

Large Bets and Stock Market Crashes

Albert S. Kyle* and Anna A. Obizhaeva[†]

First version: March 5, 2012

Most recent version: March 3, 2017

This version: March 22, 2019

Abstract

For five stock market crashes, we compare price declines with predictions from market microstructure invariance. During the 1987 crash and the 2008 sales by Société Générale, prices fell by magnitudes similar to predictions from invariance. Larger-than-predicted temporary price declines during 1987 and 2010 flash crashes suggest rapid selling exacerbates transitory price impact. Smaller-than-predicted price declines for the 1929 crash suggest slower selling stabilized prices and less integration made markets more resilient. Quantities sold in the three largest crashes indicate fatter tails or larger variance than the log-normal distribution estimated from portfolio transitions data.

JEL Codes: G01, G28, N22.

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

*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 of institutional investors.

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

After stock market crashes, rattled market participants, frustrated policymakers, and puzzled economists are typically unable to explain what happened. Even though noticeably heavy selling pressure has been often recorded during these episodes, there is no compelling quantitative explanation for the effect of selling pressure on the magnitude of crashes. It is usually believed that stock markets offer such great liquidity that these sales could not have led to such large price changes.

This paper investigates the hypothesis that stock market crashes result from the price impact of large bets from a new angle. A “bet” (or “meta-order”) is an approximately independent decision to transfer risk by buying or selling significant quantities of financial assets, implemented as a sequence of orders executed over time. Order flow moves prices, and very large bets may cause substantial market dislocations. This idea is not new; it is natural that large bets move prices in the direction of trades, as discussed by Kraus and Stoll (1972), Grinold and Kahn (1995), and Gabaix (2009). We quantify the price impact of quantities sold by applying the conceptual framework of market microstructure invariance developed by Kyle and Obizhaeva (2016). The invariance model with linear impact model implies estimates which are much larger than traditional estimates and in line with price declines during 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 dramatic 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”) documented 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 future 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 large sales of futures contracts by one entity, identified in the press as Waddell & Reed.

These were five large sell bets, which resulted either from trading by one entity or from correlated trading of multiple entities with the same motivation. 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 hard data on the size of sales which precipitated the events is not available.

Many practitioners and academics believe that selling pressure during these five market crashes was too small to induce significant price declines. Scholes (1972), Harris and Gurel (1986), and Wurgler and Zhuravskaya (2002) represent the “conventional wisdom” 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 turnover rates do not vary enormously across stocks, this belief suggests that price impact may be a function of the fraction of daily volume traded, with sales of 5% of average daily volume expected to have modest price impact. The extreme liquidity of stock index futures suggests that selling 1% of futures market daily volume will have even less price impact than selling 1% of daily volume of an individual stock.

We question this conventional intuition from the perspective of microstructure invariance. The invariance hypothesis says that all securities markets look the same, except that *business time* passes faster in more liquid markets. For example, the equity market as a whole operates much faster than less liquid markets for individual stocks. One calendar day of trading in stock index futures is therefore equivalent to *many* calendar days of trading in an individual stock. For example, if business time in the entire market passes 225 times faster than in a typical stock, then a bet of 5% of one day’s volume for stock index futures is equivalent to a bet of 1125% (not 5%) of one day’s volume for an individual stock because it is equivalent to selling of 5% of volume each day for 225 consecutive days (not one day). Both bets are expected to have similar price impacts. Thus, bets—expressed as the same fraction of average daily volume—result in greater price impact in more liquid markets than in less liquid markets. Since conventional wisdom does not make a distinction between calendar and business time, it understates price impact of bets in more liquid markets.

The invariance principle generates a universal formula for market impact. Market impact is a nonlinear function of expected dollar volume, expected returns volatility, and the dollar size of a bet as well as a couple of “invariant” parameters; Kyle and Obizhaeva (2016) calibrate these parameters using a database of about 400,000 portfolio transition trades in U.S. stocks executed during the period 2001–2005. The invariance framework implies that these invariant parameters are constant across different markets and time periods. Using the invariance principle and extrapolating these market impact estimates for individual U.S. stocks to the entire stock market—across different assets, order sizes, and time periods—we obtain impact estimates that are indeed large enough to match the size of actual stock market crashes. As discussed in section 2, Figure 1—the main figure in this paper—shows how invariance compares volume traded across assets by extrapolating along diagonal lines rather than along horizontal lines as implied by conventional wisdom.

Conventional wisdom is the alternative hypothesis because prominent academics have used it to interpret crash events. The straw-man nature of this hypothesis results from its inability to explain the size of crash events, not from inherent implausibility or lack of popularity.

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 invariance are based on volume and volatility during the month preceding the crash. The estimates implied by conventional wisdom assume that the percentage price decline 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 shows that three of the crash events involve much larger selling pressure than the other two. The 1929 crash, the 1987 crash, and the Société Générale trades of 2008 all involve 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. The price impact estimates based on conventional wisdom are minuscule in comparison to actual price changes.

In contrast, the price declines predicted by invariance are much closer in magnitude to actual price declines. Yet, there is substantial variation across events. Compared with actual price

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.

declines, the declines predicted by invariance are about twice as large for the 1929 crash, about the same for the Société Générale liquidation, about half as large for the 1987 crash, and much smaller for the 1987 Soros trades and the 2010 flash crash trades.

We conjecture that the deviations of invariance estimates from actual price declines may be mostly related to the speed of execution of bets. For the 1929 market crash, for example, the actual price decline of 25% was smaller than the predicted decline of 46.43%. This may be explained by efforts financial markets 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 quickly a few minutes later. The rapid declines and immediate reversals could be due to the unusually rapid rate at which these 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 do not affect market impact estimates. Nevertheless, 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 were so large. In 1929, market were 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; this may have reduced the size of price declines. In 1987, the actual decline of 32% was approximately double the predicted decline of 16.77%. Price declines may have been exac-

erated by breakdowns in the market mechanism documented in the Brady Report.

Since the invariance approach makes predictions about the frequency and size distribution of bets, it also has implications for how often crashes are expected to happen. Large sales of a size corresponding to the two flash crashes are 4.5-standard-deviation events, which are expected to occur several times per year. We conjecture that such large sales do not usually cause flash crashes because they are executed more slowly. The three largest crashes are approximately 6-standard-deviation events. Under the hypothesis that bets follow a log-normal distribution, these events are so rare that such large crashes would be expected to occur only once in hundreds or thousands of years. Obviously, the actual frequency of these crashes is far higher. Either the tails of the distribution are fatter than a log-normal—consistent with a power law—or the log-variance is larger than the value of 2.53 estimated from portfolio transition data by Kyle and Obizhaeva (2016).

The five bet-induced stock market crashes differ from the long-lasting macroeconomic crises catalogued by Reinhart and Rogoff (2009), who examine sovereign defaults, banking crises, exchange rate crises, and bouts of high inflation. Even after significant changes in macroeconomic policies and market regulation, it usually takes many years for the affected economies to recover from fundamental problems associated with the insolvency of financial institutions. In contrast, bet-induced stock market crashes and panics are likely to be short lived, especially if followed by appropriate government policy. 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. The wealth effect of declining equity prices may have helped trigger a recession by reducing consumption. Yet, Friedman and Schwarz (1963) describe how the Great depression of the 1930s resulted from a subsequent shift towards a deflationary monetary policy, not from the 1929 crash itself. After the liquidation of Jérôme Kerviel's rogue trades in 2008, an immediate 75-basis point interest rate cut by the Fed may have prevented this event from quickly spiraling into a deeper financial crisis, but it did not prevent the collapse of Bear Stearns a few weeks later. It was the bursting of the real estate credit bubble, not the unwinding of Kerviel's gigantic bet, that led to the deep and long lasting recession which unfolded in 2008–2009.

It is widely believed that market crashes result from irrational behavior. This “animal spirits” hypothesis of market crashes says that price fluctuations occur as a result of random changes in psychology, which may not be based on economically relevant information or rationality. Keynes (1936) said that financial decisions may be taken as the result of “animal spirits—a spon-

taneous 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.” According to this hypothesis, market crashes occur when decisions are driven by changes in mind-set based on emotions and social psychology instead of rational calculations. Promptly after the 1987 crash, for example, 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. 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 also discussed and dismissed the hypothesis that portfolio insurance sales might lead to a market meltdown. These accurate explanations did not involve animal spirits. While the sales of George Soros in 1987 may reflect the animal spirits of this one person, the rapid recoveries of prices after the event do not suggest market-wide irrationality or psychological contagion; they suggest the opposite.

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

1 Market Impact of Large Bets: Conventional Wisdom

Many prominent economists believe that stock market crashes ** do not result from selling pressure. Their views are based on the logic of perfectly competitive capital markets, the capital asset pricing model, and the efficient markets hypothesis, which we refer to as conventional intuition.

The risk premium on equities, typically assumed to be 5%–7% per year, is believed to reflect compensation for bearing risk of the entire stock market for one year. Since execution of large bets imposes on market participants risks of much smaller size than the market portfolio—and the positions are held over much shorter horizons than one year—the compensation required for absorbing these risks should be dramatically smaller than the equity risk premium. This implies an elastic demand for individual stocks and therefore small price changes in response to selling pressure. For example, using a theoretical model of competitive capital markets, Bren-

nan and Schwartz (1989) claim that portfolio insurance sales would have had an effect on prices 100 times smaller than the actual size of the 1987 crash.

Empirical studies of secondary distributions (Scholes, 1972), block trades (Kraus and Stoll, 1972), index inclusions and deletions of equities (Harris and Gurel, 1986; Wurgler and Zhuravskaya, 2002), and other events usually find that selling 1% of an individual stock’s shares outstanding has a price impact of at most 1%. Table IV (p. 603) in the latter study provides a survey of existing price elasticity estimates of demand for individual stocks; these estimates vary significantly across studies, sometimes by several orders of magnitude. By extrapolating these estimates for individual stocks to equity indexes, researchers and regulators have concluded that quantities sold during stock crashes were too small to explain the dramatic plunges in prices.

Leland and Rubinstein (1988), the academics most closely associated with portfolio insurance, echo this argument: “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.”

The conventional wisdom says that the price elasticity of demand for an equity asset is approximately one: Selling 1% of market capitalization moves prices by 1%. Mathematically, the expected log-percentage market impact $\Delta \ln P$ from buying or selling Q shares of a stock with a current stock price P dollars per share and N shares outstanding is

$$\Delta \ln P \approx \frac{\Delta P}{P} = \frac{Q}{N}. \quad (1)$$

Throughout this paper, we adopt the convention that Q is unsigned trade size and $\Delta P/P$ is expected unsigned price impact. We thus assume $Q > 0$ and $\Delta P > 0$ for both buy bets and sell bets.¹ Similar formula can be written for simple percentage impact $\Delta P/P$, where ΔP is

¹The size of market impact $\Delta \ln P$ is either the expectation of the post-trade log-price minus pre-trade log-price

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.

Price impact can be also expressed as a function of the fraction of average daily volume traded. Let V denote volume in shares per day. Assume that an asset's turnover is approximately 100% over 250 trading days per year, then V is about 0.40% per day. Since 1% of market capitalization is approximately equal to 250% of one day's volume, the conventional wisdom implies that selling a fixed fraction of average daily volume has the same price impact, regardless of the stock's capitalization:

$$\Delta \ln P \approx \frac{\Delta P}{P} = \frac{Q}{250 \cdot V}. \quad (2)$$

This equation implies that selling 5%–10% of average daily volume has price impact close to zero.

Based on this intuition, the Brady Report came to the following conclusion about daily trading volume elasticities in the 1929 crash compared 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.”

More recent models of market impact imply small price impacts of selling large quantities as well. These papers study executions of large orders by institutional investors (Chan and Lakonishok, 1995, 1997; Keim and Madhavan, 1997). The “square root model” of price impact, described by Grinold and Kahn (1995) and Torre (1997), says that the execution of an order of size Q moves price by $\sigma \cdot (Q/V)^{1/2}$. Frazzini, Israel and Moskowitz (2018) estimate a more complicated version of the square root model. Almgren et al. (2005) calibrate a related model incorporating information about execution horizon T , obtaining price impact $0.314 \cdot \sigma \cdot Q/V \cdot (N/V)^{1/4} + 0.284 \cdot \sigma \cdot (Q/(V \cdot T))^{3/5}$. All of these models imply small price impacts

for buy bets or the expectation of the pre-trade log-price minus post-trade log-price for sell bets, as in cases of market crashes.

for large bets, consistent with conventional wisdom; results for these models are reported in the Appendix B.

Since price pressure of sales was thought to be too small, 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 the price impact of trading. Yet, it is difficult to find new fundamental information shocks to which market prices would have reacted so drastically.

We disagree with the conventional wisdom. As discussed next, market microstructure invariance implies a methodology for extrapolating price impact from less liquid markets to more liquid markets in a manner than does explain how the five stock market crashes could have resulted from documented selling pressure.

2 Market Impact of Large Bets: Microstructure Invariance

Invariance does not question existing market impact estimates for individual stocks. Instead, it suggests a different methodology for extrapolating such estimates to the entire stock market.

Market Microstructure Invariance and Business Time. Market microstructure invariance is based on the simple intuition that trading in a speculative market is a game which transfers risks in business time. The speed with which business time passes varies significantly across assets. Asset-specific business time passes at a rate proportional to the rate at which bets arrive. This rate is fast in liquid markets and slow in illiquid markets.

Market microstructure invariance is the following pair 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) The expected dollar transactions costs of executing similar bets are constant across markets, *when similar bets are defined as bets transferring the same dollar risks per unit of business time*.²

These conjectures have important implications for the rate at which financial markets trans-

²Our derivation of invariance most closely follows the approach in Kyle and Obizhaeva (2016). Kyle, Obizhaeva and Wang (2014) and Kyle and Obizhaeva (2017) derive similar invariance relations in the context of one-period and dynamic equilibrium models of trading, respectively. Kyle and Obizhaeva (2018) obtain similar predictions using dimensional analysis and leverage neutrality.

fer risks. Define “trading activity” as the product of dollar volume and returns volatility:

$$W = P \cdot V \cdot \sigma. \quad (3)$$

Trading activity W is a better measure of the rate at which a market transfers risk than dollar volume $P \cdot V$ because it takes into account that trading assets with higher volatility σ transfers proportionally more risk per dollar traded. Since dollar volume $P \cdot V$ has units of dollars/day and the standard deviation of returns has units per day^{1/2}, trading activity W has units dollars/day^{3/2}.

Let H denote the length of a “business day,” which is inversely proportional to the rate at which bets arrive. The speed of business time is proportional to $1/H$. To standardize units, define the length of the business day as one calendar day for a benchmark stock with stock price $P^* = \$40$, expected trading volume $V^* = 10^6$ shares per calendar day, expected daily percentage standard deviation of returns $\sigma^* = 0.02$ 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.

Appendix A proves that the two invariance conjectures imply that H must be inversely proportional to the 2/3 power of trading activity,

$$\frac{1}{H} = \left(\frac{W}{W^*} \right)^{2/3} = \left(\frac{P \cdot V \cdot \sigma}{P^* \cdot V^* \cdot \sigma^*} \right)^{2/3}. \quad (4)$$

Business time H represents different lengths of calendar time for different assets: several weeks for thinly traded stocks, one day for the benchmark stock, a few hours for actively traded stocks, and a few minutes for the market as a whole.

The logic of invariance implies that probability distributions of random bet sizes \tilde{Q} must be the same across assets if bet size is scaled by trading volume over one business day $V \cdot H$. Let \tilde{Z}^* denote a random variable with this invariant distribution. This imposes strong restrictions on the size distribution of bets as a function of daily volume V in calendar time:

$$\frac{Q}{V \cdot H} \stackrel{d}{=} Z^* \text{ is invariant} \quad \rightarrow \quad \frac{Q}{V} = H \cdot Z^* = \left(\frac{W^*}{W} \right)^{2/3} \cdot Z^*. \quad (5)$$

Bets Q that transfer equivalent dollar risks correspond to lower fractions of daily volume in calendar time, Q/V , in more liquid markets where the length of business day H is smaller.

Let $\Delta P/P$ denote the price impact of a bet of size Q . Invariance implies that equivalent bets in different markets have equivalent price impact $\Delta P/P$ equal to the same fraction of returns

volatility over one business day $\sigma\sqrt{H}$. This implies a specific way to extrapolate market impact estimates across markets:

$$\frac{\Delta P}{P \cdot \sigma \cdot \sqrt{H}} = f(Z^*) \text{ is invariant} \quad \rightarrow \quad \frac{\Delta P}{P} = \sigma \cdot \sqrt{H} \cdot f(Z^*) \sim \sigma W^{-1/3} \cdot f(Z^*). \quad (6)$$

Bets Q that transfer equivalent dollar risks induce percentage price changes $\Delta P/P$ corresponding to lower fractions of daily volatility σ in calendar time in more liquid markets where the length of business day H is short.

The combination of these two intuitions leads to price impact functions of a particular form. If price impact function for an order of size Q is modeled as a power function $f(Z^*) = \alpha \cdot (Z^*)^\beta$ with proportionality constant α and exponent β , then it must be written as

$$\frac{\Delta P}{P} = \alpha \cdot \sigma \cdot \sqrt{H} \cdot \left(\frac{Q}{V \cdot H} \right)^\beta. \quad (7)$$

After plugging the length of business day H from equation (4), the equation (7) takes the following form for the linear model ($\beta = 1$) and the square root model ($\beta = 1/2$),

$$\frac{\Delta P}{P} = \alpha \cdot \sigma \cdot \sqrt{H} \cdot \left(\frac{Q}{V \cdot H} \right) \sim (P \cdot V)^{1/3} \cdot \sigma^{4/3} \cdot \left(\frac{Q}{V} \right) \quad \text{if } \beta = 1, \quad (8)$$

$$\frac{\Delta P}{P} = \alpha \cdot \sigma \cdot \sqrt{H} \cdot \left(\frac{Q}{V \cdot H} \right)^{1/2} \sim \sigma \cdot \left(\frac{Q}{V} \right)^{1/2} \quad \text{if } \beta = 1/2. \quad (9)$$

Linear price impact models ($\beta = 1$) are popular with finance theorists, while the square root model ($\beta = 1/2$) is popular with practitioners.

In comparison with conventional intuition (2) that bets of the same fraction of daily volume Q/V must have the same price impact, the linear specification (8) has additional terms. Invariance says that the linear impact of a bet of fraction Q/V of expected daily volume must be proportional to the product of the 1/3 power of dollar volume $V \cdot P$ and the 4/3 power of volatility σ . These powers show up due to the difference in speed of business time across markets. The square-root specification (9) is essentially the same as the Barra model.

Here is an example of bets in liquid and illiquid assets. Suppose both assets have annual turnover of 100% over 250 trading days. The first asset is a benchmark stock with $P^* = \$40$, $V^* = 10^6$, and $\sigma^* = 0.02$. The second asset is the entire U.S. stock market, which consists of both

the stock index futures market and underlying stock market. The market has daily dollar volume $V \cdot P = \$270 \cdot 10^9$ and daily returns volatility $\sigma = 0.01$, i.e., about 6,750 ($= 15^3 \cdot 2$) times the dollar volume of the benchmark stock and 1/2 of its volatility. Since business times passes at a rate proportional to $(P \cdot V \cdot \sigma)^{2/3}$, the stock market operates about 225 times faster ($= (6,750 \cdot 1/2)^{2/3}$) than the market for the benchmark stock, $H^* = 225 \cdot H$. The invariance hypothesis implies that a proper comparison must involve comparison of *one calendar day* in the entire stock market with 225 *calendar days* in the benchmark stock.

Suppose a market bet equal in size to 10% of a daily market volume is executed in the market; this would be about three times less, for example, than selling pressure during the liquidation of Kerviel's position in earlier 2008. The conventional intuition would imply that this market bet must be equivalent to the stock bet of *the same* 10% of daily stock volume, or about 4 basis points ($= 10\% \cdot (1/250)$) of stock's shares outstanding. According to existing empirical estimates for U.S. stocks and equation (1), this stock bet would be expected to have a minuscule impact of about 4 basis points on stock prices. The market bet would be predicted to have *the same* impact of 4 basis points—or even smaller—on market prices.³

By contrast, the invariance-implied extrapolation (8) leads to a very different prediction, because it implies comparison between business days, not calendar days. The same market bet of 10% of a daily market volume must be equivalent to the stock bet equal in size to 10% of volume over 225 calendar days, or 2250% of a single day stock volume! This selling would involve 9% ($= 10\% \cdot 225 \cdot 1/250$) of stock's shares outstanding and thus move a stock price by about 9% according to existing empirical estimates for U.S. stocks. This impact can be further expressed in units of standard deviation of returns in stock business time as 30% of a 225-day stock volatility equal to 30% ($= \sigma^* \cdot \sqrt{225} = 0.02 \cdot \sqrt{225}$). Invariance extrapolation implies that the market bet must have the same price impact in units of standard deviation of returns in business time, i.e., equal to 30% of daily market volatility $\sigma = 0.01$; this suggests the impact of market bet of $0.30 \cdot 0.01$, i.e., 30 basis points.

Two different extrapolation methods for determining bets “equivalent” across markets predict impacts that differ by a factor of 7.5. The impact of market bet of 10% of daily volume is equal to 4 basis points based on conventional extrapolation in calendar time and 30 basis points

³The impact of stock bet of 4 basis points is equal to 2% of daily stock volatility $\sigma^* = 0.02$. The impact of market bet of 4 basis points is equal to 4% of daily market volatility $\sigma = 0.01$, rather than the same 2% of daily market volatility, because formula (1) does not contain explicit adjustment for differences in volatility as formula (7) does.

based on invariance extrapolation in business time! To summarize, we disagree with the conventional wisdom about how to extrapolate estimates for U.S. individual stocks to the overall market, not about the empirical estimates for stocks themselves.

Evidence Based on Portfolio Transition Orders. We further confirm this invariance intuition by comparing the magnitude of selling pressure during five market crashes with the size of large institutional orders executed in U.S. equities. We examine 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. Trades converting the legacy portfolio into the new portfolio are typically handled by a professional third-party transition manager.

The distributions of these institutional orders are estimated to be symmetric about zero with unsigned order size close to the log-normal:

$$\ln\left(\frac{\tilde{Q}}{V}\right) \sim \mathcal{N}\left(-5.71 - \frac{2}{3} \cdot \ln\left(\frac{W}{W^*}\right), 2.53\right). \quad (10)$$

Thus, the empirical distribution of $\ln(\tilde{Q}/V)$ indeed has a slope of $-2/3$ with respect to change in log trading activity $\ln(W)$, as predicted by bet size invariance (5). Kyle and Obizhaeva (2016) also document evidence in favor of transaction costs invariance (6).

Figure 1 is the main figure of this paper. It shows how the largest portfolio transition orders change in magnitude across U.S. equities with different levels of trading activity W . 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 of $\ln(W/W^*) = 0$ corresponds to the benchmark stock. The figure shows that trading activity varies by a factor of about one million ($= \exp(14)$) from the least actively traded stocks with $\ln(W/W^*) = -12$ to the most actively traded stocks with $\ln(W/W^*) = 2.00$. Trading activity in the aggregate stock market, with $\ln(W/W^*) = 8.20$, can be about 500 ($= \exp(6.20)$) times larger than actively traded individual stock.

The (black) horizontal lines represent “equivalent” orders of a given percentage of calendar-day volume, thus showing the extrapolation direction as simplistically implied by the conventional intuition; for example, the horizontal line $|Q/V| = 5\%$ represents orders equal to 5% of daily volume.

The diagonal lines represent “equivalent” orders as implied by invariance estimates for the

distribution of bet sizes. The lowest (red) diagonal line has a slope of $-2/3$ and represents median bets Q/V ; this line intersects the vertical axis at -5.71 , a dot corresponding to a median bet in the benchmark stock equal to $\exp(-5.71)$, or approximately $0.33\% \cdot V$. The six parallel diagonal lines above this median line represent orders whose log-size is one through six standard deviations above the median bet size, respectively, as implied by invariance; the vertical distance between each line is the log-standard deviation of $2.53^{1/2} \approx 1.60$ from distribution (10). Each log standard deviation represents variation in bet size by a factor of $\exp(2.53^{1/2}) \approx 4.90$. For example, invariance implies that orders of five percent of daily volume ($Q/V = 5\%$) are much smaller than the median for inactive stocks, about two-standard-deviation events for active stocks, and five-standard-deviation events for the market as a whole.

For each of the 60 months from January 2001 through December 2005, each of the 400,000+ portfolio transition orders is placed into a volume bin 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 60 months. 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, confirming extrapolation along an invariance-implied line with slope of $-2/3$ explains variation in large-bet size much better than conventional-extrapolation along horizontal lines representing constant fractions of trading activity. The diamond points are certainly not on a horizontal line, as would be predicted by the conventional wisdom.

Figure 1 also shows how the largest portfolio transition orders would compare with the five crash events depicted by red round dots. Relative to the largest transition orders, the five market crashes are clearly outliers. The two flash crashes of 2.29% and 1.49% of daily volume correspond to about 4.5-standard-deviation events in log bet size. The 1929 crash of 241.52%, the 1987 crash of 66.84%, and the liquidation of Jérôme Kerviel's positions of 27.70% of daily volume correspond to about 6-standard-deviation events.

Figure 1 shows that even though crash events are small as a percentage of daily volume, they are extremely large in the context of invariance and thus expected to have significant price impact. By contrast, extrapolating along horizontal lines according to conventional wisdom implies that nothing unusual would be expected to happen during these five episodes, since even

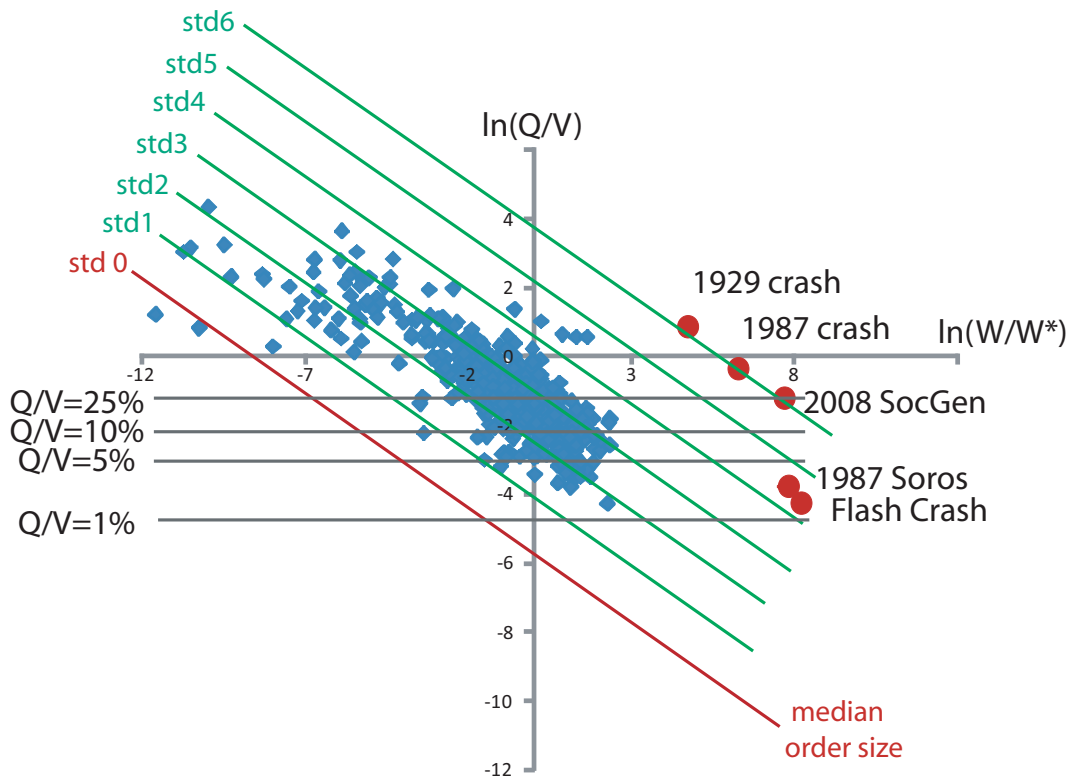


Figure 1: Largest Portfolio Transition Orders and Market Crashes.

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 $W = V \cdot P \cdot \sigma$. 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).

the largest orders such as 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.

Invariance-Implied Market Impact Formulas. We use a log-linear version of the linear impact model (8). The expected percentage price impact from buying or selling Q shares for security with share price P , expected daily volume V , and daily expected volatility σ is given by

$$\Delta \ln P = \frac{\bar{\lambda}}{10^4} \cdot \left(\frac{P \cdot V}{40 \cdot 10^6} \right)^{1/3} \cdot \left(\frac{\sigma}{0.02} \right)^{4/3} \cdot \frac{Q}{(0.01)V}. \quad (11)$$

The formula depends on the proportionality factor $\bar{\lambda}$, which is scaled to measure in basis points the market impact of trading $Q = 1\% \cdot V$ of a benchmark stock with $P^* = \$40$, $V^* = 10^6$, and $\sigma^* = 0.02$.

Invariance says that this factor $\bar{\lambda}$ is the same for all markets and time periods. Kyle and Obizhaeva (2016) estimate $\bar{\lambda}$ using portfolio transition orders. As discussed by Perold (1988), implementation shortfall is the difference between actual execution prices and prices based on transactions-cost-free “paper trading” at prices observed in the market just before the order is placed. Portfolio transition trades are ideal for using implementation shortfall to estimate transactions costs because the exogeneity of the order sizes eliminates selection bias. The inferred value of $\bar{\lambda}$ is equal to 5.00 basis points with standard errors of 0.38 basis points.⁴

We choose to consider 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 about 4.20% of daily volume and median size of 0.57% of daily volume; for these smaller orders, the distinction between continuous compounding and simple compounding is immaterial and results in the same estimates.

This is a universal formula for market impact that may be applied to different markets and time periods. Having calibrated it on portfolio transition order in the individual stocks, we apply the same formula to large market bets.

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

⁴Kyle and Obizhaeva (2016) estimate an *average* impact cost parameter of $\bar{\kappa}_I = 2.50$ basis points (standard error 0.19 basis points) for transition orders (see equation (37) on page 1400). Assuming that these orders are broken into pieces and executed at prices which tend to increase due to movement along an upward sloping supply schedule, total price impact $\bar{\lambda}$ is 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 size of standard errors suggests 2-std bounds of our estimates to be less than 20%.

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 (11) 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 (11) 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 below in Section 8.

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. Invariance alone does not explain crash events; instead, crash events are explained by applying invariance to a linear model. To make this point, “invariance” assumes a linear impact function for the rest of the paper. Due to its concavity, the square root model predicts much smaller price declines during crash events; we present these estimates along with estimates based on alternative models in the Appendix.

Sixth, while our market impact formula predicts price impact resulting from bets, the actual price changes 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 and trading by other traders. We next discuss how other factors may have influenced prices during the crashes.

3 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 seared in the memories of many after it was followed by even larger stock price declines from 1930 to 1932, bank runs, and the Great Depression.⁵

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.

⁵Our 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).

Both the stock market boom and the boom in margin lending came to an abrupt end during the last week of October 1929.

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.

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.

The Broker Loan Market. 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 be-

ginning 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).

Market Impact of Margin Selling. Liquidation of \$1.181 billion during the last week of October and \$4.348 billion over the last three months of 1929 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 (11).

In the month prior to the market crash, typical trading volume was reported to be \$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 therefore multiply reported volume by 13/10, obtaining an estimate of \$444.97 million per day. Historical volatility the month prior to October 1929 was about 2.00% per day. The total estimated margin sales of \$1.181 billion during the last week of October were approximately 265% of average daily volume in the previous month. To convert 1929 dollars to 2005

Broker Loans, Bank Loans, and DJIA, 1926-1930.

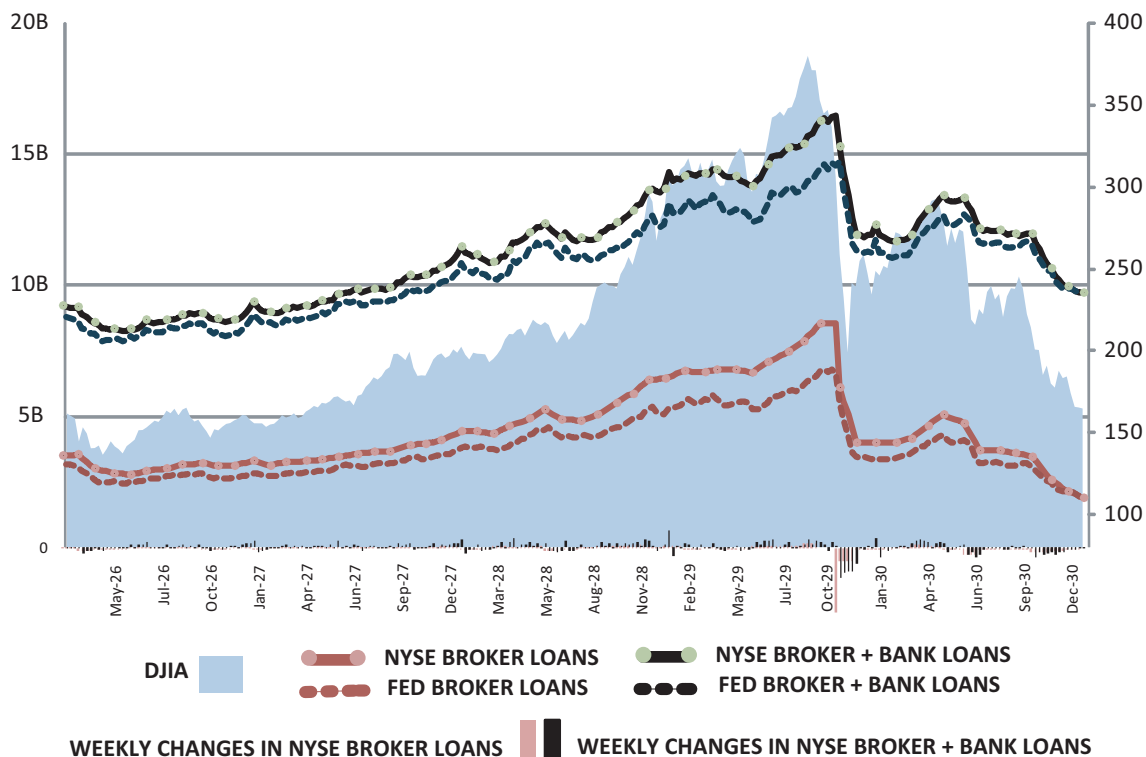


Figure 2: Broker Loans and 1929 Market Crash.

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.

dollars, we use the GDP deflator of 9.42. 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 that sample.

Equation (11) implies that margin-related sales of \$1.181 billion were estimated to trigger a price decline of 46.43%:

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

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%. The actual price change was 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 wisdom (1) implies a price change of only 1.36%. Compared with the observed price drop of 25% during the last week of October 1929, conventional wisdom predicts a much smaller decline than invariance.

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

m:	Months Preceding 24 October 1929					
	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%

The 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 given a GDP deflator adjustment which equates \$1 in 1929 to \$9.42 in 2005, 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.

We also make price impact calculations for estimated 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.

4 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 average fell 23%, and the S&P 500 futures market dropped 29%.

It has long been debated whether this dramatic decrease in prices resulted from the price impact of 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 construct estimates of sales by portfolio insurers from tables in the Brady Report, figures 13–16, pp. 197–198, obtaining results similar to Gammill and Marsh (1988). To convert 1987 dollars to 2005 dollars, we use the GDP deflator of 1.54, but report all numbers below in 1987 dollars. We consider the entire stock market to be one market; this is consistent with the Brady Report. 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. Accordingly, we estimate sales as the sum of portfolio insurance sales in the futures market and the NYSE and estimate expected daily volume as the sum of average daily volume in the futures market and the NYSE for the previous month.

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. We use the gross sales amount of \$13.75 billion in futures and stocks combined for the purpose of analyzing price impact of portfolio insurance sales.

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. In the month prior to the market crash, 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 (11) yields a predicted price decline of 16.77%:

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

Table 3 also 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%.

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.

We present in table 3 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

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. An inflation factor of 1.54 converts 1987 dollars to 2005 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.

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.

Our implied price impact is somewhat smaller than the astonishing 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 and further aggravated by break-downs in the market mechanism which disrupted index arbitrage relationships, as documented in the Brady Report. 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.

5 Trades of George Soros on October 22, 1987

George Soros is known as a philanthropist and speculator who made more than a billion dollars by shorting the British pound and “breaking the Bank of England” in 1992. On Thursday, October 22, 1987, just three days after the market crash, Soros had a bad day. He lost \$60 million in minutes by selling large numbers of S&P 500 futures contracts as prices spiked down 22% at the opening of trading. These sales has been presumably attributed to pessimistic predictions that Robert Prechter made based on “Elliott Wave Theory” and similarities between the 1929 crash and the 1987 crash. In the two years afterwards, Soros withdrew from active management of the Quantum Fund.

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 plummeted 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. 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. Remarkably, 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 this 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, 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, about \$309.60 million in 1987 dollars, was equal to 2.29% of average daily volume. Given the prior month estimates, equation (11) predicts a price decline of 6.27%:

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

Table 4 presents three sets of price impact 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).

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 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%.

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

m:	Months Preceding 22 October 1987					
	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 a GDP deflator which equates \$1 in 1987 to \$1.54 in 2005, 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%.

6 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 Author-

ity (AMF), the French stock market regulator. The AMF allowed Société Générale 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. 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,” (very early Saturday morning in Washington!), 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 could have helped the bank to 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 Kerviel's positions more expensive.

In explaining the costs of liquidating the positions to disgruntled 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. We examine whether the losses associated with price impact predicted by invariance are consistent with reported losses and observed declines in prices.

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 (11), 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% in 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 a GDP deflator of 0.92 to convert 2008 dollars into 2005 dollars.

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.

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 examine whether implied cost estimates are consistent with officially reported losses

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

m:	Months Preceding January 18, 2008					
	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, one Euro is worth \$1.4690, given an inflation adjustment of 0.92 to convert 2008 dollars to 2005 dollars. 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%.

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.⁶

Large price declines in markets where Kerviel did not hold positions suggest that the markets are well integrated as well. 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.

7 The Flash Crash of May 6, 2010

Not all market crashes happen in the United States in October, and not all of them last for a long time. The flash crash of 2010 occurred on May 6 and lasted for only twenty minutes.

During the morning of May 6, the S&P 500 declined by 3%. Rumors of a default by Greece

⁶As a robustness check, we 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 these markets as one market.

had made markets nervous in a context where there was also 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 plummeted 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 rose 5% over the next ten minutes, recovering previous losses.

Shaken market participants began a search for guilty culprits. Fat finger errors and cyber-attack theories were quickly discarded. Many accused high frequency traders of failing to provide liquidity as prices collapsed.

After the flash crash, the Staffs of the CFTC and SEC (2010a,b) issued a joint report. The report highlighted the fact 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. The period of execution corresponded precisely to 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 approximately \$4.37 billion. The joint report did not mention the name of the seller, but newspapers identified the seller as Waddell & Reed.

Many people did not believe that selling 75,000 contracts could have triggered a price drop of 5%. Indeed, the \$4.37 billion in sales represented only 3.75% of the daily trading volume of about 2,000,000 contracts per day in the S&P 500 E-mini futures market. Could the execution of such an order 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. Thus, volume in these two markets combined was \$292 billion. Not surprisingly, trading volume was much higher on the day of May 6. During the previous month, price volatility was about 1.07% per day. Given a GDP deflator of 0.90 between 2005 and 2010, equation (11) 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%:

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

Table 6 shows additional estimates based on historical volume and volatility of S&P 500 E-

mini futures contracts calculated over the preceding m months, with $m = 1, 2, 3, 4, 6, 12$. Estimates based on historical volatility range from 0.44% to 0.73%.

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% (not reported).

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

m:	Months Preceding 6 May 2010					
	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. The GDP deflator of 0.90 converts 2010 dollars to 2005 dollars. 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%.

The predicted price impact of 0.61% is smaller than the actual decline of 5.12%. As discussed later, we believe that unusually fast execution significantly increased the temporary price impact of these trades. This is consistent with the observed rapid rebound in prices. Given that the capitalization of U.S. market was about \$15.077 trillion at the end of 2009, conventional wisdom predicts an even smaller price decline of 0.03%.

8 Conclusion: Lessons Learned

Application of microstructure invariance concepts to intrinsically infrequent historical episodes requires an exercise in judgement to extract appropriate lessons learned. Nevertheless, our analysis suggests important lessons, both for policymakers interested in measuring and predicting crash events of a systemic nature and for asset managers interested in managing market impact costs associated with execution of large trades that might potentially disrupt markets.

Price Impact is Large in Liquid Markets. For the five crash events, the price declines predicted by invariance are large and much more similar in magnitude to actual price declines than predictions based on conventional wisdom and other existing market impact models calibrated for equity markets and shown in the Appendix.

The large impact estimates result from the assumption that price impact is linear in trade size and the prediction that the price impact of trading a given fraction of average daily volume (measured in volatility units) is proportional to the cube root of trading activity, define as a product of dollar volume and volatility. For example, when dollar volume $P \cdot V$ is increased by a factor of 1,000 in equation (11)—approximately consistent with dollar volume differences between a benchmark stock and stock index futures—invariance implies that market impact of a given fraction of volume will increase by a factor of $(1000)^{1/3} = 10$.

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 both for individual stocks and in more liquid markets such as markets for stock index futures. Invariance suggests that this strategy may incur much larger-than-expected transaction costs when implemented in more liquid markets.

Rapid Execution Magnifies Transitory Price Impact. Transaction costs and short-term price impacts are likely to depend on the speed of trading, while the long-term impact of bets depends on their information content.⁷ Rapid execution is likely to generate large temporary price changes associated with V-shaped price paths, in which prices plunge sharply and then

⁷Financial 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.

rapidly recover. Slowing down execution may lessen transitory price impact by signalling that the trades are not based on private information with a short half-life.

Our universal formula (11) contains parameters calibrated using the transaction costs of portfolio transition orders. Transitions are usually executed over a period of a few days, and the most complex of them are sometimes implemented over a period of a few weeks. These horizons are consistent with a prudent pace designed to keep impact costs low. In interpreting results, we make an identifying assumption that transition orders and bets during crash events have been executed at a “natural” speed. It is difficult to properly estimate the effect of execution speed on market impact.

In both the 1987 Soros episode and the 2010 flash crash, there were unusually high volatility, price changes larger than predicted by invariance, and rapid price recovery; the trades were executed unusually rapidly, over a few minutes. The Staffs of the CFTC and SEC (2010b) mentions that during the 2010 flash crash, the order was executed extremely rapidly in just 20 minutes, while two orders of similar size had been executed at previous dates over periods of 5 or 6 hours, 15 times slower. In contrast, measures implemented by bankers and the Fed in the last week of October 1929 smoothed the margin selling out over a period of five weeks rather than a few days; this appears to have lessened temporary price impact, with price declines substantially smaller than predicted by invariance.

Kyle, Obizhaeva and Wang (2014) discuss a smooth trading model which explains theoretically how significant short-term price reaction can result from unusually rapid execution of large bets. 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 quickly materialize, prices quickly rebound. As prices spike down and then recover, much of the heavy selling continues to take place. The predictions of the model are broadly consistent with the empirical patterns observed during the 2010 flash crash. Since the selling pressure during this crash occurred about 15 times faster than typically, the smooth trading model predicts that such fast selling will lead to transitory price impact 15 times greater than in the case when selling occurs at a “normal” rate. Our estimates suggest that a price decline of 0.61% would have occurred if the bet were executed at a normal speed, whereas the actual decline was 5.12%, i.e., about 9 times greater than predicted.

Policy Responses to Mitigate Crashes. Some policy responses may help to mitigate negative effects of crashes. A good example is the 1929 market crash, which may have been so well contained for several reasons.

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 in order to support prices. When their decisions were publicized, the sense of panic subsided. These meetings were not unprecedented. Similar actions, for example, were undertaken by J.P. Morgan and other bankers after a crash in 1907.

Second, the New York Fed also acted prudently in 1929. 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 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 allowed brokers to liquidate the positions of large undermargined stock investors gradually over five weeks, rather than selling collateral off at fire-sale prices over several days. They appear to have helped the market digest a large bet and reduce its temporary price impact, thus avoiding a sudden, brutal bursting of the stock market bubble.

Effect of Large Bets May Propagates Across Integrated Markets. Financial markets are integrated. Heavy selling in one market will affect correlated markets. During liquidation of Kerviel's positions in January 2008, other European markets where Société Générale did not intervene had very similar performance. The bank argued that its own market impact was therefore limited. Similar patterns were documented during the 1987 crash, when not only U.S. markets but also many major world markets experienced severe declines, despite the fact that the portfolio insurance selling was focused on U.S. stocks. According to Roll (1988), this indicates that portfolio insurance did not trigger the 1987 crash.

We disagree. Roll's argument does suggest that market impact estimates should take into account how market liquidity is shared across markets in different continents and markets of related assets, an issue we leave for future research. It supports our preferred strategy for the analysis of Société Générale's trades by assuming markets aggregated across Europe rather than focusing on isolated pools of liquidity in the market for one country's equities. It also supports our strategy for the analysis of the 1987 market crash of looking at aggregated stock and futures markets.

Efficiency versus Stability. There may be a trade-off between efficiency and stability: Less efficient trading arrangements may be more stable during volatile times.

For example, the price declines which occurred during the 1929 stock market crash were remarkably small relative to the 1987 crash given the gigantic levels of selling pressure associated with liquidation of margin loans. 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, we estimate margin related sales to be 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.

The remarkable resilience of the financial markets in 1929 may be explained by their inefficient institutional structure, which may have compartmentalized speculative capital into numerous separate silos. As a result, more capital was required to sustain orderly trading, but it also made the system more stable overall. Indeed, it may be inappropriate to assume that the stock market in 1929 was one integrated market. There were no futures markets or ETFs allowing investors to trade large baskets of stocks. Speculative trading and intermediation associated with underwriting of new stock issues often took place in "pools." These pools traded actively, used leverage, took short position, and arbitrated stocks against options, particularly when facilitating distribution of newly issued equity. The stock pools of the 1920s were typically dedicated to trading only one stock, and investors in the pools often had close connections to the company whose stock the pool traded. There were no prohibitions against insider trading and no SEC requiring firms to disclose material information to the market. When faced with massive liquidations of margin loans, the market may have therefore found that it had more

speculative capital available to stabilize markets than in a more “efficiently” leveraged system in which institutional investors can spread their capital efficiently across markets by trading hundreds of stocks simultaneously.

Invariance provides a way to assess quantitatively the effect of the compartmentalization of the 1929 stock market. How would our price impact estimates change if it were not considered to be one large market, but rather many small markets for different stocks? As a hypothetical illustration, suppose the 1929 stock market consists of 125 separate markets for 125 different stocks, and assume all of them are of the same size and turnover. Comparing to an integrated market, 125 small markets absorb same shocks $125^{2/3} = 25$ times more slowly and their market impacts are $125^{1/3} = 5$ smaller, as implied by equations (4) and (11).

Early Warning Systems May Be Useful and Practical. For all five crash events, policymakers or stock market participants had in hand the information market microstructure invariance requires to quantify the price impact and thus the systemic risks resulting from sudden liquidations of large stock market exposures.

Contrary to beliefs, market crashes in 1929 and 1987 were not completely unpredictable. In both cases, data was publicly available before the crash event. Data on broker loans was published by the Federal Reserve System and the NYSE. Estimates of assets under management by portfolio insurers were available before the 1987 crash. In both cases, 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 trading contributed to the 1987 market crash started long before the crash itself occurred as well. 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 a cascade scenario induced by portfolio insurance sales. After describing in some detail a potential meltdown scenario which closely resembled the subsequent crash in October 1987, the study dismissed the risk of such an event as a remote possibility, in agreement with conventional Wall Street wisdom at the time. On the day the 1987 crash occurred, academics were holding a conference on a topic of potential “market meltdown” induced by portfolio insurance sales!

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 market microstructure 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). At that time, 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 a direction opposite from portfolio insurance, invariance implies that potential portfolio insurance sales were already on the verge of triggering precisely the cascade meltdown scenario.

A similar analysis may be warranted to assess the effect of leveraged ETFs on financial markets. For example, Tuzun (2012) 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 would lead to further price declines of about 1% and potentially trigger a downward spiral.

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 examination of five case studies should not be interpreted as a regression with five data points. Instead, we think that examining them leads to useful insights about why stock market crashes occur, how to prevent them if possible, and how to respond to them appropriately 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 Business*, 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, Securities 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.
- 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., and Anna A. Obizhaeva.** 2016. "Market Microstructure Invariance: Empirical Hypotheses." *Econometrica*, 84(4): 1345–1404.

- Kyle, Albert S., and Anna A. Obizhaeva.** 2017. "Market Microstructure Invariance: A Dynamic Equilibrium Model." Working Paper. Available at SSRN: <https://ssrn.com/abstract=2749531> or <http://dx.doi.org/10.2139/ssrn.2749531>.
- Kyle, Albert S., and Anna A. Obizhaeva.** 2018. "Dimensional Analysis, Leverage Neutrality, and Market Microstructure Invariance." Working Paper, University of Maryland.
- 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.

9 Appendix A: Proof of the Invariance Formula for H .

We outline here the proof of equation (4), which relates the length H of an invariance-implied business day to observable dollar volume $V \cdot P$ and return volatility σ . The proof is a simplified version of the proof in Kyle and Obizhaeva (2016). It is based on two simple equations.

Suppose that trading of a risky asset, on average, involves γ bets per calendar day, and average bet size is \bar{Q} shares. Business time is determined by the expected arrival rate of bets, so the length of a business day is

$$H = 1/\gamma. \quad (12)$$

The faster is the arrival rate of bets, the faster the market transfers risks among traders, and the shorter is the business day.

First, since $V = \gamma \cdot \bar{Q}$, trading activity $W := P \cdot V \cdot \sigma$ in equation (3) can be written as

$$W = \gamma \cdot \bar{Q} \cdot P \cdot \sigma. \quad (13)$$

Second, the invariance conjecture says that dollar risk $P \cdot \sigma$ transferred by a bet of \bar{Q} per units of business time H is invariant across assets and time,

$$Q \cdot P \cdot \sigma \cdot \sqrt{H} = \frac{Q \cdot P \cdot \sigma}{\sqrt{\gamma}} = \text{const.} \quad (14)$$

Define $a := Q \cdot P \cdot \sigma$ and $b := \gamma$. Then the two equations (13) and (14) define a system of two equations in two unknowns a and b :

$$a \cdot b = W \quad \text{and} \quad a \cdot b^{-1/2} = \text{const.} \quad (15)$$

The solution for a and b is

$$a = \text{const} \cdot \left(\frac{W}{\text{const}} \right)^{1/3} \quad \text{and} \quad b = \left(\frac{W}{\text{const}} \right)^{2/3}. \quad (16)$$

Hence, average bet size \bar{Q} and the expected arrival rate of bets $\gamma = 1/H$ satisfy

$$\bar{Q} \cdot P \cdot \sigma = a \sim W^{1/3} \quad \text{and} \quad \gamma = 1/H = b \sim W^{2/3}. \quad (17)$$

After adjustment of the proportionality constant for the length of business day for a benchmark stock, the last equation leads directly to equation (4). More active and more liquid markets have more bets of larger sizes. For example, if trading activity W increases by a factor of 8, then invariance predicts that there will be $8^{2/3} = 4$ times more bets and these bets will be $8^{2/3} = 2$ times larger. Trading activity W is also directly related to measures of liquidity, as discussed in Kyle and Obizhaeva (2018).

10 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):

$$\frac{\Delta P}{P} = \frac{2 \cdot 2.50}{10^4} \cdot \left(\frac{P \cdot V \cdot d_{\text{gdp}}}{40 \cdot 10^6} \right)^{1/3} \cdot \left(\frac{\sigma}{0.02} \right)^{4/3} \cdot \frac{Q}{(0.01) \cdot V}. \quad (18)$$

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

$$\frac{\Delta P}{P} = \frac{2 \cdot 12.08}{10^4} \cdot \left(\frac{\sigma}{0.02} \right) \cdot \left(\frac{Q}{(0.01) \cdot V} \right)^{1/2}. \quad (19)$$

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

$$\frac{\Delta P}{P} = \frac{Q}{N}. \quad (20)$$

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

$$\frac{\Delta P}{P} = \frac{Q}{250 \cdot V}. \quad (21)$$

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

$$\frac{\Delta P}{P} = \sigma \cdot \left(\frac{Q}{V} \right)^{1/2}. \quad (22)$$

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

$$\frac{\Delta P}{P} = 0.314 \cdot \sigma \cdot \frac{Q}{V} \cdot \left(\frac{N}{V} \right)^{1/4} + 2 \cdot 0.142 \cdot \sigma \cdot \left(\frac{Q}{V \cdot T} \right)^{3/5}. \quad (23)$$

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

$$\frac{\Delta P}{P} = \left(-0.2 \cdot \ln(1 + N \cdot 10^{-9} \cdot d_{\text{gdp}}) + 0.35 \cdot \frac{Q}{0.01 \cdot V} + 9.32 \cdot \left(\frac{Q}{0.01 \cdot V} \right)^{1/2} + 0.13 \cdot \sigma \cdot \sqrt{252} \cdot 100 \right) \cdot \frac{2}{10^4}. \quad (24)$$

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 (23) 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 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.

Table 7: Alternative Models.

<i>Panel A: Simple percentage impact $\Delta P/P$.</i>								
	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%

<i>Panel B: Log-percentage impact $\Delta \ln P$.</i>								
	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%

The table 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 $\Delta \ln P$.

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 (18)–(24), 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 in-

stead; 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%.

11 Appendix C: 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 \tilde{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

$$\ln\left(\frac{\tilde{X}}{V}\right) = -5.71 - \frac{2}{3} \cdot \ln\left(\frac{\sigma \cdot P \cdot V}{(0.02)(40)(10^6)}\right) + \sqrt{2.53} \cdot \tilde{Z}, \quad (25)$$

where $\tilde{Z} \sim \mathcal{N}(0, 1)$. 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

$$\ln(\gamma) = \ln(85) + \frac{2}{3} \cdot \ln\left(\frac{\sigma \cdot P \cdot V}{(0.02)(40)(10^6)}\right). \quad (26)$$

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 (25) and (26) 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.