2015

Essays In Asset Bubbles And Financial Crises

Harlan Michael Holt
University of Mississippi

University of Mississippi

Doctoral Thesis

Essays in Asset Bubbles and Financial Crises

Author:

Harlan Holt

Supervisor:

Dr. John Conlon

A thesis submitted in fulfilment of the requirements

for the degree of Doctor of Philosophy

in the

Department of Economics

August 2015

UNIVERSITY OF MISSISSIPPI

Abstract

Department of Economics

Doctor of Philosophy

Essays in Asset Bubbles and Financial Crises

by Harlan Holt

The first essay examines insider trading before, during, and after the October, 2008 stock market crash. I show that inside traders did not appear to predict the crash in the aggregate by selling big before the crash. They also were not able to predict whether their stocks would lose especially badly during the crash individually at either long or short horizons. However, insiders bought in very large numbers immediately after the crash, and were especially active buyers in small firms, high book-to-market firms, and high-beta firms. The insiders who bought during this period successfully predicted post-crash returns that are substantially larger and longer-lasting than the prior literature on insider trading shows for normal trading periods. These results are similar to those in a previous study (Seyhun (1990)) on the 1987 crash. This new evidence of no pre-crash prediction followed by large post-crash returns to insider trading portfolios, combined with Seyhun's (1990) results, suggest this may be a pattern in widescale market crashes.

The second essay presents a model of rational greater-fool asset bubbles that, unlike previous models, includes risk averse agents. Risk aversion allows us to more realistically examine the welfare implications of using central bank policy to prevent speculative bubbles. This improves on previous models in Allen, Morris, and Postlewaite (1993) and Conlon (2015) because they assume risk neutrality among the agents in the market, and therefore cannot comment on the implications of risk aversion for the welfare effects of bubble policy. This paper's results should better inform policy makers on the full costs and benefits of trying to calm swings in asset markets. Comparative dynamics results suggest that a policy of general deflation of overpriced assets is welfare diminishing for buyers, but has ambiguous effects for sellers. We go on to prove that under the reasonable assumption that utility is diminishing in absolute risk aversion, welfare for buyers in the asset market will always be harmed by the general deflation policy.

ABBREVIATIONS

Acronym	**Definition**
ARA	Absolute Risk Aversion
CAPM	Capital Asset Pricing Model
CARA	Constant Absolute Risk Aversion
CRRA	Constant Relative Risk Aversion
DARA	Decrasing Absolute Risk Aversion
EMH	Efficient Markets Hypothesis
FF	Fama-French 3-factor model
IARA	Increasing Absolute Risk Aversion

ACKNOWLEDGEMENTS

I thank my committee members Josh Hendrickson, Mark van Boening, and Bonnie van Ness for their constructive comments on previous drafts of both chapters. I also thank the attendees of the Department of Economics Seminar in which I presented the first chapter for their helpful suggestions and discussions. Lastly, I thank the tireless efforts of my dissertation adviser, John Conlon, whose (many) criticisms substantially improved both chapters.

Contents

List of Figures

List of Tables

CHAPTER 1

INTRODUCTION

For as long as there have been asset markets, there have been asset market panics and crashes. The earliest western capital markets, those for the debt and equity issued by the first joint-stock companies such as the South Sea Company (1720) and the Mississippi Company (also 1720), have shown a propensity to vacillate wildly causing great harm to both those within the market, and those without. These wild fluctuations in the price of asset markets has persisted through the decades to the present day. Many of these instances, such as the Japanese asset market crash (late 1980s), and the global subprime mortgage crisis (2000-2008) have been blamed for widespread economic consequences beyond the asset markets themselves. However, despite the long and infamous history of asset market panics and crashes, we still know precious little about them. What causes wild swings in prices in asset markets? To what extent do they damage the economy? Whose information is most important, and how does the market internalize this information? What can be done to prevent these wild swings in prices (if anything at all)? All of these questions are of first-order importance for the field of economics to explain. This thesis is an attempt to contribute to the existing body of knowledge of financial crises.

The first essay details the stock market crash of 2008. Specifically, we focus on the information communicated by inside traders, those traders with a significant informational advantage about a single firm. We ask whether or not inside traders were able to predict the crash of 2008 by selling early, and whether or not they could predict which stocks would recover the best after the crash.

The answer to the first question, whether or not inside traders could have predicted the crash, is a tentative “no”. We test the hypothesis by gathering information on inside traders leading up to the stock market crash in the second week of October, 2008, then matching the level of insider activity with their stocks’ returns during the crash. We compare the returns of those stocks whose insiders bought most actively with those whose insiders sold most actively. No apparent pattern is detected. Inside traders seemed not to be able to forecast the crash in the aggregate, nor could they individually forecast if their stocks would do better or worse during the crash. While these results suggest that inside traders did not predict the stock market crash of 2008, we cannot say for sure that they were completely unable to do so. Recent developments in finance now make it possible to identify which inside traders are trading on information, and which are trading for routine portfolio rebalancing purposes. We leave these concerns to future research.

The answer to the second question, whether or not inside traders could predict whether their own stocks would outperform the market in the post-crash period, is a resounding “yes”. Inside traders who bought their stock most actively in the weeks immediately after the crash outperformed those who sold their stocks most actively by a wide margin - much wider than the literature suggests exists during normal trading periods. That this result appears without any filtering of routine inside traders out of the sample is remarkable. The pattern is also robust to controlling for the usual factors as well. Comparing stocks on the basis of market capitalization, leverage, and risk does not alter the results. We add some discussion about whether or not the existence of these returns constitute some evidence against the Efficient Markets Hypothesis of Fama (1968). We conclude that this is not likely given that traders in 2008 had very little prior experience with widespread market crashes, and thus had no way to know where to look for profit, and thus, could not arbitrage it away.

The second essay examines the hypothesis that a central bank can protect investors by calming down wild swings in the prices of assets. We begin by recreating a market that develops an asset bubble with fully rational agents with heterogeneous preferences and uncertainty. In certain states of the world, there exist "bad sellers" who know that the asset is overpriced. However, there are other states where there exist honest "good sellers" who think the asset has value. Buyers, however, cannot distinguish between the two types of sellers, and thus, may wind up buying from a bad seller. An asset bubble develops because in certain states of the world, two bad sellers might try to sell to each other, not knowing that the other is also a bad seller. Thus we wind up with a state of the world where everybody knows the asset has a value of zero, and yet the asset curiously has a positive price.

From here we assume the existence of a central bank which knows for sure when at least one agent knows the asset is overpriced. In these states of the world, the central bank announces this fact which pushes the price of the asset down to zero, preventing bad sellers from taking advantage of naive buyers. However, this policy causes an interesting unintended consequence. In the states of the world where the central bank *does not* announce that the asset is overpriced, buyers are encouraged to offer higher prices since they know for sure that they are no longer buying an asset from a bad seller. So while the buyer benefits by being prevented from purchasing an overpriced asset in some states of the world, he is harmed by paying a higher price in other states. Thus, the *a priori* expectation of the effect of the policy on buyers is ambiguous.

We first perform some numerical comparative dynamics experiments to shed some light on whether or not the policy is helpful to buyers in the asset market using numerical examples both derived analytically, and through fixed-point methods. Afterwards, we are able to prove that, for a general class of utility functions, the central bank's efforts to eliminate overpricing in the market will always harm the buyer more by increasing the price in states of the world where the central bank does not announce, than help him by preventing him from buying an asset that is known to be overpriced.

This result has profound implications for the current debate over how central banks should treat asset market bubbles. Namely, it is the first paper to examine this particular unintended consequence the buyers experience. Not only does it suggest that central banks cannot help protect buyers by deflating overpriced assets (or, at least, it harms them on average by doing so), it suggests that even if the central bank knew *with certainty* that assets were overpriced, it still could not improve the outcome for buyers of the asset. Since such a perfect condition is unlikely to ever hold in the real world, we can definitively say that central bank policy should not be used to protect buyers in asset markets from unscrupulous sellers, or from ups and downs in asset market prices.

CHAPTER 2

INSIDER TRADING AND THE 2008 STOCK MARKET CRASH

2.1 Introduction

The ability of insiders to forecast future price movements in their own stocks is well documented. A vast literature has been devoted to uncovering the extent to which inside traders can predict future returns [See Lakonishok and Lee (2001), Finnerty (1976), Lorie and Niederhoffer (1968), Jaffe (1974), Chowdhury et al. (1993), and Seyhun (1986) for examples]. The current paper adds to this body of work by examining insider trading during the October 2008 stock market crash. This particular time period is interesting because of the extreme movements in the stock market aggregates at the time,[1] and the obviously pertinent information flows that come along with widespread market crashes. We have only one other similar period in recent US stock market history, namely the October 19, 1987 "Black Monday" stock market crash. That crash is documented in detail by Seyhun (1990, 1999). Along with the 1987 crash, the 2008 stock market crash presents an excellent opportunity to study how insiders trade before, during and after a widespread, traumatic market event, and whether the market processes information differently during such periods.

[1]For instance, the Dow Jones Industrial Average fell by nearly 20% the week of October 6-10, the worst single week in its history. Similar declines were experienced at the same time by the NASDAQ composite, S&P, Nikkei, Hang Seng, and FTSE indices.

Specifically, I address two questions. First, I test whether or not inside traders in firms that performed especially poorly during the crash predicted the crash by selling their own stocks before the crash. Any evidence of pre-crash selling would imply that insiders had some sort of ability to foresee the crash. Given that previous studies establish insiders' ability to predict future movements in their own stocks, it may have been possible for them to see the 2008 crash coming. Secondly, one of the many interesting results from Seyhun (1990, 1999) is that insiders bought in very large numbers immediately after the 1987 crash. I follow Seyhun by examining whether or not insiders bought in large numbers after the 2008 crash. Additionally, I also examine the post-crash returns to stocks sorted on the basis of insider trading, as in Seyhun (1990), to determine the extent to which these buyers could predict which stocks would recover most rapidly after the crash.

While the prior literature is unambiguous in showing that inside traders can forecast future stock returns, most of these returns are small enough to be wiped out by transactions and/or research costs (Rozeff and Zaman, 1988), and thus the persistent existence of these returns is not thought to conflict with the weak or semi-weak forms of the Efficient Markets Hypothesis (EMH) of Fama (1965, 1970).

Curiously, the returns reported by Seyhun (1990) after the 1987 stock market crash are a great deal larger than most of the returns documented in the prior literature on this topic. To quote Seyhun (1990), p. 1386, "...the difference in the 14-month post-crash returns between [the most heavily bought and most heavily sold stocks] reaches 18.5%. ...This difference substantially exceeds the previously reported stock price movements following insiders' transactions." Even the most optimistic of the other papers generally report no more than 6-8% abnormal returns over the next year for stocks that are bought heavily by insiders. However studies on insider trading and future stock returns are exclusively conducted using data during normal (non-crash) trading periods, whereas Seyhun studies a period with a major market crash. Given this discrepancy, the evidence for the 1987 crash, presented by Seyhun (1990, 1999), begs for a follow up study on returns to stocks purchased by insiders after market crashes. Widespread market crashes don't happen that often, and thus the data generated by these events is exceedingly rare. At this point, the crashes of 1987 and 2008 are the only two major U.S. stock market crashes for which data on insider trading is readily available.

I show that, like the 1987 crash, inside traders did not forecast the stock market crash of 2008 at either long or short horizons by selling early. Not only did they seem not to predict the crash in the aggregate, they also did not predict whether their own individual stocks would do better or worse during the crash, and an outside observer could not have profited by selling the same securities that insiders sold heavily before the crash. Also like the 1987 crash, I show that there was a large burst of insider buying as soon as the crash began, that lasted through the end of 2008. The stocks that were most heavily bought by insiders during that period earned substantial abnormal returns over the next 14 months after the crash. These returns were similar in magnitude to those reported by Seyhun (1990). An outside observer mimicking the behavior of the heaviest of insider buyers could have earned abnormal returns of 26-32%. In addition these returns were very long-lasting, being well spread out over the next 14 months. These returns continued to be large well into the second and third quarters of 2009, about the same time the stock market began recovering.

The primary results in this paper strongly support the observations made by Seyhun (1990) about the 1987 crash: insiders appear not to predict the crash ahead of time, but buy in large numbers afterwards, and insiders are able to take advantage of the market by picking winners as a group. Once again, the returns generated post-crash are significantly larger than even the most optimistic prior literature report on insider trading during normal market periods (See Seyhun (1986), Chowdhury, et al. (1993), and Lakonishok and Lee (2001)). Since Seyhun (1990) studies the only other market crash for which we have adequate data, his work together with the current paper, suggests that returns to mimicking insider movements after market crashes are larger than during normal periods.

The fact that the returns are large and long lasting may seem to be a major violation of the efficient markets hypothesis (EMH). However, because market crashes are such exceedingly rare events, investors likely do not yet know how strongly to react to insider signals during major stock market crashes.

Finally, as in the previous literature (Lakonishok and Lee (2001)), I find insiders were more likely to purchase their own stock if their firms were small firms, and high book-to-market firms. I additionally show that insiders also tended to favor high-beta stocks. Each of these behavioral patterns should imply that portfolios sorted by the intensity of insider purchases should earn higher returns due to the size effect (Banz (1981)), value effect (Fama and French (1992)), and high beta. The abnormal returns generated by portfolios sorted by insider trading patterns, however, continue to be present after inclusion of these factors as control variables.

2.1.1 Literature Review

The prior insider trading literature primarily asks whether insider trades can predict future returns during normal trading periods (Again, see Jaffe (1974), Seyhun (1986), and Lakonishok and Lee (2001) for examples). If insiders can generate significant returns, then it is possible that outsiders mimicking the actions of insiders can potentially profit from the information on insider trades, and thus the existence of persistently large abnormal returns represents evidence against the EMH. The extant literature generally concludes that insiders possess superior information, and are able to profit from this information (Jaffe (1974), Seyhun (1986), Rozeff and Zaman (1988), and Lakonishok and Lee (2001)). Table 2.1 summarizes some of the more highly cited examples in this particular field, in chronological order, along with the predicted returns to portfolios built around insider trading patterns.[2]

[2] In general, the calculation and methods of estimation of these returns vary between studies, but for the most part, they can all be interpreted as abnormal market returns generated by constructing a portfolio based on some measure of insider trading. The measure of insider trading, the holding period length, and the degree of insider trading included in each portfolio varies from study to study. These differences are noted in the table where necessary. With the exception of Seyhun (1990), however, these studies arrive at the same qualitative conclusions.

Table 2.1: Summary of Some Highly Cited Insider Trading Papers

Summary of some of the related literature on the topic. These are a few of the most highly cited empirical studies that examine the predictive power of insider trading. The table is not intended to be a detailed list of all of the papers that study insider trading.

Study	Sample used in study	Size of returns after insider trades	Notes
Lorie and Niederhoffer (1968)	Stratified random sample of NYSE stocks, 1960-1970	-	Lorie and Niederhoffer (1968) shows that insider trades can forecast price changes of 8% or more within six months, but provides no evidence on how large the returns are.
Jaffe (1974)	Large cap CRSP stocks from 1962-68	1%-7.4% over 8 months	
Finnerty (1976)	All NYSE stocks, 1969-1972	8.3% for buys over 11 months, -4.8% for sales	The returns reported to the left probably overstate Finnerty's results somewhat. Finnerty calculates his returns for each month after the trading event. The numbers to the left are simply the sums of all the monthly returns estimated by Finnerty. The results are presented this way to remain consistent with the existing literature that cites it (Seyhun (1986)). Returns for stocks are strongest one month after the insider trading event, but dissipate quickly, mostly disappearing by the third month. Returns after month five are generally very imprecisely estimated with p-values near or above 0.1.
Seyhun (1986)	Stratified random sample of NYSE and AMEX stocks, 1975-1981	3.1%-4.3% for buys, -2.2% for sales over approximately 10 months	Seyhun (1986) also shows that most of the abnormal returns occur within the first 100 days after an insider trading event, similar to Finnerty (1976).
Rozeff and Zaman (1988)	All NYSE stocks, 1973-1982	$\sim$ 6% over 12 months, 3-3.5% with transactions costs included.	
Seyhun (1990)	**All NYSE, NASDAQ, and AMEX stocks 1987-88 in response to the October 1987 stock market crash.**	**10.1%-18.5% additional returns over 14 months for stocks bought after the crash compared to stocks with no inside trading.**	

Continued on next page

Table 2.1 – *Continued from previous page*

Study	Sample used in study	Size of returns after insider trades	Notes
Chowdhury, Howe, and Lin (1993)	All NYSE and AMEX stocks, 1975-1986	0%-3% for buys, 0% for sales over 5 months	Predicted returns are very small and disappear after the first week of trading after an unexpected increase in insider purchases. Nearly no effect is detected for sales.
Lakonishok and Lee (2001)	All NYSE, NASDAQ, and AMEX stocks, 1975-1995	$\sim$ 0%-7.2% for buys, -3.7%-0% for sales over the next 12 months	L&L (2001) generally shows returns of about 3.5% for purchases for all firms and under 2% for sales, but the returns are more extreme for small firms.

From Table 2.1, two general conclusions from the literature are immediately apparent. First, the informational content of purchases is well established and sales seem to be less informative than buys. Constructing a portfolio based on insider purchases is able to generate positive market returns, but sales generate smaller returns, or even no returns at all.[3] This general trend is documented by Finnerty (1976), Seyhun (1986), Chowdhury, Lowe, and Lin (1993), and Lakonishok and Lee (2001).

Secondly, the returns to mimicking inside traders are usually so small that they are likely either severely reduced or completely eliminated by transactions and research costs, hence the persistent existence of these returns need not necessarily conflict with the EMH (Rozeff and Zaman (1988)). In general, the returns for the next 12 months after the insider trading event do not exceed 8%, with the bulk of those returns coming within the first few months after the insider trading event.[4]

The big outlier in Table 2.1 however is Seyhun's (1990) study of the 1987 stock market crash. Seyhun finds much larger and more economically significant returns to portfolios built on insider trading behavior in the year following the 1987 stock market crash.[5] Seyhun finds that portfolios built around insider purchases post-crash did far better than the portfolio of stocks that had no insider trading at all, accumulating between 10-17% higher returns. This figure is higher than even the most optimistic

[3] Marin and Olivier (2008) find that insider sales can predict large downward stock movements on the individual level, but only at long lag lengths.

[4] A potential exception, not listed here, is Cohen, et al. (2012), who find returns between 10% and 20% for portfolios based of "opportunistic" inside traders after screening out "routine" inside trades. However since previous studies do not attempt to screen out these routine trades, the Cohen, et al. study is not directly comparable to the studies included in Table 2.1.

[5] Seyhun (1999) also looks at the 1987 crash. However, that study is largely just a restatement of the findings in Seyhun (1990) and consequently is not listed here.

studies on normal trading period returns, which generally found no more than 8% abnormal returns over the year following portfolio construction.

The disparity between Seyhun's results and the rest of the literature raises several questions. Do Seyhun's results simply reflect statistical noise? Is this a pattern that is common to market crashes? If so, what makes market crashes different, and why did investors not arbitrage these gains away?

I demonstrate in this paper that Seyhun's results were not a statistical fluke. As in 1987, inside traders were able to pick big winners after the 2008 crash, and revealed their information by buying winners in large numbers. Thus, not only were Seyhun's results not a fluke, they seem to be part of a pattern of large post-crash returns to insider trading portfolios.

While the returns to post-crash insider trading portfolios are quite large, these results probably do not violate the EMH. Because crashes are so rare, investors could not have had much knowledge about how to interpret insider trades during the market crash. The US stock market has had only two such crashes for which we have sufficient data to study: 1987 and 2008. Thus, the average US investor after the 2008 crash has likely experienced only one major market crash during his career. We should therefore not be surprised to find that these investors do not know how to interpret the informational content of trades by insiders before, during, and after a market crash, so the subsequent returns signaled by those trades should behave very differently from insider trades during ordinary times.

Seyhun's other principal finding is that insiders did not predict the 1987 crash ahead of time by selling early.[6] The lack of early selling by insiders recalls much of the literature which finds that sales have less informational content than buys. This also presents a hypothesis that can be applied to the 2008 crash, allowing us to further test if these market crashes are similar to each other.

2.2 Insiders

In the United States, an insider in any particular company is defined by the SEC act of 1934 as a senior officer, a person on the board of directors, or a holder of 10% or more of the company's outstanding shares. A central feature of insider regulation is disclosure. Insiders are required to report to the SEC their trades of stock issued by the company they are considered insiders for. Whenever an insider trades their stock, she or he must file a "Form 4" with the SEC, which details the time and date of the trade, whether it was a purchase or a sale, the number of shares traded, and the price of those shares. As soon as they are filed, these forms are made available to the public by the SEC. Prior to 2002, insiders were required to file a Form 4 with the SEC within 10 days after any change in ownership. However in 2002, the Sarbanes-Oxley Act shortened the reporting time to within two business days. Additionally, insiders are not allowed to short-sell their own stock and most companies prohibit or strongly discourage their insiders from trading around major company events such as earnings or dividends announcements. Lastly, short-swing profits from trades made by an insider reversing his or her position within six months are recoverable by either the corporation or the SEC.

[6] Marin and Olivier (2008) anecdotally show a curious increase in insider selling approximately 10 months before the 1987 crash. However their study does not examine the 1987 crash in detail. In light of the findings in Seyhun (1990) and this paper, their anecdote is likely mere happenstance, though their major finding that insiders can successfully predict *individual* stock crashes remains important.

This effectively prohibits insiders from reversing their positions within six months. This restriction likely causes insiders to be much more thoughtful and deliberate about their actions. Since any position cannot be reversed for six months, the informed insider must be sure of his trades and potentially must be able to accurately predict events that may be far off in the future.

2.3 Data and Methodology

Insider trading data are made public by the SEC's Ownership Reporting System (ORS), which is now a part of the Electronic Data Gathering, Analysis, and Retrieval (EDGAR) system. Insider trades are then relayed by other media platforms to the general public. Insider trading data for this paper is collected from the reporting site SECform4.com. Among the data reported on Form 4 are: the date and time of the trade, the date on which the trade was reported,[7] whether the trade was a buy or sell, the size of the trade, the current price of the stock, the identity of the insider, and his/her current holdings of the company's stock. The sample consists of US exchange traded common stock. The sample period for inside trades runs from the week starting January 1, 2007 to December 31, 2008. Table 2.2 presents the descriptive statistics for the insider trading sample.

The sample contains a total of 163,795 insider transactions over the period, from 7,073 firms listed on the NYSE and NASDAQ exchanges. The first line in Table 2.2

[7] As mentioned above, insiders may report their trades at any point within two days after the trade is executed.

TABLE 2.2: Summary Insider Transaction Statistics, Pre- and Post-Crash

Sample statistics for insider trading during the period 1/1/2007 - 12/31/2008. The top line presents the statistics on the average purchases, sales, and transactions per month with a month defined as 22 trading days, as well as the average buy ratio for the entire period, and the total number of firms who had insider trades during this period. The next four lines are sub-samples. First, the pre-crash period designated as January, 2007 through August, 2008. September, 2008 is shown alone to see if insiders in the aggregate increased their selling behavior significantly just before the crash. The crash period is designated as between the dates 10/2/2008 and 10/10/2008. The post-crash period starts at the end of the crash period and continues through the end of 2008.

Subsample	Purchases (per month)	Sales (per month)	Transactions (per month)	Buy Ratio	Firms
Total (1/1/2007 - 12/31/2009)	2,489.0	4,335.8	6,824.8	0.36	7,073
Pre-Crash (1/1/2007 - 8/30/2008)	2,291.0	4,619.9	6,910.9	0.33	6,784
Sept 2008 (Month before the crash)	2,097.0	3,139.0	5,236.0	0.40	1,801
Crash Period (10/2/2008 - 10/10/2008)	2665.1	1,474.0	4,139.1	0.64	624
Post-Crash (10/13/2008 - 12/31/2008)	4,556.2	2,615.9	7,172.1	0.63	2,768

shows statistics on insider trading (per month) for the entire sample period. The next two lines show the descriptive statistics in the pre-crash period separated into two subperiods. The second line presents the data from the beginning of 2007 to August 2008, and the third line shows the only the month of September (the month leading up to the crash). The fourth line reports the figures for the crash period, defined as October 2 through Oct. 10, 2008. The final line shows insider trades after the crash through the end of 2008. Before the crash, insider sales consistently outnumber buys about 2 to 1, but selling lightly decelerates in September to a ratio of about about 3 to 2. In the post-crash period, the ratio reverses completely with buys now outnumbering sales. After October 10, purchases vastly outnumber sales through the end of the year.

To test whether or not insiders could predict the crash, and whether they were able to earn abnormal profits after the crash, returns were collected from the Center for Research in Security Prices (CRSP) over two horizons. Any firms with missing information in the CRSP database were dropped from the sample. First, daily returns from October 2 through Oct. 10, 2008 were used to calculate holding period returns for the crash itself. These returns are matched by firm with insider trades in the pre-crash period to assess insiders' ability to predict the crash. To test whether insiders were able to predict returns *after* the crisis, I collect monthly returns for the next year from Jan. 2009, to Dec. 2009, and match those returns by firm with insider trades in the post-crash period.

Returns for both the crash and post-crash subperiods are risk-adjusted using the typical CAPM and Fama-French (FF) regressions shown in Equations (2.1) and (2.2):

$$ret_{i,t} - rf_t = \alpha_{i,t} + \beta_i(mkt_t - rf_t) \quad (2.1)$$

$$ret_{i,t} - rf_t = \alpha_{i,t,FF} + (\beta_M)_i(mkt_t - rf_t) + (\beta_H)_i HML_t + (\beta_S)_i SMB_t \quad (2.2)$$

Here, ret_i is the raw return for stock i, rf is the (risk-free) rate of return on a 1-month U.S. treasury bill, and $\alpha_{i,t}$ and $\alpha_{i,t,FF}$ are the abnormal returns (hereafter the "alphas") for each stock. The independent variables are the factors for the different specifications. The CAPM model only adjusts for exposure to the market excess. mkt_t is the value-weighted market return in month t. The FF model also includes the book-to-market factor, HML (*high-minus-low*), and the size factor, SMB (*small-minus-big*). Factor data are obtained from the Fama French factor database provided by Wharton Research Data Services (WRDS).

Using the above regressions, I estimate the beta coefficients for each stock in the sample using both the CAPM and Fama-French specifications. Equations (2.1) and (2.2) are estimated using 60 months of past data (Jan. 2002-Dec. 2006) to avoid any overlap with the pre-crash portfolio formation period. The betas from this estimation are saved, and then used to adjust the returns in the relevant sample periods. The risk adjustment is performed by subtracting from the raw return, each of the factors multiplied by the corresponding betas to obtain the alpha for each stock. As a robustness check, I also consider the raw returns to each stock adjusted by subtracting the value weighted NYSE/Nasdaq/AMEX index return.

Lastly, to test whether the behavior of returns is sensitive to firm size and book-to-market (B/M) value in the post-crash period, market capitalization and B/M data are taken from Compustat. Book value (per share) is calculated as the total value of assets minus the total value of liabilities divided by the number of common shares outstanding. B/M is simply the book value per share divided by the market price.

As is common in cross-sectional studies (see: Lakonishok and Lee (2001) and Conrad and Kaul (1993)) I drop all stocks with a price of less than $1 at the end of December, 2008.[8] I also drop all stocks with any missing information from the CRSP tapes or from Compustat. The final sample size for the post-crash period is 1,352 firms.

2.3.1 Insider Trading Metrics for Individual Stocks

I first define the *Net Insider Purchases* of insider trading,

$$NIP_i = numbuys_i - numsells_i$$

NIP_i can be thought of as an insider trading differential for each firm. A positive NIP_i indicates that the stock was a net buy for insiders (the stock was bought more than it was sold), and a negative NIP_i indicates a net sell for insiders. The larger (smaller) the

[8]There were fewer than 10 firms in the sample that were highly insolvent (Liabilities > Assets) at the end of December, 2008. Accordingly, these firms also had very low stock prices, which means a small number of firms in the sample have very large negative B/M figures. Though these firms were small in number, they had B/M values sufficiently large negative enough to cause the mean of B/M to become negative. The problem is solved by simply eliminating these stocks from the sample. This does not cause any of the primary results in this paper to change greatly.

NIP_i, the more that stock was bought (sold) by insiders. Once each stock is assigned an NIP, the stocks are then partitioned into quartiles. Q1 is designated the "high sell" portfolio. It contains the stocks that were most sold during the portfolio formation period (so they have the smallest, i.e. the biggest negative, NIP_i's). Conversely, the "high buy" portfolio is Q4. It contains the stocks that were most bought (have the biggest positive NIP_i's).

NIP is a basic measure of insider trading intensity, but in some cases, it may not be completely informative. For instance, if a stock has an NIP of 20, it could have 20 buys and no sales, or conversely, it could have 120 buys and 100 sells. In the first case, the insider trades are exclusively buys. In the second, there are only 20% more buys than sells. However, the NIP measure will give these two outcomes the same weight in any analysis. For this reason, I also calculate the *Net Insider Purchase Ratio* as in Seyhun (1990) and Lakonishok and Lee (2001), defined as

$$NIPR_i = \frac{numbuys_i - numsells_i}{numbuys_i + numsells_i}$$

This measure gives the appropriate weight in terms of the ratio by which buys exceed sales, and vice-versa. Once again a positive value indicates a stock was a net insider buyer, and a negative value means it was a net inisder seller. The ratio, however, also has a potential drawback. A ratio measure would assign the same value to a firm with one insider trade that was a buy during the portfolio formation period as a firm that had 20 insider trades that were all buys. The previous measure, NIP, therefore remains useful to distinguish between these cases.

Finally, counting transactions may be problematic because a given transaction could be large or small. The two metrics described above treat all transactions the same, regardless of how large the transaction is. Therefore, I also define share-weighted values for the insider trading metrics. The *Net Insider Share Purchase Ratio* is the share-weighted analogue to $NIPR_i$:

$$NISPR_i = \frac{(sharesbought)_i - (sharessold)_i}{(sharesbought)_i + (sharessold)_i}$$

Summary statistics for the insider trading variables and other selected variables appear in Table 2.3 separated by the two different subperiods that this paper studies. The pre-crash period is defined from Oct. 1, 2007 to Sept. 30, 2008, and the post-crash period is defined as Oct. 13, 2008 to Dec. 31, 2008.

2.3.2 Aggregate Insider Trading

Marin and Olivier (2008) provides evidence that insiders can predict crashes in their own individual stocks during normal trading periods, but only for very long lead times.[9] It is important to stress that their results apply for *normal* trading periods for individual stocks that experience a crash. During a period with many stocks crashing at once, it is conceivable that the information signaled by insiders is very different than during normal trading periods. In this subsection, I examine how aggregate insider trading patterns changed around the crash.

[9] Marin and Olivier (2008) also presents some anecdotal evidence of a noticeable increase in insider selling approximately ten months before the 1987 stock market crash. See Footnote 7.

Table 2.3: Cross-sectional Stock Summary Statistics

Selected summary statistics for stocks during the pre-crash period (Oct. 1, 2007 - Sept. 30, 2008) post-crash period (Oct. 13, 2008 - Dec. 31, 2008). All statistics are presented cross-sectionally among firms in the sample. The collection period for the insider trading variables is Oct. 13, 2008 to Dec. 31, 2008. *Size* and B/M data are the values for each stock at the end of Dec., 2008. Since they are not used for the pre-crash analysis, they are omitted in Panel A. The betas were estimated for each stock from equations (2.1) and (2.2) using monthly data from the period Jan. 1, 2002 to Dec. 31, 2006. The insider trading metrics are listed at the bottom of each panel. In order, they are the *Net Insider Purchases* (NIP), *Net Insider Purchase Ratio* ($NIPR$), and *Net Insider Share Purchase Ratio* ($NISPR$).

Variable	Mean	St. Dev.	Min	Max
Panel A - Pre-crash data: Oct. 1, 2007- Sept. 30, 2008				
β (CAPM)	1.143	0.873	-4.042	8.428
β_M (FF)	1.007	0.901	-11.680	7.500
β_H	0.190	1.303	-9.180	10.390
β_S	0.679	1.159	-8.560	12.910
NIP	-3.500	49.500	-579	2638
$NIPR$	-0.053	0.845	-1.00	1.00
$NISPR$	-0.117	0.891	-1.00	1.00
Panel B - Post-crash data: Oct. 13, 2008 - Dec. 31, 2008				
Size (billions of $)	2.474	11.86	0.0007	218.91
B/M	0.394	0.573	-3.48	5.590
β (CAPM)	1.287	0.655	-1.276	3.981
β_M (FF)	1.156	0.683	-1.568	6.576
β_H	-0.106	1.347	-7.511	7.359
β_S	0.722	1.347	-14.04	7.864
NIP	0.820	7.277	-58	84
$NIPR$	0.131	0.915	-1.00	1.00
$NISPR$	0.097	0.950	-1.00	1.00

As in Seyhun (1999), define the aggregate proportion of insider buys as:

$$BuyProportion_t = \frac{\sum_{j=1}^{n_t} BUY_{j,t}}{\sum_{j=1}^{n_t} BUY_{j,t} + \sum_{j=1}^{n_t} SALE_{j,t}} \tag{2.3}$$

where j represents transaction j during week t. The ratio represents the proportion of insider transactions in any particular week that are buys. This measure is an aggregate measure of *transaction level* insider trading, similar to $NIPR$ in Section 2.3.1. However it measures insider buys as a percentage of total transactions *for all stocks* in each week. Figure 2.1 plots the results from computing the transaction proportion on a weekly basis for the two years from the beginning of 2007 to the end of 2008. The vertical lines denote the crash (Oct. 2 - Oct. 7).

FIGURE 2.1: Proportion of insider transactions that are buys. Vertical Lines denote dates of the stock market crash.

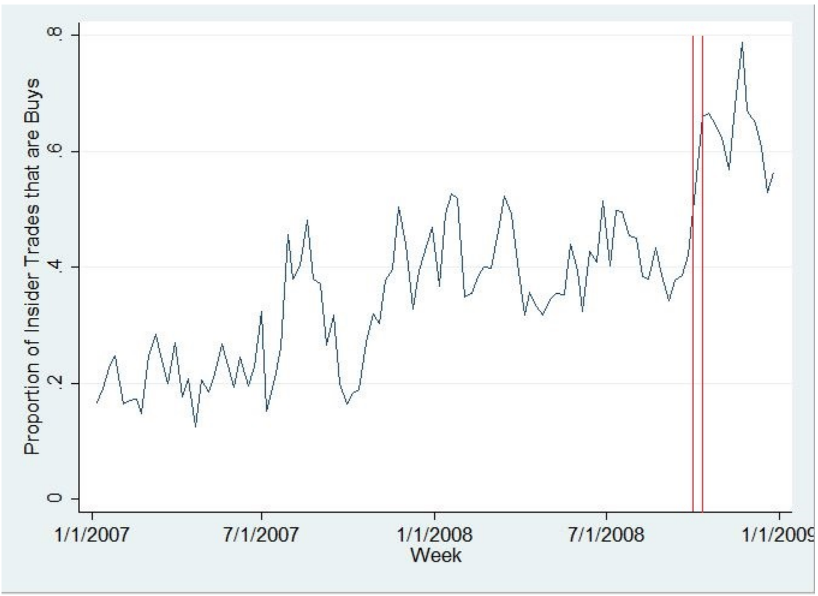

Unlike the 1987 crash, where there was a significant sell-off by insiders around 10-months before the crash (Marin and Olivier 2008), there is not an obvious sign of an abnormally large dip in the insider buying ratio at any point in the two years prior to the crash. In general, the buy-ratio begins at a low level in early 2007, and remains low through most of 2008.

What is immediately noticeable in Figure 2.1 is the extremely large increase in the buy ratio immediately during the week of the crash. Taking the mean value of the simple insider transaction proportion for the five weeks leading up to the crash (0.40) and the five weeks after (0.63), the difference of 23% is economically significant, and statistically significant with a p-value $< .001$. It's worth noting that this jump is not the only significant jump in aggregate insider trading during the sample. For instance, measuring peak-to-peak, there appears to be a large shift upwards in the trading metric in July, 2007. The jump in late 2007 is probably inconsequential for the crash however, since we still get a large peak-to-peak increase at the exact time the crash occurs, and insiders do not shift from being net sellers to net buyers. No noticeable sell off, comparing trough-to-trough, appears at any point in 2008, or in the months leading up to the crash, as compared to what occurred in 1987 (Marin and Olivier (2008)). For a more detailed test of whether insiders predicted the crash, I next construct a test similar to that of Marin and Olivier (2008) to see if insiders individually predicted abnormally large crashes in their own stocks at both long and short horizons.

2.4 Analysis and Results

2.4.1 Were insiders able to predict the crash?

From the data in Section 2.3.3, it appears that insiders (in the aggregate) did not signal the crash with any noticeable increase in selling before the crash. However, it is still possible that they predicted a crash in their individual stocks by selling off the biggest losers before the crash. To further examine this hypothesis, I employ a statistical test similar to the one used in Marin and Olivier (2008) to identify whether insiders could predict crashes in *their stock* far in advance of the crash itself.

I calculate NIP, $NIPR$, and $NISPR$ for each firm for two portfolio formation periods: first for the month of September (one full month before the crash), and then for the months from October 2007 to August 2008 (the period between the 12th and 2nd months before the crash).[10] I then match each stock with its risk adjusted

[10] These portfolio formation periods are parallel to the ones used by Marin and Olivier (2008) so this paper approximately tests the same hypothesis: that insider sales predicted the 2008 crash.

abnormal holding return during the crash, i.e., from Thursday, Oct. 2 through Friday, Oct. 10. Any stocks that were not matched due to missing returns or reporting errors were eliminated from the data.

Table 2.4 presents the abnormal returns (alphas) along with the Market Adjusted returns for portfolios organized using each of the three insider trading metrics presented in the previous section. For the NIP measure the stocks are sorted into quartile portfolios, and the results are presented in Panel A. For the $NIPR$ and $NISPR$ measures, the quartile approach is insufficient because there are many stocks that had only buys or only sells during the portfolio formation period, so I sort the stocks according to the value of $NIPR$ and $NISPR$. The top portfolio in Panels B and C in Table 2.4 correspond to stocks who had no buys (only sells) during the portfolio formation period (or $NIPR = NISPR = -1$). The next portfolio includes stocks with more sells than buys ($NIPR, NISPR \in (-1, 0]$). The third portfolio includes stocks with more buys than sells ($NIPR, NISPR \in (0, 1)$). Finally the last portfolio includes stocks who had no sells (only buys) during the portfolio formation period. Because Sept, 2008 is such a short time period, less than 50 stocks had a mixture of buys and sales during the portfolio formation period, so results are only shown for the "only sell" and "only buy" portfolios. Abnormal returns were calculated using the CAPM, and the Fama-French three-factor methods as in Fama and French (1993). I also include the raw returns for each portfolio minus the value-weighted NYSE/Nasdaq/AMEX market return for comparison under the column labled "Market Adjusted".

Table 2.4: Prediction of Crash Returns

Abnormal returns during the 2008 stock market crash based on a holding period from Oct. 2 to Oct. 10, 2008 (t-statistics in parentheses). The portfolios are formed on the basis of insider trading activity calculated by each stock's Net Insider Purchases (NIP), Net Insider Purchase Ratio ($NIPR$), $andNetInsiderSharePurchaseRatio$(NISPR) during the portfolio formation period. The portfolio formation periods are parallel to those used by Marin and Olivier (2008), Oct. 2007 - Aug. 2008 is the period twelve to two months before the crash. Sept. 2008 is the month immediately preceding the crash. Abnormal returns are reported for both the CAPM and FF 3-factor specifications. "Market Adjusted" returns are the portfolio raw returns minus the value-weighted NYSE/Nasdaq/AMEX return and are included for comparison and robustness. Panel A presents the results by sorting stocks into quartiles according to the Net Insider Purchases (NIP) measure of insider trading. Panel B presents the results using the Net Insider Purchase Ratio, and panel C presents the results for the Net Insider Share Purchase Ratio sorted into portfolios according to the criteria in the first column.

	Oct. 2007 - Aug. 2008 Portfolios			Sept. 2008 Portfolios		
Portfolio	Market Adjusted	CAPM alpha	FF 3-factor alpha	Market Adjusted	CAPM alpha	FF 3-factor alpha
Panel A - NIP						
Quartile 1	-0.0068**	0.0664***	0.0074	0.0064	0.0207	-0.0004
(Most sold)	(-2.06)	(9.09)	(1.01)	(0.80)	(1.80)	(-0.02)
Quartile 2	-0.0003	0.0474***	0.0002	0.0095	0.0402**	-0.0304*
	(-0.06)	(4.78)	(0.02)	(1.00)	(2.26)	(-1.89)
Quartile 3	-0.0068	0.0096	-0.0219**	-0.0026	0.0043	-0.0217
	(-1.51)	(1.15)	(-2.41)	(0.22)	(0.23)	(-1.02)
Quartile 4	-0.0030	0.0209**	-0.0060	0.0031	0.0509**	0.0163
(Most bought)	(-0.55)	(2.11)	(-0.61)	(1.23)	(2.63)	(0.82)
Panel B - NIPR						
Only sells	0.0002	0.0667***	-0.0011	0.0031	0.0272***	-0.0106
	(0.05)	(7.74)	(-0.13)	(0.50)	(2.62)	(-1.28)
$NIPR \in (-1, 0]$	-0.0007	0.0488***	0.0091	-	-	-
	(-0.18)	(6.01)	(1.02)	-	-	-
$NIPR \in (0, 1)$	-0.0009	0.0245**	0.0107	-	-	-
	(-0.16)	(2.53)	(0.97)	-	-	-
Only buys	-0.0133***	0.0013	-0.0329***	-0.0024	0.0219	-0.0069
	(-2.97)	(0.14)	(-3.64)	(-0.25)	(1.54)	(-0.47)
Panel C - NISPR						
Only sells	0.0002	0.0667***	-0.0011	0.0031	0.0272***	-0.0106
	(0.05)	(7.74)	(-0.13)	(0.50)	(2.62)	(-1.28)
$NISPR \in (-1, 0]$	0.0025	0.0376***	0.0059	-	-	-
	(0.64)	(5.02)	(0.72)	-	-	-
$NISPR \in (0, 1)$	-0.0067	0.0427***	0.0166	-	-	-
	(-1.09)	(3.85)	(1.28)	-	-	-
Only buys	-0.0133***	0.0013	-0.0329***	-0.0024	0.0219	-0.0069
	(-2.97)	(0.14)	(-3.64)	(-0.25)	(1.54) ‘	(-0.47)

As can be seen from the Table 2.4, no significant patterns emerge based on insider trading from either time period. Most importantly, quartiles 1 and 2 do not display any consistently negative abnormal returns, indicating that insiders are not forecasting losers. If anything the results seem to show some weak evidence of *negative* prediciton, that is the most sold stocks are doing better than the most bought ones. This primarily appears in the CAPM results for the long-horizon (Oct. 2007-Aug. 2008) subperiod where the returns to the portfolio of most sold stocks is larger than the returns to the portfolio of most bought stocks. It also appears weakly in the long-horizon Market Adjusted returns and the 3-factor alphas. These results are not robust to all measures of insider trading, even for the Oct. 2007-Dec. 2008 subperiod. They're also economically small, being between 1 and 6%.

The CAPM results for the short-horizon in Panel A provide some very weak evidence for prediction. The high-buy quartile does about 3% better than the low-buy portfolio. This result however is not robust to the 3-factor returns, nor does it appear for either of the other metrics in Panels B and C. Overall, these results combined with that in Section 2.3.3 suggest we can soundly reject the hypothesis that insiders predicted the stock by selling early. This is consistent with the results from Seyhun (1990) who also did not observe appreciable early insider selling before the 1987 crash, nor does he observe any evidence of insiders predicting the crash individually.

It is important to stress that, while this paper tests approximately the same hypothesis as Marin and Olivier (2008) these results do not invalidate or contradict those presented in Marin and Olivier (2008). Table 2.4 above simply shows that their results do not apply to this period where many stocks crashed at the same time. One obvious explanation is that insiders simply may not have been able to predict the crash ahead of time. This could be the case if enough stocks with good fundamentals were crashing just as badly as stocks with bad fundamentals, much in the same way that both solvent and insolvent banks suffer during a bank run. A second possibility is that insiders may have been aware that their firms would do especially poorly if the market were to crash, but were unwilling to signal their pessimism to investors by selling. This might have been the case if insiders were more concerned about the confidence of outside investors in their firms than value of their own stock portfolios.

2.4.2 Post-Crash Predictions

The ability of insiders to predict returns is well documented in the literature. However as mentioned above, most of the classic studies on insider trading show that these gains, while statistically significant, are small in magnitude, and almost completely exhausted within a couple of months (Finnerty 1976, Rozeff and Zaman 1988, Chowhdury et al. 1993). Whether or not prices display this pattern after a major stock market crash is another question, and one that has little empirical documentation. I examine this question by forming portfolios based on the NIP, $NIPR$, and $NISPR$ of insider trading for the period immediately following the crash. This portfolio formation period runs from Oct. 13 to Dec. 31, 2008. I then match each stock's NIP, $NIPR$, and $NISPR$ with it's monthly abnormal returns from Jan. 1 to Dec. 31, 2009. The portfolio's abnormal return is defined as the equally-weighted average of the abnormal returns for each stock over the holding period.[11] As in Section 2.4.1, stocks are divided

[11] Seyhun (1990) forms his insider portfolios by measuring the change in the net purchase ratio (purchases divided by transactions) during the week of the 1987 crash. Therefore, he defines insider trading to be the degree to which insiders reversed their positions during the crash. This method captures

into approximately equally sized quartile portfolios based on NIP. I also form portfolios based on $NIPR$, and $NISPR$. For each of these measure, I divide the stocks into four portfolios based on the ratio: A high-sell portfolio where $NIPR$ or $NISPR = -1$ (all sales), a sell portfolio where $NIPR$ or $NISPR \in (-1, 0]$ (more sales than buys), a buy portfolio where $NIPR$ or $NISPR \in (0, 1)$ (more buys than sales), and a high-buy portfolio, where $NIPR$ or $NISPR = 1$ (all buys).

The results are presented in Table 2.5 for both CAPM and Fama-French 3-factor treatments. Panels A, B, and C present the returns for NIP, $NIPR$, and $NISPR$ portfolios respectively. The Market Adjusted returns (raw returns minus total value-weighted market return) are again presented side-by-side for comparison and robustness only.

What's striking in Table 2.5 is how large the abnormal returns are to the high-buy portfolios. The high-buy portfolio for all three measures show returns between 26% and 32% for 2009. The high-buy portfolio for all three measures show returns between 26% and 32% for 2009. The difference between the high-buy and high-sell portfolios range between 22.90% (3-factor alphas for $NIPR$ and $NISPR$) and 27.47% (CAPM alphas for NIP). These returns are slightly larger, but similar to those reported by Seyhun

similar information, however the measures are different from the current paper. The current paper also includes a longer portfolio formation period than Seyhun (1990). This is appropriate for this paper for two reasons. First, the 2008 stock market crash occurred in the context of a much wider economic event: the Subprime Mortgage Lending Crisis of (2007-2008). The market crash occurred in October, 2008, but the market began to decline late in 2007, and continued to decline until well into 2009. Second, the short swing rule (see section 2.2) may prevent an insider from trading within a short period of time even if he wants to trade. Therefore, it makes sense to use a longer formation period for portfolios based on insider trades after the crash. These differences, however, are not likely to seriously affect a qualitative comparison of the two papers.

TABLE 2.5: Prediction of Post-crash Returns - Jan. 2009 - Dec. 2009

Abnormal returns during the post-crash holding period from Jan. 1, 2009 to Dec. 31, 2009 (standard errors in parentheses). The portfolios are formed on the basis of insider trading activity calculated by each stock's Net Insider Purchases (NIP), Net Insider Purchase Ratio ($NIPR$), and Net Insider Share Purchase Ratio ($NISPR$) during the portfolio formation period from Oct. 13, 2008 to Dec. 31, 2008. Abnormal returns are reported for both the CAPM and FF 3-factor specifications. "Market Adjusted" returns are the portfolio raw returns minus the value-weighted NYSE/Nasdaq/AMEX return and are included for comparison and robustness. Panel A presents the results by sorting stocks into quartiles according to the Net Insider Purchases (NIP) measure of insider trading. Panel B presents the results using the Net Insider Purchase Ratio, and panel C presents the results for the Net Insider Share Purchase Ratio sorted according to the criteria in the first column.

Portfolio	Market Adjusted	CAPM alpha	FF 3-factor alpha
Panel A - NIP			
Quartile 1	0.0759***	0.0488**	0.0235
(Most sold)	(2.51)	(1.97)	(0.81)
Quartile 2	0.2073***	0.1368***	0.1105***
	(7.24)	(5.67)	(4.65)
Quartile 3	0.4502***	0.3455**	0.3096**
	(2.89)	(2.36)	(2.08)
Quartile 4	0.4885***	0.3235***	0.2723***
(Most bought)	(9.07)	(7.74)	(6.86)
Panel B - NIPR			
Only sells	0.0909***	0.0566***	0.0310
	(3.89)	(2.84)	(1.53)
$NIPR \in (-1, 0]$	0.1158***	0.0723***	0.0486**
	(4.81)	(3.60)	(2.24)
$NIPR \in (0, 1)$	0.3022***	0.1997***	0.1698***
	(5.24)	(4.12)	(3.46)
Only buys	0.4213***	0.3014***	0.2600***
	(6.07)	(4.72)	(4.02)
Panel C - NISPR			
Only sells	0.0909***	0.0566***	0.0310
	(3.89)	(2.84)	(1.53)
$NISPR \in (-1, 0]$	0.1241***	0.0798***	0.0579***
	(5.08)	(3.92)	(2.62)
$NISPR \in (0, 1)$	0.2541***	0.1566***	0.1166***
	(4.74)	(3.47)	(2.61)
Only buys	0.4213***	0.3014***	0.2600***
	(6.07)	(4.72)	(4.02)

(1990) for the 1987 crash (10-19%). The difference between the results in this paper and Seyhun's may simply reflect the timing of the 2008 crash, which occurred as part of a much larger financial event, or it may reflect the longer portfolio formation period, or both. For comparison, Finnerty (1976) constructs portfolios of stocks weighted by how heavily those stocks are bought by insiders during normal trading periods, and reports only a 4% abnormal return in the first month for a portfolio of insider buys, 1% after the second month, and mere basis points thereafter, completely disappearing by the sixth month. Thus it seems that the advantage to following the heaviest of inside buyers produces much larger returns after a crisis than is typical during normal trading periods.

Interestingly, the CAPM abnormal returns to the high-sell portfolios are positive. However, this may not be surprising since insider sales do not tend to forecast future negative returns very well (Lakonishok and Lee 2001, Chowdhury et al. 1993). They are also not robust to the 3-factor adjustment which shows no significant returns above zero for the high-sell porfolios.

The returns are all monotonically increasing in the insider trading measure, except for the third and fourth quartile portfolios for the NIP measure, where the 3rd quartile portfolio has a higher CAPM and 3-factor return than the 4th quartile portfolio. However, the third quartile returns are well within the 95% confidence interval for the fourth quartile returns for both the CAPM (.2415, .4056), and the 3-factor (.1944, .3502) models, so we cannot say whether they are statistically different from each other.

The time series of monthly returns to the portfolios based on the NIP measure is presented in Table 2.6. I include the abnormal CAPM and 3-factor return for each month of 2009 for the high-buy and high-sell portfolios. I also report the difference between them in the third column with t-statistics in parentheses.

Table 2.6: Post-crash abnormal returns to insider stock portfolios

Monthly returns for quartile portfolios based on Net Insider Purchases (NIP) for the portfolio formation period Oct. 13, 2008 to Jan. 2, 2009. Returns are presented for both the CAPM and FF 3-factor specifications. Q1 alpha is the abnormal return to the portfolio of the top 25% of stocks most sold by insiders, while Q4 alpha is the abnormal return to the portfolio of stocks most bought by insiders. Difference is the spread between the Q4 and Q1 abnormal returns.
***- Returns to Q4 are larger than returns to Q1 at the 1% significance level, **- at the 5% significance level, *- at the 10% significance level.

	CAPM alphas			FF 3-Factor alphas		
Month	Q1 alpha (Most sold)	Q4 alpha (Most bought)	Difference	Q1 alpha (Most sold)	Q4 alpha (Most bought)	Difference
Jan 2009	0.0037	0.0994	0.0957***	0.0054	0.1195	0.1141***
			(6.20)			(7.13)
Feb 09	-0.011	-0.0013	0.0097	-0.0190	0.0086	0.0272**
			(0.87)			(2.33)
Mar 09	0.0135	0.0202	0.0067	0.0189	0.0194	0.0004
			(0.39)			(0.02)
Apr 09	0.0524	0.1451	0.0927***	0.0258	0.1098	0.0840***
			(5.52)			(4.92)
May 09	0.0032	0.0522	0.0491***	0.0280	0.0726	0.0446***
			(3.41)			(3.08)
Jun 09	-0.0255	0.0194	-0.0061	0.0084	0.0097	0.0013
			(0.62)			(-0.44)
Jul 09	0.0086	-0.0044	-0.0130	-0.0009	-0.0224	-0.0215**
			(-1.22)			(-1.98)
Aug 09	0.0160	0.0705	0.0545***	0.0208	0.0593	0.0385**
			(2.96)			(2.00)
Sep 09	0.0194	0.0302	0.0108	0.0069	0.0165	0.0096
			(1.07)			(0.93)
Oct 09	-0.0354	-0.0332	0.0021	-0.0077	-0.0005	0.0072
			(0.27)			(0.83)
Nov 09	-0.0289	-0.0615	-0.032***	-0.0035	-0.0400	-0.0364***
			(-3.70)			(-3.911)
Dec 09	0.0353	0.0339	-0.0019	-0.0013	-0.0001	0.0013
			(0.18)			(-0.12)

The spread between the top and bottom portfolios, is listed under the column labeled *Difference*. This captures the extent to which insiders were able to distinguish between stocks which were going to rebound well, and those which were not.The difference between the returns to the high-buy and high-sell portfolios behaves as expected, with the largest returns occurring closest to the crash. However, the high-buy portfolio continues to significantly outperform the high-sell portfolio well into the second and third quarters of 2009, six months after the crash.

The spread between the top and bottom portfolios, is listed under the column labeled *Difference*. This captures the extent to which insiders were able to distinguish between stocks which were going to rebound well, and those which were not. The

difference between the returns to the high-buy and high-sell portfolios behaves as expected, with the largest returns occurring closest to the crash. However, the high-buy portfolio continues to significantly outperform the high-sell portfolio well into the second and third quarters of 2009, six months after the crash. This delayed response in the returns is curious when compared with most of the literature, which finds that insider trades lose their predictive power very quickly.[12] The monthly returns for all four quartiles are represented graphically in Figures 2.2 and 2.3, while the Q4-Q1 spread appears in Figures 2.4 and 2.5.

As we can again see, the high-buy portfolio exhibits large gains over the next 12-months, much larger than would be implied by even the most optimistic studies of insider trading during normal periods (Finnerty 1976), and vastly more than conservative estimates (Seyhun 1986, Chowdhury et al. 1993). The results do, again, by and large, agree with Seyhun's (1990) study of the 1987 crash, which found between 10.1% and 18.5% abnormal returns for a buy portfolio versus a sell portfolio. When compared with the results from Seyhun (1990), it seems that the informational content of inside trades after the 2008 stock market crash was similar to, or perhaps even greater than for the 1987 stock market crash.

[12] For instance, Finnerty (1976) reports almost no abnormal returns after the fourth month, Chowdhury et al. (1993) reports most returns disappear after the first couple of weeks, and Seyhun (1986) shows that most abnormal returns occur within the first 100 days after the insider trade. Seyhun (1990), which studies the 1987 crash, has a very limited treatment of the length of the abnormal returns. He only presents his results for a 14-month holding period from November 1987 to December 1988, and a 12-month holding period which consists of January 1988 to December 1988 alone. The returns that exclude the last two months in 1987 are significantly positive for his buy portfolios, but less so than the full 14 month sample.

Figure 2.2: Monthly CAPM abnormal returns to insider intensity portfolios

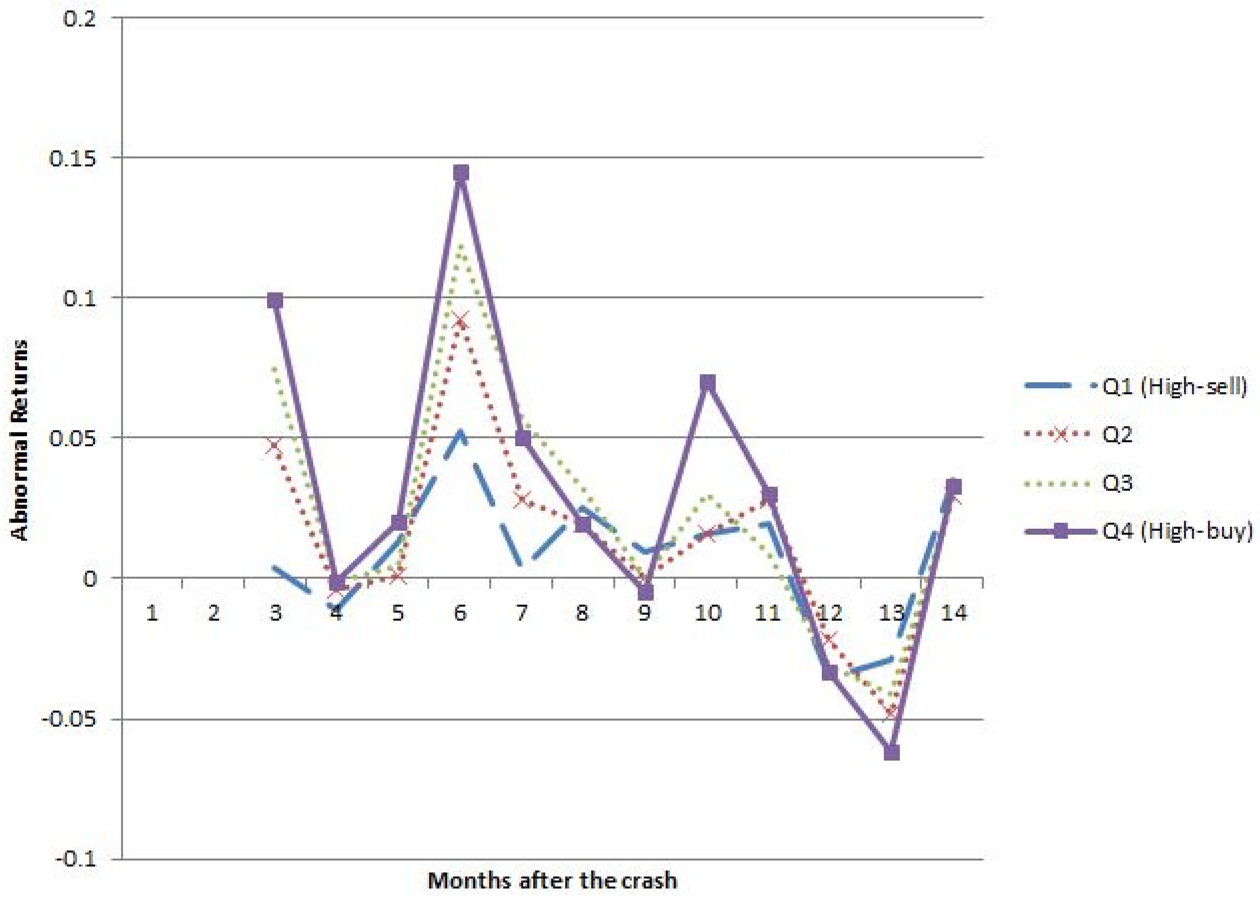

Figure 2.3: Monthly FF abnormal returns to insider intensity portfolios

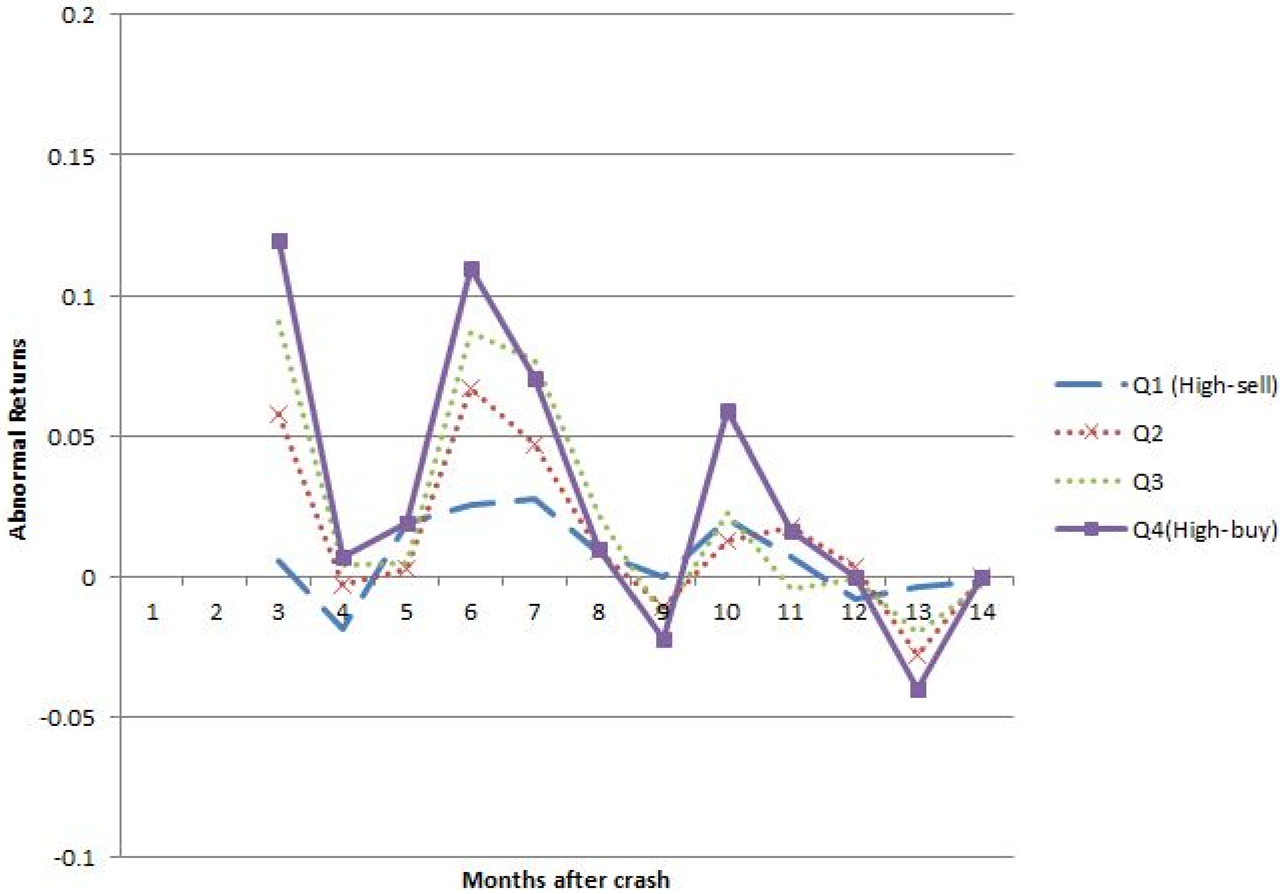

FIGURE 2.4: The difference in CAPM returns between holding the high-buy (Q4) portfolio over the high-sell portfolio (Q1) by month

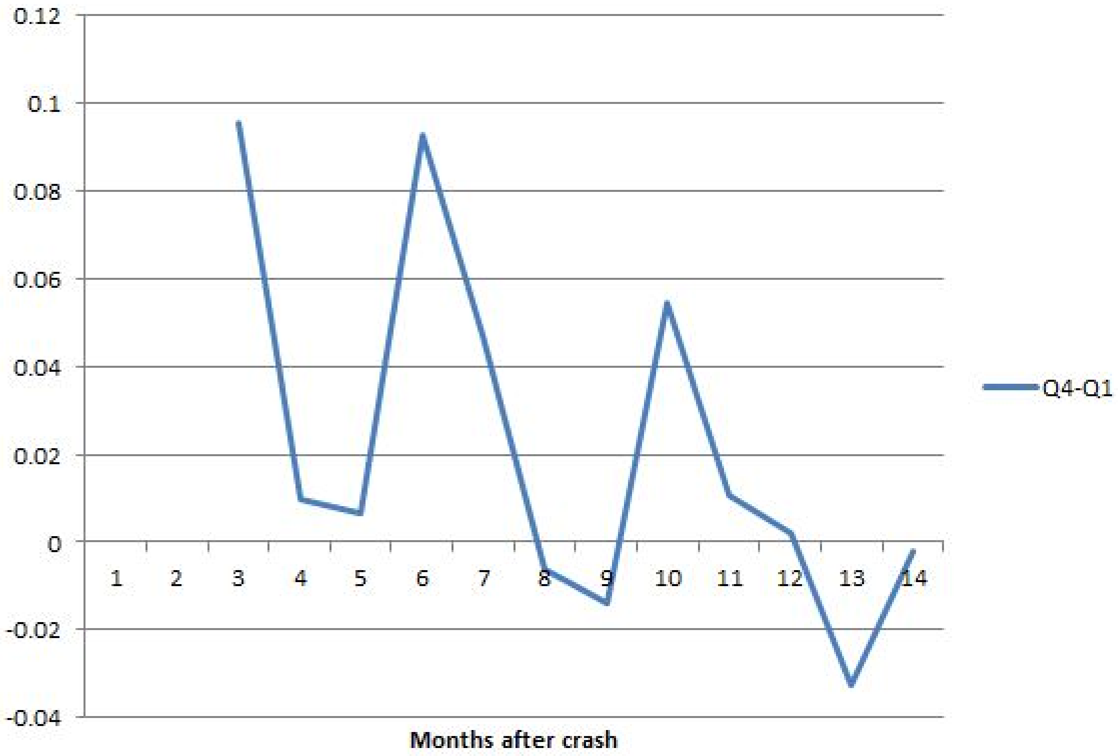

FIGURE 2.5: The difference in FF 3-factor returns between holding the high-buy (Q4) portfolio over the high-sell portfolio (Q1) by monthh

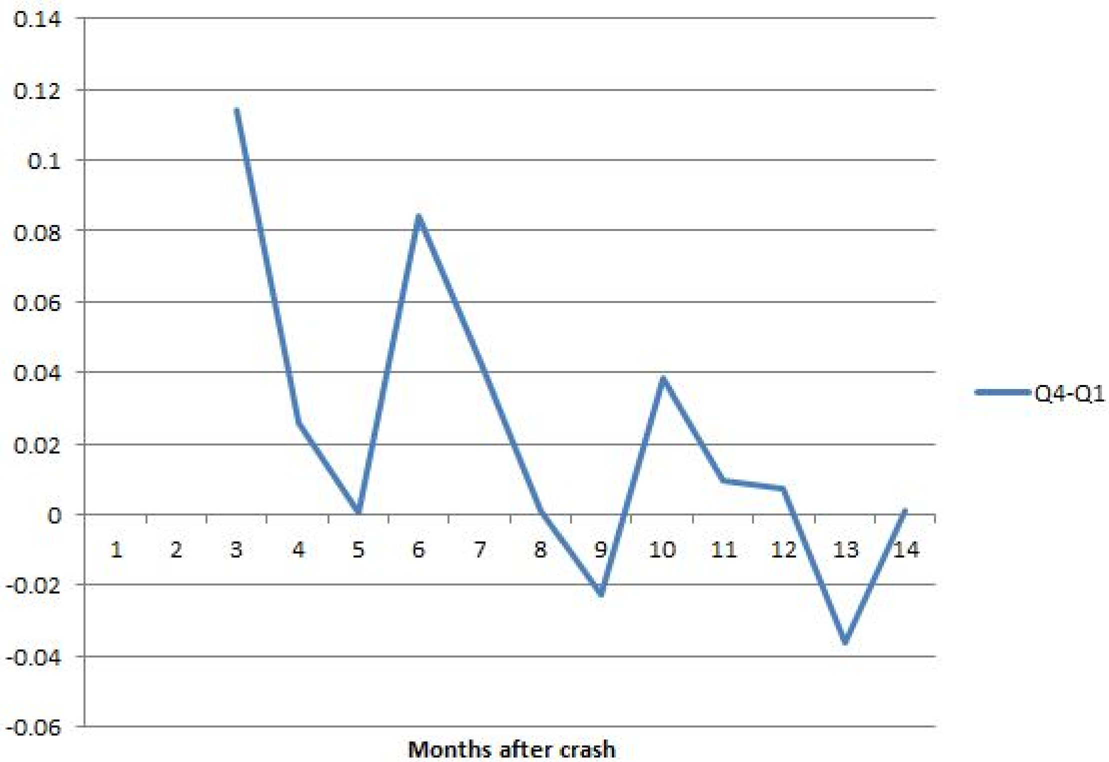

2.4.3 Which Insiders Are Trading Their Stocks After the Crash?

One curious observation from Table 2.5 is that the "market adjusted" raw returns are quite a bit higher than either the CAPM or FF adjusted returns. Since overall stock prices were rising very rapidly during this period, these high "market adjusted" raw returns could be explained if insiders are trading predominantly in high-beta stocks. In this case, the CAPM and FF models will more fully adjust for the unusually high market returns, compared to the "market adjusted" returns, which *de facto* assume that the beta for each stock in the sample is one. In fact, the average (CAPM) market beta for the sample of post-crash insider traded stocks is 1.29. The average beta for the high-sell (Q1) portfolio is 1.09, and the average beta for the

high-buy (Q4) portfolio is 1.50. Since the average beta for both portfolios is larger than one, the CAPM-adjusted returns will be smaller than the "market adjusted" raw returns in both portfolios. Since the average beta is higher for Q4 than for Q1, the *difference* between the returns to Q4 and Q1 is also larger. This implies that insiders of high-beta firms are more likely to buy than insiders of low-beta firms.

To further explore the behavior of inside traders, Table 2.7 presents the results for the following regression:

$$Insiderbuying_i = \delta_0 + \delta_1 Size_i + \delta_2 (B/M)_i + \delta_3 (\beta_M)_i + \varepsilon_i \tag{2.4}$$

The first line of Table 2.7 presents the results with *Net Insider Purchase* (NIP) for each stock on the left hand side. The next two use the *Net Insider Purchase Ratio*

TABLE 2.7: Insider behavior

Regression results for equation (2.4). The dependent variable is the insider trading metric. The independent variables are market capitalization ($Size$), the book-to-market (B/M) ratio, and each stock's market beta ($(\beta_M)_i$) as estimated from (2.2). The first line shows the results with *Net Insider Purchases* (NIP) as the dependent variable, the second for *Net Insider Purchase Ratio* ($NIPR$), the third for *Net Insider Share Purchase Ratio* ($NISPR$). Standard t-statistics are included in parentheses. ***-significant at the 1% level. **- 5% significance level.

$$Insiderbuying_i = \delta_0 + \delta_1 Size_i + \delta_2 (B/M)_i + \delta_3 (\beta_M)_i + \varepsilon_i$$

Dependent variable	$Size$	B/M	β_M
NIP	-0.043***	1.521***	1.911***
	(-2.69)	(4.45)	(4.29)
$NIPR$	-0.048**	0.258***	0.241***
	(-2.40)	(6.00)	(6.18)
$NISPR$	-0.005**	0.287***	0.225***
	(-2.51)	(6.52)	(5.55)

($NIPR$), and the *Net Insider Share Purchase Ratio* ($NISPR$), respectively.[13] As can be seen, the results are similar in sign and significance regardless of which measure of insider trading is used. The consistently negative and significant coefficient for $Size$ indicates that insiders are more likely to buy if their firm is a small cap firm. The consistently positive coefficients for B/M and β indicate that insiders are also more likely to buy if their firm has a high B/M value or a high beta. The first two of these results are largely consistent with those of Lakonishok and Lee (2001), who also found that insiders tended to buy small firms and high B/M firms during normal trading periods. The negative coefficient for $Size$ could indicate that insiders are aware that the market for small stocks is less efficient. It could also mean that insiders for large

[13]The regression was also performed with the total insider sales of each firm on the left hand side. The signs were consistent with those reported in Table 2.7. Not only were small stocks, high B/M stocks, and high-beta stocks more likely to be bought by insiders, they were also were less likely to be sold by insiders.

companies are more heavily policed by firm management, which makes it harder for them to capitalize on inside information. For example, many large companies, especially banks, prohibited their insiders from trading at all in September, 2008.

2.4.4 Cross-Sectional Regression Analysis of Returns to Insider Trading

If insiders truly are predicting returns, then we should observe a relationship between the future excess returns to a stock and how frequently that stock was bought or sold by insiders. I consider the following multivariate relationship, similar to the regressions performed by Lakonishok and Lee (2001) and Seyhun (1990):

$$RET_i = \gamma_0 + \gamma_{NIP} NIP_i + \gamma_S Size_i + \gamma_{BM}(B/M)_i + \gamma_\beta(\beta_M)_i + \varepsilon_i \qquad (2.5)$$

where RET_i is the holding-period raw return to stock i minus the value weighted NYSE/AMEX/ Nasdaq market return.[14] The regression also includes the net insider purchases for each stock for the last two months of 2008 (NIP_i), market capitalization in billions ($Size_i$), and the book-to-market value ($(B/M)_i$). The market betas ($(\beta_M)_i$) are also included in the regression. Recall these were estimated using pre-sample period data from equations (2.1) and (2.2). The results presented use the beta for each stock generated from the Fama-French 3-factor model in (2.2) as a right hand side variable to control for market risk. Using the CAPM betas estimated from (2.1) instead does not substantially change any results. Using either the CAPM or Fama-French corrected

[14] This is the same measure used in much of the insider trading literature, including Lakonishok and Lee (2001).

abnormal returns in place of the raw returns also does not appreciably change any of the primary results in terms of the coefficients on the insider trading variable.[15]

Lakonishok and Lee (2008) suggests that the relationship between insider trading and future returns, during normal trading periods, is sensitive to firm size. The most likely reason for this is that small stocks carry significantly higher transactions costs (Chan and Lakonishok, 1995). This suggests that the market for small cap stocks is less efficient than the one for large cap stocks. Thus we would expect that returns to smaller firms should be more sensitive to the insider trading variable, NIP. To capture this interaction between size and the degree of insider trading, I split the sample into deciles based on market capitalization. Small cap stocks are those in the lowest three deciles of $Size$, while medium caps are those in the four middle deciles, and large caps are those in the three highest deciles. I also reorganize the post-crash return sample into the four quarters of 2009 and perform the regression for each quarter separately.

The results for these regressions are in Table 2.8. The panel labeled "All" presents the regression results for all stocks in the sample. The coefficient on NIP is positive and significant through the third quarter of 2009, which suggests that insiders are successfully predicting returns through the third quarter of 2009 (six to eight months after the crash). These results indicate that insiders are predicting returns over much longer time horizons than those reported by either Finnerty (1976) (three months) or Chowdhury et al. (1993) (two months or less). Insiders do the best job of

[15]However, since these returns are already corrected for size, B/M and the market excess return, the control variables mostly become less significant as might be expected.

Table 2.8: Returns to Insider Purchases by Size and Quarter - NIP

Regression results for equation (2.5) on the effect of insider trading on abnormal returns. The left hand side variable is raw returns minus the value weighted NYSE/AMEX/Nasdaq market return. The measure of insider trading for each stock are the Net Insider Purchases (NIP_i). Heteroskedasticity-robust t-statistics are shown in parentheses below estimates. The right hand side variables are the insider trading variable NIP, market capitalization ($Size$), book-to-market (B/M), and each stock's market beta as estimated from (2). The regression is performed for each quarter in 2009. The first subset of results shows the effect for all stocks. The second, for stocks with small market capitalizations (less than $95m), the third for medium size market caps (between $95m and $880m), and the fourth for large market caps (above $880m). ***- significant at the 1% level **- 5% significance level. *- 10% significance level.

$$RET_i = \gamma_0 + \gamma_{NIP} NIP_i + \gamma_S Size_i + \gamma_{BM}(B/M)_i + \gamma_\beta(\beta_M)_i + \varepsilon_i$$

Size	2009 Quarter	NIP	$Size$	B/M	β_M
All	1	0.0028**	0.0004	0.0350**	0.0096
		(2.00)	(1.33)	(2.01)	(0.63)
	2	0.0064***	-0.0013***	-0.0090	0.1756***
		(4.27)	(-3.25)	(-0.19)	(9.49)
	3	0.0031***	-0.0003	0.1026***	0.0812***
		(3.10)	(-0.67)	(3.48)	(5.45)
	4	-0.0008	0.0004**	0.0027	0.0043
		(-0.80)	(2.00)	(0.19)	(0.11)
Small	1	0.0053*	-1.248	-0.0026	0.0210
		(1.82)	(-1.32)	(-0.10)	(0.57)
	2	0.0033	0.9423	-0.0518*	0.1408***
		(1.38)	(0.77)	(-1.72)	(3.99)
	3	0.0033	0.614	0.1312***	0.0948***
		(1.32)	(0.60)	(2.82)	(3.41)
	4	-0.0020	0.2146	0.0019	-0.0012
		(1.11)	(0.35)	(0.09)	(0.03)
Medium	1	0.0040*	0.0607	0.0265	0.0212
		(1.67)	(1.11)	(0.88)	(1.05)
	2	0.0097***	-0.1151	-0.0282	0.1697***
		(3.59)	(-1.59)	(-0.18)	(5.23)
	3	0.0050**	0.0012	0.0569	0.0676***
		(2.27)	(0.02)	(1.44)	(2.47)
	4	0.0027	0.0415	0.0113	0.0031
		(1.29)	(1.06)	(0.54)	(0.19)
Large	1	-0.0001	0.0001	0.0429	0.0165
		(-0.07)	(0.33)	(1.13)	(0.75)
	2	-0.0001	-0.0002	-0.0343	0.2150***
		(-0.10)	(-1.00)	(0.74)	(8.14)
	3	0.0014*	-0.0003**	0.0522	0.0962***
		(1.75)	(-1.83)	(1.33)	(5.70)
	4	-0.0007	0.0001	-0.0473**	0.0230*
		(-0.84)	(0.54)	(-2.03)	(1.70)

predicting returns for the second quarter of 2009 (January-March 2009), or 3-5 months after the crash.

The next three panels of Table 2.8 perform the same regressions for the different size groups. As expected, insider trades in small and medium cap stocks are doing the best job of predicting returns, with positive, statistically significant coefficients in the first two to three quarters. Insider trades in large cap stocks do appear to predict returns in the third quarter of 2009, but the magnitude of these returns are quite small compared to those for the small and medium cap stocks. Overall, it appears that the small and medium cap stocks are mostly, if not wholly responsible for the large spread in returns to insider trading portfolios. Table 2.8 is replicated with $NIPR$, and $NISPR$ in place of NIP as the insider trading variable. The results are included in Tables 2.9 and 2.10, respectively. Most of the important results for insider trading are unchanged in terms of sign and significance, but different measures do seem to influence the timing of the estimated returns. For instance, when $NIPR$ is used on the right hand side, the positive returns for the first quarter seem to weaken, though the positive returns for the second and third quarters of 2009 remain significant. Most of the statistically significant returns appear for the medium stocks. This is likely because returns to the small stocks are imprecisely estimated. The coefficients for the insider trading variable in the small and medium portfolios are similar to each other, but large compared to the large stocks. This seems to suggest that here, as in Lakonishok and Lee (2001), and Chan and Lakonishok (1995), the market for large firms' stock is more efficient than that of smaller firms. The other variables are included strictly for control. The coefficient on the β for each stock is reliably positive and significant, indicating

TABLE 2.9: Returns to Insider Purchases by Size and Quarter - NIPR

Regression results for equation (2.5) on the effect of insider trading on abnormal returns. The left hand side variable is raw returns minus the value weighted NYSE/AMEX/Nasdaq market return. The measure of insider trading for each stock is the Net Insider Purchase Ratio ($NIPR_i$). Heteroskedasticity-robust t-statistics are shown in parentheses below estimates. The right hand side variables are market capitalization ($Size$), book-to-market (B/M), and each stock's market beta as estimated from (2). The regression is performed for each quarter in 2009. The first subset of results shows the effect for all stocks. The second, for stocks with small market capitalizations (less than \$95m), the third for medium size market caps (between \$95m and \$880m), and the fourth for large market caps (above \$880m). ***- null hypothesis of $\gamma = 0$ rejected at the 1% significance level **- rejected at 5% significance level. *- rejected at 10% significance level.

$$RET_i = \gamma_0 + \gamma_{NIP} NIP_i + \gamma_S Size_i + \gamma_{BM}(B/M)_i + \gamma_\beta(\beta_M)_i + \varepsilon_i$$

Size	2009 Quarter	$NIPR$	$Size$	B/M	β_M
All	1	0.0142	0.0003	-0.0355**	0.0115
		(1.41)	(1.09)	(-2.00)	(0.74)
	2	0.0620***	-0.0013***	-0.0153	0.1729***
		(5.31)	(-3.03)	(-0.33)	(9.53)
	3	0.0198**	-0.0003	0.1022***	0.0824***
		(2.51)	(-1.58)	(3.43)	(5.58)
	4	0.0086	0.0005**	-0.0061	0.0008
		(1.35)	(2.40)	(-0.47)	(0.07)
Small	1	0.0224	-1.201	0.0032	0.0257
		(0.55)	(-1.25)	(0.12)	(0.69)
	2	0.0241	0.9935	-0.0521*	0.1435***
		(0.70)	(0.80)	(-1.70)	(4.09)
	3	0.0283	0.6755	0.1307***	0.0975***
		(1.07)	(0.66)	(2.82)	(3.41)
	4	0.0211	0.2622	0.0171	-0.0206
		(1.08)	(0.43)	(0.84)	(-0.84)
Medium	1	0.0066	0.0618	0.0259	0.0214
		(0.51)	(1.14)	(0.85)	(1.07)
	2	0.0643***	-0.1073	-0.0341	0.1734***
		(3.74)	(-1.52)	(-0.21)	(5.46)
	3	0.0275**	0.0007	0.0548	0.0709***
		(2.02)	(0.01)	(1.35)	(2.83)
	4	0.0192*	0.0448	0.0094	0.0037
		(1.85)	(1.16)	(0.45)	(0.26)
Large	1	-0.0046	0.0002	0.0430	0.0184
		(-0.39)	(0.62)	(1.15)	(0.82)
	2	0.0215*	-0.0001	-0.0349	0.2057***
		(1.89)	(-0.75)	(-0.76)	(7.86)
	3	0.0075	-0.0003*	0.0506	0.0968***
		(0.89)	(-1.88)	(1.31)	(5.71)
	4	-0.0083	0.0001	-0.0468**	0.0174
		(-1.20)	(0.61)	(-2.25)	(1.25)

Table 2.10: Returns to Insider Purchases by Size and Quarter - NISPR

Regression results for equation (2.5) on the effect of insider trading on abnormal returns. The left hand side variable is raw returns minus the value weighted NYSE/AMEX/Nasdaq market return. The measure of insider trading for each stock is the Net Insider Share Purchase Ratio ($NISPR_i$). Heteroskedasticity-robust t-statistics are shown in parentheses below estimates. The right hand side variables are market capitalization ($Size$), book-to-market (B/M), and each stock's market beta as estimated from (2). The regression is performed for each quarter in 2009. The first subset of results shows the effect for all stocks. The second, for stocks with small market capitalizations (less than $95m), the third for medium size market caps (between $95m and $880m), and the fourth for large market caps (above $880m). ***- null hypothesis of $\gamma = 0$ rejected at the 1% significance level **- rejected at 5% significance level. *- rejected at 10% significance level.

$$RET_i = \gamma_0 + \gamma_{NIP} NIP_i + \gamma_S Size_i + \gamma_{BM}(B/M)_i + \gamma_\beta(\beta_M)_i + \varepsilon_i$$

Size	2009 Quarter	$NISPR$	$Size$	B/M	β_M
All	1	0.0202**	0.0003	0.0334**	0.0104
		(2.08)	(1.19)	(1.88)	(0.67)
	2	0.0500***	-0.0013***	-0.0136	0.1766***
		(4.70)	(-3.01)	(-0.29)	(9.54)
	3	0.0138*	-0.0003*	0.1034***	0.0841***
		(1.72)	(-1.68)	(3.46)	(5.67)
	4	0.0088	0.0005**	-0.0064	0.0009
		(1.40)	(2.43)	(-0.49)	(0.08)
Small	1	0.0290	-1.210	0.0020	0.0255
		(0.79)	(-1.27)	(0.07)	(0.69)
	2	-0.0037	0.9535	-0.0501	0.1441***
		(-0.09)	(0.75)	(-1.64)	(4.08)
	3	0.0301	0.6550	0.1298***	0.0974***
		(1.26)	(0.64)	(2.80)	(3.41)
	4	0.0179	0.2405	0.0168	-0.0206
		(0.96)	(0.39)	(0.83)	(-0.84)
Medium	1	0.0149	0.0676	0.0239	0.0196
		(1.17)	(1.25)	(0.79)	(0.99)
	2	0.0514***	-0.1201	-0.0337	0.1782***
		(3.57)	(-1.63)	(-0.21)	(5.40)
	3	0.0158	-0.0091	0.0563	0.0745***
		(1.08)	(-0.16)	(1.38)	(2.95)
	4	0.0186*	0.0434	0.0089	0.0044
		(1.82)	(1.13)	(0.42)	(0.29)
Large	1	0.0027	0.0002	0.0427	0.0152
		(0.23)	(0.63)	(1.13)	(0.68)
	2	0.0195*	-0.0002	-0.0358	0.2072***
		(1.78)	(-0.74)	(-0.77)	(7.99)
	3	0.0031	-0.0003*	0.0506	0.0989***
		(0.37)	(-1.89)	(1.31)	(5.77)
	4	0.0098	0.0001	-0.0473**	0.0171
		(1.53)	(0.62)	(-2.27)	(1.25)

Table 2.11: Returns to Insider Purchases by B/M and Quarter

Regression results for equation (2.5) on the effect of insider trading on abnormal returns. The left hand side variable is raw returns minus the value weighted NYSE/AMEX/Nasdaq market return. In the first column, the measure of insider trading for each stock are Net Insider Purchases (NIP_i). In the second column, the measure of insider trading for each stock is the Net Insider Purchase Ratio ($NIPR_i$). In the last column, the measure of insider trading for each stock is the Net Insider Share Purchase Ratio ($NISPR_i$). Heteroskedasticity-robust t-statistics are shown in parentheses below estimates. The other independent variables are, market capitalization ($Size$), book-to-market (B/M), and each stock's market beta as estimated from (2). The regression is performed for each quarter in 2009. The first subset of results shows the effect for all stocks. The second, for stocks with low book-to-market values (less than 0.11), the third for middle book-to-market values (between 0.11 and 0.48), and the fourth for high book-to-market (above 0.48). ***- null hypothesis of $\gamma = 0$ rejected at the 1% significance level **- rejected at 5% significance level. *- rejected at 10% significance level.

$$RET_i = \gamma_0 + \gamma InsiderTrading_i + \gamma_S Size_i + \gamma_{BM}(B/M)_i + \gamma_\beta(\beta_M)_i + \varepsilon_i$$

B/M Group	2009 Quarter	NIP	$NIPR$	$NISPR$
All	1	0.0028**	0.0142	0.0202**
		(2.04)	(1.41)	(2.08)
	2	0.0064***	0.0620***	0.0500***
		(4.22)	(5.31)	(4.72)
	3	0.0031***	0.0197**	0.0138
		(3.04)	(2.51)	(1.72)
	4	-0.0008	0.0086	0.0088
		(-0.78)	(1.35)	(1.40)
Low B/M	1	0.0029	0.0530***	0.0544***
		(1.52)	(2.94)	(3.12)
	2	0.0038	0.0385**	0.0338*
		(1.44)	(2.12)	(1.91)
	3	0.0010	0.0100	0.0115
		(0.76)	(0.66)	(0.07)
	4	-0.0001	0.0073	0.0112
		(-0.44)	(0.58)	(0.92)
Middle B/M	1	0.0028**	0.0085	0.0028*
		(1.88)	(0.62)	(1.88)
	2	0.0046**	0.0324**	0.0047**
		(2.50)	(2.55)	(2.50)
	3	0.0026**	0.0146	0.0264**
		(2.39)	(1.48)	(2.39)
	4	0.0007	0.0094	0.0064
		(0.46)	(0.97)	(0.68)
High B/M	1	0.0037	-0.0100	-0.0078
		(1.18)	(-0.40)	(-0.33)
	2	0.0041*	0.0518**	0.0317
		(1.73)	(2.24)	(1.23)
	3	0.0044*	0.0231	0.0236
		(1.69)	(1.27)	(1.36)
	4	-0.0023	0.0149	0.0152
		(-1.28)	(1.18)	(1.31)

that the higher beta stocks in each sub-sample are predictably earning higher returns. The coefficients for $Size$ and B/M are mostly insignificant.

I perform the same regressions for B/M groups in Table 2.11, with the controls suppressed for brevity. Even though insiders tended to buy high B/M stocks during the portfolio formation period (see section 2.4.3) no real pattern of predictive ability emerges among the B/M groups. For NIP as the insider trading variable, the medium and high B/M groups do a better job of predicting returns than the low B/M group, but this is not robust to the other measures of insider trading.

2.5 Conclusion

The predictive power of insider trading has been well established in the literature. However, I find no evidence that insiders sold heavily before the crash of 2008, nor can I find any support for the hypothesis that insiders predicted the crash by selling their stocks prior to the crash if they did especially poorly during the crash. There was, however, a large increase in insider buys during and immediately after the crash, and the future returns to the stocks that were most bought during the period after the crash did very well relative to the market. Portfolios comprised of stocks whose insiders bought heavily after the crash earned an additional 26-32% abnormal returns compared to the CAPM and Fama-French models over a 12-month holding period from January to December, 2009. This was far above the abnormal returns reported by Finnerty (1976), Seyhun (1986), Rozeff and Zaman (1988), and most of the literature on insider trading during normal periods. These returns lasted a very long time, persisting at least until the third quarter of 2009, 9-11 months after the crash, much longer than during normal trading periods as shown by Finnerty (1976), Seyhun (1986), and Chowdhury, et al. (1993). As expected, most of the returns were generated by small and medium cap stocks, consistent with the conjecture that the market for large cap stocks is more informationally efficient.

These results compare favorably with Seyhun (1990) and Seyhun (1999), which study the 1987 market crash. More research is likely needed in this area to understand the behavior of informed traders during market panics versus during normal trading periods, and whether this empirical pattern extends to other asset markets and other stock market crashes.

The presence of large returns seems to violate the Efficient Markets Hypothesis, but this need not necessarily be true. If investors have not had sufficient exposure to market crashes, they will not know where to look for the strongest signals. If insider information is carrying stronger signals during a crash than during normal trading periods, it will take time for investors to learn this. The general pattern of "learning by repetition" is almost universally observed in the repeated asset market experiments of Smith, Suchanek and Williams (1988), where human subjects routinely produce bubbles and crashes in simulated asset markets. This "bubble and crash" behavior persists for the first couple of trials for a group of subjects before the price converges to fundamentals by the third or fourth repetition. It is likely that investors learn from past experiences in asset markets in the real world as well. Since market crashes do not occur very often in the lifetime of a typical investor, it is likely that they do not know to look for such high returns following post-crash insider trades. It is also true that comprehensive data reporting on insider trading has only been available for the past few decades, and even then, only for asset markets in the most highly developed economies. If this is true, then this research, combined with Seyhun (1990, 1999) may prompt traders to look for large returns to portfolios based on insider purchases after market crashes. It is also likely that we should see lower returns to insider trading in markets that crash more often. Therefore, while the same pattern is observed after the 1987 and 2008 market crashes, we might not expect to find these large returns in future crashes since investors may learn to arbitrage these gains away.

CHAPTER 3

SHOULD CENTRAL BANKS INTERVENE IN ASSET MARKETS?

3.1 Introduction

Should central banks react to asset bubbles? This issue has been hotly debated within the economics profession, particularly since the late 2000's housing bubble and subsequent global financial crisis. One side of the argument (see Cecchetti et al (2000), and Roubini (2006)) contends that central bank policy can and should respond to fluctuations in asset prices since wild swings in asset markets may be accompanied by large fluctuations in the output gap and employment. The counterargument suggests that it is not even possible to target asset prices because deviations from fundamentals are unobservable, and even if it was possible, the outcome might be destabilizing (see Bernanke and Gertler (1999, 2001), Posen (2006)).

Comparatively little attention, however, has been paid to the microeconomic aspects of bubble bursting policies by central banks. It may not be a stretch to suggest that arguments in favor of bubble bursting presuppose the idea that bubble bursting is welfare improving for the individual agents in asset markets. If the central bank can identify when assets are overpriced, it can protect naive buyers from malicious sellers who may have more information than they do. The existence of these "bad sellers" also creates a lemons problem. Buyers cannot know beforehand if they are buying from a “good seller” who thinks the asset might be worth something, and a bad seller who knows the asset is worthless. Thus, the good sellers are also harmed because the buyers will bid down the price if they believe that there are unscrupulous sellers trying to take advantage of naive buyers (Conlon (2015)). We can eliminate the problem of asymmetric information by having the central bank announce when it thinks we have entered a state of the world where at least some sellers know the asset is overpriced. This is a tempting argument to make, and it does have some merit.

One simple response is that the central bank cannot know with certainty when assets are overpriced, and therefore, by having a policy of deterring bad sellers, might cause perfectly good, wealth creating transactions to go unconsummated and therefore hurt buyers and good sellers (see Kai and Conlon (2008) for a detailed exploration of this argument).

However, there exists another, more subtle issue. Even if the central bank can accurately idenfity when assets are overpriced, the argument that the central bank should deflate the price of those assets does not address the following unforeseen consequence of such a policy. Specifically, eliminating the lemons problem will cause buyers to raise their prices in the states of the world where they are now encouraged because bad sellers do not exist. The increase in prices could cause the buyers to be *worse* off than before if the increase in price is large enough to overcome the good that was done for them by eliminating those transactions from which the buyers bought from bad sellers.

The purpose of this paper is to address exactly this point. We do not ask whether the central bank can identify states of the world where overpriced assets exist.

Instead we imagine a world where the central bank knows with certainty whether anyone knows that assets are overpriced, and ask, “Can the elimination of overpricing in asset markets make the agents in the world better off?” As we will explain below, this is a question that has gone only partially answered because of the lack of attention that the microeconomic issues of bubble policy have recieved. This paper attempts to fill the lacuna in the literature by building on work done by Allen, Morris, and Postlewaite (1993) (hereafter AMP), Conlon (2004), and most importantly, Conlon (2015).

AMP presents a simple asset market model with three heterogeneous individuals and asymmetric information. In certain of the states of the world, one or more of the sellers receive a private signal that the asset will pay no dividend. As a result, the agent (a “bad seller”) will try to sell the worthless asset to an unsuspecting buyer. However, the buyer does not know whether or not anyone has received the negative signal, so from the buyer’s perspective, sellers might potentially be “good sellers” willing to sell to him. Thus, the buyers are willing to pay a positive price for the asset, even if there’s a chance that the seller knows they are selling a worthless asset. In the states of the world where this occurs, we say that the asset is “overpriced”. In these states there is at least one agent who knows that the true value of the asset is zero, but the market price is positive.

AMP go on to show that even if all three agents receive the same bad signal (everyone realizes the asset is worthless), the price can still be positive. This is because even though each agent knows the true value of the asset is zero, none of the agents realize that the other agents also know this. Since, from the point of view of each seller, there is a possibility that the other agents did not recieve the signal, he hopes to sell the asset in the market for a positive price. AMP call this state a *strong bubble* because each of the agents realizes (in equilibrium!) that the asset is overpriced.

Conlon (2004) simplifies AMP by showing that their central results can be expressed in a model with just two agents, making it vastly easier to construct numerical examples and, as this paper will show, conduct comparative dynamics experiments on the welfare of the agents in the model.

Finally, Conlon (2015) incorporates a broad policy of general deflation of overpriced assets into a simple rational greater fool bubble model like the one presented in AMP In this paper, the central bank announces any time it observes a state of the world where at least one agent has recieved the private signal that the asset is worthless. In other words, if only one agent knows the true value, the central bank also knows, and announces this fact without fail.

Conlon (2015) performs comparative dynamics experiments to determine how bubble policy affects the agents' welfare in the model. However, his model, like AMP, assumes constant marginal utilities, so buyers have perfectly elastic demands for the asset, and therefore have zero consumer surplus in both the bubble equilibrium and the post-policy equilibrium. Thus, buyers are not affected by the policy, so his welfare analysis is restricted to the effects on sellers only. Conlon shows that bubble policy has an ambiguous effect on sellers' welfare. Sellers lose out in the states where they were able to sell an overpriced asset to naive buyers (the greater fools), but gain in the remaining states because inaction by the central bank provides an implicit endorsement of the asset's price (guaranteeing that there are no bad sellers) and encourages buyers to offer higher prices.

This paper builds on the discussion begun by Conlon (2015), by adding risk averse agents to the model. Downward sloping marginal utility schedules produce downward sloping demand curves, so buyers receive positive consumer surplus. As a result, the welfare effects now extend to buyers as well as sellers. This can inform the discussion on bubble bursting or general deflation policies because it can show us whether the policy accomplishes its goal of protecting naive buyers from malicious sellers. As we will soon see, this is not a trivial extension. Concave utility functions require us to keep track of the effect of wealth on the marginal utilities in each state, and many utility functions do not yield analytic solutions.

The central result in this paper is that buyers in the market are made worse off by the policy because the cost of paying higher prices in the absence of the lemons problem outweighs the benefit of not buying assets known to be worthless. We show this by first conducting numerical comparative dynamics experiments with several classes of utility functions. It is always the case that the higher prices generated by the policy make buyers worse off regardless of which utility function is assumed. This suggests that bubble bursting policy can only make things worse for the buyers. In addition to this, the sellers also seem to be hurt by bubble policy relative to a risk neutral baseline case.

Some general theorems are proven to support these results. Theorem 1 shows that as long as utility is decreasing in absolute risk aversion (DARA), the general deflation policy will raise non-announcement prices more when the agents become more risk averse. This result clears up some conceptual difficulties about the type of utility we want the agents to have to display the kind of behavior we expect. It is also essential for proving Theorem 2 which shows that, if absolute risk aversion is decreasing, then buyers will always be hurt by a policy that deflates overpriced assets. The effect on sellers is ambiguous.

3.1.1 Bubbles

Several high-profile episodes suggest that large destructive asset bubbles, while uncommon, occur with some regularity. Historical examples include the Dutch Tulipmania (1634-1637), the Mississippi Company boom/bust (1720), the South Sea Company boom/bust (1720), the United States stock market boom and crash (1920-1929), the Japanese real estate crash (1985-89), the Swedish housing crash (1990-1992), the NASDAQ tech stock boom/bust (1999), and the 2000's global housing boom/bust (2000-2008). In all of these cases, asset prices rose dramatically over a relatively short period, and then collapsed quickly, though the exact timing varies for each instance. Many of these examples have coincided with destructive events beyond the asset market itself. The Japanese asset market crash is frequently cited as one cause of Japan's Lost Decade, the US stock market crash in 1929 is often credited for beginning the Great Depression, and the housing bubble of the late 2000's is often blamed for the Global Financial Crisis of 2008.[1]

It is, in general, difficult to disentangle which of these historical examples represents a bubble, and which are merely boom/busts in demand or abrupt revisions in expectations by investors. For instance, suppose that investors feel optimistic that a firm will be highly profitable in the future. As a result, the price of that firm's stock will rise with investor optimism. At some later date, suppose expectations of future profitability

[1]See Kindleberger and Aliber (2011) for more examples, and see Garber (2001) for an alternative interpretation of many historical asset bubbles.

are revised downward.[2] This will cause the price to crash. But can this simple thought experiment be accurately described a bubble? Or was the market pricing the stocks appropriately given the available information, which was later updated to induce a large negative shock to expectations? Regardless of which argument the reader is inclined to believe, our current understanding of asset bubbles is best described as incomplete.

In some sense, the question of how an asset bubble might occur is, itself, curious. There is no easy way to see how a reasonably intelligent person might hold an over-priced asset when placed in an environment with other reasonably intelligent people. If we know that there is no one else foolish enough to hold an over-priced asset, how could we ever expect to sell to someone at an overly high price? We therefore arrive at the result that the more people understand that a tradeable asset is over-priced, the less likely the asset is to be over-priced.

Modelers therefore have an interesting challenge to overcome. How can a model produce over-priced assets in a population of intelligent traders?

Even if we suggest that bubbles do indeed exist, the question of how to model bubble behavior has not been fully addressed. One solution has been to abandon the idea that traders are fully rational (Harrison and Kreps (1978), Scheinkman and Xiong (2003), Abreu and Brunnermeier (2003)). While some of these models, especially Abreu and Brunnermeier (2003), are quite sophisticated, it is conceptually very easy to generate bubbles if certain traders are systematically irrational. Bubbles can be easily constructed by having smart traders selling bad assets to overconfident or noise traders at high prices.

[2] A possible example of this may be the late 90's "tech stock bubble".

If all traders are rational, however, then this presents a more difficult problem, especially if each trader knows all of the other traders in the market are similarly rational. Rational traders learn from price signals. In Milgrom and Stokey (1983), if an agent is selling a risky asset at a high price, the rational buyer will ask, “Why is this investor selling the asset?” The buyer deduces that the seller is sending a negative signal about the value of the asset. As a result the buyer lowers his willingness to pay until the seller is no longer willing to sell. Thus, no assets are ever traded, and any possible overpricing disappears. Hence, with rational actors, the more people believe the asset is overpriced, the less likely it is that the asset will continue to be overpriced.

As stated already, AMP present a rational heterogeneous-agent bubble model with three agents and three periods. The agents are prompted to trade with each other by having different marginal utilities in relevant states of the world, so there are at least some states in the world where sellers wish to sell to rebalance their portfolios, for example, and not simply selling on private information. AMP then impose a specific asymmetric information structure, in which there exist states where one or more agents receive private information that the asset will pay a dividend of zero. However, since potential buyers believe that they may be in a state of the world where the other agents are selling for portfolio rebalancing purposes, these potential buyers are willing to buy, so a bubble is possible.

AMP’s most important contribution, is that they generate a state in which each agent values the asset at a price of zero, but the asset has a market price greater than zero. This occurs because even though each agent realizes the asset is worthless, they *don’t know* that the other agents know the asset is worthless. Therefore, they all hope to sell to each other at a price higher than zero. Bubble models like this unfortunately only produce very short-lived bubbles. Once the agents try to sell to each other, they realize they both received the same information, and the price crashes to zero in the subsequent period. They also require short-sale constraints, or else the bad sellers would sell so much of the asset that it would push the price down to zero.

3.2 General model

The model very closely follows the ones in AMP and Conlon (2015). There exists a competitive market where rational, hetergeneous agents trade money for an asset that pays a stochastic dividend. We construct the market in a way that produces a *strong bubble* (AMP), where each of the agents knows the asset is worthless (will pay a dividend of zero), yet the market price is above zero. The model achieves this by generating a state where each seller knows the asset is worthless, but does *not* know that the other agent(s) also know that the asset is worthless.

First, we will generate an example similar to the one presented Conlon (2015). We generate a strong bubble for a specific state of the world where each agent knows with certainty that the asset will pay a dividend of zero, but the price is positive.

We then depart from AMP and Conlon (2015) by supposing that the marginal utilities used to generate the bubble are not given exogenously, but instead are generated from risk averse agents.[3] We then work backwards, from the marginal

[3] In both of these earlier papers, the agents' marginal utilities are constant for each state of the world. This is useful for simplifying the analysis, but renders the model incapable of delivering confident qualitative predictions about the welfare effects on the agents. AMP briefly discusses downward sloping marginal utilities but performs no welfare analysis.

utilities supporting the initial bubble equilibrium, to calculate the final wealth belonging to each agent necessary to reproduce these marginal utilities for a given concave utility function, as well as the corresponding initial wealths. We can then calculate the utility generated by these final wealth holdings, given the bubble.

In the final step, we assume the existence of a central bank which can perfectly observe states of the world where at least one agent observes that the asset is worthless. The central bank can protect buyers in the market by announcing the state of the world whenever it observes these conditions. We will call this a policy of "general deflation of overpriced assets" as in Conlon (2015).[4] We then recalculate the equilibrium prices and trades in the asset market after the policy announcement, calculate the final wealth for each of the agents in each state, calculate their resulting expected utility, and then compare to the original bubble equilibrium. By doing this, we can determine the qualitative (and quantitative) welfare effects of the general deflation policy.

The welfare implications of either a bubble bursting policy or a policy of general deflation, even when the central bank can perfectly identify asset bubbles, are not obvious. Suppose that the central bank actively protects the market against sellers who know the asset is overpriced, by announcing when it observes states where at least some of the agents know the asset is overpriced. When the central bank then turns out to *not* make such an announcement, the central bank *implicitly endorses* the price of the asset. The implicit endorsement of these asset prices can raise prices in the market, which may

[4]The general deflation policy should not be confused with a "bubble bursting" policy, in which the central bank only announces when it observes a bubble. This may lead to different conclusions about the agents' welfare, and will be discussed in a later paper.

reduce the utility of buyers. Other implications (notably, those of a central bank that imperfectly observes bubbles) are of potential importance, but are not discussed here.

3.2.1 Assumptions and setup

There are two agents in the model, who can be thought of as being heterogeneous traders in an asset market. These agents maximize their utility across all states of the world. We will call the two agents Ellen (E) and Frank (F), denoted by $i \in \{E, F\}$. The market lasts for three periods $t = 1, 2, 3$. After period 3, the asset pays a stochastic liquifying dividend, the market closes, and agents consume their wealth. There are a finite number of states of the world ω, where ω is an element of a state space Ω. The agents have a common, prior probability distribution over Ω, given by $\pi(\omega)$ which is known to everyone.

Agents are endowed with an initial amount of both a risky asset, $a_{i,0}(\omega)$, and a riskless asset (money) that can be used to acquire the risky asset or can be consumed. An agent's holdings of the risky asset at time t, are given by $a_{i,t}(\omega)$. The risky asset cannot be used in consumption, but can be sold for money, or held in hopes of recieving the dividend d at time $t = 3$.

There are no short-sales of the risky asset, so $a_{i,t}(\omega) \geq 0$. As in AMP and Conlon (2004, 2015) we have to rule out short-sales of the asset because in the presence of an overpriced asset, sellers will short sell until the asset is no longer overpriced.

3.2.1.1 The information structure

There are nine possible states of the world, e_1^G, e_2^G, e_3^G, e^B, f_1^G, f_2^G, f_3^G, f^B, b. Each state is equally likely to occur, so $\pi(\omega) = 1/9$ for all ω. The risky asset pays a dividend, d, only in states e_3^G and f_3^G.

In states e_1^G, e_2^G, and e_3^G, Ellen is a "good seller". That is, she wishes to sell but thinks the asset might be worth something. In states b and e^B, Ellen is a "bad seller". In these states, Ellen knows the asset will pay no dividend, so she clearly wants to sell to Frank. Finally, Ellen is a buyer in states f_1^G, f_2^G, f_3^G, and f^B. Symmetrically, there are states of the world in which Frank is a bad seller (f^B and b), where he is a good seller (f_1^G, f_2^G, and f_3^G), and where he is a buyer (e_1^G, e_2^G, e_3^G, e^B). Note that in state b, both agents are "bad sellers" who know the asset is worthless.

Lastly, we assume that agents are endowed with a single unit of the asset when they are sellers. The buyers are endowed with none of the asset.

Accordingly, we can write Ellen's period 1 information partition as:

$$E_1^{BS} = \{b, e^B\}, E_1^{GS} = \{e_1^G, e_2^G, e_3^G\}, \text{ and } E_1^{Buyer} = \{f^B, f_1^G, f_2^G, f_3^G\}$$

Frank's period 1 information partition is likewise:

$$F_1^{BS} = \{b, f^B\}, F_1^{GS} = \{f_1^G, f_2^G, f_3^G\}, \text{ and } F_1^{Buyer} = \{e^B, e_1^G, e_2^G, e_3^G\}$$

The information partitions for period 1 are represented graphically in Figure 1.

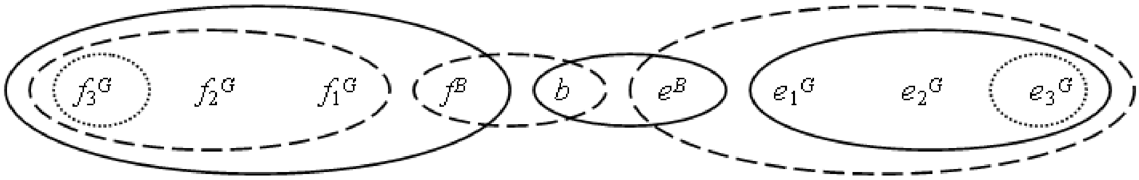

FIGURE 3.1: Period 1 information partitions. Solid ovals - Ellen, Dashed ovals - Frank, Dotted ovals - dividend paying states

In period 2, the agents learn whether the state is b, e_1^G, or f_1^G. Thus the period 2 information partition for Ellen is

$$E_2^1 = \{b\}, E_2^2 = \{e_1^G\}, E_2^3 = \{f_1^G\},$$
$$E_2^{BS} = \{e^B\}, E_2^{GS} = \{e_2^G, e_3^G\}, \text{ and } E_2^{Buyer} = \{f^B, f_2^G, f_3^G\}$$

and likewise for Frank. The nontrivial part of the new information structure is displayed in Figure 2 with states b, e_1^G and f_1^G omitted.

In period 3, all information is revealed. The agents learn the true state of the world, the dividend is paid, and agents consume their wealth.

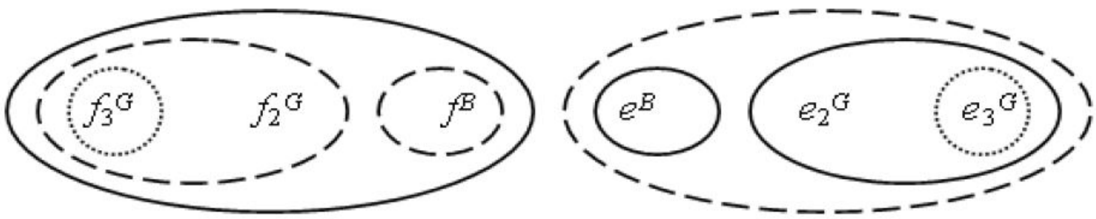

FIGURE 3.2: Period 2 information partitions. Solid ovals - Ellen, Dashed ovals - Frank, Dotted ovals - dividend paying states

3.2.2 The bubble equilibrium

We begin by generating a bubble equilibrium where the final marginal utilities for each state of the world are as in Table 1, where x, y, and z are all greater than zero. These marginal utilities will generate certainty equivalent prices for the risky asset in each state of the world. They are also arranged in such a way that changing x, y, and z do not affect the asset's price. They simply control the weight of importance that the agents give to different states of the world.[5] In states $\{b, e^B, e_1^G, e_2^G, e_3^G\}$ Ellen is endowed with one unit of the risky asset at the beginning of the world. In states $\{b, f^B, f_1^G, f_2^G, f_3^G\}$, Frank is endowed with one unit of the asset.

TABLE 3.1: Assumed marginal utilities in the bubble equilibrium

State	b	e^B	f^B	e_1^G	e_2^G	e_3^G	f_1^G	f_2^G	f_3^G
Ellen	x	x	z	$2y$	y	y	$5z$	z	$2z$
Frank	x	z	x	$5z$	z	$2z$	$2y$	y	y

Working backwards, at $t = 3$, all information has been revealed and the agents know the true state of the world, thus $p_3(\omega) = d(\omega)$. For each preceding time period, the certainty equivalent willingness to pay is defined as

[5] At this point, these marginal utilities are simply assumed. We will work backwards from this example to derive the implied wealths in each state necessary to produce this particular bubble equilibrium.

$$WTP_i(\omega) = \frac{E\left[MU_i(\omega)p_{t+1}(\omega)\right]}{E\left[MU_i(\omega\right]} \quad \text{for} \quad \omega \in P_{i,t} \tag{3.1}$$

where $P_{i,t}$ is any given information set for agent i at time t. The market is competitive with short sale constraints, so the price, $p_t(\omega)$, is just the greater of the two willingnesses to pay. As an illustration, we suppose the asset will pay a dividend of $d = 4$ in states e_3^G and f_3^G, and a dividend of $d = 0$ in all other states (as in Figure 1). For F_2^{Buyer}, in period 2 Frank believes the state could be e^B, e_2^G, or e_3^G. Applying (3.1) to obtain Frank's period 2 willingness to pay, using the marginal utilities from Table 1, we have, for $\omega \in F_2^{Buyer}$

$$WTP_F(\omega) = \frac{E\left[MU_F(\omega)p_3(\omega)|\omega \in F_2^{Buyer}\right]}{E\left[MU_F(\omega)|\omega \in F_2^{Buyer}\right]} = \frac{(z)(0)(1/9) + (z)(0)(1/9) + (2z)(4)(1/9)}{(z)(1/9) + (z)(1/9) + (2z)(1/9)} = 2$$

Ellen's problem is similar. Within the set F^{Buyer}, Ellen has two possible information sets: $E_2^{BS} = \{e^B\}$, and $E_2^{GS} = \{e_2^G, e_3^G\}$. For $\omega \in E_2^{BS}$,

$$WTP_E(\omega) = \frac{(x)(0)(1/9)}{(1)} = 0$$

and, for $\omega \in E^{GS}$,

$$WTP_E(\omega) = \frac{(y)(0)(1/9) + (y)(4)(1/9)}{(y) + (y)} = 2$$

Therefore, in state e^B, Ellen sells to Frank at a price of 2, and in states e_2^G, and e_3^G, Ellen is indifferent between selling, and not selling at the price of 2. Similar calculations can be made for $\omega \in E_2^{Buyer}$, and also for the period 1 information sets for both agents.

Applying similar calculations to the remaining states in both periods 1 and 2, we arrive at an equilibrium set of prices and net sales for each state of the world. The prices are shown below in Table 2, and the pattern of trade is shown in Table 3. In period 2, Ellen sells to Frank in information sets $E_2^{BS} = \{e^B\}$ and $E_2^{GS} = \{e_2^G, e_3^G\}$, while Frank sells to Ellen in information sets $F_2^{BS} = \{f^B\}$ and $F_2^{GS} = \{f_2^G, f_3^G\}$.[6]

TABLE 3.2: Bubble Equilibrium Prices

State	b	e^B	f^B	e_1^G	e_2^G	e_3^G	f_1^G	f_2^G	f_3^G
Period 1	1	1	1	1	1	1	1	1	1
Period 2	0	2	2	0	2	2	0	2	2
Period 3	0	0	0	0	0	4	0	0	4

TABLE 3.3: Sales from Ellen to Frank - Bubble Equilibrium

State	b	e^B	f^B	e_1^G	e_2^G	e_3^G	f_1^G	f_2^G	f_3^G
Period 1	0	0	0	0	0	0	0	0	0
Period 2	0	1	-1	0	1	1	0	-1	-1
Period 3	0	0	0	0	0	0	0	0	0

This exercise generates three states of the world where the asset is overpriced. In states b, e^B, and f^B, at least one agent knows the asset pays a final dividend of zero, yet the price is above zero in period 1. In state b, *both* agents realize the asset is worthless, but the price is positive. This is the familiar *strong bubble* result from AMP. In state b, Ellen knows the asset is worthless, but does not know that Frank also knows the asset is worthless. From her point of view, she may be in e^B, in which case Frank will be in

[6] It is not an accident that we have chosen the parameters such that trade occurs only in period 2. This greatly simplifies the analysis.

F_1^{Buyer}. This means that she believes that Frank might believe that the true state could be e_3^G. If this is the case, Frank will be willing to pay a positive price for the asset. Frank, meanwhile, is thinking the exact same thing about Ellen in state b, so in period 1, both agents are hoping to sell a worthless asset to the other in the next period. In period 2, both agents learn that since they're trying to sell to each other, the true state must be b, and the bubble crashes to a price of 0 in period 2.

3.2.3 Bubble policy and the baseline scenario with constant marginal utility

Now that we have a well defined equilibrium that generates a bubble, we can examine what happens when a market authority, e.g., a central bank, uses policy to prevent overpricing. We assume the existence of a central bank that can identify when there are overpriced assets in the market (i.e., if the true state of the world is e^B, f^B, or b). Since the bank knows the true state of the world when there are bad sellers in the market, the central bank can pursue a policy of "general deflation of overpriced assets" by announcing to the public when it observes these states.

We assume that the central bank observes bad sellers with complete accuracy. This has two important functions. First, it greatly simplifies the analysis. Second, it allows us to study what the effects of the central bank's policy could be *in the most perfect of all worlds.* In other words, if the central bank cannot improve welfare even in the case where it knows everything, it surely cannot improve welfare in the vastly more realistic case where the central bank has limited information.[7]

[7]See Kai and Conlon (2008) for a model, with constant marginal utilities, where the central bank observes overpricing imperfectly.

Following Conlon (2015), we examine what happens to the market when the central bank follows the general deflation policy as described above. In period 1, the central bank announces whether it knows if the world is in states b, e^B, or f^B. The policy forces the price to zero in these states because it is now common knowledge among the agents that the asset is worthless. Each agent knows the asset is worthless, each agent knows that the other agent knows the asset is worthless, and each agent knows that the other agent knows that they know the asset is worthless, etc.[8] Ellen's new information partition in period 1 is

$$E_1^{1P} = \{b\}, E_1^{2P} = \{e^B\}, E_1^{3P} = \{f^B\},$$

$$E_1^{BSP} = \varnothing, E_1^{GSP} = \{e_1^G, e_2^G, e_3^G\}, \text{ and } E_1^{BuyerP} = \{f_1^G, f_2^G, f_3^G\}$$

and similarly for Frank, where the superscript P indicates the information sets under the general deflation policy. As before, assume that in period 2 the agents learn whether the state is e_1^G or f_1^G. Ellen's period 2 information partition is then

[8]See Allen, Postlewaite, and Shin (1995), Conlon (2004), and Allen, Morris, and Shin (2006) for more discussion of these interesting types of repetitive information structures involving depth of knowledge problems.

$$E_2^{1P} = \{b\}, E_2^{2P} = \{e^B\}, E_2^{3P} = \{f^B\}, E_2^{4P} = \{e_1^G\}, E_2^{5P} = \{f_1^G\}$$

$$E_2^{BSP} = \varnothing, E_1^{GSP} = \{e_2^G, e_3^G\}, \text{ and } E_1^{BuyerP} = \{f_2^G, f_3^G\}$$

and similarly for Frank.

Re-applying the calculations in Equation (3.1) to the new information structure, holding the marginal utilities from Table 3.1 constant, we arrive at a new set of prices and trades shown in Tables 3.4 and 3.5.

TABLE 3.4: Risk Neutral Prices Under Bubble Bursting Policy

State	b	e^B	f^B	e_1^G	e_2^G	e_3^G	f_1^G	f_2^G	f_3^G
Period 1	0	0	0	1.33	1.33	1.33	1.33	1.33	1.33
Period 2	0	0	0	0	2.67	2.67	0	2.67	2.67
Period 3	0	0	0	0	0	4	0	0	4

TABLE 3.5: Sales from Ellen to Frank - Bubble Equilibrium

State	b	e^B	f^B	e_1^G	e_2^G	e_3^G	f_1^G	f_2^G	f_3^G
Period 1	0	0	0	0	0	0	0	0	0
Period 2	0	0	0	0	1	1	0	-1	-1
Period 3	0	0	0	0	0	0	0	0	0

The most important thing to notice in the general deflation equilibrium is that prices have risen in the remaining states where the agents trade. Why did this happen? Under the general deflation policy, the central bank announces when it observes overpriced assets, but, as in Conlon (2015), the more important implication is that when the central bank *does not* announce, it is providing an implicit endorsement of asset prices to the market. By shutting down trade in the bad seller states, buyers no longer have to worry about buying an asset that the seller definitely knows is overpriced. As a result, buyers are encouraged to bid up prices because they are buying an asset that is more likely to pay a dividend.

This can also be interpreted as the elimination of a lemons problem. Because buyers in the bubble equilibrium have no guarantee that they are buying from someone who thinks the asset may pay a dividend, they reduce their willingness to pay in all states of the world. So, the bad sellers also inflict an externality on good sellers by reducing the price good sellers get for the asset. Elimination of the overpricing in the bad seller states results in the good sellers receiving higher prices.

3.2.3.1 Seller Welfare

In the risk neutral case we have presented here, buyers have no consumer surplus. This is because the marginal utilities assumed in Table 3.1 are constant. Therefore, the demand curve for the asset is perfectly elastic, and welfare comparisons for the buyer are trivial since consumer surplus is zero in both equilibria. We can, however, make some general statements about the sellers' welfare in the market. Because the policy eliminates overpricing in the market, we may be tempted *a priori* to conclude that this policy damages sellers in terms of their expected welfare. This is not necessarily the case however. Recall that the policy caused the price of the asset to rise in the remaining states of the world, so sellers are better off in those states. Thus, the overall effect on the sellers' welfare is ambiguous and depends on their marginal utilities in the relevant states of the world. If the sellers are are poorer in the states where they are bad sellers (x is large), the policy could make them worse off. If sellers are poorer in the states where they are good sellers (y is large), the policy will make them better off.

To begin with, we know that a seller in the bubble equilibrium could be in any one of five possible states of the world. Ellen, for example, could be in b, e^B, e^G_1, e^G_2, or e^G_3 (Frank's problem is similar). Her expected welfare would be her expected proceeds from her period 2 sale of assets in states e^B, e^G_2, and e^G_3, weighted by her marginal utility in each state

$$EU_{Seller} = \frac{1}{9}(x + y + y)\Pi(p_2)$$

where $\Pi(p_2) = a_{i,2}p_2$ is the profit from selling $a_{i,2}$ at price p_2. In the context of the above example, since Ellen has only one unit of the asset to sell, and the price in the bubble equilibrium derived above is $p_2 = 2$, we can simplify Ellen's expected welfare to

$$EU_{Seller} = \frac{1}{9}(x + 2y)(2)$$

When the central bank pursues the general deflation policy, it pushes the price to zero in the bad seller states. Ellen can no longer sell in these states, but the price rises in the remaining states. Her expected welfare under the policy is therefore

$$EU^P_{Seller} = \frac{1}{9}(2y)(2.6667)$$

Ellen will be hurt by the policy if her expected welfare in the bubble equilibrium is higher than her expected welfare under the general deflation policy, or if

$$\frac{1}{9}(x+2y)(2) > \frac{1}{9}(2y)(2.6667), \text{ or } y > \frac{3}{2}x$$

Recall that from Table 3.1, Ellen's marginal utilities contain y in the good seller states, and x in the bad seller states. The larger y is, the higher the marginal utility in the good seller states (e.g., Ellen may be poorer in these states), and thus Ellen ascribes more importance to these states. If y is larger than x by a factor of 3/2, then Ellen will be helped by the policy because she is poorer in the states of the world where the price is rising. The converse is true when x is relatively big compared to y. It easily follows from this line of analysis that the point where the policy has no effect on the welfare of sellers (in the risk neutral case) occurs when

$$\text{Baseline: } y = \frac{3}{2}x$$

3.3 An Example with Quadratic Utility

The above example can generate some nontrivial welfare effects for the sellers in the market. However, the buyers are unaffected by this policy. This is because buyers have constant marginal utilities, so their demand curves are perfectly elastic, and they enjoy no consumer surplus. While this greatly simplifies the analysis, it is has the distinct drawback of restricting our welfare discussions of the policy to the sellers. A commonly stated goal of a bubble bursting policy described above is to protect the buyers of assets (the greater fools) from malicious sellers who know the assets are overpriced. If this type of model is to provide a complete picture of how bubble bursting policy affects the agents trading in asset markets, it is important that these models be able to seriously address the welfare of buyers.

In the following section, I provide a model that remedies these criticisms by introducing agents with risk averse utility functions. Risk aversion implies that the agents' marginal utility functions are downward sloping functions of wealth. This will then yield nonzero consumer surplus, and allow us to perform a nontrivial welfare analysis for the buyers before and after the policy takes effect.

It is important to realize that while the goal of the policy is to protect naive buyers from sellers who might know more about the asset then the buyers, there is no guarantee that the policy will be able to achieve this goal. Recall from the above experiment in the risk neutral case, that prices rise in some states of the world as a result of the general deflation policy. If the buyers are poor enough in these states then they will have high marginal utilities of wealth in these states, so paying higher prices for the asset could harm them, *even if* we prevent them from buying in states of the world where the sellers know the asset to be worthless.

What follows is an example of the above model with a simple quadratic utility function. The marginal utility curves in this case are downward sloping but linear. Thus while marginal utility falls as a function of wealth, the rate by which it falls does not depend on wealth (i.e., $U''(w)$ is constant in wealth). We consider this case because quadratic utility produces an analytic solution for the price, making it relatively straightforward to compute a numerical example. We extend the analysis to more realistic utility functions in Section 4.

We begin this section by supposing that agents with the marginal utilities in Table 3.1 are risk averse. Then we ask, for what initial wealths in the different states will, if combined with the trades, prices and dividends, yield the marginal utilities in Table 1? To answer this, we must first give the agents a specific utility function. We first assume the agents' preferences are represented by a quadratic utility function that depends only on the agent's final wealth, W:

$$U_i(W) = AW - \frac{\lambda}{2}W^2 \tag{3.2}$$

This function produces downward sloping marginal utility curves given by $MU_i(W_i) = \frac{\partial U_i}{\partial W_i} = A_i - \lambda W_i$. The parameter λ controls the slope of the marginal utility curve. As $\lambda \to 0$, the marginal utility curve flattens out, and the consumer becomes more risk neutral. To complete the welfare analysis of the agents, we must determine what wealth they would have needed to have to generate the bubble equilibrium described in Section 2.

3.3.1 Wealth

Consider the final marginal utility (after the agents have traded) $MU(\omega) = A - \lambda W_{i,t=3}$ derived frim equation (3.2) above, where $W_{i,t=3} = W_{i,3}$ is the agent i's final wealth. We can rearrange this as

$$W_{i,3}(\omega) = (A - MU(\omega))/\lambda. \tag{3.3}$$

If we take $MU(\omega)$ from Table 3.1, then this equation gives the final wealth necessary in each state for *risk averse* agents to wind up with these final marginal utilities. The initial money holdings of each agent can then be backed out by reversing the trades and dividend payouts in each state of the world. Recalling that exchange takes place only in period 2, final wealth is given by

$$W_{i,3}(\omega) = M_{i,0}(\omega) + p_2(\omega)X_{i,2}(\omega) + (a_{i,0}(\omega) - X_{i,2}(\omega))\, d(\omega) \tag{3.4}$$

where $p_2(\omega)$ is the price of the asset in period 2, $M_{i,0}(\omega)$ is the initial money holdings of agent i, $X_{i,2}(\omega)$ is the sales of the asset in period 2, and $d(\omega)$ is the dividend the asset pays. Equation (3.4) states that the final wealth for agent i is given by initial money holdings, $M_{i,0}(\omega)$, plus the value of any sales of the asset at time 2, $p_2(\omega)X_{i,2}(\omega)$, plus the amount of the asset held in period 3 times the dividend it pays, where $a_{i,0}(\omega)$ is the initial asset holdings, and $X_{i,2}(\omega)$ is the total sales of the asset at time 2.

If we denote p_2^B and $X_{i,2}^B$ as the period 2 prices and net sales from the bubble equilibrium in Tables 2 and 3 above, equation (3.4) becomes $W_{i,3}(\omega) = M_{i,0}(\omega) + p_2^B(\omega)X_{i,2}^B(\omega) + (a_{i,0}(\omega) - X_{i,2}^B(\omega)d(\omega))$. Therefore, to yield final, equilibrium marginal utility $MU(\omega)$ in state ω, money holdings in that state would need to be given by

$$M_{i,0}(\omega) = (A - MU(\omega))/\lambda - p_2^B(\omega)X_{i,2}^B(\omega) - (a_{i,0}(\omega) - X_{i,2}^B)d(\omega) \tag{3.5}$$

These calculations can be illustrated with a simple numerical example. First, let $y = z = 1$ and $x = 2/3$. This ensures that in the risk neutral case, the deflation policy will not change the seller's expected utility. Also let $\lambda = 0.5$, $d = 4$, and $A = 15$. Plugging into equation (3.3), Frank's final wealth in state e_3^G, for example, is given by

$$W_{F,3}(e_3^G) = \frac{(A - MU(e_2^G))}{\lambda} = (15 - 2)/0.5 = 26$$

These calculations are carried out for both agents in all states and the results are presented in Table 3.6.

TABLE 3.6: Implied final wealths for risk averse agents - bubble equilibrium

State	b	e^B	f^B	e_1^G	e_2^G	e_3^G	f_1^G	f_2^G	f_3^G
Ellen	28.67	28.67	28	26	28	28	20	28	26
Frank	28.67	28	28.67	20	28	26	26	28	28

We can now use (3.5) to calculate the agents' initial wealths by undoing the trades, prices, and dividend payments. Again for example, Frank in state e_3^G bought one unit of the asset from Ellen at a price of 2, so $X_{F,2}^B(e_2^G) = -1$ and $p_2(e_3^G) = 2$, while the dividend paid is 4. Therefore his implied initial wealth *before* the trade and dividend payment was

$$W_{F,0}(e_3^G) = W_{F,3}(e_3^G) - p(e_3^G)(-1) - d(e_3^G) = 26 + 2 - 4 = 24$$

Carrying out these calculations for both agents in all states of the world give us Table 3.7.

Table 3.7: Implied initial wealths for risk averse agents - bubble equilibrium

State	b	e^B	f^B	e_1^G	e_2^G	e_3^G	f_1^G	f_2^G	f_3^G
Ellen	28.67	26.67	30	26	26	26	20	30	24
Frank	28.67	30	26.67	20	30	24	26	26	26

Table 3.7 tells us what wealths each agent would need to have entered the market with, had they had utility given by equation (3.2), in order for the equilibrium in Tables 2 and 3 to give them the marginal utilities given in Table 3.1. Close examination of Table 3.7 reveals something interesting about the agents. (For the moment, ignore states e_1^G and f_1^G; the agents learn after period 1 that the dividend will be zero, and no trade occurs rendering these states uninteresting for now.) The dividend paying state is the one where the buyer's wealth is at its lowest (e_3^G for Frank, and f_3^G for Ellen). Hence, a risk averse buyer should fear this state. However the asset pays a dividend in that very same state, so it winds up being a good hedge against the buyer's low wealth state. Therefore, we can think of the asset as a form of insurance.

For the moment, examine Frank's situation. In period 2 he knows that three states of the world could be possible: e^B, e_2^G, and e_3^G, in which his initial wealth will be 30, 30, and 24, respectively. As a risk averse agent, Frank fears the risk of landing in state e_3^G because his wealth will be at its lowest. Therefore, he will want to spread this wealth out more evenly, sacrificing some of his wealth in states e^B and e_2^G to gain more in e_3^G. Comparing Table 3.7 with Table 3.6, we can see the asset does a good job of smoothing out his wealth. Frank's final wealth in states e^B, e_2^G, and e_3^G is now much more even: 28, 28, and 26 respectively. Thus, the asset is a good hedge against the buyer's risk. Much of the subsequent analysis of welfare, especially that in Section 5, relies on this intuition.

3.3.2 Asset demand under risk aversion

The model in Conlon (2015) is sufficient to produce a strong bubble, but unfortunately it cannot say much about the welfare implications for buyers in a market where a central bank deflates overpriced assets. Since the model in Conlon (2015) assumes that each agent has a constant marginal utility in each state of the world, demand curves are perfectly elastic, so there is no consumer surplus, and the discussion of the welfare effects of bubble policy is necessarily limited to analysis of the sellers.

To provide a definite qualitative prediction about the welfare effects of general deflation policy in the previous example, we recall the utility function presented in equation (3.2). Each agent's objective is to maximize the expected value of (3.2), given his or her information. We have the following maximization problem

$$\max_{X_{i,t}} E\left[U_i\right] = \sum_{\omega \in P_{i,2}} \pi_(\omega) \left[AW_{i,3}(\omega) - \frac{\lambda}{2} W_{i,3}(\omega)^2\right] \tag{3.6}$$

where $W_{i,3}(\omega)$ is the final wealth of each agent, and again is given by expression (3.4).

We need to maximize (3.6) cell by cell in the information partition. Since agents cannot tell which state exists inside a given cell, $P_{i,t}$ in their information partition, $X_{i,t}(\omega)$ is constant over all $\omega \in P_{i,t}$, so we'll call it $X_{i,t}$. Carrying out the maximization problem in (3.6) for a given information set $P_{i,2}$ gives the relevant first order condition

$$\sum_{\omega \in P_{i,2}} \pi(\omega) \left[A_i \frac{\partial W_{i,3}(\omega)}{\partial X_{i,2}} - \lambda W_{i,3}(\omega) \frac{\partial W_{i,3}(\omega)}{\partial X_{i,2}} \right] = 0 \tag{3.7}$$

where $X_{i,2}$ is constant over $P_{i,2}$.

Using (3.4), we can evaluate the derivative inside the brackets of (3.7). Doing this gives us the agent's sensitivity of final wealth to his/her sales: $\frac{\partial W_{i,3}(\omega)}{\partial X_{i,2}} = p_2(\omega) - d(\omega)$. Notice that this result is intuitive. If an agent sells one more unit of the asset, s/he recieves the market price for the asset, but loses the dividend paid in state ω.

By way of illustration, consider a first-order condition with just two states of the world, so $\omega \in \{1, 2\}$, though it is trivial to extend the analysis to more states. Again, agents cannot distinguish between $\omega = 1$ and $\omega = 2$ if they are in the same information set, so $p_2(1) = p_2(2) = p_2$. Then, we have

$$\pi(1)\Big[A(p_2 - d(1)) - \lambda W_{i,3}(1)(p_2 - d(1))\Big] + \pi(2)\Big[A(p_2 - d(2)) - \lambda W_{i,3}(2)(p_2 - d(2))\Big] = 0 \tag{3.8}$$

Substituting in equation (3.4) gives

$$\pi(1)\Big[A[p_2 - d(1)] - \lambda(M_{i,0}(1) + p_2 X_{i,2} + (a_{i,0}(1) - X_{i,2})d(1))[p_2 - d(1)]\Big]$$
$$+ \pi(2)\Big[A[p_2 - d(2)] - \lambda(M_{i,0}(2) + p_2 X_{i,2} + (a_{i,0}(2) - X_{i,2})d(2))[p_2 - d(2)]\Big] = 0$$

Let $M_{i,0}(\omega) + a_{i,0}(\omega)d(\omega) = N(\omega)$. This is the total wealth of each agent if they do not trade, which henceforth will be referred to as the agents' "no-trade wealth". After some simplification, the above equation becomes

$$\pi(1)\Big[[p_2-d(1)](A-\lambda N(1))-\lambda[p_2-d(1)]^2X_{i,2}\Big]+\pi(2)\Big[[p_2-d(2)](A-\lambda N(2))-\lambda[p_2-d(2)]^2X_{i,2}\Big]=0$$

Moving the terms with $X_{i,2}$ to the right side gives

$$\pi(1)\Big(p_2-d(1)\Big)\Big(A-\lambda N(1)\Big)+\pi(2)\Big(p_2-d(2)\Big)\Big(A-\lambda N(2)\Big)=\pi(1)\lambda\Big(p_2-d(1)\Big)^2X_{i,2}+\pi(2)\lambda\Big(p_2-d(2)\Big)^2X$$

Finally, factoring out $X_{i,2}$ and solving produces a "net demand curve" given by

$$X_{i,2}(p_t)=\frac{\pi(1)\Big[(A_i-\lambda N(1))(p_2-d(1))\Big]+\pi(2)\Big[(A_i-\lambda N(2))(p_2-d(2))\Big]}{\pi(1)\Big(p_2-d(1)\Big)^2+\pi(2)\Big(p_2-d(2)\Big)^2} \tag{3.9}$$

where $X_{i,2}(p_t)$ represents demand for the asset as a function of the price prevailing in the market. Since the demand curve is nonlinear in price, we must be careful that we do not end up on a part of the demand curve that bends backwards. This problem appears because the utility function (3.2) is quadratic. This means that there is a bliss point. If the buyer, for example, faces a price that is sufficiently low, he or she may be able to buy so much of the asset that it may put him or her on the downward sloping portion of the utility function. Diminishing marginal utility will then cause the demand curve to slope upwards since a falling price will cause buyers to buy less in an attempt to stay near the bliss point (an obviously unrealistic situation for the agents). We consider more realistic utility functions in Section 4 below, for which this is not a problem.

3.3.3 The short-sale constraint, market clearing, and the price

Since Ellen and Frank are short sale constrained, they can only sell up to the amount of the asset that they're endowed with. This means that they can only sell up to their asset holdings in that period, so $X_{i,t} \leq a_{i,t}$. If the constraint holds with equality, solving (3.9) for p_t gives a polynomial of order 2 in p_t,

$$a_{i,t}\lambda p_t^2 - \left(2\lambda\Big[\pi_1 d_1 + \pi_2 d_2\Big] + \Big[\pi_1(A - \lambda N(1)) + \pi_2(A - \lambda N(2))\Big]\right)p_t$$
$$+ \lambda\Big[\pi_1 d_1^2 + \pi_2 d_2^2\Big] + \pi_1\Big(A - \lambda N(1)\Big)d_1 + \pi_2\Big(A - \lambda N(2)\Big)d_2 = 0 \quad (3.10)$$

The demand curve in equation (3.9) is shown graphically in Figure 2. The solutions to equation (3.10) occur at the intersections between the buyer's demand curve and the seller's short sale constraint, labeled "SS constraint". As can be seen, equation (3.10) produces two solutions for the price of the asset, one where the demand curve is downward sloping, and one where it is upward sloping. The intersection on the upward sloping portion occurs when the buyer is to the right of his/her bliss point in at least one state (on the downward sloping portion of his/her utility function). At this point the trader is hurt by lower prices since they push him/her further away from the bliss point. The trader can then make him or herself better off by buying less of the cheap asset to stay close to the bliss point. Thus, this solution is unrealistic as an equilibrium. The only sensible equilibrium point will lie where the demand curve is downward sloping, at solution 1.

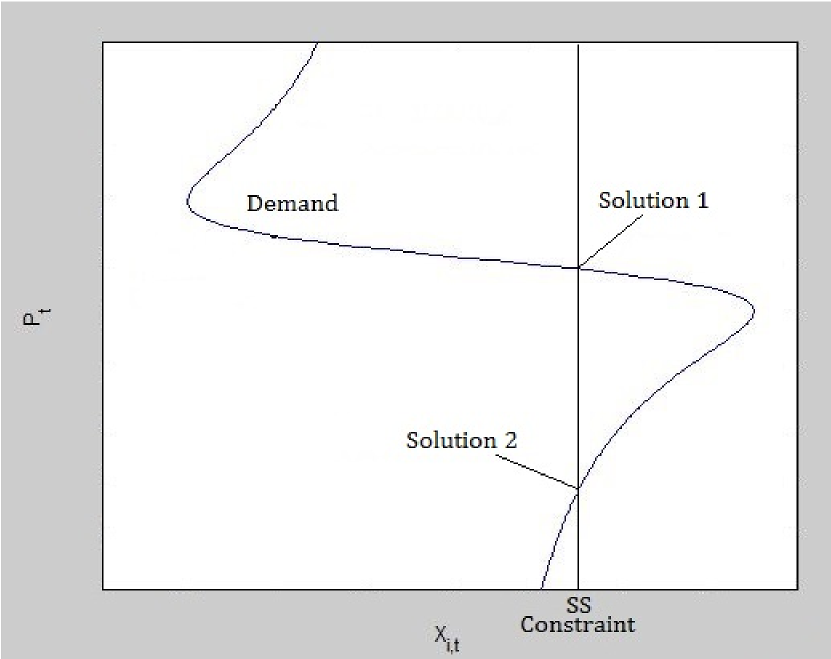

FIGURE 3.3: An example demand curve

3.3.4 Bubble Policy

As in Section 2.3, we now posit the existence of a central bank which can credibly and accurately identify when assets are overpriced. It reveals this information by announcing when it observes overpriced assets (states b, e^B, and f^B), so the price of the asset falls to zero in these states. The period 2 prices will now be positive only in states e_2^G and e_3^G, where Frank will buy from Ellen, and f_2^G and f_3^G where Ellen will buy from Frank.

The equilibrium prices in period 2, resulting from this new information partition, are calculated using the pricing expression (3.10), and presented in Table 5. Since all trade occurs in period 2, the period 1 prices are just the certainty equivalent of the period 2 prices using equation (3.1), but including states e_1^G, and f_1^G, and the final marginal utilities in from the policy equilibrium. Again, $A = 15$, and $d = 4$.

Once again, the elimination of overpricing in the market makes the buyers feel

Table 3.8: Prices as Risk Aversion Rises

λ	p_1	p_2
0.0001	1.33	2.67
0.1	1.27	2.64
0.5	1.07	2.56
1.00	0.83	2.50

more secure. In other words, when the central bank has a general deflation policy, but turns out not to announce that it observes overpriced assets, it is providing an *implicit endorsement* of the asset's price. As a result, the asset prices are higher in each period, for each of the remaining states of the world. The pattern of trade remains the same as in the bubble equilibrium, with the exception of states e^B and f^B where the price has been pushed down to zero by the central bank.

As the buyer becomes more risk averse, the price still rises, but by less and less. In the risk netural case ($\lambda \to 0$), the price rises by \$0.67 as a result of the general deflation policy, as in Section 2.3 (recall the price was \$2.00 in the bubble equilibrium). However when buyers are modestly risk averse ($\lambda = 1$), the price rises by only \$0.50. This is surprising given our intuition about risk aversion. As the degree of risk aversion rises, the buyer should fear low income states more. Since the asset is acting as insurance against the low-income state, we expect the price to rise when risk aversion gets higher. This awkward result is addressed in Section 3.6.

Although the period 1 prices are positive, the seller's willingness to pay in period 1 still exceeds the buyer's in all states. Thus in period 1, Ellen wants to buy from Frank in states e_1^G, e_2^G, and e_3^G. However, recalling the assumption made in Section 2.1, Frank cannot sell to Ellen because Frank does not have any of the asset to sell (and vice versa for Ellen in states f_1^G, f_2^G, and f_3^G). Since the sellers are short-sale constrained, no trade takes place in period 1.

3.3.5 Welfare Effects

The new, richer model presented above now allows us to relax the risk neutrality assumption in AMP and Conlon (2015). With downward sloping marginal utility curves, we can now generate welfare effects of the bubble policy by examining how utility changes for buyers as well as sellers.

To calculate final utility from the general deflation policy, we must first calculate final wealth. As an example, in state e_3^G, frank buys from Ellen and receives the dividend. If $\lambda = 0.5$, then the period 2 price is 2.56, so his final wealth in this state is

$$W_{F,3}^P(e_3^G) = 24 - 2.56 + 4 = 25.44$$

The remaining calculations proceed likewise. Table 3.9 shows the final wealths for both agents in all states for $\lambda = 0.5$.

TABLE 3.9: Final wealths for risk averse agents after general deflation policy

State	b	e^B	f^B	e_1^G	e_2^G	e_3^G	f_1^G	f_2^G	f_3^G
Ellen	28.67	26	30	26	28.56	28.56	20	27.44	25.44
Frank	28.67	30	26	20	27.44	25.44	26	28.56	28.56

As before, the asset is serving as a good hedge against the buyer's income. Where Frank's initial wealth in states e_2^G and e_3^G was initially 30 and 24 respectively, it is 27.44 and 25.44 after the market.

We can examine welfare effects of the policy by calculating utility in both the bubble equilibrium and the bursting equilibrium. Utility in the bubble equilibrium is calculated using the final wealths in Table 3.6, recalling that $\lambda = 0.5$, and $A = 15$. Since $\pi(\omega)$ is equal for all states, expected utility, EU_i is just the arithmetic mean of the utilities. E^{Buyer} contains four states f^B, f_1^G, f_2^G, and f_3^G.

$$EU_{Buyer} = [U_E(28) + U_E(20) + U_E(28) + U_E(26)]/4 = 217.5$$

$E^{Seller} = E^{GS} \cup E^{BS}$ contains the other five states, b, e^B, e_1^G, e_2^G, and e_3^G. Her conditional expected final utility as a seller in the bubble equilibrium is

$$EU_{Seller} = [U_E(28.67) + U_E(28.67) + U_E(26) + U_E(28) + U_E(28)]/5 = 223.62 \tag{3.11}$$

Similar calculations, using Table 3.9 for the general deflation equilibrium, yield

$$EU^P_{Buyer} = 217.04$$

and

$$EU^P_{Seller} = 223.34$$

The welfare effects of the deflation policy can be found by taking the difference between the conditional expected utilities from the two equilibria

$$\Delta EU_{Buyer} = EU^{P}_{Buyer} - EU_{Buyer} = -0.21 \tag{3.12}$$

$$\Delta EU_{Seller} = EU^{P}_{Seller} - EU_{Seller} = -0.27 \tag{3.13}$$

Calculations like those above for different levels of risk aversion are shown in Table 3.10. As expected, in the near-risk neutral case there are no effects on welfare. We calibrated the model such that this would be true for the sellers. For the buyers in the risk neutral case, the marginal utility curves are flat, which implies that their demand curves are perfectly elastic. In this case there is no consumer surplus before or after the policy, so there are no changes in welfare either positive or negative.

Table 3.10: Changes in Utility as Risk Aversion Rises

λ	ΔEU_{Buyer}	ΔEU_{Seller}
0.0001	0.000	0.000
0.1	-0.040	-0.059
0.5	-0.209	-0.274
1.00	-0.437	-0.517

As risk aversion rises, however, both the buyer and seller become increasingly and monotonically worse off.

3.3.6 Discussion

The state in which the dividend is paid winds up being the state where the buyer's wealth is the lowest. As risk aversion rises, the buyer should fear the risk of the low income state more and more. We expect that the buyer will have a higher willingness to pay to smooth his/her wealth out, and bid a higher price for the insurance. Curiously however, this is not what happens. The price the buyer is willing to pay actually falls as the agents get more risk averse.

In period 2, the buyer knows that only two states of the world are possible. Frank, for example, observes $F^{BuyerP} = \{e_2^G, e_3^G\}$. If Frank is more risk averse he should be willing to pay more to switch wealth from the high wealth state, e_2^G, to the low wealth state, e_3^G, using the asset which pays a dividend in e_3^G. In other words, we would expect the shadow value of switching wealth to state e_3^G to be increasing, so the marginal rate of substitution between states 2 and 3 should also be increasing. Final marginal utilities generated by the wealths in Table 3.9 are shown in Table 3.11.

TABLE 3.11: Final marginal utilities for risk averse agents after general deflation policy

State	b	e^B	f^B	e_1^G	e_2^G	e_3^G	f_1^G	f_2^G	f_3^G
Ellen	1	2	0	2	0.72	0.72	4	1.28	2.28
Frank	1	0	2	4	1.28	2.28	2	0.72	0.72

The marginal rate of substitution is ratio of the marginal utilities in the two states, e_2^G, and e_3^G:

$$MRS_{e_2^G, e_3^G} = \frac{MU_{e_3^G}}{MU_{e_2^G}}$$

Suppose $\lambda = 0.1$, and $d(e_3^G) = 4$. Then in the general deflation policy equilibrium, the marginal rate of substitution is

$$MRS = \frac{2.06}{1.06} = 1.94 \tag{3.14}$$

If the degree of risk aversion rises to 0.5 (marginal utility curves get steeper), the MRS becomes

$$MRS = \frac{2.28}{1.28} = 1.78 \tag{3.15}$$

Contrary to what we expect, as risk aversion rises, the marginal rate of substitution falls! This means that the shadow value of moving wealth from state e_2^G to state e_3^G is also falling, so the price of the asset, which is serving as insurance against state e_3^G, is falling as well when Frank becomes more risk averse. This is an awkward result, and one that is contrary to our prior intuition about how the price of insurance should behave when buyers become more risk averse. But we will show in Section 4 that this result does not generalize to all utility functions, and is an artifact of choosing the quadratic utility function to illustrate this particular example. The awkward result arises from the fact that the quadratic utility function has increasing absolute risk aversion. When wealth falls due to the policy, the agents become less risk averse, and thus feel less inclined to shift wealth to the low-wealth state of the world. As a result the price of the asset which is serving as insurance against the low-wealth state of the world falls.

3.4 Generalizing

In general, the numerical exercise above cannot be repeated for most utility functions. This is because more general utility functions do not lend themselves to analytic solutions. However, it is possible to derive numerical solutions using fixed point methods.

3.4.1 The fixed point method

A constant x^* is a fixed point of $f(x)$ if and only if $f(x^*) = x^*$. All functions that pass through the 45-degree line are guaranteed to have at least one fixed point. Fixed points can be attractive or non-attractive. Attractive fixed points can be found using iterative strategies. Suppose a function has an attractive fixed point at x^* and we choose a value x sufficiently nearby. Then the sequence

$$x, f(x), f(f(x)), f(f(f(x))), ... \tag{3.16}$$

will converge to the fixed point, x^*. As an example, the function in Figure 3.4 has two fixed points. An attractive fixed point exists at x^*, and a non-attractive fixed point at the origin. For any choice of $x > 0$, the sequence (3.16) will converge to x^*. The non-attractive fixed point can only be found by setting the initial value to $x = 0$.

Lastly, not all functions have a fixed point. The function in Figure 3.5 does not cross the 45-degree line. Thus, for any choice of initial value x, the sequence (3.16) will not converge.

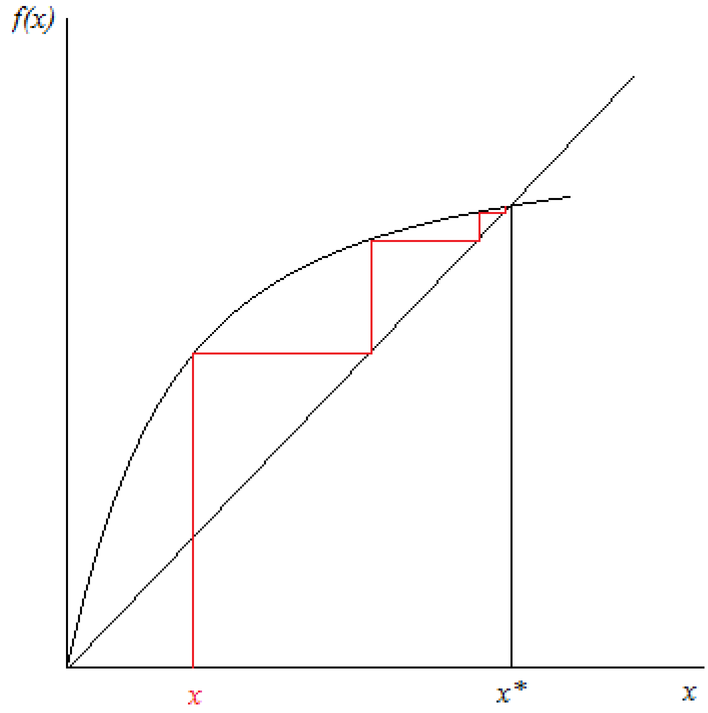

FIGURE 3.4: Finding a fixed point through iteration

3.4.2 Finding solutions for more general utility functions using the fixed point method

We can use the iterative method of finding attractive fixed points for more general classes of utility. In this section we wish to consider whether the results in Sections 3.4 and 3.5 hold for more general utility functions. Here we examine a Constant Absolute Risk Aversion (CARA) utility function,

$$U(w) = -e^{-\lambda w} \tag{3.17}$$

and a Constant Relative Risk Aversion (CRRA) utility function

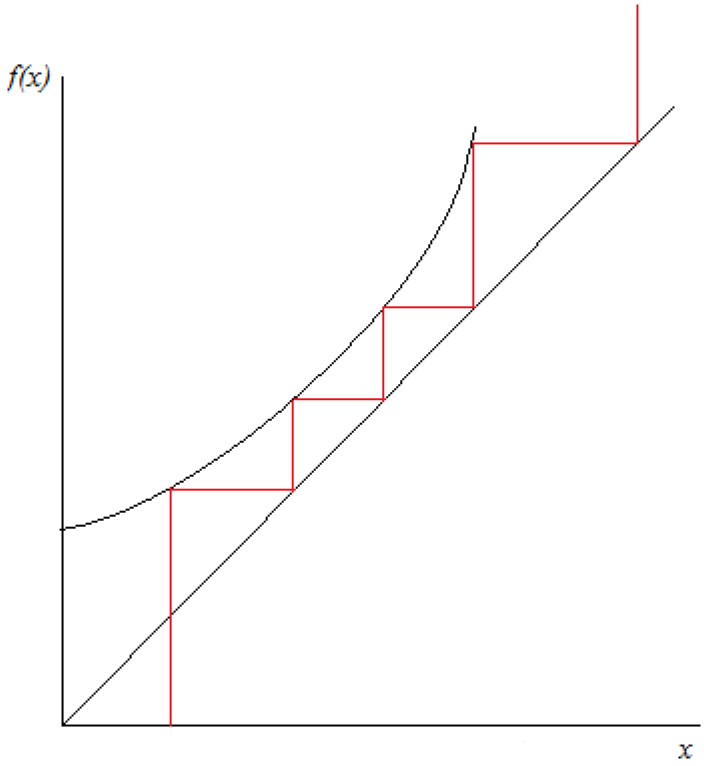

FIGURE 3.5: A function with no fixed points.

$$U(w) = \frac{w^{1-\lambda}}{1-\lambda} \tag{3.18}$$

We can apply the fixed point method described above to these new utility functions by using sequence (3.16). The procedure is roughly the same as Section 3.3, up to Section 3.2. We first back out the implied final bubble equilibrium wealths, and the implied initial wealths using the assumed marginal utilities in Table 3.1. This step does not require any fixed point calculation.

Next, to calculate prices under the deflation policy, we do not derive an analytic solution for the price. Instead, we make a guess at the price and use this guess, together with the initial wealths above, to derive the final wealths obtained in a manner similar to Section 3.1, and final marginal utilities in the policy equilibrium. Using equation (3.1), we then calculate the new certainty equivalent price given the final marginal utilities just calculated. We then iterate this process as in sequence (3.16), for the price by using the previous solution for equation (3.1) as the new guess. This process is repeated until the sequence of certainty equivalent prices converge, yielding x such that $x \approx f(x)$. The results are presented in Table 3.12 along side the previous results for quadratic utility. Once again we use the target bubble equilibrium marginal utilities from Table 3.1 with $z = y = 1$ and $x = 2/3$, and the dividend is set to 4.[9]

Notice that the pattern described in Table 3.8 changes when we introduce new utility functions. In the quadratic case, the price rises by less as the agents become more risk averse because the marginal rate of substitution between states 2 and 3 is getting smaller. But this doesn't seem to make much sense. Recall that the buyer's wealth is very low in the dividend paying state (state e_3^G for Frank, and f_3^G for Ellen), and the asset pays off in those same states, so it serves as insurance against the low wealth states. But the quadratic utility results suggest that the buyer's willingness to pay for the asset, which is serving as insurance against the low wealth state, falls as the buyer becomes more risk averse.

By contrast, for the CARA function, the price is invariant to risk aversion, and the CRRA function produces the results we expect, where prices increase with risk aversion. However, for all three functions, buyer welfare decreases, while seller welfare is decreasing for the quadratic and CARA cases, but increasing for the CRRA case.

[9]The CRRA function is undefined when $\lambda = 1$, so the results are shown up to $\lambda = 0.99$.

Table 3.12: Period 2 prices and utility loss for different utility functions

	Quadratic				CARA				CRRA			
λ	p_2	MRS	ΔEU_{Buyer}	ΔEU_{Seller}	p_2	MRS	ΔEU_{Buyer}	ΔEU_{Seller}	p_2	MRS	ΔEU_{Buyer}	ΔEU_{Seller}
0.0001	2.67	2.00	0.00	0.00	2.67	2.00	0.00	0.00	2.67	2.00	0.00	0.00
0.1	2.64	1.94	-.04	-.06	2.67	2.00	-.06	-.04	2.81	2.35	-.02	+.01
0.5	2.56	1.78	-.21	-.27	2.67	2.00	-.27	-.23	3.32	4.90	-.11	+.05
0.99	2.50	1.67	-.43	-.52	2.67	2.00	-.49	-.66	3.83	22.73	-.35	+.06

The puzzling result for the price is explained by the fact that the buyer's degree of *absolute* risk aversion is changing as the buyer's wealth falls due to the policy. We can use the Arrow-Pratt measure of absolute risk aversion to help us understand what is happening here.

Define the Arrow-Pratt coefficient of absolute risk aversion as

$$ARA = -\frac{U''(w)}{U'(w)} \tag{3.19}$$

The coefficient of absolute risk aversion for the quadratic utility function (equation (3.2)) is

$$ARA_{quadratic} = \frac{\lambda}{A - \lambda w}$$

The sensitivity of the ARA to an increase in wealth is

$$ARA'(w)_{quadratic} = \left(\frac{\lambda}{A - \lambda w}\right)^2 > 0$$

Thus, risk aversion *rises* (*falls*) as wealth gets higher (lower). Recall that the policy reduces wealth in period 2 for both states the buyer actually buys in. Thus, the buyer's risk aversion *falls* as wealth falls, meaning that the buyer cares less about the risk of the low wealth state, and therefore reduces his willingness to pay for the insurance. This is an awkward result, but it suggests that the behavior of price, as risk aversion rises, depends on the form of the utility function we assume. Perhaps other utility functions yield more intuitive results.

When CARA utility is assumed, the coefficient of absolute risk aversion is invariant to changes in wealth. So in this case, as can be seen from Table 3.12, the marginal rate of substitution and the price do not change as λ rises. This too, however is a difficult result to reconcile with practical intuition. Finally, under CRRA utility, $ARA'(w)_{CRRA} < 0$. Therefore, in this case, as wealth falls, risk aversion increases, the marginal rate of substitution increases (the buyer fears the risk of the low wealth state more), and the price paid for the asset rises.[10]

The factor that distinguishes each type of utility function is how the degree of absolute risk aversion changes when wealth changes. Since the price goes up, buyers' wealth falls in the states where trade occurs. How the degree of risk aversion changes as the agents move along the marginal utility curve determines what happens to the marginal rate of substitution between states 2 and 3, and therefore what happens to the price. Following this line of thought, we have the following theorem.

Theorem 3.1. *The marginal rate of substitution between buyers' states under the policy equilibrium will be decreasing in wealth if ARA is increasing in w, constant in w if ARA*

[10] Similar calculations to those above give us

$$ARA_{CARA} = -\frac{-\lambda^2 e^{-\lambda w}}{\lambda e^{-\lambda w}} = \lambda$$

Thus, $ARA'_{CARA}(w) = 0$, so as expected, there is no effect of wealth on the degree of absolute risk aversion, and the price rises the same for any choice of λ.

For CRRA utility,

$$ARA_{CRRA} = \frac{-\lambda w^{-\lambda-1}}{w^{-}\lambda} = \frac{\lambda}{w}$$

so $ARA'_{CRRA}(w) = -\frac{\lambda}{w^2} < 0$, and CRRA utility exhibits decreasing absolute risk aversion.

is constant in w, and increasing in w if ARA is decreasing in w. Thus, as agents become more risk averse, the price will decrease if ARA is increasing in w, remain constant if ARA is constant in w, and incrase if ARA is decreasing in w.

Proof. Recall that $MRS_{2,3} = \frac{U'(w_3)}{U'(w_2)}$, and that wealth is decreasing as a result of the policy in both states of the buyer's information set (for example $F^{BuyerP} = \{e_2^G, e_3^G\}$). The amount by which wealth falls is the same in both states because the buyer pays the same price in both states. Call this change in wealth Δw^P, which, of course, is negative. Differentiating MRS by the change in wealth,

$$\frac{\partial MRS}{\partial \Delta w^P} = \frac{U'(w_2 + \Delta w^P)U''(w_3 + \Delta w^P) - U'(w_3 + \Delta w^P)U''(w_2 + \Delta w^P)}{[U'(w_2 + \Delta w^P)]^2}$$

The denominator must be positive, so the sign of $\frac{\partial MRS}{\partial \Delta w^P}$ is controlled by the two terms in the numerator.

If, for example, $U'(w_2 + \Delta w^P)U''(w_3 + \Delta w^P) < U'(w_3 + \Delta w^P)U''(w_2 + \Delta w^P)$, then $\frac{\partial MRS}{\partial \Delta w^P} < 0$. Dividing both sides of the inequality by $U'(\cdot)$ gives us

$$\frac{U''(w_3 + \Delta w^P)}{U'(w_3 + \Delta w^P)} < \frac{U''(w_2 + \Delta w^P)}{U'(w_2 + \Delta w^P)}$$

Using the definition of Arrow-Pratt absolute risk aversion above, this can be restated as

$$ARA(w_3 + \Delta w^P) > ARA(w_2 + \Delta w^P) \tag{3.20}$$

Since we know that $w_2 > w_3$, it must be the case that (3.20) implies that ARA is falling with wealth (or conversely rising when wealth falls). Thus, for a utility function where ARA is decreasing in wealth, it must be the case that a decrease in wealth causes the marginal rate of substitution to rise, and therefore the price of the asset rises. The cases where $ARA(w_3 + \Delta w^P) = ARA(w_2 + \Delta w^P)$, and $ARA(w_3 + \Delta w^P) < ARA(w_2 + \Delta w^P)$ follow similarly.

■

Given the undesirable properties of Increasing Absolute Risk Aversion (IARA) and CARA utility functions, we focus the rest of the discussion on the case where utility displays Decreasing Absolute Risk Aversion (DARA). This leads us to Theorem 2.

Theorem 3.2. *A general deflation policy will be welfare reducing for the buyer if absolute risk aversion falls with wealth.*

Proof. Suppose Frank is the buyer, and Ellen the seller. We know that in the bad seller state, the general deflation policy causes Frank's wealth to rise because the central bank prevents him from buying a worthless asset. Define the gain in utility arising from not purchasing the worthless asset to be equal to

$$U(w_{e^B}^P) - U(w_{e^B}) = \Delta U_B^{Frank}$$

If marginal utility is constant, we know that the utility curve is linear, so utility will rise by exactly

$$\Delta U_B^{Frank} = (p_2 - p_2^P(e^B))MU_{e^B}$$

where p_2 is the price in the bubble equilibrium, and $p_2^P(e^B)$ is the general deflation policy price in e^B. We know that $p_2^P(e^B) = 0$ so this becomes

$$\Delta U_B^{Frank} = p_2 MU_{e^B}$$

Frank's wealth will fall in Ellen's good seller states because the policy causes the price to rise. The expected utility loss from the increase in price is, similarly,

$$\Delta U_G^{Frank} = (p_2^P(e^G) - p_2)MU_{e_2^G} + (p_2^P(e^G) - p_2)MU_{e_3^G}$$

where $p_2^P(e^G)$ is the price in e_2^G and e_3^G under the deflation policy (the price must be the same in both states).

In the risk neutral case, the buyer's utility does not change. Thus $\Delta U_B^{Frank} = -\Delta U_G^{Frank}$, or

$$p_2 MU_{e^B} = (p_2^P(e^G) - p_2)MU_{e_2^G} + (p_2^P(e^G) - p_2)MU_{e_3^G} \tag{3.21}$$

Rearranging terms and simplifying the notation for price, we get

$$\frac{p_2^P - p_2}{p_2} = \frac{MU_{e^B}}{MU_{e_2^G} + MU_{e_3^G}}$$

The left hand side of this equation is just the percentage increase in the price of the asset due to the policy in the good seller states of the world, so this simplifies to

$$\frac{\Delta p_2}{p_2} = \frac{MU_{e^B}}{MU_{e_2^G} + MU_{e_3^G}} \tag{3.22}$$

As Frank becomes more risk averse, then when the policy is passed, his wealth rises in state e^B, and falls in states e_2^G and e_3^G. Therefore his marginal utility falls in state e^B, and rises in states e_2^G and e_3^G. Thus, risk aversion causes the numerator of the fraction $\frac{MU_{e^B}}{MU_{e_2^G}+MU_{e_3^G}}$ to fall, and the denominator to rise, causing the overall value of the fraction to fall. Meanwhile, from Theorem 1, the policy guarantees that $\frac{\Delta p_2}{p_2}$ gets larger if and only if the utility function is decreasing in absolute risk aversion (DARA). Therefore for a DARA utility function, it must be the case that

$$\frac{\Delta p_2}{p_2} > \frac{MU_{e^B}}{MU_{e_2^G} + MU_{e_3^G}} \tag{3.23}$$

or, that

$$p_2 MU_{e^B} < (p_2^P(e^G) - p_2)MU_{e_2^G} + (p_2^P(e^G) - p_2)MU_{e_3^G} \tag{3.24}$$

or in other words, the utility gain from the policy preventing Frank from buying an overpriced asset in state e^B is overwhelmed by the loss in utility from the price increasing in states e_2^G and e_3^G. Therefore, when the utility function is decreasing in absolute risk aversion, it must be the case that $\Delta U_B^{Frank} < \Delta U_G^{Frank}$, and the policy is welfare diminishing for buyers. ■

3.5 Conclusion

This model extends the basic model from AMP and Conlon (2015) to include risk aversion. After deriving a bubble equilibrium and reverse-engineering it to back out the initial wealths a risk averse agent would need to have in order to generate that particular equilibrium, we use these initial wealths to find the corresponding equilibrium where the central bank has deflated the overpriced assets by announcing when bad sellers are in the market. While this policy will protect buyers from purchasing overpriced assets, it harms them by causing them to pay higher prices in the states where no announcement is made. We show, under the generally plausible assumption of decreasing absolute risk aversion, that this policy must be welfare diminishing for buyers in the asset market.

This result has serious implications for the debate surrounding central bank policy concerning asset markets. In essence we have shown that even if the central bank had good enough information to know precisely when assets in the market were overpriced, it cannot help buyers in the market by protecting them from overpriced assets. A policy of general deflation of overpriced assets helps buyers in states where they would have bought an asset that was known to be overpriced, but hurts them in states where the price rises due to the central bank's elimination of the lemons problem. Since buyers' expected utility falls, it suggests this policy cannot, and should not be employed by central banks.

The policy's welfare effects on the sellers are ambiguous. Under DARA utility, it is possible that the sellers could be helped by the policy. However, ostensibly, this is contrary to the policy's stated goals since the policy still hurts buyers, and therefore should not be argued in its defense.

Interestingly, the unintended consequences of bubble policy by central banks presented in this paper seem to have been largely ignored by those arguing in favor of these policies. The analysis in this paper suggests that these effects are large enough to cause serious doubts about the policy's ability to acheive its stated goals.

CHAPTER 4

CONCLUSION

For as long as asset markets have existed, they have been a source of intense curiosity for the field of Economics. We believe that the results in this thesis have contributed to the body of knowledge that currently exists.

We have shown that inside traders either did not or could not predict the 2008 stock market crash. This combined with previous results by Seyhun (1990) and Marin and Olivier (2008) suggests that informational signals sent by inside traders behave differently in wide spread financial crises than during normal times. This result deserves closer examination in future work. Namely, it would be informative to see if the same result holds once we have filtered routine traders out of the sample. We have also shown that inside traders performed very well picking winners in the wake of the 2008 stock market crash. This result, supported again by Seyhun (1990) suggests that insiders are much more able to pick stocks after a market crash than during normal times. The existence of large returns in the post-crash data suggests that investors may not be used to dealing with crises, as the EMH suggests they should easily be able to arbitrage these gains away. More research on precisely why these signals seem to be so much stronger after a crash, and why investors do not seem to be able to eliminate these gains is needed to say anything definite. However this essay has provided an empirical base of knowledge from which to work on.

We have also shown that asset market policy by central banks is more difficult than it first seems. Even if a central bank has perfect information about whether an asset is overpriced, it cannot help buyers by protecting them from the overpricing. This suggests that buyers would prefer a world where the central bank ignores swings in asset markets, to one where the central bank deflates overpriced assets. Further research on this topic includes formalizing the effects of the general deflation policy on sellers in the market, and what happens when we relax short sale constraints.

BIBLIOGRAPHY

[1] Abreu, D, and M. Brunnermeier, "Bubbles and Crashes", 2003, *Econometrica*, Vol. **71**, No. 1, 173-204.

[2] Allen, F., and D. Gale, 2000, Bubbles and Crises, "The Economic Journal", Vol. 110, 236-255.

[3] Allen, F., S. Morris, and A. Postlewaite, 1993, "Finite Bubbles with Short Sale Constraints and Asymmetric Information," *Journal of Economic Theory*, Vol. 61, 206-229.

[4] Allen, F., S. Morris, and H. Shin, 2006, "Beauty Contests and Iterated Expectations in Asset Markets", *Review of Financial Studies*, Vol. 19, No. 3, 719-752.

[5] Allen, F., A. Postlewaite, and H. Shin, 1995, "The Depth of Knowledge and the Effect of Higher Order Uncertainty", *Economic Theory*, Vol. 6, 453-467.

[6] Banz, R., 1981, "The Relationship between Return and the Market Value of Common Stocks," *Journal of Financial Economics*, Vol. 9, 3-18.

[7] Cecchetti, S., H. Genberg, J. Lipsky, and S. Wadhwany, 2000, *Asset Prices and Central Bank Policy*, Geneva: International Center for Monetary and Banking Studies.

[8] Chan, L. and J. Lakonishok, 1995, "The Behavior of Stock Prices Around Institutional Trades," *Journal of Finance*, Vol. 50, 1147-1174.

[9] Chowdhury, Howe, and Lin, 1993, "The Relation between Aggregate Insider Transactions and Stock Market Returns", *The Journal of Financial and Quantitative Analysis*, Vol. 28, No. 3, 431, 437.

[10] Cohen, L., C. Malloy, and L. Pomorski., 2012, "Decoding inside information." *Journal of Finance*, Vol. 67, no. 3, 1009-1043.

[11] Conlon, J. R., "Simple Finite Horizon Bubbles Robust to Higher Order Knowledge," 2004, *Econometrica*, Vol. Vol. 72, No. 3, 927-936.

[12] Conlon, J. R., 2015 "Should Central Banks Burst Bubbles? Some Microeconomic Considerations" *The Economic Journal*, Vol. 125, No. 582, 141-161.

[13] Conrad, J., and G. Kaul, 1993, "Long Term Market Overreaction or Biases in Computed Returns?" *Journal of Finance*, Vol. 48, no. 1, 39-63.

[14] Doblas-Madrid, A., "A Robust Model of Bubbles with Multidimensional Uncertainty," 2012, *Econometrica*, Vol. 80, No. 5, 1845-1893.

[15] Fama, E., 1965, "The Behavior of Stock Market Prices," *The Journal of Business*, Vol. 38, No. 1, 34-105.

[16] Fama, E., 1970, "Efficient Capital Markets: A Review of Theoretical and Empirical Work," *Journal of Finance* Vol XXV, No. 2, 383-417.

[17] Fama, E., and K. French, 1992, "The cross section of expected stock returns." *Journal of Finance* 47, no. 2, 427-465.

[18] Fama, E., and K. French, 1993, "Common risk factors in the returns on stocks and bonds." *Journal of financial economics* 33, no. 1, 3-56.

[19] Fama, E., and J. MacBeth, 1973, Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.

[20] Finnerty, J., 1976, "Insiders and Market Efficiency," *Journal of Finance* Vol XXXI, No. 4, 1141-1148.

[21] Garber, P. 1990, "Famous First Bubbles", *Journal of Economic Perspectives*, Vol. 4, No. 2, 35-54.

[22] Gorton, G. *Slapped by the Invisible Hand*, Oxford University Press, New York, NY, 2010.

[23] Garber, P. *Famous First Bubbles*, MIT University Press, Cambridge, MA, 2001.

[24] Harrison, J., and D. Kreps, 1978, "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics* Vol. 92, 323-336.

[25] Jaffe, J., 1974, "Special Information and Insider Trading," *Journal of Business*, Vol. 47, 410-428.

[26] Jiang, X., and M. Zaman, 2010, "Aggregate insider trading: Contrarian beliefs or superior information?", *Journal of Banking and Finance*, Vol. 34, 1225-1236.

[27] Kai, G. and J. Conlon, 2008, "Why Bubble Bursting is Unpredictable: Welfare Effects of Anti-bubble Policy when Central Banks Make Mistakes", working paper.

[28] Kindleberger, C. and R. Aliber, 2011, "Manias, Panics, and Crashes: A History of Financial Crises", Palgrave MacMillan, Basingstoke, UK.

[29] Lakonishok, J., and I. Lee., 2001, "Are Insider Trades Informative?" *Review of Financial Studies*, Vol. 14, No. 1, 79-111.

[30] Lorie, J, and V. Niederhoffer, 1968, "Predictive and Statistical Properties of Insider Trading," *Journal of Law and Economics*, Vol. XI, 35-51.

[31] Marin, J, and J. Olivier, 2008,"The Dog that Did Not Bark: Insider Trading and Crashes," *Journal of Finance*, Vol. LXIII, No. 5, 2429-2476.

[32] McKeon R., and J. Netter, "What Caused the 1987 Stock Market Crash and Lessons for the 2008 Crash," *Review of Accounting and Finance*, 2009, 123-137

[33] Milgrom, P., and N. Stokey, 1982, "Information, Trade and Common Knowledge", *Journal of Economic Theory*, Vol. 26, 17-27.

[34] Posen, A., 2006, "Why Central Banks Should Not Burst Bubbles", *International Finance*, vol. 9, no. 1, 109-24.

[35] Morris, S., A. Postlewaite, and H. S. Shin, "Depth of Knowledge and the Effect of Higher Order Uncertainty," 1995, *Economic Theory*, Vol. 6, 453-467.

[36] Rozeff, M. and M. Zaman, 1988, "Market Efficiency and Insider Trading: New Evidence", *The Journal of Business*, Vol. 61, No. 1, 25-44

[37] Santos, M., and M. Woodford, "Rational Asset Pricing Bubbles", 1997, *Econometrica*, Vol. 65, 19-57.

[38] Scheinkman, J., and W. Xiong, "Overconfidence and Speculative Bubbles," *Journal of Political Economy* Vol. 64, 467-482.

[39] Seyhun, N., 1986, "Insiders' Profits, Costs of Trading and Market Efficiency," *Journal of Financial Economics*, Vol. 16, 1-24.

[40] Seyhun, N., 1990, "Overreaction or Fundamentals: Some Lessons from Insiders' Response to the Market Crash of 1987," *Journal of Finance*, Vol. XLV, No. 5, 1363-1388.

[41] Seyhun, N., 1988 "The Information Content of Aggregate Insider Trading," *Journal of Business*, Vol. 61, 1-24.

[42] Seyhun, N., 1992, "Why Does Aggregate Insider Trading Predict Future Stock Returns?", *Quarterly Journal of Economics*, 1303-1331.

[43] Seyhun, N. *Investment Intelligence: From Insider Trading*, MIT Press, Cambridge, MA, 1999.

[44] Shiller, R. J., *Irrational Exuberance*, Princeton University Press, Princeton, NJ, 2000.

[45] Simon, H., 1955, "A behavioral model of rational choice." *Quarterly Journal of Economics* 69, no. 1, 99-118.

[46] Sims, Christopher A., 2003, "Implications of rational inattention." *Journal of Monetary Economics* 50, no. 3, 665-690.

[47] Smith, V., G. Suchanek, and A. Williams, 1988, “Bubbles, Crashes, and Endogenous Expectations in Spot Asset Markets,” *Econometrica*, Vol. 56, 1119-1151.

[48] Takavoli, M. D. McMillan, and P. McKnight, 2012, “Insider Trading and Stock Prices,” *International Review of Economics and Finance*, Vol. 22, 254-266.

VITA

Harlan M. Holt

Department of Economics
Holman 341
University, MS 38677
(224) 392-3854
hmholt22@gmail.com

PERSONAL *Website*: https://sites.google.com/site/harlanmholt/home
Citizenship: USA

EDUCATION *Ph.D.,* Economics
University of Mississippi, University, MS, August 2015
M.A., Economics
University of Mississippi, University, MS, May 2011
B.S., Economics
Texas A&M University, College Station, TX, December 2009

RESEARCH AREAS Applied Macroeconomics, Financial Economics, Asset Markets, Search Models, Asset Bubbles

POSITIONS *Visiting Assistant Professor* Aug. 2015-
Northern Illinois University, DeKalb, IL

RESEARCH *Working Papers*

- "Turning Pink Slips into Red Tape: The Unintended Effects of Employment Protection Legislation" (with Josh Hendrickson), Working Paper Link **(Revise and resubmit: *Contemporary Economic Policy*)**

- "Insider Trading and the 2008 Stock Market Crash", Working Paper Link
- "Money, Credit, and Growth" (with Josh Hendrickson), Working Paper Link
- "Rational Greater Fool Bubbles with Risk Averse Agents"

Works in Progress

- "Why Don't Banks Hold More Capital?" (with Josh Hendrickson and Mark V. van Boening)
- "Risk Aversion in Experimental Asset Markets" (with Tim Carle and Ken Yamada)
- "Information, Insiders, and Asset Markets: An Experiment"

PRESENTATIONS

- "Did Insiders Predict the 2008 Stock Market Crash?", University of Mississippi Finance department seminar, Fall 2012, University, MS.
- "Did Insiders Predict the 2008 Stock Market Crash?", Midwest Finance Association Annual Meeting, Spring 2013, Chicago, IL.
- "Insider Trading and the 2008 Stock Market Crash", University of Mississippi Economics department seminar, Fall 2014, University, MS.
- "Money, Credit, and Growth", Southern Economic Association Annual Meeting, Fall 2014, Atlanta, GA.

OTHER *Workshops*

- Summer 2013 - Barcelona LeeX Experimental Economics Summer School in Macroeconomics in Universitat Pompeu Fabra. June 13-18, 2013

TEACHING EXPERIENCE *Instructor* 2010-2015

- Principles of Microeconomics: *Spring 2011*, *Fall 2011*, *Fall 2012*, *Fall 2013*, *Fall 2014*
- Principles of Macroeconomics: *Summer 2011*, *Spring 2012*, *Summer 2012*, *Spring 2013*, *Summer 2013*, *Spring 2014*, *Summer 2014*, *Spring 2015*