers in the Brady Report abandoned their portfolio insurance strategies as prices crashed and switch to buying securities. Even though we believe that for the purpose of analyzing the price impact of portfolio insurance sales it is better to use the gross sales amount, we also report estimates for net sales of \$11.11 billion of futures contracts and stocks combined (\$9.51 billion in futures and \$1.60 billion in stocks). Their predicted impact ranges from 9.71% to 13.78%.

Second, we show implied estimates if we treat markets for futures contracts and NYSE stocks separately. To avoid radically different price impacts in two markets, we adjust quantities sold in both markets by the NYSE's estimate of net NYSE index-arbitrage sales of \$3.27 billion (Brady Report, figures 13–14). We add this number to portfolio insurance sales in NYSE stocks and subtract the same amount from portfolio insurance sales in the futures market because arbitrageurs transferred some price pressure from futures to stocks. This results in net sales of \$7.21 billion in the futures market with impact ranging from 10.00% to 14.11% and \$6.54 billion in NYSE stocks with impact ranging from 9.09% to 13.00%. The fact that index arbitrage sales make price impact estimates similar in both markets is consistent with the interpretation that portfolio insurance sales were indeed driving price dynamics in both markets.

Estimation Details for Liquidation of Kerviel's Trades in 2008

We also examine whether implied cost estimates are consistent with officially reported losses of €6.30 billion. We assume that average impact cost is equal to half of predicted price impact since—assuming no leakage of information about the trades—a trader can theoretically walk the demand curve, trading only the last contracts at the worst expected prices. Accounting for compounding, invariance predicts that the total cost of unwinding Kerviel's position is equal to 5.55% of the initial €50 billion position, i.e., €2.77 billion.

Officially reported losses also include mark-to-market losses sustained by hidden naked long positions as markets fell from the end of the previous reporting period on December 31, 2007, to the decision to liquidate the positions when the market re-opened after January 18, 2008. From December 28, 2007, to January 18, 2008, the Euro STOXX 50 fell by 9.18%, DAX futures fell by 9.40%, and FTSE futures fell by 8.68%. If we assume that Kerviel held a constant long position from December 31, 2007, to January 18, 2008, then these positions would have sustained €4.62 billion in mark-to-market losses during that period. Société Générale reported, however, that Kerviel acquired his hidden long position gradually over the month of January. If we assume that Kerviel acquired his position gradually by purchasing equal quantities of futures contracts at each lower tick level from the end-of-year 2007 close to January 18 close, we estimate that such positions would be under water by only half as much, i.e., €2.31 billion, at the close of January 18.

Table 5 reports that the estimated market impact costs of liquidating the rogue position range from €2.72 billion to €3.34 billion under different assumptions about expected volume and volatility. Adding mark-to-market losses sustained prior to liquidation leads to estimated losses ranging (A) from €5.03 billion to €5.65 billion if positions were acquired gradually and (B) from €7.34 billion to €7.96 billion if positions were held from the end of 2007. These estimates are similar in magnitude to losses of €6.30 billion reported by the bank.

As a robustness check, we also estimate market impact under the assumption that the Euro STOXX 50, the DAX, and the FTSE 100 futures markets are distinct markets, not components of one bigger market.

In the month preceding January 18, 2008, historical volatility per day was 98 basis points for futures on the Euro STOXX 50, 100 basis points for futures on the DAX, and 109 basis points for futures on the FTSE 100. Average daily volume was €55.19 billion for Euro STOXX 50 futures, €32.40 billion for DAX futures, and £7.34 billion for FTSE 100 futures. Kerviel's positions of €30 billion in Euro STOXX 50 futures, €18 billion in DAX futures, and €2 billion in FTSE 100 futures represented about 54%, 56%, and 20% of daily trading volume in these contracts, respectively. We use an exchange rate of €1.3440 for £1 on January 17.

Our calculations estimate a price impact of 12.08% for liquidation of Kerviel's position, 10.77% for liquidation of his DAX futures position, and 4.12% for liquidation of his FTSE futures position. Indeed, from the close on January 18 to the close on January 23, Euro STOXX 50 futures fell by 10.50%, DAX futures fell by 11.91%, and FTSE 100 futures fell by 4.65%. Note that from the close on January 18 to the lowest point during January 21 through January 23, Euro STOXX 50 futures fell by 11.67%, DAX futures fell by 12.71%, and FTSE 100 futures fell by 9.54%. The similarity of actual price declines for the STOXX 50, DAX and FTSE suggests substantial integration of European markets, consistent with our strategy of thinking about them as one market.

From the close on January 18 to low points on January 22, the Spanish IBEX 35, the Italian FTSE MIB, the Swedish OMX, the French CAC 40, the Dutch AEX and the Swiss Market Index fell by 12.99%, 10.11%, 8.63%, 11.53%, 10.80%, and 9.63%, respectively. By January 24, all of these markets had largely reversed these losses. Euro Stoxx 50 and FTSE reversed losses as well, but DAX recovered only partially. Large price declines in markets where Kerviel did not hold positions suggest that the markets are well integrated as well.

Estimation Details for the Flash Crash of May 6, 2010

Since the price drop in the morning may have reset market expectations about volatility, as a robustness check, we also report results for expected volatility of 2.00% per day; they range from 1.39% to 1.65%.

If we do not treat the cash market and the futures market as one market but focus only on the futures market, then the estimates range from 0.76% to 1.29% for historical volatility and from 2.35% to 2.91% for volatility of 2%.

Appendix D: The Frequency of Market Crashes

Market microstructure invariance can be used to quantify the frequency of crash events, including both the size of selling pressure and the resulting price impact.

Using portfolio transitions orders as proxies for bets, Kyle and Obizhaeva (2016) find that the invariant distributions of buy and sell bet sizes can be closely approximated by a log-normal. The distribution of unsigned bet size $\widetilde{X}$ of a stock with expected daily volume of $P \cdot V$ dollars and expected daily returns volatility σ can be approximated as a log-normal

$$\ln\left(\frac{\widetilde{X}}{V}\right) = -5.71 - \frac{2}{3} \cdot \ln\left(\frac{\sigma \cdot P \cdot V}{(0.02)(40)(10^6)}\right) + \sqrt{2.53} \cdot \widetilde{Z}, \tag{23}$$

where $\widetilde{Z} \sim \mathcal{N}(0,1)$. Under the assumption that there is one unit of intermediation trade volume for every bet, the bet arrival rate γ per day is given by

$$\ln(\gamma) = \ln(85) + \frac{2}{3} \cdot \ln\left(\frac{\sigma \cdot P \cdot V}{(0.02)(40)(10^6)}\right). \tag{24}$$

These equations have the following implications for a benchmark stock with dollar volume of \$40 million per day and volatility 2% per day$^{1/2}$. The estimated mean of -5.71 implies a median bet size of approximately \$132,500, or 0.33% of daily volume. The estimated log-variance of 2.53 implies that a one-standard-deviation increase in bet size is a factor of about 4.91. The implied average bet size is \$469,500 and a four-standard-deviation bet is about \$77 million, or 1.17% and 192% of daily volume, respectively ($0.33\% \cdot \exp(2.53/2)$ and $0.33\% \cdot \exp(2.53 \cdot 4)$). There are 85 bets per day. The standard deviation of daily order imbalances is equal to 38% of daily volume ($85^{1/2} \exp(-5.71 + 2.53)$). Half the variance in returns results from fewer than 0.10% of bets and suggests significant kurtosis in returns.

Now let us extrapolate these estimates to the entire market, where volume is the sum of the volume of CME S&P 500 futures contracts and all individual stocks. Using convenient round numbers based on the 2010 flash crash, the volume for the entire market is about \$270 billion per day, or 6,750 times the volume of a benchmark stock. The volatility of the index is about 1% per day, or half of 2% volatility of a benchmark stock. With 6,750 conveniently equal to $15^3 \cdot 2$, invariance implies that market volume consists of 19,125 bets ($85 \cdot 15^2$) with the median bet of about \$4 million ($\$132,500 \cdot 15 \cdot 2$), or 0.0014% of daily volume. The implied average bet size is \$14 million, or 0.0052% of daily volume, and a four-standard-deviation bet is \$2.310 billion ($\$469,500 \cdot 15 \cdot 2$ and $\$77 \cdot 10^6 \cdot 15 \cdot 2$), or 0.86% of daily volume. The implied standard deviation of cumulative order imbalances is 2.55% of daily volume (38%/15).

Equations (23) and (24) can be used to predict how frequently crash events occur. The three large crash events—the 1929 crash, the 1987 crash, and the 2008 Société Générale trades—are much rarer events than the two smaller crashes—the 1987 Soros trades and the 2010 flash crash.

We estimate the 1929 crash, the 1987 crash, and the 2008 liquidation of Kerviel's positions to be 6.15, 5.97, and 6.19 standard deviation bet events, respectively. Given corresponding estimated bet arrival rates of 1,887 bets, 5,606 bets, and 19,059 bets per day, such events would be expected to occur only once every 5,516 years, 597 years, and 674 years, respectively. Obviously, either the far right tail of the distribution estimated from portfolio transitions is fatter than a log-normal or the log-variance estimated from portfolio transition data is too small. In the far right tail of the distribution of the log-size of portfolio transition orders in the most actively traded stocks, Kyle and Obizhaeva (2016) do observe a larger number observations than implied by a normal distribution. It is also possible that portfolio transition orders are not representative of bets in general. If the true standard deviation of log bet size is 10% larger than implied by portfolio transition orders, then 6.0 standard deviation events become 5.4 standard deviation events, which are expected to occur about 34 times more frequently.

We estimate the 1987 Soros trades and the 2010 flash crash trades to be 4.45 and 4.63 standard deviation bet events, respectively. Given estimated bet arrival rates of 14,579 bets and

29,012 bets per day, respectively, bets of this size are expected to occur multiple times per year. We believe it likely that large bets of this magnitude do indeed occur multiple times per year, but execution of such large bets typically does not lead to flash crashes because such large bets would normally be executed more slowly and therefore have less transitory price impact.