

Experimental economics

Lecture X - Behavioral Finance

Matej Lorko

matej.lorko@euba.sk

Student resources: www.lorko.sk

References:

- Cartwright, E. (2018). *Behavioral economics*. Routledge.

Life cycle hypothesis

- One of the important and controversial ideas in economics, is the life cycle hypothesis. The life cycle hypothesis comes in various different guises but the basic idea is relatively simple: a person should take a long-term view about their future income and smooth consumption over their lifetime to maximize expected utility.
- Income goes through large variation over the life cycle, as people typically earn little when being students, have increased income later and then earn little income again when they retire. Presumably people will want their standard of living and consumption to be more stable than this. Consequently, they may smooth consumption by borrowing in student years, saving during working years and using savings during retirement.
- When we look at actual spending data, we see that people undeniably do smooth consumption over their lifetime by saving for retirement, borrowing when a student or borrowing to buy a first home. However, they do not smooth consumption as much as we might expect. Instead, consumption seems to track current income more closely and be more sensitive to macroeconomic cycles than would seem optimal. We need to try and understand why this is the case. One potential contributory factor is mental accounting.




Figure 2.30 An illustration of consumption smoothing. Income varies a lot with age, but consumption can be completely smoothed, as in schedule B, or partly smoothed, as in schedule A.




Figure 2.31 An inverse U-shaped consumption and income pattern derived from UK couples with a husband born between 1936 and 1943. Cyclical and growth effects are removed from the data.

Source: Browning and Crossley (2001).

- Imagine that you have decided to see a play for which admission costs \$10 per ticket. As you enter the theater you discover that you have lost a \$10 bill. Would you pay \$10 to watch the play?
- Imagine that you have bought the ticket to see a play for which admission costs \$10. As you enter the theater, you discover you have lost the ticket and there is no way to recover it. Would you pay \$10 for another ticket?

Mental accounting

- Tversky and Kahneman (1981) report that 88 percent of subjects asked question one said they would pay, but only 46 percent asked question two said they would pay. In both cases the loss is \$10, so where the loss comes from clearly matters.
- To capture such things we can introduce mental accounting. Mental accounting is the process of coding, categorizing and evaluating choices and outcomes. The primary component of mental accounting is to put any spending or income into separate accounts for specific purposes.
- People use mental accounts to keep track of income and spending. Outcomes are perceived and experienced relative to the particular account that is brought to mind.

Mental accounting and the life cycle hypothesis

- Mental accounting is very important in thinking about the life cycle hypothesis. That's because the life cycle hypothesis assumes fungibility, meaning that all money is treated the same, no matter where it came from. If people keep mental accounts, then money is likely to be segregated according to how it was obtained, and so would not be fungible.
- To illustrate, consider these three scenarios.
 - You have been given a special bonus at work, meaning you will receive \$200 a month over the next year.
 - You have been given a special bonus at work, meaning you will receive a lump sum payment of \$2,400.
 - You have been told of a distant relative who has left you an after-tax inheritance of \$2,400, but you will not receive the money for five years.
- In all three scenarios your wealth will increase by \$2,400. Fungibility requires that the use to which you put this extra money should not depend on where it came from. You should just think how you will spend the \$2,400 over the rest of your lifetime. In reality, things seem different, as the results of a study by Shefrin and Thaler (1988) illustrate. There is a stark difference in how much of the extra \$2,400 people said they would consume over the following year. If the money came from the regular payment, people planned to consume half of it in one year, but, if it came from a future inheritance, would spend none of it.

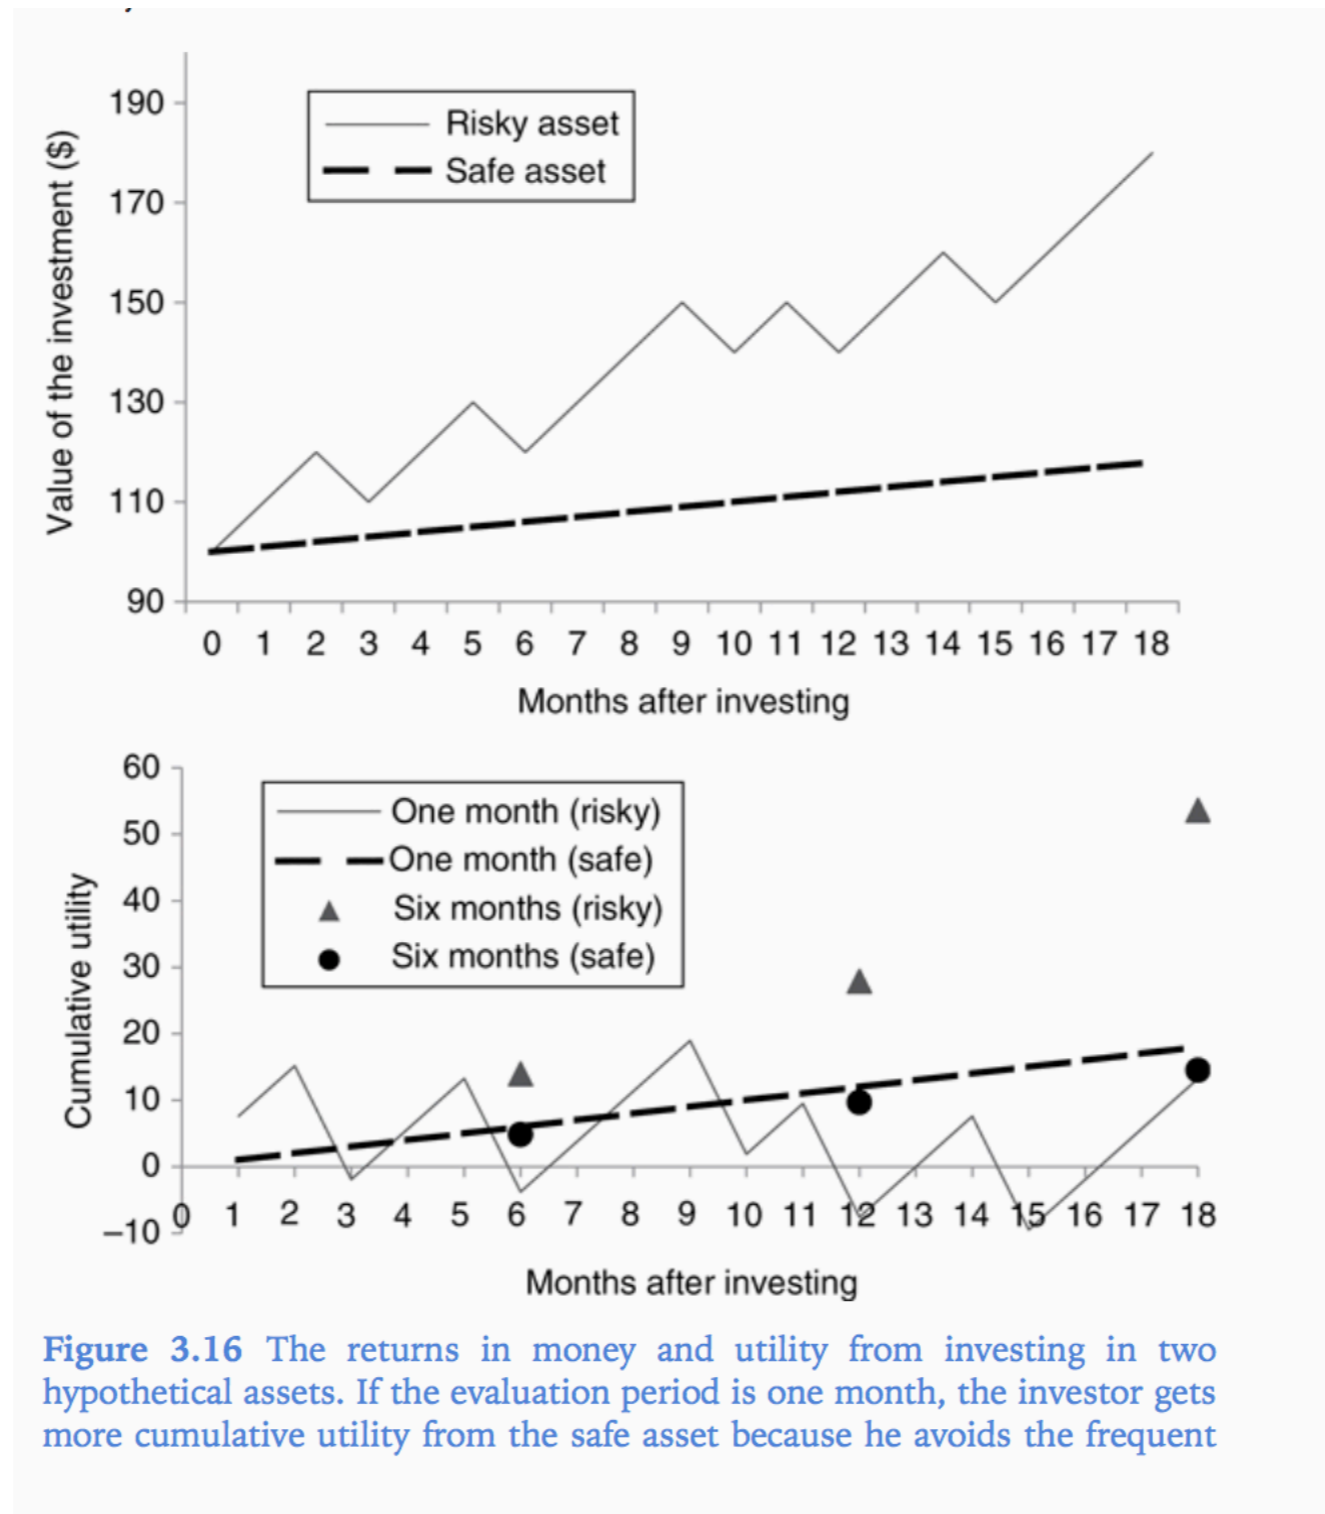
Mental accounting and the life cycle hypothesis

- Broadly speaking, we can think that households keep in mind three basic mental accounts: a current income account, an asset account and a future income account.
 - The current income account is for day-to-day spending. We would expect that the marginal propensity to consume from this account, or the proportion of the money spent, is relatively high, because the money is there to be spent.
 - The asset account is for saving and investing for things such as retirement or a new house. We would expect a much lower marginal propensity to consume from this account. Indeed, its purpose may be to commit the household to save adequately.
 - Finally, the future income account includes predictable future income such as from a pension, inheritance or expected salary increases. The marginal propensity to consume from this account is likely to be very low if people are reluctant to spend income they do not have.
- All in all, households treat money differently depending on where it comes from, how much it was and the purpose for which it was given. If people keep mental accounts, then it is natural that money should not easily flow from one account to the next. The question we need to get back to is what implications this has for the life cycle hypothesis.
- One thing we have seen is that people seem reluctant to spend future unearned income. Only when they got the money did households spend the tax rebate and restitution payment. We know, more generally, that people do spend unearned income. Student debt and families taking out a mortgage to buy a house are ample evidence of this. The point is, however, that people may spend less of their unearned income than the life cycle hypothesis might suggest they should. Given that most people get increasing amounts of money through their lifetime, this would mean spending tracks income more than would be predicted by the life cycle hypothesis. This is what we observe.

Investing and the equity premium puzzle

- The equity premium puzzle is that the real rate of return on stocks is much higher than the return on 'safer assets' such as Treasury bills. The differences in annual return can multiply to huge differences over a few years. For example, \$1 invested in stocks in the United States in 1926 would have been worth over \$2,500 by 2000, while the same amount invested in Treasury bills would have been worth less than \$20. Given such numbers, why would anyone invest in Treasury bills?
- The obvious answer would be that investors are risk-averse and so invest only in stocks, the riskier of the two, if the average return is higher. The equity premium that we observe could not be explained solely by risk aversion, however, because to do so would require extremely high coefficient of relative risk aversion. We need a different explanation. One explanation that does work is loss aversion and prospect theory. To make it work, however, we are going to have to think about a complicating factor: time.
- The prices of stocks are changing minute by minute during a working day, and in principle an investor could change his investments at any time. There is, therefore, no definite moment in time when the outcome of investing in the stock market is known. Instead, an investor must choose how long to leave his investments before he evaluates how they are doing. We call this length of time the evaluation period. It seems sensible to assume that the investor gets utility when he evaluates his investments at the end of each evaluation period.

Equity premium puzzle



Equity premium puzzle

- The evaluation period is going to be crucial. To explain why, it is easiest to work through an example. In Figure, you can see the value over time of \$100 invested in two hypothetical assets. In any month the risky asset could fall or rise by \$10 while the safe asset grows by a steady \$1. If we just look at the value of the investment, in the top half of the figure, then the risky asset looks a clear winner. Things are not so simple, however, when we look at the investor's utility.
- What we see in this example is how the volatility of a risky asset can make it undesirable for someone who is loss-averse, even if the overall returns are high. For every month that the investment goes down in value there need to be three months of gain to compensate the investor for the disutility of the bad month. This effect is lessened by lengthening the evaluation period. The longer the evaluation period, the less chance the investment will have decreased in value, and so the less chance that the investor feels a loss. This means there should be an evaluation period when he is relatively indifferent between the risky and safe asset.
- Benartzi and Thaler (1995) asked how long the evaluation period would have to be to explain the equity premium that we do observe. That is, how long the evaluation period would have to be to make the expected utility on safe assets the same as on stocks. They evaluated this by simulating the utility an investor would have got from investing in US stocks, bonds or Treasury bills for various time intervals between 1926 and 1990. They find that, with an evaluation period of one month, estimated utility would be higher from investing in safe assets. With an evaluation period of 18 months, it would be higher from stocks. The point where the estimated utility from safe assets and stocks are equal is an evaluation period of 12 months.
- The estimate of 12 months is remarkably plausible! It does seem intuitive that ordinary investors might use a one-year period to evaluate returns on investments. Managers of pension and investment funds might also be evaluated on their year-by-year performance. We have, therefore, an explanation for an equity premium. This explanation is based on people being both loss-averse and myopic, in the sense that the evaluation period is shorter than it could be.

Demand, supply and markets

- Markets are at the heart of economics, so they seem an apt place to begin looking at applications of findings in behavioral economics. It's easiest to start with a laboratory market. The standard way to create a market in the experimental lab begins by randomly allocating subjects to the role of buyer or seller. Each buyer is given a value to her of buying one unit of a fictitious good that will be traded in the experiment. Each seller is given a value to her of selling one unit of the good. Different subjects are given different values, and only they know their value.
- Buyers and sellers are then able to interact and do a deal. The profit of a buyer is her value minus the price she paid. The profit of a seller is the price she sold for minus her value. So, in principle, buyers want to buy as cheap as possible and sellers to sell as dear as possible.
- If we know the values of all sellers and buyers we can derive a demand and supply curve. To derive the supply curve we have to look at all possible prices and ask how many sellers would sell at each price. To derive the demand curve we have to look at all possible prices and ask how many buyers would buy at each price.

Demand, supply and markets




Figure 2.17 The demand and supply curves before and after a shift in demand.

- Demand and supply curves are one of the most basic ideas in economics, and these curves are no different from those you may be familiar with. For instance, the demand curve must slope down and the supply curve slope up. (You might be used to smooth curves rather than the stepped ones drawn here, but if you have enough buyers and sellers the curve will start to look smooth.) The market equilibrium price and quantity are where demand equals supply. One remarkable thing about market equilibrium is that total profit is maximized at the equilibrium.

Demand, supply and markets

- This stylized example illustrates why demand and supply analysis is so fundamentally important in economics: it is a tool to predict what may happen in markets, and should tell us what the most efficient outcome is. One thing we have not done is specify how people interact in order to make trades. Three of the more common ways we observe are the following.
 - negotiated price institution, in which buyers and sellers are free to talk to each other and try to do a deal. This is similar to bartering in a local market.
 - double-auction institution, in which buyers submit bids of what they are willing to pay and sellers submit asks of what they are willing to sell for. All bids and asks are displayed on a screen for all to see. A buyer can buy by accepting the lowest ask price. A seller can sell by accepting the highest bid price. This institution is used in most financial, commodity and currency markets.
 - posted offer institution, in which each seller displays a ‘take it or leave it’ price. Buyers can go to any seller and agree to trade at the displayed price. This institution is what we are familiar with when we go to the grocery store or most shops.


Double-auction markets in the laboratory


- In one of the earliest economic experimental studies, Vernon Smith (1962) ran ten experimental double-auction markets. Since then, the experiment has been replicated thousands of times. The typical results are: the average price and quantity observed in experimental double-auction markets almost always get close to the market equilibrium with repeated trading, and typically start close. The price and quantity also change as predicted by changes in demand and supply.
- This is a fantastic result. It suggests that demand and supply are reliable predictors of what will happen in a market, which is good news for economists. It also suggests that markets are a great way to allocate resources, which is good news for everyone.
- The efficiency of double-auction markets is arguably the most important lesson so far from behavioral economics. The reason that double-auction markets work so well is that subjects need use only simple heuristics. Recall that the only thing a subject knows at the start is her own value to buy or sell. So she simply needs to buy for less or sell for more than her value to make a profit, and this is something very familiar to most of us.

Posted offer markets and market power

- Having given the good news, it is time for some of the bad. Markets are not always so efficient.
- For example, in posted offer markets, prices stay consistently above the equilibrium and quantity below the equilibrium, and there is no evidence that this would change if the market were to be repeated more times.
- It is not too much of a surprise that we do not always see convergence to the market equilibrium in posted offer markets. This is because suppliers have an incentive to try and keep the price high.
- Sellers can keep prices high because they have the power to set prices. They can signal to each other from one period to the next through the prices they offer and potentially collude on higher prices. Prices are, therefore, higher because sellers have market power. More generally, market power can move the price away from the market equilibrium.
- The equilibrium is also often not reached due to search costs. Because of search costs, we observe variation in prices for an identical good that is sold at more than one location.
- In addition, when searching, people usually focus on relative instead of absolute savings. While saving 5 EUR on a 20 EUR product seems like a great deal, saving 5 EUR on 1000 EUR product does not impress. However, choosing to search more on the first item and less on the second is irrational, if both searches take similar time.

But what if you can be both buyer and a seller?








Reservation prices


1 2 3 4 5 6 7 8 9 10 11 12






Reservation prices

1 2 3 4 5 6 7 8 9 10 11 12




Pair – Think - Share



Reservation prices

1 2 3 4 5 6 7 8 9 10 11 12





What happens in the market assuming that a person can only eat 1 cheesecake?

- who is going to be a seller and who is going to be a buyer?
- how many trades would you expect to happen?
- what would be the market clearing (median) price?




Reservation prices






Reservation prices

1 2 3 4 5 6 7 8 9 10 11 12




Sellers

Buyers



Reservation prices


1 2 3 4 5 6 7 8 9 10 11 12



Sellers


Buyers

Expected market outcome



Reservation prices

1 2 3 4 5 6 7 8 9 10 11 12





Sellers

Buyers

Expected number of trades: 3

Expected clearing (median) price: 6.5




I am happy to buy a cake for 7 EUR.
I am happy to sell a cake for 7 EUR.

WTP vs. WTA

- The standard assumptions of economic theory imply that differences between an individual's maximum willingness to pay (WTP) for a good and minimum compensation demanded for the same entitlement (willingness to accept [WTA]) should be negligible
- The assumption that entitlements do not affect value contrasts sharply with empirical observations of significantly higher selling than buying prices.
- The endowment effect
- Thaler (1980) labeled the increased value of a good to an individual when the good becomes part of the individual's endowment the "endowment effect." This effect is a manifestation of "loss aversion".
- If a good is evaluated as a loss when it is given up and as a gain when it is acquired, loss aversion will, on average, induce a higher dollar value for owners than for potential buyers, reducing the set of mutually acceptable trades.

1 2 3 4 5 6 7 8 9 10 11 12



? ? ? ? ? ? ? ? ? ? ? ?



TABLE 2
RESULTS OF EXPERIMENT 1
INDUCED-VALUE MARKETS

Trial	Actual Trades	Expected Trades	Price	Expected Price
1	12	11	3.75	3.75
2	11	11	4.75	4.75
3	10	11	4.25	4.25

CONSUMPTION GOODS MARKETS				
Trial	Trades	Price	Median Buyer Reservation Price	Median Seller Reservation Price
Mugs (Expected Trades = 11)				
4	4	4.25	2.75	5.25
5	1	4.75	2.25	5.25
6	2	4.50	2.25	5.25
7	2	4.25	2.25	5.25
Pens (Expected Trades = 11)				
8	4	1.25	.75	2.50
9	5	1.25	.75	1.75
10	4	1.25	.75	2.25
11	5	1.25	.75	1.75

The Endowment Effect in Choices between Goods

- Undertrading happens not only in exchanges of money and consumption goods, but also in the exchange between two goods.
- Participants were either offered of a choice between chocolate and a mug, or were given a mug, or were given a chocolate. All of them had an opportunity to change their endowment at the end of the class
- For most participants a mug was more valuable than the chocolate when the mug had to be given up but less valuable when the chocolate had to be given up.

- Peter Lynch (Fidelity Magellan Fund)
- 1977 - 1989 (13 years career)
- Managed to beat the market in 11 out of those 13 years.
- His portfolio grown from USD 20 million to USD 14 billion
- Would you invest money to his mutual fund?

- Imagine a coin toss competition
- There are 500 players, each tosses the coin 13 times
- The winner is the one who tosses HEADS most times

Fidelity Magellan's Alpha Relative to IFA Large Blend Mix*

50 Years (1/1/1964 - 12/31/2013)

Minimum Track Record to Indicate Skill (t-stat > 2): **24 Years**

Average Alpha: **7.30%**
Standard Deviation of Alpha: **17.90%**

t-Statistic: **2.88**
(a t-stat of >2 = 95% confidence of skill)


FUND: Fidelity Magellan | BENCHMARK: IFA Large Blend Mix*

Full Fund Period
1964 - 2013

Before Peter Lynch
1964 - 1976

Peter Lynch at the Helm
1977 - 1990

After Peter Lynch
1991 - 2013



Fidelity Magellan's alpha is benchmarked relative to IFA Chosen Benchmark Mix. The benchmark mix is calculated by running 3 factor regression of Fidelity Magellan fund. The benchmark mix was created to achieve a similar size and value exposure as the Fidelity Magellan Fund over the whole period.

*IFA Large Blend Mix

IFA Large Company Index	50%	IFA Small Growth Index	10%
IFA Large Growth Index	10%	IFA Small Cap Index	30%

Stock markets specifics

- Bubbles and crashes are familiar in many markets, such as the housing and financial markets. It is headline news when markets crash, and bad news for many. This is why crashes are often etched into the history books as a black day: 'black Friday' on September 24, 1869, 'black Thursday' and 'black Tuesday' on October 24 and 29, 1929, 'black Monday' on October 19, 1987, 'black Wednesday' on September 16, 1992, and 'black week' beginning October 6, 2008! Why we get such extreme swings in the prices?
- We will primarily focus on stock markets. The efficient markets hypothesis says that the price of stock should be equal to its fundamental value. The fundamental value is, basically, the expected future returns from the stock, in the form of dividends. We cannot predict for certain what these future returns will be, and so the fundamental value is uncertain. The share price should, however, take into account all the information available at the time. It follows that the share price should change only if there is some new information about fundamental value.
- So, information is key. Given, therefore, that new information is likely arriving all the time, is it any surprise that prices fluctuate? This turns out to be a tricky question. The consensus, however, is that the fluctuations we observe in share prices seem too large to be caused solely by new information.
- To explain: mathematical logic says that the fluctuations in the fundamental value must exceed fluctuations in the expectation of the fundamental value. This means that share prices should be less volatile than fundamental value. The intuition for this is that, if people do not know all the relevant information, the share price should vary less, on average, than it would have done if they had known everything. What we observe is the opposite. Prices clearly fluctuate more than the fundamental value.
- We observe bubbles, when price exceeds fundamental value, busts, when price is below fundamental value, as well as sudden rises and falls in prices.

Stock markets




Figure 5.7 Price of the S&P 500 index compared to the estimated fundamental value. The fundamental value at a particular time is calculated by looking at what the dividends were in the following years.

Source: Shiller (2003).

Bubbles in the laboratory

- Earlier we saw that double-auction markets appeared to work well, in that the price and quantity traded was close to the efficient amount with supply equaling demand. In those experiments, subjects were either a buyer or a seller, knew their value for the good and made their trades one at a time. This is a good representation of many markets, such as the labor market or new car market, but is not so good a representation of many other markets, such as the stock market. In a stock market we observe traders who sometimes buy and sometimes sell, do not know the value of the good they are trading and trade over consecutive periods. What happens when we add such things to experimental markets?
- To illustrate, we shall look at some results of a study by Smith, Suchanek and Williams (1988), in which subjects took part in an experimental asset market. There are three subjects or traders. At the start of the experiment each trader is given an initial amount of money and some units of the asset. In each period traders can sell any units of the asset they hold, or buy units if they have sufficient money. This buying and selling takes place via a double auction. At the end of each period the asset yields a randomly determined amount of money called a dividend. Traders receive this dividend and the next period begins. Thus, traders can accumulate money and assets over the periods. At the end of the last period, traders take home any money they have accumulated, and the assets become worthless.
- Given that assets will be worthless by the end of the experiment, the primary reason to hold an asset is the dividend it may yield. For example, if the maximum dividend is \$0.40 per period, and so the maximum value (fundamental value) of holding the asset in period one is (number of periods \times \$0.40).
- Do traders trade in accordance with the fundamental value? Not at all! What we see is a spectacular bubble and crash! The price rises in the first few periods with relatively large amounts of trading. It reaches a price far higher than can be justified by the value of the asset, even on the most optimistic hopes of future dividends. Then, after next few periods, the volume of trade slows down and the price crashes towards zero.

Bubbles in the laboratory

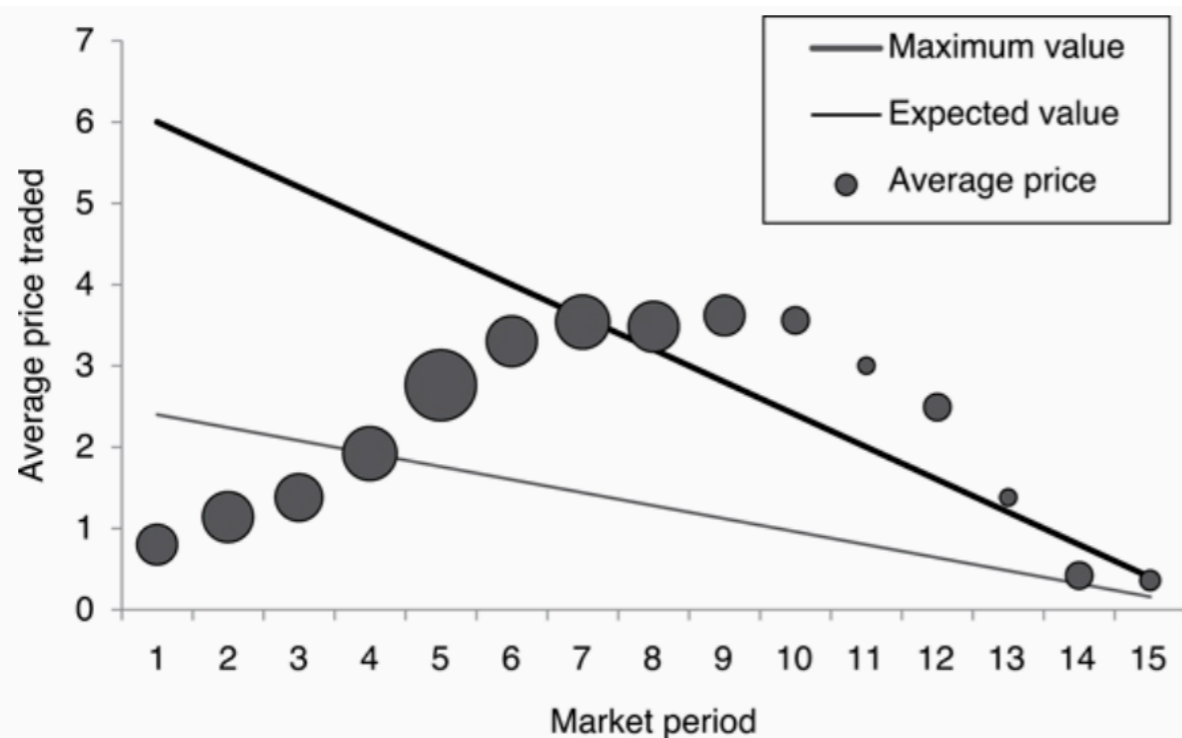


Figure 5.9 A bubble in an experimental asset market. The size of the average price circle indicates the number of trades done. The efficient markets hypothesis would suggest average price should track expected value or, at least, not go above maximum value.

Source: Smith, Suchanek and Williams (1988).

- What we observe in Figure is startling, and one of the most fascinating insights from behavioral economics to date. This is a very simple market to understand, with a minimal amount of uncertainty, yet we observe trading at prices well below and well above the fundamental value of the asset. If we get such bubble and bust in so simple an environment then no wonder we see bubble and bust in more complicated financial and goods markets.
- You might think that bubbles would stop happening if we change the market institution. For example, what if we allow short selling, when a trader can sell assets they do not own? Or what about futures trading, when traders can trade, say, period eight assets ahead of schedule? The basic conclusion seems to be that institutions sometimes dampen bubbles, but they do not stop them happening. Bubbles, it seems, are hard to stop, even in these simple experimental markets.
- Even though experience in the same environment reduces the duration of bubbles it does not reduce their amplitude. Sometimes this is enough to stop bubbles happening, but not always. Experience, therefore, does not stop bubbles happening. Indeed, nothing seems guaranteed to stop bubbles happening.

Explaining bubbles

- Having seen bubbles, we now need to try and work out why they happen. There are good reasons to think that bubbles are caused in part by biases in interpreting new information, however. The law of small numbers would suggest that, if the price has gone up three or four periods in a row, it should keep on rising. Confirmatory bias would suggest that, if people expect prices to rise, they will interpret a slight fluctuation in price as a signal that this is going to happen. Putting these and other biases together gives something called feedback theory. The basic idea of feedback theory is that prices rise because people expect them to rise, and they expect them to rise because they have been rising! A rising price translates, therefore, into a rising price.
- Feedback theory can explain bubbles, and comes about because people are biased by what they see happening to the price. What it cannot do is explain the crash. Why do prices not keep on rising? Well, there are only so many investors, and so prices cannot keep on going higher and higher, because eventually no one is left to buy. In the laboratory, a crash is usually preceded by a period with fewer trades. Basically, therefore, bubbles run out of steam. In real markets there is a less pronounced drop-off in volume than in the laboratory, presumably because real markets are not going to end after a certain number of periods, but the logic still holds.
- To give a bit more substance to the feedback theory, I want to finish with an interesting study by Barber and Odean (2008). They looked at data from individual and institutional investors to see whether the news matters. The news might matter because of various reasons, but one is the availability heuristic. When choosing where to invest, people might remember seeing a stock in the news and choose that one. If stocks get into the news because their price has increased this can cause a feedback of prices rising because prices are rising.
- One key to Barber and Odean's empirical strategy is the difference between buying and selling. When investing, the investor has thousands of stocks to choose from, but when he comes to sell, he has only those stocks he has invested in to look at, and these probably number at most 100. So, we should expect investors to be more influenced by the news when they buy than when they sell. For institutional investors this is less likely to matter, because they can short-sell stocks they do not own. More generally, we might expect professional, institutional investors to be less influenced by the news.
- The authors indeed observe that there is a large difference in the buy-sell imbalance depending on whether the stock was in the news or not that day. For institutional investors, there is no such difference. It would appear, therefore, that people are influenced by what stocks have appeared on the news when deciding where to invest. Interestingly, this effect becomes even more pronounced if the stock went down in value that day (i.e. the news was likely to be bad) than if it went up! That people are more likely to buy stocks that have appeared on the news brings us nicely on to something called the beauty contest.

Disposition Effect

Game 1: you are given \$30,000. It's yours to keep. Then you are asked to choose between the following two possibilities:

- (A) Receive an additional \$10,000 for sure
- (B) Toss a coin: if it comes up Heads, you get an additional \$20,000; if tails, you get nothing.

Game 2: you are given \$50,000. Then you are asked to choose between the following two possibilities:

- (C) A guaranteed *loss* of \$10,000
- (D) Toss a coin: if it comes up heads, you lose \$20,000; if tails, you lose nothing.

If you picked (A) in the first game, and (D) in the second, you are in very good company: this pair is the most commonly selected combination

Note, however, that the outcomes in the combination (A) and (C) are identical: in both cases, you walk away \$40,000 richer.

Similarly, (B) and (D) are identical: together, they generate a 50% chance of either \$30,000 or \$50,000.

This preference “switch” is known as the **Disposition Effect**. Why do so many people “flip” their preferences? If they selected (A) in Game 1, why not stick with (C) (which has identical outcomes in all scenarios) in Game 2?

Prospect Theory

Reference Points

