Essays in Risk

DISSERTATION

zur Erlangung des akademischen Grades
doctor rerum politicarum
(Dr. rer. pol.)
im Fach Volkswirtschaft

eingereicht an der
Wirtschaftswissenschaftlichen Fakultät I
Humboldt-Universität zu Berlin

von
Frau Dipl. Kfm. Astrid Matthey
geboren am 21.02.1978 in Jena

Präsident der Humboldt-Universität zu Berlin:
Prof. Dr. Christoph Markschies

Dekan der Wirtschaftswissenschaftlichen Fakultät I:
Prof. Dr. Oliver Günther

Gutachter:

1. Prof. Dr. Dorothea Kübler
2. Priv.-Doz. Dr. Franz Hubert

Tag der mündlichen Prüfung: 30. Januar 2007

Abstract

The dissertation consists of three chapters.

The first chapter considers a novel component of individual utility, which I term "adjustment utility". In a classroom experiment, I first show that this component of utility exists. I then develop a model to show when and in what way adjustment utility affects overall utility and economic decision making. Data on HIV infections and use of condoms in Germany shows the relevance of the results.

For the second chapter I conducted an experiment, which shows that individuals imitate intentionally, even in settings where they cannot learn anything by doing so. This complements previous experimental research, which could show that individual behavior is consistent with imitation motives, but where behavior could also be explained by learning motives. In addition, the results show that when subjects choose whom to imitate, they consider the results of other players over several periods, rather than only of the last period, as assumed in previous work.

Finally, in the third chapter, I analyze the question whether state-owned banks have a competitive advantage over private banks due to a state guarantee on their deposits. State-owned banks face a restriction of their business strategies, which is due to their mandate of "supporting economic development". As a consequence, state-owned banks cannot publicly declare to liquidate all borrowers in financial distress. This offers private banks the opportunity to separate borrowers by self-selection, enter the market and make profits in equilibrium.

Keywords:
utility from expectations, reference-dependent preferences, imitation in experiments, competition between public and private banks

Zusammenfassung

Die Dissertation besteht aus drei Kapiteln.

Im ersten Kapitel wird unter der Bezeichnung „Adjustment Utility“ eine neue Komponente individuellen Nutzens eingeführt. Mit einem Experiment, dass ich mit Studenten durchgeführt habe, zeige ich erst, dass diese Nutzenkomponente existiert. Dann entwickele ich ein Modell, welches aufzeigt, wann und in welcher Weise Adjustment Utility den Gesamtnutzen von Individuen sowie ihr ökonomisch relevantes Entscheidungsverhalten beeinflusst. Daten zu HIV Infektionen und der Verwendung von Kondomen in Deutschland zeigen die Relevanz der Modellergebnisse.

Das zweite Kapitel betrachtet ein weiteres Experiment, welches zeigt, dass Individuen absichtlich imitieren, auch in Situationen, in denen sie durch Imitation nichts lernen können. Das ergänzt die bisherige experimentelle Forschung, die zwar zeigen konnte, dass individuelles Verhalten mit Imitationsmotiven konsistent ist, bei der das beobachtete Verhalten sich jedoch auch mit genuinen Lernmotiven erklären liess. Darüber hinaus zeigen die Ergebnisse des Experiments, dass bei der Wahl dessen, den die Individuen imitieren, sie die Ergebnisse ihrer Mitspieler über mehrere Runden berücksichtigen, statt nur das Ergebnis der letzten Runde, wie in der Literatur meist angenommen.

Abschliessend analysiere ich im dritten Kapitel die Frage, ob staatliche Banken aufgrund der staatlichen Einlagengarantie einen Wettbewerbsvorteil gegenüber privaten Banken haben. Staatliche Banken unterliegen Einschränkungen ihrer Geschäftsstrategie, die durch ihr Mandat begründet sind, die wirtschaftliche Entwicklung zu unterstützen. Das heisst, dass staatliche Banken nicht öffentlich erklären können, alle Kreditnehmer, welche sich in finanziellen Schwierigkeiten befinden, dem Konkurs zu überlassen. Diese Einschränkung gibt privaten Banken die Möglichkeiten, Kreditnehmer durch Selbstselektion zu separieren, in den Markt einzutreten und sogar im Gleichgewicht Gewinne zu erwirtschaften.

Schlagwörter:
Nutzen aus Erwartungen, referenzabhängige Präferenzen, Imitation im Experiment, Wettbewerb zwischen staatlichen und privaten Banken

Widmung

Meinen Eltern

als Dank für die Wurzeln

und die Flügel.

Contents

Chapter 1

Introduction

Risk is an integral part of our life. Every day we face risks to our health, our social relations, our career, our financial situation etc. The way we deal with these risks, our strategies to reduce them, the influence they have on our wellbeing, and the opportunities they can provide for economic activity are the subjects of this thesis.

In the second chapter I use surveys to show that people get used to their expectations in the same way as they get used to their consumption, income etc. If we have a certain expectation of what our future will look like, and this expectation persists for some time, it becomes our reference point. This means that we compare other expectations to this reference expectation. In particular, if new information induces us to form new expectations, we evaluate them not only according to their absolute value, but also relative to the reference expectation as an improvement or a deterioration. This mechanism also affects our attitude towards risks. Risks induce unpleasant expectations. Getting used to these unpleasant expectations means that we get used to bearing the risk. Compared to a new risk, known risks are less frightening and we accept lower costs in order to avoid them.

This result has implications for the regulation of health and environmental risks. Since the acceptance of a new risk, e.g., the consumption of an unfamiliar, risky asset, must be expected to increase over time, regulatory measures should not be based on consumption patterns observed at the time of the introduction of the asset. Rather, they should target predicted future consumption that takes into account the change in behavior implied by people getting used to the risk.

The concept of reference dependence that the model in this chapter is based on is not new. It was introduced into economics by Kahneman and Tversky a long time ago. They used simple experiments to demonstrate that the utility people derive from a certain outcome does not only depend on its absolute size, but on the reference outcome it is compared to. But the surveys described in this essay show that reference formation is a more general concept than has been considered so far. And the extension they imply is not trivial. The effects on utility from expectations that are analyzed in the model have the potential to influence behavior and wellbeing just as much as the effects on utility from experienced outcomes like income and consumption. However, the essay also shows that the question regarding the reach of reference dependence and the exact mechanics of reference state formation is still open. The model can only contribute another stone to the mosaic that forms the answer.

The third chapter considers one of the strategies people have developed to deal with the risks they face in their daily life: imitation. Imitation is a widely observed behavior with both animals and humans. On one hand, imitation can help to reduce risks if by imitation one learns how to handle unfamiliar situations. On the other hand, imitation can help to reduce the risk of bad outcomes *while* one learns how to deal with a situation, e.g., by observing others.
The experiment described in this chapter analyzes the behavior of people in situations where imitation can help them avoid bad outcomes, but is not required for learning. Participants faced a simple task where they had to choose cells from a table. Each cell returned points, but which cell returned how many points changed according to a regular pattern. There were two types of participants: informed and uninformed, where the informed types knew more about the pattern of changes than the uninformed types. In addition to choosing cells themselves, players could also imitate other players by following their cell choice. They would then obtain the number of points that the imitated player obtained, minus a fee.
The experiment is designed to mimic market situations where some players are better informed than others, but no one can learn anything through imitation that he could not learn by pure observation. The results show that players imitate to a significant degree, and do so intentionally. They start imitating once they have found out whom they should imitate, i.e., which players receive high numbers of points. They stop imitating once they have learnt the pattern of changes themselves. When choosing a player to imitate, they do not only consider the points of the last period, but of at least four periods prior to the decision. This means that only if a player has high outcomes over several periods do the other players believe that the risk of bad outcomes is low when imitating him.

Finally, in the fourth chapter I analyze how the risk involved in lending offers private banks the chance to enter a market that would otherwise be dominated by a public (i.e., state-owned) bank. I consider the situation where there is an incumbent public bank that holds a state guarantee on its deposits, and a number of private banks that want to enter the market but face higher funding costs than the public bank. The public bank has the mandate to support the economy, while private banks are unrestricted in setting their policies. The borrowing firms have a certain risk of entering into financial distress during

the duration of the loan. There are two types of firms in the market: safe firms with a low probability of distress, and risky firms with a higher probability of distress. Some of the firms that are in distress are viable and produce output if their loan is extended. The rest of the distressed firms is unable to recover and produce output, such that extending their loans is inefficient.
I show that private banks can credibly commit to liquidating all borrowers in financial distress, and separate the borrower pool by offering a lower interest rate than the public bank. They can successfully enter the market and obtain positive profits in equilibrium. This strategy does not require that the public bank has the true objective to support the economic development. It suffices that firms have the *perception* that they will not be liquidated by the public bank at the onset of financial distress.
The mandate of the public bank helps private banks to overcome the competitive disadvantage of higher funding costs. However, the entry of the private banks, which would be expected to increase competition, may lead to higher interest rates for all firms.

I thank Dorothea Kübler for valuable comments and advice, and Franz Hubert for insightful comments and for patiently listening to all my ideas - the reasonable ones and the not so reasonable ones. The essays in this thesis also benefitted from the comments of participants in seminars at Humboldt-Universität zu Berlin, Technische Universität Berlin, Wissenschaftszentrum Berlin and UC Berkeley, as well as participants of many conferences, work shops and summerschools, e.g., the 21st Annual Congress of the European Economic Association, the Verein für Socialpolitik Annual Meeting 2006, the IAREP/SABE Congress 2006, the International Workshop on Behavioral Game Theory and Experiments 2006 in Capua, and the 14th International Congress of the International Economic Association.

Last but not least I want to thank my boyfriend and my family for being the most lively discussants of my ideas and the pilot subjects of my experiments, and for making my life the pleasure it is.

Chapter 2

Getting Used to Expectations

2.1 Introduction

Feeling disappointed when expectations change for the worse and delighted when they change for the better are emotions that are well known to most people. However, they cannot be explained with the existing models in either expected utility theory or behavioral economics. In this paper I introduce a novel concept, *adjustment utility*, which can explain the observed preferences, and yields new predictions for consumer behavior.

Consider the following two situations. In the first situation, Ms. Smith learns on Friday, February 4, that she will receive a wage increase in April, which she did not expect. However, on Monday, February 7, she learns that Payroll made a mistake and she will not receive this wage increase. In April she receives her usual wage. In the second situation, Ms. Smith is not told anything about a wage increase in February and receives her usual wage in April. In which situation is Ms. Smith happier on Monday, February 7?[1]

According to expected utility theory and behavioral economics, she is equally happy in both situations. She expects the same future outcomes, and the time until their realization is long enough such that the reference state for her wage in April can adjust to her expectations. However, most people would expect Ms. Smith to be happier on Monday if she were never told about the wage increase. In a classroom survey with 47 students, 44 ($> 90\%$) expected Ms. Smith to be happier in the second situation.[2] This means that past expectations have an influence on current utility from expectations, i.e., preferences

[1] Assume that she does not undertake consumption during the weekend which she would not have done without the wage increase.

[2] I conducted this survey in January 2006 with participants of a Master course in Game Theory at the Technical University Berlin. Most of the students in the course study industrial engineering. One student expected Ms. Smith to be happier in the first situation, for two she was equally happy in both. See the translated questionnaire in appendix B for details.

are path-dependent with respect to expectations. The state of Ms. Smith's expectations over the weekend influences the utility she derives from her expectations on Monday.

I account for such preferences by introducing reference-dependent utility from expected future outcomes. Since this component of utility is caused by the adjustment of expectations, I term it *adjustment utility*. It explains individual behavior like that observed in the survey.

The model focuses on situations where the formation of expectations over future outcomes and the actual realization of those outcomes are temporally separated. This applies, first, to situations in which people form expectations over a consumption decision prior to the actual decision. For example, people form expectations over the purchase of seasonal goods like summer shoes prior to actually buying them, and over the consumption of goods bought in catalogues prior to actually receiving them. Second, it applies to situations in which the consumption of the asset takes place prior to the realization of outcomes. This includes assets and actions whose consumption causes future health risks, like the risk of an HIV infection through unsafe sex, or the risk of adverse health effects from eating genetically-modified food. It also includes assets and actions which cause delayed benefits, like consuming healthy food and exercising.

The survey shows that past changes in expectations are relevant for current utility, which means that adjusting expectations is not utility-neutral. This implies that the reference state for expectations can never be fully endogenized. The reason is that for any economically relevant asset, there is at least one moment in which the individual's expectation regarding the asset changes. Even in the case of immediate full information that remains the same until final consumption, expectations change from initial unawareness to awareness of the future acts of consumption. After information is received, during the adjustment of the reference state from past to current expectations, the reference state is (partly) exogenous, since initial unawareness is exogenous.

Regarding the formation of the reference state, the survey shows the relevance of one's own expected future state. A second survey shows a similar influence of the social comparison component. I again asked students to evaluate two situations. In the first situation, Ms. Smith is head of a department at her company and knows already in August 2006 that she will earn 50.000 EUR in 2007. She expects the other heads of department at her company to earn 40.000 EUR on average. On October 1, she learns that the other heads will also earn 50.000 EUR on average. In the second situation, she expects the other heads of department to earn 60.000 EUR on average, but then also learns that they will earn 50.000 EUR. In which situation is Ms. Smith happier on October 1? Again, according to received theory, there is no difference in her utility, since she expects the same future outcomes and her reference state has plenty of time to adjust to her expectations until outcomes are realized. However, in the survey 61 out of 78 students (78%) expected Ms. Smith to be happier in the second situation.[3]

These results allow me to qualitatively specify the reference state for expectations to depend positively on the individual's past expectations regarding her own future state, and her expectations regarding the future state of other members of her reference group.[4] Even though these two qualitative relations do not characterize the reference formation process in as much detail as one would hope, they are all one can take from the surveys. They are also sufficient to derive results on people's utility, and their behavior towards assets and actions that induce future outcomes.

To derive the model's predictions, I analyze the equilibrium strategies in the spirit of the personal equilibrium that Koszegi and Rabin (2006) derive for realized outcomes. It is shown that when considering adjustment utility, an individual's strategy can in equilibrium lead to choices opposite to those that would obtain in the personal equilibrium without adjustment utility. In particular, in situations where the individual would have chosen not to consume a certain asset in the personal equilibrium without utility from expectations, with adjustment utility it may be her unique equilibrium strategy to consume the asset.

In general, reference formation regarding expectations leads to a change in the utility from assets that induce outcomes in the future, even if no change in the parameters of the assets occurs. For example, people are less frightened by the risk of adverse health effects from genetically-modified food once they get used to consuming such food, i.e., once their reference state adjusts towards the effects that may result from its consumption. As a result, sophisticated individuals who are able to predict their future reference states derive higher utility from consuming an asset that induces negative expected outcomes than naive individuals who are unable to predict their reference states. For both types of individuals, there evolves a tendency of increased risk acceptance in the population without a change in the exogenous parameters.

People getting used to their expectations and forming reference states regarding them offers an explanation, e.g., for the increasing acceptance of the risk of an HIV infection in Germany since the late 1990s. The data shows that people are at least as well informed about the risk of an HIV infection and the protective effect of condoms as before.[5]

[3]The survey was conducted at the Technical University Berlin in May 2006. Again, most of the students are industrial engineers. No student participated in both surveys. Note that it cannot be determined for either survey whether the students interpreted Ms. Smith to believe in the information or her expectation to 100% or less. This does not, however, effect the argument.

[4]I define the reference group of an individual endogenously as the group of people who's utility is affected by the choices of this individual, and who's choices affect the utility of the individual.

[5]Data from the Robert Koch Institute Berlin and Bundeszentrale für gesundheitliche Aufklärung. See

However, the use of condoms has decreased slightly but significantly between 1997 and 2004, which provides one reason for the increase in HIV infections in recent years. Since the data show that lack of care due to the availability of new therapies cannot explain this effect, gradual reference point adjustment towards the risk of an infection offers a plausible explanation. Over several years, people may have themselves accepted the risk of an infection occasionally and have learnt of others accepting the risk, e.g., through the media. This means that the risk has - to a small but significant degree - been included in people's reference states. Accordingly, the disutility from accepting the risk has decreased (the utility has increased). This means that less (non-monetary) costs are accepted in order to avoid the risk, leading to a decrease in the use of condoms.

In addition to explaining observed preferences and attitudes towards risks, the model yields implications for policy. It shows that the release of positive but dubious information may decrease welfare. Contrary to existing models, when accounting for adjustment utility, the release and subsequent correction of positive information may lead to negative net utility. This is the case if the initial delight after the release of the information is overcompensated by the disappointment caused by the correction. The more loss averse individuals are, the more likely is a negative overall effect.

Further, if regulatory measures of newly introduced assets that involve future health or environmental risks are based on current rate of consumption estimations without accounting for changes in the reference state, they will underestimate the extent of consumption in the medium to long-term. Accordingly, either will production and consumption of the asset exceed the socially optimal level, or regulation will have to be adapted later on, which may be costly and/or hard to implement.

In section 2.2, I briefly describe reference-dependent preferences and utility from expectations. The concept of adjustment utility, the definition of the reference state for expectations, and the expectation personal equilibrium are developed in section 2.3. I derive the change in utility that an individual experiences when repeatedly consuming a risky asset in section 2.4. The individual's choice between the risky and a comparable safe asset is analyzed in the context of her overall consumption optimization in section 2.5. Section 2.6 derives the value of information, while section 2.7 shows that the reference state cannot be endogenized in general. Finally, section 2.8 derives further implications and concludes.

table 1 in appendix A for details.

2.2 Reference-dependent preferences and utility from expectations

Consider an asset which induces a set of outcomes $x \in X$ with distribution function $f(x)$ some time after consumption.[6] The asset is consumed in period t and outcomes are realized in period T, with $T > t$. The asset may also induce outcomes immediately upon consumption, but since they are not relevant for the analysis, they are normalized to zero.[7] Where no confusion can arise, I will abbreviate the asset with F.

Individuals derive reference-dependent utility from the outcomes induced by asset F. Utility in T consists of the reference-independent part $u^x(x)$, and the reference-dependent part $v^x(x|r^x)$ with r^x as the reference state for realized outcomes.[8] I will call the utility U^x that the individual derives from the realization of x *outcome utility*: $U^x(x) = u^x(x) + v^x(x|r^x)$.

The reference-dependent outcome utility fulfills Kahneman and Tversky (1979) conditions for the relative value function, i.e., $v(x|r)$ is defined as a finitely valued function $v : X \times X \to \mathbf{R}$, and

A1: $v(x|x) = 0$ for all x.

A2: $v(x|r^x)$ is continuous for all x and twice differentiable for $x \neq r^x$;

A3: $v'(x|r^x) > 0$ for all x given r^x.

$v(x|r^x)$ is to be interpreted as a measure of the value of x relative to r^x when both are viewed from the reference point r^x. In any period $s < T$, the individual then expects the utility $U^x(f(x), r^x) = u^x(f(x)) + v^x(f(x)|r^x)$ to be experienced in period T, with r^x as the expected reference state for outcomes in T.

There exists both experimental and empirical evidence for the reference dependence of individuals' outcome utility. It shows that the majority of people evaluate outcomes relative to a reference point in many spheres of utility, e.g., for consumption levels and wages (see, e.g., Kahneman et al. (1991), for experiments on the endowment effect, or Benartzi and Thaler (1993), for an explanation of the *equity premium puzzle* for loss averse investors).

However, individuals do not only derive outcome utility from the consumption of asset F. Rather, from expecting $f(x)$ to be realized in T they derive utility in each period s with $t \leq s < T$, i.e., in each period between the consumption of asset F and just prior to the realization of outcomes. Utility from expectations is experienced whenever an individual

[6] The degenerate lottery with $f(x_0) = 1$ for a certain outcome x_0 is included.

[7] Since I do not consider preference reversals over time here, e.g., quasi-hyperbolic discounting, having non-monetary outcomes at two moments in time would not add to the analysis.

[8] I use the terms reference state and reference point interchangeably.

looks forward to expected future events (see, e.g., Loewenstein (1987) for the utility from anxiety and pleasant anticipation). In an economic context, utility from expectations has been considered by, e.g., Caplin and Leahy (2001), Koszegi (2005) and Brunnermeier and Parker (95). There, *anticipatory utility* reflects the reference-independent utility people derive from expecting future outcomes. These models show that anticipatory feelings influence economic behavior and are relevant for utility.

I define the anticipatory utility that is experienced in all periods s with $t \leq s < T$ as $u_s^a = u_s^a(f(x), r^x)$, i.e., it depends on the utility individuals expect to experience in T. Accordingly, it depends on the outcome distribution and the reference state they expect in T. $u^a(f(x), r^x)$ is assumed to be increasing in $f(x)$, decreasing in r^x, continuous and twice differentiable in both dimensions. For simplicity, I assume that $u^a(f(x), r^x) > 0 \Leftrightarrow U^x(f(x), r^x) > 0$ and $u^a(f(x), r^x) < 0 \Leftrightarrow U^x(f(x), r^x) < 0$, i.e., assets which induce negative expected utility also induce negative anticipatory utility, while assets which induce positive expected utility induce positive anticipatory utility.

Combining the concepts of reference dependence and anticipatory utility then yields the expected utility from consuming asset F as

$$\hat{U}(f(x)) = \sum_{s=t}^{T-1} u_s^a(f(x), r^x) + u^x(f(x)) + v^x(f(x)|r^x) \tag{2.1}$$

where anticipatory utility u^a is experienced in each period $t \leq s < T$ and outcome utility $u^x + v^x$ is experienced only in T. For simplicity, the discount factor is set to $\delta = 1$.

An alternative scenario to the one considered here is where the individual in period t only *expects* to consume asset F in T, instead of consuming it. Upon consumption in T, however, she experiences outcomes immediately. In this case, the delay is in consumption, rather than in the realization of outcomes.[9] Formally, these two scenarios are identical for the purposes of this model, and all the results hold for both cases. For reasons of expositional simplicity, I will frame the model only in terms of the first scenario.

2.3 Adjustment utility

2.3.1 Definition

Reference dependence and anticipatory utility are not always able to explain observed preferences. To see this, consider again the two situations in the first survey. Regarding anticipatory utility, Ms. Smith is expecting the same future wage in both situations, such

[9] This is the setting in which Koszegi and Rabin (2006) develop their model of reference-dependent outcome utility.

that the same utility results.[10] Regarding reference-dependent outcome utility, in the first situation she expects a higher wage between Friday and Sunday, but then expects the lower wage from Monday onward. Reference-dependent outcome utility would only yield a difference in utility if these three days in early February had a lasting effect on Ms. Smith' reference state in April. This seems implausible. It also contrasts with the literature. For example, if expectations form the reference state for outcome utility as in the model of Koszegi and Rabin (2006), one would expect that in the two months that remain until April Ms. Smith' reference point for her wage has sufficient time to adjust to her wage expectation. Accordingly, she derives the same expected reference-dependent outcome utility from her wage in April as if she had always expected her usual wage.[11]

With reference dependent utility from expected outcomes, preferences as expressed in the survey can be explained. For simplicity, in the further analysis I will follow Koszegi and Rabin (2006) and assume that for all individuals $E(r^x) = E(x)$, i.e., all individuals assume that between periods t and T their reference point for outcomes in T adjusts to their expectations regarding those outcomes. This allows me to abstract from the dependence of the utility from expectations on the expected future reference state r^x, i.e., to write $u^a(f(x))$ rather than $u^a(f(x), r^x)$.

Definition 1 *Let $f(x)$ denote a lottery over the set of possible outcomes X induced by asset F and $\boldsymbol{f}(x)$ the set of possible lotteries over X.*
The adjustment utility *that an individual derives from expecting lottery $f(x) \in \boldsymbol{f}(x)$ given her reference lottery $r^a \in \boldsymbol{f}(x)$ is defined as $v^a(f(x)|r^a)$, where v^a denotes a finitely valued function $v^a : \boldsymbol{f}(x) \times \boldsymbol{f}(x) \to \boldsymbol{R}$ and*

A4: *$v^a(f(x)|f(x)) = 0$ for all $f(x)$.*

A5: *$v^a(f(x)|r^a)$ is continuous for all $f(x)$ and twice differentiable for $f(x) \neq r^a$;*

A6: *$v^{a\prime}(f(x)|r^a) > 0$ for all $f(x)$ given r^a.*

In analogy to Kahneman and Tversky's (1978) value function, adjustment utility $v^a(f(x)|r^a)$ is to be interpreted as a measure of the value of expecting lottery $f(x)$ relative to expecting the reference lottery r^a, when both are viewed from the expectation of the reference lottery r^a.

The utility from expecting the lottery implied by asset F then consists of anticipatory utility $u^a(f(x))$ and adjustment utility $v^a(f(x)|r^a)$[12]:

$$U^a(f(x)) = u^a(f(x)) + v^a(f(x)|r^a) \qquad .$$

[10] This also holds if anticipatory utility depends on expected future reference states, see the argument below.

[11] A more detailed analysis of the argument that Ms. Smith simply fails to correctly predict her reference state, which may seem to explain the data at first sight, is relegated to appendix C.

[12] To be correct, one would have to write $v^a(f(x)|r^x, r^a)$ since adjustment utility also depends on the

Adjustment utility captures the observation that the pleasure or pain people derive from expecting future outcomes depends on what they compare them to, not only on their absolute amount. If you expected a vacation in the Himalayas next summer, and then learn that you will only go to the Alps, you will initially derive different utility from anticipating this vacation than if you had expected to stay at home, even if you are aware that your reference state for outcome utility will adjust to the vacation in the Alps in both cases until next summer.

Instead of the utility in (2.1), the individual's utility from consuming asset F is therefore given as

$$U(f(x)) = \sum_{s=t}^{T-1} [u^a_s(f(x)) + v^a_s(f(x)|r^a)] + u^x(f(x)) + v^x(f(x)|r^x). \tag{2.2}$$

Anticipatory utility $u^a(f(x))$ and adjustment utility $v^a(f(x)|r^a)$ are experienced in all periods between the consumption of the asset and the period just prior to the realization of outcomes. When outcomes are realized, reference-independent outcome utility u^x and reference-dependent outcome utility v^x are experienced. Note again the important difference between the two reference states: r^x is the individual's reference for the outcomes she experiences in T, i.e., a reference *outcome*. r^a is the reference for the expectations she experiences in all periods up to $T-1$, i.e., a reference *expectation*. $E(r^x)$ is the reference outcome the individual predicts for T, while $E(r^a)$ is the reference expectation she predicts to have at some time prior to T. Even though there is probably a connection between the two, they are never relevant for experienced utility at the same time.

The sequence of events in each period is as follows. If there is new information regarding the asset it becomes available first. Then the person experiences utility (from either expected or realized outcomes, depending on the period) given the available information and her current reference state. Next she observes others' behavior. Finally, her reference state for the next period is formed.

2.3.2 Reference state formation

From the surveys there emerge two factors that influence the reference state. The first is individuals' past expectation of their own state: Ms. Smith' expected April wage at the weekend influenced her reference wage expectation on Monday. The second is the expectation of others' states: Ms. Smith's expectation in September 2006 of her colleagues' wages in 2007 influenced her reference expectation regarding their wages on October 1, 2006.

---

expected utility in T, which depends on r^x. However, since I assume $r^x = E(x)$ throughout, I drop the additional notation.

This parallels the findings for outcome utility. There the reference state has been found to be influenced by the individual's own past consumption level, wage, investment return etc. (e.g., Kahneman et al. (1990), Campbell and Cochrane (1999) and by relevant others' consumption levels, wages etc. (e.g., Abel (1990), Constantinides (1990).

The surveys do not, however, provide sufficient insight to infer a quantitative relationship between these two factors, or even to quantify their influence on the reference state.[13] Accordingly, I will not specify an explicit functional form of the reference state, but limit the specification to qualitative relations. I also take account of the fact that events further in the past are less present in people's minds than the same events in more recent periods. Accordingly, they must be expected to be less relevant for the formation of the reference state and will receive a "mental discount".

Let $f_{i,t}(x)$ denote the expectation person i has in period t of her own future state regarding asset F, and $\mathbf{f}_{-i,t}(x)$ denote the set of expectations she has of the future states of the members of her reference group. β denotes the mental discount factor, which may differ across assets and individuals: $0 \leq \beta \leq 1$. Considering a sequence of past expectations and weigthting it with a discount factor smoothes the reference formation process. It takes account of the fact that reference formation in most cases must be expected to be a gradual process. After a change in expectations, the reference state may take some time to adjust to the new situation. In contrast, considering only last period's expectations would lead to jumps in the process, i.e., ad hoc adjustments of the reference state. In addition, through adjusting β properly, the process does not depend on the definition of the length of a period.

Another principle that I assume the reference formation process to observe is *convergence*. If expectations regarding one's own future state and the states of relevant others are constant for a long time at f_i and $\mathbf{f}_{-i}$, the reference state for adjustment utility converges towards a state that reflects only these expectations, i.e., towards the reference state that would result if it only depended on last period's expectations f_i and $\mathbf{f}_{-i}$.

Definition 2 *The reference state for the adjustment utility of individual i in period t is defined as*

$$r_{i,t}^{a} = r_{i,t}^{a}\left(\beta_{t-s}, f_{i,s}(x), \boldsymbol{f}_{-i,s}(x)\right)_{s<t} \tag{2.3}$$

with

$$\frac{\partial r_{i,t}^{a}}{\partial f_{i,s}(x)} > 0 \qquad \text{and} \qquad \frac{\partial r_{i,t}^{a}}{\partial \boldsymbol{f}_{-i,s}(x)} > 0 \tag{2.4}$$

[13] This is, unfortunately, a problem this research shares with many others' on the issue of reference state formation.

$$\frac{\partial \beta_{t-s}}{\partial s} > 0 \qquad and \qquad \beta_1 = 1$$

$$r^a(\beta, 0, 0) \quad = \quad 0 \tag{2.5}$$

$$lim_{t-s\rightarrow\infty}\beta_{t-s} \quad = \quad 0 \tag{2.6}$$

$$lim_{t-s\rightarrow\infty} \; r^a_{i,t}\left(\beta_{t-s}, f_i, \boldsymbol{f}_{-i}\right)_{s<t} = r^a_{i,t}\left(f_i, \boldsymbol{f}_{-i}\right) \tag{2.7}$$

The term $r^a_{i,t}\left(\beta_{t-1-s}, f_{i,s}(x), \mathbf{f}_{-i,s}(x)\right)_{s<t}$ captures that the expectations of all periods prior to period t are taken into account, weighted with their respective discount factor. Note that (2.3) and (2.4) are applied to each line of the probability vector separately.

An example of a functional form of the reference state that complies with this definition is

$$r^a_{i,t} = q\left(\sum_{s<t}\frac{\beta^{t-s}}{\sum_{s<t}\beta^{t-s}} f_{i,s}(x)\right) + (1-q)\left(\sum_{s<t}\frac{\beta^{t-s}}{\sum_{s<t}\beta^{t-s}}\mathbf{f}_{-i,s}(x)\right)$$

with $0 < q < 1$ as the weighting factor for individual vs. social comparison. To illustrate the formation of the reference state, consider again the situations in the surveys. Equation (2.3) and the inequalities in (2.4) imply that the reference state to which an individual compares her wage expectation increases if she expects a higher wage and if others expect a higher wage. If no one ever expects to receive a wage, the individual's reference state for her wage expectation is zero (equation (2.5)). Equation (2.6) captures the property that the wage expectations the individual and the members of her reference group had a long time ago have a negligible influence on her current reference state. Finally, with convergence (2.7), if the individual and relevant others expect the same wage for a long time, the individual's reference state converges towards the reference state that would result if only last period's wage expectations were considered.

2.3.3 Expectation Personal Equilibrium

In this section, I apply the concept of Koszegi (2005) *personal equilibrium* (PE) to the model of adjustment utility and consider the predictions this yields. Consider an individual who has rational expectations and is aware of her reference formation process. At any time t, she forms expectations regarding the asset's outcomes in T given her behavior in t. From these expectations she can predict her future reference states r^a_s for all periods $t < s < T$ according to (2.3).[14] Knowing her reference states she can derive her adjustment utility. Considering all four components of the individual's overall utility as defined in (2.2), she then chooses her optimal behavior in t. If this behavior differs from what she assumed initially, it induces new expectations regarding her outcomes in T, her reference

[14] I omit the indices i and s where no confusion can arise.

state r_s^a, her expected utility, and hence her (updated) optimal behavior. This process continues until it reaches a state where the behavior from which she derives her expectations and the optimal behavior that result from them are identical. Depending on what behavior she assumed initially, this fixed point problem may yield several equilibria. It is similar to the optimization problem in Koszegi and Rabin (2006) for outcome utility. It includes two loops. The first one is the feedback loop from the individual's behavior into her reference state and from her reference state into her behavior. This loop includes the (potentially) exogenous r^a and the subsequent reference formation process. The second is the social comparison feedback loop from the individual's behavior to the reference state and behavior of others, and from the behavior of others back to the individual's reference state and behavior. Solving the fixed point problem, therefore, requires the assumption of common knowledge of the reference formation process and utility of all members of the reference group. Even though this may seem demanding, it is common in the literature and useful in order to be able to compare my results with those of earlier models. It shows where the inclusion of adjustment utility yields novel predictions. I briefly consider the case when this assumption is violated, i.e., naivety, in the next paragraph.

Definition 3 *Let $r^a(f(x))$ and $r^x(f(x))$ denote the reference states for adjustment utility and outcome utility, respectively, that result from choosing the lottery $f(x)$. A choice $f(x) \in \boldsymbol{f}(x)$ forms an* expectation personal equilibrium *if for all possible choices $f'(x) \in \boldsymbol{f}(x)$*

$$\sum_{s=t}^{T-1}[u^a(f(x)) + v^a(f(x)|r_s^a(f(x)))] + u^x(f(x)) + v^x(f(x)|r^x(f(x))) \geq$$
$$\sum_{s=t}^{T-1}[u^a(f'(x)) + v^a(f'(x)|r_s^a(f'(x)))] + u^x(f'(x)) + v^x(f'(x)|r^x(f'(x)))$$

In words, an individual's choice forms an expectation personal equilibrium if, given her reference state for adjustment utility and the reference formation process of her as well as of others given this choice, she derives utility from this choice which is at least as high as that derived from any other possible choice and the reference formation processes that result from this choice.

When does an expectation PE yield predictions different from Koszegi and Rabin (2006) outcome PE? For ease of comparison, consider their example of a shoe purchase. Assume that an individual is used to buying new shoes in spring. Towards the end of winter, she receives the catalogue for shoes that will be available in spring. But this year she realizes from studying the catalogue that shoe prices have increased and new shoes may be beyond her means. Assume that if the individual considers only outcome utility, it is her unique outcome personal equilibrium strategy not to buy the shoes in spring. She makes this decision predicting that her reference state for the purchase will adjust to her expectation of not buying the shoes until spring actually comes, such that not buying shoes is not felt as a loss then. Now add the utility from expectations. First, taking into account anticipatory utility makes the purchase of the shoes more attractive, since the individual can look forward to the purchase and derive utility from that. Second, and more important here, including adjustment utility makes the purchase of the shoes more attractive since it avoids the disappointment the individual would experience from adjusting her expectations towards expecting not to buy the shoes since she initially expected - without detailed knowledge of the market - to buy shoes as every year. Her reference state in winter, which included the expectation to buy shoes in spring, makes the expectation of not buying the shoes be felt as a loss. This is different from the utility in spring, for which the individual has time to get used to the idea of not buying the shoes and may not feel disappointed. In summary, if the reference state in winter is the expectation to buy the shoes, e.g., because the individual is used to buying shoes in spring, adjustment utility increases the utility from buying the shoes. This means that there exist cases where it is the individual's unique strategy in Kőszegi and Rabin's outcome personal equilibrium not to purchase the shoes, but her unique strategy in the expectation personal equilibrium to purchase the shoes.

Proposition 1 *Consider an individual who has the choice between $f(x) \in \boldsymbol{f}(x)$ and $f'(x) \in \boldsymbol{f}(x)$. Then, if*

$$\sum_{s=t}^{T-1} [v^a(f(x)|r_s^a(f(x))) - v^a(f'(x)|r_s^a(f'(x)))] <$$

$$\sum_{s=t}^{T-1} [u^a(f'(x)) - u^a(f(x))] + u^x(f'(x)) - u^x(f(x)) + v^x(f'(x)|r^x(f'(x)) - v^x(f(x)|r^x(f(x)))$$

the individual's strategy in expectation personal equilibrium is $f'(x)$, while the strategy in outcome personal equilibrium is $f(x)$.

Proposition 1 results directly from the definition of an expectation PE. It shows that if the difference in aggregated adjustment utility between two choices is larger than the difference in anticipatory and outcome utility and of the opposite sign, the choice predicted by the expectation PE is opposite to the one predicted by the outcome PE, even if one accounts for anticipatory utility in both equilibria.

2.3.4 Naivety vs. Sophistication

An individual that is aware of her reference formation process knows that the decision to consume the asset F involves a change in her reference state for expectations, which induces a change in utility. When assessing the utility she derives from the consumption of the asset, she takes account of this effect. In particular, if she consumes an asset which induces negative expected future utility, she knows that her reference state will decrease and she will derive higher utility from expecting the asset's outcomes in the periods up to their realization. Similarly, if she consumes an asset which induces positive expected future utility, she knows that her reference state will increase and she will derive lower utility from expecting the asset's outcomes.

An individual that is not aware of her reference formation process does not predict the effect of the reference point change on her utility and hence assumes that she will experience the same utility in each period up to T.
Accordingly, a sophisticated individual assigns a higher overall utility to the consumption of an asset with negative expected future utility than a naive individual, but a lower overall utility to an asset with positive expected future utility.

2.4 Repeated consumption

So far I have analyzed the consumption of the asset only in one period. In this section, I derive the change in the individual's net utility from asset F when she considers its consumption repeatedly. I begin in section 2.4.1 with analyzing the utility of consuming F at different stages of the reference formation process. For the same reference formation process, in section 2.4.2 I derive the utility of *not* consuming F, i.e., the utility of not expecting its future outcomes. The comparison of the two utilities shows the change in the attraction of F when the reference state changes.

Finally, rather than analyzing F in isolation as in sections 2.3 and 2.4, in section 2.5 I introduce prices and extend the analysis to the individual's overall consumption optimization. In addition to F, the individual can now choose to consume an alternative asset, which does not induce future outcomes but differs from F in its price. This means that since the choice of F also influences the individual's ability to consume other goods, her reference state for these goods has to be taken into account. This finally allows me to derive the individual's decision regarding F from her overall consumption optimization program.

2.4.1 Utility of consuming the asset

Consider now the repeated consumption of asset F. For example, an individual may repeatedly eat genetically-modified food and after a while get used to expecting possible adverse health effects from this in the future. I analyze utility starting in period t, where

t is the period such that before period t no member of the reference group has ever consumed the asset, but in t at least one individual consumes it. However, the argument that is developed below holds for all situations where r^a has not yet fully converged to the state that reflects permanent consumption of F by all members of the reference group.

The results of the previous sections were derived assuming only that people's preferences were reference-dependent and decreasing in the reference state. In order to obtain quantitative results, I now add the assumption that individuals are loss averse for both realized and expected outcomes, i.e., their relative value function for realized outcomes and their adjustment utility function are steeper for losses than for gains.

Compare the individual's utility from consuming the asset for the first time in period t to consuming it the second time in period $t+1$. If consumed in t the asset induces outcomes in T, while if consumed in $t+1$, it induces outcomes in $T+1$. I do not allow for any new information regarding the asset to become available in $t+1$, i.e., I keep $f(x)$ constant. This means that $u^a_s(f(x))$, $u^x(f(x))$ and $v^x(f(x)|r^x)$ are constant in all periods until $T+1$.

Note that, in order to account for the accumulation of expected outcomes, r^a_s now explicitly refers to all relevant expected outcomes of asset F. Considering only decisions in period t and $t+1$, this includes outcomes in T and $T+1$. The number of outcomes could, however, be extended to include any number of relevant expectations. Formally, this means that[15]

$$\begin{aligned} r^a_{i,t} &= r^a_{i,t}\left(\beta_{t-s}, f_{i,s}(x_T), \mathbf{f}_{-i,s}(x_T), f_{i,s}(x_{T+1}), \mathbf{f}_{-i,s}(x_{T+1})\right)_{s<t} \quad \text{and} \\ r^a_{i,t+1} &= r^a_{i,t+1}\left(\beta_{t+1-s}, f_{i,s}(x_T), \mathbf{f}_{-i,s}(x_T), f_{i,s}(x_{T+1}), \mathbf{f}_{-i,s}(x_{T+1})\right)_{s<t+1} \quad . \end{aligned}$$

Consider the individual in isolation first, i.e., ignore the effects of social comparison. For simplicity, and since it would not add any new insights, I again abstract from discounting ($\delta = 1$). Since I only deal with one individual here, I drop the index i. The difference in utility is then given as

$$\begin{aligned} U_{t+1}(f(x)) - U_t(f(x)) &= \sum_{s=t+1}^{T}[u^a_s(f(x)) + v^a_s(f(x)|r^a_s)] + u^x_{T+1}(f(x)) + v^x_{T+1}(f(x)|r^x) \\ &- \sum_{s=t}^{T-1}[u^a_s(f(x)) + v^a_s(f(x)|r^a_s)] - u^x_T(f(x)) - v^x_T(f(x)|r^x) \quad . \end{aligned}$$

Since $u^x(f(x))$ and $v^x(f(x)|r^x)$ are the same in T and $T+1$, $u^x_T + v^x_T$ and $u^x_{T+1} + v^x_{T+1}$ offset each other:

$$U_{t+1}(f(x)) - U_t(f(x)) = \sum_{s=t+1}^{T}[u^a_s(f(x)) + v^a_s(f(x)|r^a_s)] - \sum_{s=t}^{T-1}[u^a_s(f(x)) + v^a_s(f(x)|r^a_s)] \quad .$$

[15] For the general case when n periods are considered, this extends to $r^a_{i,t} = r^a_{i,t}\left(\beta_{t-s}, \sum_n f_{i,s}(x_{T+n}), \sum_n \mathbf{f}_{-i,s}(x_{T+n})\right)_{s<t}$ where the interval $[T, T+n]$ includes all periods with relevant expected outcomes.

Both terms above contain $(T-t-1)$ terms $u^a_s(f(x))$, which, given that $u^a_s(f(x))$ is constant in all periods, is the same such that

$$U_{t+1}(f(x)) - U_t(f(x)) \;=\; \sum_{s=t+1}^{T} v^a_s(f(x)|r^a_s) - \sum_{s=t}^{T-1} v^a_s(f(x)|r^a_s) \quad .$$

Since r^a_s is relevant for all assets in period s, i.e., the reference state is the same for all assets F, reference-dependent anticipatory utilities of the two assets between $t+1$ and $T-1$ drop out, yielding that

$$U_{t+1}(f(x)) - U_t(f(x)) \;=\; v^a_T(f(x)|r^a_T) - v^a_t(f(x)|r^a_t) \quad .$$

The individual did not consume asset F before t. Hence, $f(x)_{s<t} = \mathbf{0}$ and, abstracting from social comparison effects, $r^a_{s\leq t} = \mathbf{0}$.[16] After consuming the asset in t, $f(x)_{t\leq s<T} \neq \mathbf{0}$ and r^a adjusts towards $f(x)$ according to (2.3). The more time between t and s, i.e., the further the reference formation process has proceeded, the further away from zero is r^a, since for $\beta < 1$ the weight of $f(x)_{s<t} = \mathbf{0}$ decreases while the weight of $f(x)_{s\geq t} \neq \mathbf{0}$ increases.

Hence, if the individual consumes the asset repeatedly, her r^a decreases over time if the asset induces negative expectations and increases if it induces positive expectations. This means that:

$$\begin{aligned} U_{t+1}(f(x)) \;&>\; U_t(f(x)) \quad \text{if } f(x) \prec \mathbf{0} \\ U_{t+1}(f(x)) \;&<\; U_t(f(x)) \quad \text{if } f(x) \succ \mathbf{0} \quad . \end{aligned}$$

with the initial reference state $r^a = \mathbf{0}$. In general, $U_{t+1}(f(x)) > U_t(f(x))$ if $f(x) \prec r^a$ and $U_{t+1}(f(x)) < U_t(f(x))$ if $f(x) \succ r^a$.

Consider now the process for individuals with social comparison preferences, where $\mathbf{f}_{-\cdot}(x)$ denotes the expectations of the respective individual regarding the outcomes of all individuals of her reference group. As before, $r^a_t = \mathbf{0}$ at time t for all individuals in the reference group, and at least one individual consumes F in t. This leads to $r^a_{t+1} \prec r^a_t$ if the asset induces $f(x) \prec 0$ and $r^a_{t+1} \succ r^a_t$ if it induces $f(x) \succ 0$ for all individuals in the group. Since for all individuals $\mathbf{f}_{-\cdot}(x) \neq 0$ in all periods up to $T-1$, $r^a_T \neq r^a_t = \mathbf{0}$, i.e. $r^a_T \prec r^a_t$ if $\mathbf{f}_{-\cdot}(x) \prec 0$ and $r^a_T \succ r^a_t$ if $\mathbf{f}_{-\cdot}(x) \succ 0$, for all individuals in the reference group. Then, if an individual considers consuming F in $t+1$, independently of the predicted future behavior of herself and relevant others, it results that

$$\begin{aligned} U_{t+1}(f(x)) \;&>\; U_t(f(x)) \quad \text{if } f(x) \prec \mathbf{0} \\ U_{t+1}(f(x)) \;&<\; U_t(f(x)) \quad \text{if } f(x) \succ \mathbf{0} \end{aligned}$$

for all individuals in the reference group. These results are summarized in

[16]Since $f(x)$ and r^a are vectors, $\mathbf{0}$ denotes the zero vector.

Proposition 2 *Consider an individual that in period t starts consuming an asset which induces future expected outcomes. The utility that this individual and the other members of her reference group derive from consuming the asset in period $t+1$ increases if the asset induces negative expected utility, but decreases if it induces positive expected utility.*

For example, the utility people derive from consuming genetically modified food increases the more often they consume it. In contrast, the utility they derive from organic food decreases the more they get used to it.

2.4.2 Utility of not consuming the asset

For the decision of whether to consume the asset or not, the effect of reference state formation on the alternative utility of not consuming the asset is equally important, since eventually the difference in utility is decisive. This utility is affected by r^a in the same way as is the utility of consuming the asset. Consider an asset which induces negative expected future outcomes, $f(x) \prec 0$. The utility of not consuming the asset increases if it was consumed in the past, since with a lower r^a a utility of zero feels like a higher gain. Starting from $r_t^a = \mathbf{0}$, a marginal reference point change to $r_{t+1}^a \prec 0$ leads to a state where not consuming the asset is felt as a gain (from neither gain nor loss for $r_t^a = \mathbf{0}$) and the consumption of the asset is felt as a marginally smaller loss. Considering the entire reference formation process, for loss averse individuals a change in the reference point from $r^a = 0$ to $r^a = f(x) \prec 0$ leads to a higher increase in the utility of consuming the asset ($\Delta_C = v^a(f(x)|r_2^a) - v^a(f(x)|r_1^a)$ than in the utility of not consuming it ($\Delta_{NC} = v^a(0|r_2^a) - v^a(0|r_1^a)$, see figure 2.1). Hence, the consumption of the asset becomes more attractive.[17]

For assets which yield positive expected outcomes, e.g., health benefits from exercising, healthy lifestyle etc., the utility of not consuming the asset decreases with an increasing reference point, and since it is felt as a loss, for loss averse individuals it decreases faster initially than the utility of consuming F. Accordingly, the relative utility of consuming the asset compared to not consuming the asset increases. Hence, even though section 2.4.1 showed that the utility from the asset decreases, consuming the asset nevertheless becomes *more attractive* relative to not consuming it.[18]

[17]Note that if $f(x)$ causes large losses, the initial increase in utility of not consuming the asset may exceed the increase for the consumption of the asset. This is the case if v^a at $0 > r^a$ is steeper than at $f(x) < r^a$. Then, for a certain part of the reference formation process, the consumption of the asset becomes relatively less attractive after a decrease in r^a.

[18]Only if the final loss that is derived from an expectation of zero, given a reference point $r^a = f(x) \succ 0$, is large, the utility from consuming the asset may temporally decrease faster for an increase in the reference state than that from not consuming the asset.

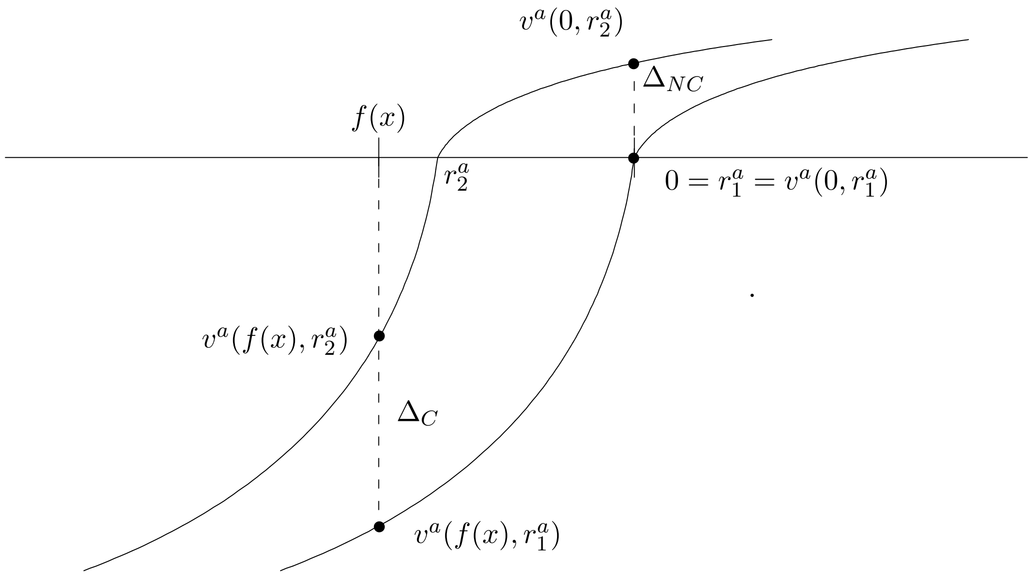

Figure 2.1: Effect of reference state change with $f(x) \prec 0$

If the reference lottery moves from r_1^a to r_2^a, the adjustment utility of the degenerate lottery that yields zero in all states (no consumption of F) increases from 0 to $v^a(0, r_2^a)$. The adjustment utility of the lottery $f(x)$ induced by the consumption of asset F increases from $v^a(f(x), r_1^a)$ to $v^a(f(x), r_2^a)$. With v^a steeper for losses than for gains, the net effect is an increase in utility from $f(x)$ relative to the zero lottery.

2.5 Consumption optimization

2.5.1 General case

As the final step of the analysis I now add prices, and take into consideration the individual's reference formation regarding her overall consumption. I compare her utility from consuming asset F with that of consuming an alternative asset which induces the same current utility as F but no future outcomes. For ease of exposition, I restrict the analysis to assets F which yield negative expected outcomes, i.e., "risks". In addition, I focus on assets that induce non-monetary outcomes like health effects. This allows me to simplify the analysis by assuming that monetary and non-monetary utility are separable.

In period t, the individual receives monetary income Y_t and can choose between the risky asset F and the alternative safe asset S, which yields the same outcomes at present but no outcomes in T. I normalize present outcomes of assets F and S to zero.[19] The price for the safe asset S is normalized to $P_S := 0$. The risky asset F has a discount

[19]Since present outcomes are the same, reference point effects for those outcomes are also the same and can be ignored in a relative analysis.

π, such that $P_F = -\pi$. For example, the individual has the choice between consuming genetically-modified food and organic food, where organic food is safe but more expensive (and I assume that both taste the same at present).

$r^a(f(x))$ as before denotes the reference state that results if F is consumed, while $r^a(S)$ denotes the reference state that results if S is consumed, i.e., if there are neither bad nor good outcomes in T. In addition to F and S, individuals can invest into monetary consumption c_t, and utility from monetary consumption is denoted with U^c. To keep the model simple, I abstract from savings, which means that monetary consumption in periods $s > t$ has no influence on the optimization.[20]

Let the parameter ξ denote the decision, where $\xi = 0$ if the safe asset is consumed and $\xi = 1$ if the risky asset is consumed. The individual's optimization problem is then given by

$$\begin{aligned} max_\xi \quad U^c(c_t) \quad &+ \quad \xi \left[\sum_{s=t}^{T-1} (u_s^a(f(x)) + v_s^a(f(x)|r_s^a(x))) + u_T(f(x)) + v_T(f(x)|r^x) \right] \\ &+ \quad (1-\xi) \sum_{s=t}^{T-1} [u_s^a(S) + v_s^a(S|r_s^a(S))] \\ \text{s.t.} \quad c_t \quad &= \quad Y_t + \xi\pi \quad . \end{aligned}$$

This yields the first order condition

$$U'(Y_t + \xi\pi)\pi = \sum_t^{T-1} [u_s^a(S) + v_s^a(S|r_s^a(S)) - u_s^a(f(x)) - v_s^a(f(x)|r_s^a(x))] - u_T(f(x)) - v_T(f(x)|r^x)$$

Since S is safe, i.e., it does not induce health effects in the future, $u_T(S) = v_T(S|r^x) = u^a(S) = 0$, where I use the fact that if the individual does not consume the asset, she expects no health effect, i.e., $r^x = 0$. However, $u^a(S) = 0$ is considered a gain for $r^a < 0$, in which case $v^a(S|r^a(S)) > 0$. I then obtain for the optimum that

$$U'(Y_t + \xi\pi)\pi = \sum_t^{T-1} [v_s^a(S|r_s^a(S)) - u_s^a(f(x)) - v_s^a(f(x)|r_s^a(x))] - u^x(f(x)) - v^x(f(x)|r^x)$$

In words, the risky asset F is preferred over the safe asset S if the utility increase due to the lower price of F exceeds the utility decrease due to the lower utility from expectations and realized outcomes caused by a possible adverse outcome in T. Note that this general trade-off is not affected by introducing saving opportunities into the model.

[20]This simplification does not affect the main results of the model, since I assume utility from monetary and non-monetary assets to be separable. This is common in the literature and seems at least a good approximation(see, e.g., Feldman and Dowd, 1991, who assume separability between medical care consumption and non-medical care consumption).

2.5.2 Indivisible assets

To keep the analysis tractable, I now limit it to assets that are indivisible, $\xi \in \{0, 1\}$. This means that asset F can be consumed or not, but cannot be partly consumed. This applies, e.g., to decisions like whether to eat gm-food or not, practice safe or unsafe sex etc. I account explicitly for reference effects regarding consumption of monetary goods, c_t. In contrast to the non-monetary outcomes induced by $f(x)$ (e.g., health effects), these are goods that can be bought, like food, clothing, cars etc. r^c denotes the reference state for the outcome utility from these goods. Then, the risky asset is preferred over the safe asset if $U_{\xi=1} > U_{\xi=0}$, i.e., iff:

$$
\begin{aligned}
& u^c(Y_t + \pi) + v^c(Y_t + \pi | r^c) - u^c(Y_t) - v^c(Y_t | r^c) \quad > \\
& \sum_{s=t}^{T-1} [v_s^a(S | r_s^a(S)) - u_s^a(f(x)) - v_s^a(f(x) | r_s^a(x))] \quad - \quad u^x(f(x)) - v^x(f(x) | r^x)
\end{aligned}
\tag{2.8}
$$

In this condition, expressions $r^a_{s>t}(f(x))$, $r^a_{s>t}(S)$ and r^x are endogenous, while r^a_t and Y_t are exogenous in t.

Consider first the influence of the reference formation process regarding adjustment utility on the optimization. The reference state r^a only affects the term $\sum_{s=t}^{T-1}[v_s^a(S|r_s^a(S)) - v_s^a(f(x)|r_s^a(f(x)))]$ in (2.8). Assume that F induces negative outcomes. Then, no matter whether the individual consumes S or F, the social comparison component of the reference state leads to $r^a < 0$ if somebody else in her reference group accepts the risk. However, with the individual herself accepting the risky asset, the reference state decreases faster according to (2.3), such that $r^a_{s>t}(f(x)) < r^a_{s>t}(S)$. Hence, with t denoting the time when the first member of the reference group accepts the risk, $v^a_{s>t}(S|r^a_s(S)) > 0$ and $v^a_{s>t}(f(x)|r^a_s(f(x))) < 0$, that is, the utility from expectations that arises from not accepting the risk always exceeds that from accepting the risk, given the respective paths for the reference state. Put differently, even though reference states for $s > t$ are determined endogenously by the individual's decision of whether to accept the risk or not, the utility from non-monetary goods that is expected for periods $t < s \leq T$ is always higher when consuming the safe asset than when consuming the risky asset.

With decreasing r^a, consuming the safe asset becomes a larger gain, while consuming the risky asset becomes a smaller loss. This may initially increase the attractiveness of the safe asset if losses are large. However, as shown in section 2.4.2, the net effect of reference state formation is a decrease in $\sum_{s=t}^{T-1}[v_s^a(S|r_s^a(S)) - v_s^a(f(x)|r_s^a(f(x)))]$ and hence an increase in the attractiveness of the risky asset relative to the safe alternative. Accordingly, if the individual chooses between F and S later in the reference formation process, she derives a higher utility from the risky asset relative to the safe alternative than earlier in the process. This does not depend on whether individuals are aware of their preferences being reference dependent, or whether they are able to correctly predict reference state formation.

How is this result affected by changes in the reference point for monetary consumption? An increase in r^c from Y_t towards $Y_t + \pi$ makes paying for the risky asset felt as less high a gain and paying for the safe asset as higher a loss. Beginning from $r^c = Y_t$ when the risky asset is first consumed, $v^c(Y_t|r^c)$ decreases fast initially, while $v^c(Y_t+\pi|r^c)$ decreases only moderately. Accordingly, paying for the risky asset becomes more attractive. For the overall process, r^a decreases from 0 to $f(x)$, and the utility from the risky asset F increases more than the utility from the safe asset S.[21] In contrast, with r^c increasing from Y_t to $Y_t+\pi$, the overall utility from consuming Y_t (and the safe alternative) decreases more than the utility from consuming $Y_t + \pi$ (and asset F). This yields

Proposition 3 *Consider an individual that has the choice between consuming the indivisible, safe asset S and the risky but cheaper indivisible asset F. Then, the net effect of the reference formation process regarding both the utility from expectations and the utility from realized outcomes is an increase in the utility from buying the risky asset relative to the utility from buying the safe asset.*

Accordingly, the probability that the individual consumes the risky asset weakly increases if she gets used to expecting the asset's outcomes.

2.6 Value of information

Consider an individual who in t expects an outcome distribution $f(x)$ to be realized in T. At time $s > t$ she receives new information and updates her expectation to $f'(x)$. At time z with $s < z < T$ the information is corrected and expectations are adjusted back to $f(x)$. In T, $f(x)$ is realized as expected.

Compared to the case where no information is obtained between s and T, i.e., expectations are constant at $f(x)$, the net effect of this (false) information is

$$\Delta U = \underbrace{\sum_{s}^{z-1}[u^a(f'(x)) - u^a(f(x))]}_{I} + \underbrace{\sum_{s}^{z-1} v^a(f'(x)|r^a)}_{II} + \underbrace{\sum_{z}^{T-1} v^a(f(x)|r^a)}_{III} + \underbrace{v_T^x(f(x)|r^x)}_{IV}$$

Term I denotes the aggregated effect on anticipatory utility for all periods between the announcement and the correction of information. Term II denotes the aggregated adjustment utility that is experienced after the initial announcement of the information

[21] Only for reference states close to $Y_t + \pi$ and large losses from paying for the safe asset may this be reversed, namely if v_a gets steeper for gains than for losses at the respective locations.

but before the correction, while term III denotes the aggregated adjustment utility after the correction of the information but before outcomes are realized. Finally, term IV denotes the effect on reference-dependent outcome utility in case the reference state has not adjusted to the corrected expectations until T.

Consider first the case when negative information is released, i.e., $f'(x) \prec f(x)$. Then,

$$\begin{aligned} \Delta U > 0 \quad & \text{if} \quad |I + II| < III + IV \quad \text{while} \\ \Delta U < 0 \quad & \text{if} \quad |I + II| > III + IV. \end{aligned}$$

Terms I and II are negative, such that adjustment utility after the initial announcement of the information is experienced as a loss. Terms III and IV are positive, that is, adjustment utility after the correction and reference-dependent outcome utility are experienced as gains.

For simplicity, and since I am not interested in outcome utility here, consider as before the case when the reference state r^x has fully adjusted to expectations until T, such that term IV is zero. Adjustment utilities II and III refer to the same difference in utility, only with different sign. The experimental literature suggests that reference states do not change faster for losses than for gains, and losses are felt no less than gains of equal size. Assuming this to hold for r^a, one obtains that $\Delta U < 0$. This is consistent with the implication one would derive if considering only anticipatory utility. Applying loss aversion to adjustment utility strengthens the expected effect.

If positive information is released, $f'(x) \succ f(x)$,

$$\begin{aligned} \Delta U > 0 \quad & \text{if} \quad I + II > |III + IV| \quad \text{while} \\ \Delta U < 0 \quad & \text{if} \quad I + II < |III + IV|. \end{aligned} \tag{2.9}$$

Now, terms I and II are positive, such that adjustment utility after the initial announcement of the information is experienced as a gain. Terms III and IV are negative, i.e., adjustment utility after the correction and reference-dependent outcome utility are experienced as losses. Consider again the case where the reference state r^x has fully adjusted to expectations until T, such that term IV is zero. Accounting only for anticipatory utility (term I) would then imply an unambiguous positive effect. Including adjustment utility, however, shows that positive false information, if believed by the individual, is no "free lunch". Rather, if the disutility from the disappointment after the correction of the information is large relative to the initial elation, the overall effect may be negative. Whether the positive or negative effect prevails depends on the functional forms of u^a, and v^a given r^a, as well as on the formation process of r^a.[22]

[22]If r^x has not fully adjusted to $f(x)$ until T, it also depends on the functional form of v^x given r^x and the formation process of r^x.

In summary, (2.9) shows that in contrast to the implications of anticipatory utility and reference-dependent outcome utility, spreading good news that later prove to be false may have a negative overall effect on utility, even if it is corrected in time for r^x to adjust to the individual's expectations until outcomes are realized.

2.7 Exogeneity of the reference state

A theoretical implication of adjustment utility is that the reference state for expectations cannot be derived endogenously in general. The reason is that changes in the reference state are gradual and always affect utility. To see this, consider again the framework developed in section 2.3.1. For simplicity, timing in the model is discrete. However, in reality people experience pleasure and pain in any moment in time. This means that both utility and the reference state are experienced continuously. Hence, the time Δ between two periods in the model should be thought of as very short. With $\Delta \to 0$, it is sensible to assume that the reference state adjusts to changes in the determining factors over several periods rather than from one period to the next, i.e., that reference formation is a smooth process without jumps.

Consider a period when information arrives that induces a change in expectations. The new expectation is then gradually included in the reference state r^a. But until the adjustment of the reference state is complete, the reference state is also partly determined by the former state of expectations, i.e., the former reference state. These former expectations are exogenous to the new adjustment process. Only in the periods when the reference point has fully adjusted to expectations and no new information arrives is the reference point determined endogenously.

One may argue that if there is never a change in expectations, the reference state is endogenous. However, for any asset that is relevant for utility, an individual's initial unawareness of the asset has to be overcome in order for her to analyze the asset's utility or consider its consumption. She only becomes aware of the asset if the necessary information is supplied, which then affects her expectations. Hence, at least when leaving initial unawareness, the reference point is exogenous.

The exogeneity of the reference state for expectations differs from the results of Koszegi and Rabin (2006) for realized outcomes. In their model, the individual's expectation before T over her consumption decision in T is interpreted as the reference state for outcomes in T. And the adjustment of the reference state to new expectations is completed *before* the reference state affects utility. Accordingly, during its adjustment the reference state does not cause any utility effects. In contrast, with individuals continuously experiencing utility from expectations, reference states regarding expectations are affecting utility at any time. This means that these reference states cannot be formed before they are relevant, and hence are not endogenous.

Consider as an example a person who buys a wedding dress. Until marriage crossed her mind, she was unaware of the market for wedding dresses and the utility of wearing them. Hence, she had no expectation of buying and wearing a wedding dress, such that her reference state regarding this expectation was zero. Now she searches a catalogue for wedding dresses and picks one that she likes. After choosing the dress, she decides not to look into wedding catalogues again. The expectation of wearing the chosen dress then gradually becomes her reference state for her expectation of what to wear at her wedding. If she does not reconsider her decision, her choice together with her initial unawareness determine the reference states for her expectations for all periods until adjustment is complete or the wedding takes place. Accordingly, until complete adjustment, the reference state has an exogenous and an endogenous component.

Alternatively, this person may have attended previous weddings and have already formed an expectation about what kind of dress to wear. This expectation forms her reference state before she starts searching the market and it is exogenous at the time when she obtains detailed market information. When finally searching the catalogues she may find out that her budget does not allow her to buy any dress of the kind she had in mind. If she decides for an affordable dress, this then forms her expectation about what to wear at her wedding and how much to spend on it. If she never reconsiders her decision, her future reference states for her expectation are determined by this new expectation together with her prior, uninformed expectation. However, the change from the initial uninformed expectation to the informed expectation will already have caused her disappointment.

2.8 Implications and conclusion

Adjustment utility also has implications for the forecasting of consumption behavior. Proposition 3 implies that consumers' disposition to buy assets which induce adverse future outcomes is not constant, even if people's preferences and the assets' characteristics are. Rather, if these assets get increasingly included in people's reference states, this disposition increases. Further, the surveys described in the introduction imply that the inclusion of such assets (their expected outcomes) in the reference state increases if the individual herself or relevant others have consumed the asset in the past. Hence, once some individuals start consuming these assets, their disposition to buy them in the future increases, and so does the disposition of other people in the reference group. Accordingly, an individual's current consumption behavior is not a good predictor of her future consumption behavior if some individual in her reference group, including herself, consumes the asset. This is relevant, e.g., when a new source of risk is about to be released into the market and regulatory measures are discussed. In particular, if a higher level of consumption requires a different regulatory policy, or if consumption is not wanted to exceed a certain level, regulation should not be based on current or currently planned consumption behavior but on behavior updated with reference formation regarding expectations. This is particularly relevant if later changes in the regulation are expensive or hard to implement.

Another implication of the model regards the behavior of firms. If a firm wants to introduce an asset into the market that is perceived as being risky, e.g., genetically modified food, it will focus on increasing the number of individuals that consume the asset. This decreases consumers' average reference states regarding the risk induced by this asset, which increases their utility from consuming the asset and their willingness to buy it.

Accordingly, selling a product initially at below-cost prices is profitable in the long run if market penetration is high enough to substantially decrease people's reference states regarding the involved risks. At a later stage, the firm can then slowly raise prices to profitable levels without a marked reduction in sales volume, since reference states continue to decrease until the asset's risk is fully included in consumers' reference states.

It should finally be noted that in order to quantify the predictions of the model, more research on the formation of the reference state is needed, i.e., on the relative importance of the influencing factors, their interrelation, the time period that is considered (the size of β), and the characteristics of the reference group.

Appendix A

Table 1: HIV infections and use of condoms in Germany, 1996-2004

Indicator	1996	1997	1998	1999	2000	2001	2002	2003	2004
Number of new infections with HIV in Germany	1871	2070	1923	1752	1690	1425	1635	1827	2058
People who know that HIV is transmitted through unsafe sex (%)	98	98	98	99	99	99	100	99	100
People who believe that condoms reduce the risk of an HIV infection (%)	-	-	-	-	86	85	88	87	89
People who intend to use condoms (%)	91	90	93	92	92	93	94	91	93
People who possess condoms (%)	64	62	62	61	64	65	64	67	67
People who actually use condoms (%)	72	73	72	72	70	70	69	71	69
People who actually use condoms (new relationships, %)	72	70	73	72	78	77	75	73	70
People who actually use condoms (more than one sexual partner, %)	81	80	78	80	79	83	82	78	77
People who are informed about new therapies and use condoms (%)	-	74	-	-	-	-	-	74	78
People who are not informed about new therapies and use condoms (%)	-	75	-	-	-	-	-	77	72

Source: Robert-Koch-Institute, Berlin, Germany (line 1) and Bundeszentrale für gesundheitliche Aufklärung (BZgA), Köln, Germany (lines 2-10).

Appendix B

Questionnaire 1

This questionnaire is part of a research project of the chair of Microeconomics, Prof. Dr. Dorothea Kuebler. Please answer the questions and return the questionnaire to the lecturer. Thank you for your participation!

Consider the following situations:

Situation 1

Ms. Schulz is told on Friday, February 3, that she will receive a 5% wage increase from April 1. Ms. Schulz did not expect this wage increase. On Monday, February 6, in the morning, she is told that the payroll department made a mistake and that she will not receive the wage increase.
On April 1 she receives her usual wage.

Situation 2
Ms. Schulz does not receive any information about a wage increase. On April 1 she receives her usual wage.

Question (Please tick the situation in which you think Ms. Schulz is happier.):
In which situation is Ms. Schulz happier on February 6 at noon?
Situation 1 o Situation 2 o Equally happy in both o

Thank you for your participation!

Questionnaire 2

This questionnaire is part of a research project of the chair of Microeconomics, Prof. Dr. Dorothea Kuebler. Please answer the questions and return the questionnaire to the lecturer. Thank you for your participation!

Consider the following situations:

Situation 1

Mr. Meier is department manager at his company. He knows already in August 2006 that he will receive a salary of EUR 50,000 in 2007. He does not know the salaries of the other department managers at his company, and the other managers do not know his salary. However, Mr. Meier believes that the other managers will receive an average salary of EUR 40,000 in 2007. On October 1, 2006, Mr. Meier learns by coincidence that the other department managers will also earn EUR 50,000 on average in 2007. The other managers do not get to know his salary. Assume that the average salary of the other managers does not have an impact on Mr. Meier's future salaries or career perspectives.

Situation 2

Mr. Meier is department manager at his company. He knows already in August 2006 that he will receive a salary of EUR 50,000 in 2007. He does not know the salaries of the other department managers at his company, and the other managers do not know his salary. However, Mr. Meier believes that the other managers will receive an average salary of EUR 60,000 in 2007. On October 1, 2006, Mr. Meier learns by coincidence that the other department managers will also earn EUR 50,000 on average in 2007. The other managers do not get to know his salary. Assume again that the average salary of the other managers does not have an impact on Mr. Meier's future salaries or career perspectives.

Question (Please tick the situation in which you think Mr. Meier is happier.):

In which situation is Mr. Meier happier on October 1?

Situation 1 o Situation 2 o Equally happy in both o

Thank you for your participation!

Appendix C

In this appendix I analyze the argument that Ms. Smith in the survey question simply fails to correctly predict her reference state.

Let r^x_{Apr} denote the reference state for the outcome utility Ms. Smith experiences when receiving her wage for April. If Ms. Smith' higher expected April wage on Feb. 4-6 means that on Feb. 7 her expectation of this reference state is higher than her expectation of her wage in April, $E_{Feb7}(r^x_{Apr}) > E_{Feb7}(x_{Apr})$, and we assume that anticipatory utility takes into account future reference states, we get the following:

$$u^a(f(x), E_{Feb7}(r^x_{Apr})) \quad < \quad u^a(f(x), E_{Feb7}(x_{Apr}))$$

This difference in anticipatory utility would explain the difference in happiness Ms. Smith is expected to experience according to the survey. So how then would this difference arise? $E_{Feb7}(r^x_{Apr}) > E_{Feb7}(x_{Apr})$ means that even though Ms. Smith correctly expects her wage in April, she expects that her reference state will not adjust to this expectation during the two months up to the realization of outcomes. She expects her future self to be disappointed when receiving her April wage. This does not necessarily contrast with the literature (e.g., Koszegi and Rabin, 2005), since *expecting* to be disappointed in April does not necessarily mean that she will indeed *feel* disappointed when April arrives. She may simply predict her reference formation process incorrectly and underestimate adjustment. However, it seems highly counterintuitive for Ms. Smith to overestimate the duration of her disappointment in the future, without feeling disappointed at present. But feeling disappointed at present is what cannot be explained with the existing models.

In summary, some individuals may indeed be unhappy in the first situation because they do not expect their reference state to adjust to the new wage expectation until April. But these individuals will feel disappointed also at present, i.e., the argument made in the paper will apply to them as a precondition for the argument above.

The same argumentation applies to the second survey. Ms. Smith will most likely not overestimate the duration of her disappointment or delight onto the moment when she receives her wage if she does not feel disappointed or delighted in February.

In general, however, one would not expect 90% or 78% of the respondents to predict their reference state to not adjust to their expectations within two months, unless one assumes that reference states indeed *do not* adjust to expectations over time.

Bibliography

A. Abel. Asset prices under habit formation and catching up with the joneses. *The American Economic Review, Papers and Proceedings of the Hundred and Second Annual Meeting of the American Economic Association*, 80(2):38–42, 1990.

S. Benartzi and R. Thaler. Myopic loss aversion and the equity premium puzzle. *NBER Working Paper #4369*, 1993.

M. Brunnermeier and J. Parker. Optimal expectations. *American Economic Review*, 2005 (4):1092–1118, 95.

J. Campbell and J. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, April 1999.

A. Caplin and J. Leahy. Psychological expected utility and anticipatory feelings. *Quarterly Journal of Economics*, 116(1):55–80, 2001.

G. Constantinides. Habit formation: A resolution of the eep. *Journal of Political Economy*, 98:519–543, 1990.

D. Kahneman and A. Tversky. Prospect theory. *Econometrica*, 42:33–35, 1979.

D. Kahneman, J. Knetsch, and R. Thaler. Experimental tests of the endowment effect. *Journal of Political Economy*, 98:1325–1348, 1990.

D. Kahneman, J. Knetsch, and R. Thaler. The endowment effect, loss aversion and status quo bias. *Journal of Economic Persepectives*, 5:193–206, 1991.

B. Koszegi. Utility from anticipation and personal equilibrium. *Working Paper, University of California, Berkeley*, 2005.

B. Koszegi and M. Rabin. A model of reference-dependent preferences. *Quarterly Journal of Economics, forthcoming*, 2006.

G. Loewenstein. Anticipation and the valuation of delayed consumption. *Economic Journal*, 97:666–684, 1987.

Chapter 3

Imitation with Intention and Memory: an Experiment

3.1 Introduction

There are two reasons for imitation. First, imitation can be necessary for the development of skills. This applies to tasks that require practice, like eating with chop sticks or walking on one's hands. Second, imitation can be useful to achieve high outcomes, but no information or skill is acquired. For example, buying the same stocks as André Kostolany would not have meant to learn why he picked which stocks, but to profit from his knowledge without acquiring it oneself. In such cases, the necessary information is received *before* the imitation, when one learns what person is worthwhile imitating. Skill is not required.

The experiment described in this paper focuses on the second kind of imitation. It is designed to mimic situations where players make decisions in a changing environment, but cannot learn from imitation in the sense of acquiring skill or information through imitating others. It is particularly relevant under circumstances where players have (or think they have) different levels of information, and can observe others' current behavior. The experiment shows that people imitate even if they cannot obtain information or acquire skill by doing so. Overall, 39% of all decisions in the experiment show an intention to imitate, but are consistent neither with genuine learning nor with randomization strategies. This result complements previous experimental research (Apesteguia et al., 2006), (Selten and Apesteguia, 2002) that focussed on showing that behavior was *consistent* with an imitation strategy, but could not show whether it was actually driven by imitation motives.

The second point of the paper concerns the choice of imitation examples, that is, whom or what players imitate.[1] In the literature, people are usually assumed to repeat an action that was successful either in the period just before the decision (Vega-Redondo, 1997) (Selten and Ostmann, 2001) or in some period during the time the player can recall (Alos-Ferrer, 2004), (Josephson and Matros, 2004). In experiments, players only have the opportunity to imitate past strategies (Huck et al., 1999). In contrast, the hypothesis in this experiment is that players imitate other players, rather than repeating particular actions. The main difference between the two concepts is that successful players have long-term strategies that make them successful in changing environments. One-shot actions, on the other hand, may be successful in one period, but not necessarily in another. To distinguish the two concepts, repeating others' actions is costless in the experiment, while imitating players is costly. Nevertheless, participants frequently imitated successful players.

Finally, the experiment sheds some light on the use of memory in individual decisions. Even though some theoretical models account for multi-period memory (Alos-Ferrer, 2004), the experimental literature so far usually considered the case where only last period's payoffs can be recalled and used in the decision (Huck et al., 1999), (Altavilla et al., 2006).[2] In this experiment, unlimited memory of all previous periods is induced by providing information on all actions and payoffs to all participants throughout the experiment. The results suggest that players indeed use this information when making their decisions. Although 85% of all imitation decisions are consistent with considering only last period's payoffs and imitating the player with the highest payoff there, over 80% of the decisions in the relevant periods are also consistent with considering the payoffs of the last 2 to 17 periods. Considering all past periods and weighting them with a discount factor also supports the hypothesis that results of more than one past period are relevant for players' decisions. Finally, a regression shows that the payoffs of up to four periods prior to the decision have a significant impact on the choice of imitation examples.
The results of the experiment are consistent with the behavioral rules suggested by Alos-Ferrer (2004) and Bergin and Bernhardt (1999), with the restriction that the experimental design is not well suited for a rigorous test of these rules. In contrast, the win-stay lose-shift rule suggested by Dixon (2000) and Oechssler (2002) receives only weak support in this context.

The situations that are mimicked in the experiment, where the environment changes and some players may be better informed than others, are common in real market situations. For example, consider a atomistic market where a new firm enters. This firm

[1] I choose the term "imitation examples" following Selten and Ostmann (2001) "success examples".

[2] It should be noted that under this assumption it becomes undistinguishable whether actions or players are imitated.

may follow the marketing strategy of a firm that proved successful in the market, without understanding yet why this strategy is actually profitable. Assuming that the strategy and payoffs of the successful firm are observable, the new firm does not learn anything from imitation that it could not learn from pure observation. However, by imitating it may ensure high payoffs until it has finally learned to successfully market its products itself.

Similar examples exist for individual decision making. Assume an investor who wants to invest in stocks and does not have much knowledge about the market. But he happens to know that George Soros recently bought shares of the publicly traded company X. By buying such stocks, the investor will not learn anything new about the market, since stock prices are observable. He may nevertheless imitate Soros, believing that this increases his payoffs until he has finally learned to make profitable investment decisions himself. In both cases, the relevant increase in information occurs *before* the imitation. The new firm learns by observation which firms are successful in the market, and the newcomer learns who are successful investors. Practicing skills is not necessary.

Imitation processes have been analyzed theoretically by Vega-Redondo (1997), Schlag (1998), Schlag (1999), Selten and Ostmann (2001), and more recently by Fudenberg and Imhof (2006) and Levine and Pesendorfer (2006). In these models, any behavior is defined as imitation which leads to the individual choosing an action which was played last period. Only the exact action the individual should choose is determined differently. While Vega-Redondo and Selten/Ostmann use an imitate-the-best rule, Schlag proposes a proportional rule, where each strategy which yields a higher outcome than one's own is imitated with a probability proportional to the difference in outcomes.

These models' focus on last period's outcomes applies well to evolutionary contexts, where periods can be interpreted as generations, and individuals enter and drop out of the population each period. For individual learning within relatively short time (either by people or firms), however, such a limited memory seems a restrictive assumption. Alos-Ferrer (2004) and Josephson and Matros (2004) relax this restriction and include multi-period memory in their models. Alos-Ferrer modifies the imitate-the-best rule, such that players remember the payoffs of $K+1 > 1$ periods, and imitate the strategy which yielded the highest payoff in memory. He notes, however, that more sophisticated strategies are imaginable, i.e., to imitate the strategy with the highest recalled average payoff, the average population payoff and the highest within-period payoff.

The experimental literature so far followed the evolutionary, no-memory approach of the earlier imitation models. For example, Apesteguia et al. (2006) set up a Cournot market game and test the different predictions that the theoretical models yield for this market. On the aggregate level, different informational settings allow them to distinguish between the models according to the quantities players choose, i.e., the equilibrium the market converges to. On the individual level, the authors find that many players' behavior is consistent with imitation in the majority of periods. However, given the experimental setting, it seems impossible to distinguish explicitly between behavior that looks like imitation but is the result of other strategies, and behavior which is actually driven by imitation motives.[3] This problem equally applies to the experiments of Selten and Apesteguia (2002), Huck et al. (1999), Altavilla et al. (2006). It is described in more detail below.

Consider a situation where in an oligopoly market game player 1 sets price x and obtains payoff a. Player 2 sets price y and obtains payoff $b > a$. In the next period, player 1 also sets price y. In the experiments mentioned above, this would be interpreted as evidence for imitation in individual behavior. This inference, however, is based on the very broad interpretation of "imitation" as any behavior which makes an individual choose an action that was successful in the past. In particular, it does not require an intention to imitate. How, then, does learning occur which is not imitation? If somebody learns by observing others that in chess the player wins who keeps his king until the end, and then tries to keep her king when playing herself - does she imitate? Or did she just learn the rules of the game, and now plays by them?[4]

In this paper, I choose a narrower definition of imitation, in order to be able to distinguish between behavior which is imitating in the colloquial sense of the word (copying, mimicking) on one hand, and behavior which involves observation and genuine learning on the other hand.[5] Imitation in the sense of practicing to acquire skills is not considered. According to this narrower definition of imitation, players' behavior is imitating if they have an intention to do what somebody else does or did, and do not have sufficient understanding of the situation to choose a successful action themselves. This means that player 1 may truly have wanted to imitate player 2. However, he may as well have learnt the structure of the market from observation in the first period, and then used this knowledge to pick the action with the higher expected payoff in the second.

In the next section I explain the design of the experiment. The results are presented and discussed in section 3.3. Some concluding comments are provided at the end.

[3]The authors note that a questionnaire that the participants filled in after the experiment gives hints that some players indeed intended to imitate.

[4]Even trial and error learning could be classified as imitation with this definition, namely imitation of one's own past successful actions.

[5]I define as *genuine learning* all strategies that lead to an expected increase in relevant knowledge about the game, i.e., its structure and rules. This includes complex strategies like Bayesian updating from the observation of the strategies and payoffs of others, but also simple strategies like trial and error.

3.2 Experimental design

Participants in the experiment had to choose cells from a table. Each cell returned points. The pattern according to which cells were matched to points changed, but not necessarily in each round. Participants knew that this change occurred according to a regular, non-trivial rhythm. They also knew that other players might have more or less information than they themselves. However, they did not know for sure whether there were others with a different amount of information, nor whether they were better or worse informed than others, nor of what the difference in information consisted. I will call the better informed players "informed", and the less informed players "uninformed".

Informed players saw tables like the one in Figure 3.1.

a1	b2	c3
b3	a3	a2
c2	c1	b1

Figure 3.1: Board shown to informed players

They knew that a-cells were always best, b-cells were medium and c-cells were always worst. The position of the cells in the table changed, but players could observe this. However, the ranking of the three cells of a particular letter also changed. In some period, a1 could be best, in another a2, and in yet another a3. The rhythm of this change was not known to the informed players but could be learnt during the experiment. So the task for informed players was in each period to choose an a-cell and find out during the experiment in which period which a-cell returned the highest number of points.[6]

Uninformed players saw the table in figure 3.2 in each period:

a	b	c
d	e	f
g	h	i

Figure 3.2: Board shown to uninformed players

This table did not contain any payoff-relevant information. Players knew that each cell contained points, and that these point might change every period. But they did not

[6]It was not explicitly stated that they should only choose from a-cells. But it was clearly stated that a-cells were always best.

see the a-cells. Hence, it was their task to find out among all cells - instead of just the a-cells - which cells yielded the highest points in which period.

Instead of choosing a cell themselves, players could name another player, whose cell choice was then also valid for them. They had to make this decision before they saw other players' current choices, that is, when making the decision to imitate another player, they did not know his cell choice.[7] For example, player 3 might have decided in period 5 to imitate player 7. Player 7 chose, say, cell c and received 14 points. Only afterwards did player 3 learn about the choice of player 7, and the number of points he received.

After each round, all players' cell choices and points were made visible to all participants, and this information about all rounds was available on screen throughout the experiment. For consistency, the choices of one type of player were translated into the strategy space of the other. For example, an $a1$ choice of an informed player appeared as a b, e etc. choice on the screen of uninformed players, depending on its position in the relevant period. Similarly, a b, f etc. choice of an uninformed player appeared as a $b1$, $c3$ etc. choice on the screen of informed players, depending on the positions in the relevant period. This means that all information that was generated during the experiment and was relevant for learning how to obtain a high number of points was available to all players at all times. Imitating another player did not increase the players' information. Rather, since by imitating others players did not generate their own data points (no cell choice), imitation reduced overall information compared to choosing one's own cell.

Players could earn between 1 and 15 points by choosing a cell themselves. If they chose to imitate a player, they received the number of points of the imitated player, minus a fee. There was a "cheap" treatment, in which this fee was one point, and an "expensive" treatment, where the fee was three points.

The experiment lasted for 30 periods, without any practice periods. In each session there were 2 informed and 12 uninformed players (except for session 5, where there were only 10 uninformed players for technical reasons). Four sessions were conducted for the "cheap" treatment, and three for the "expensive" treatment. In sum, there were 48 uninformed players in the "cheap" treatment, and 34 in the "expensive" treatment. Treatments differed only in the size of the fee. Participants were informed that after the 30 periods one period was chosen randomly, of which payoffs were paid out immediately. Points were exchanged into EUR according to the rate 1 point = 0.7 EUR for uninformed players and 1 point = 0.5 EUR for informed players. The different exchange rates were chosen in order to account for the difference in information and give players roughly equal expected payoffs. At the end of the instructions test questions ensured that all participants under-

[7] The words "imitate" or "imitation" were not used anywhere in the experiment. See also the translated instructions in the appendix.

stood the rules of the experiment.

Participants were students of Humboldt Universität zu Berlin and Technische Universität Berlin, with mostly non-economics majors. The experiments took place at Technische Universität Berlin in June 2006. The experiments were computerized using z-tree (Fischbacher, 1999). Sessions lasted about 60 min. The average amount earned was 9.8 EUR.

3.3 Results

3.3.1 Imitation intention

The first question the experiment addresses is whether players have an intention to imitate. Imitation is defined as players following the cell choices of other players, although they do not know yet what this choice is.

Overall, uninformed players imitated in 39% of the decisions they made (959 of 2460). In the cheap treatment they imitated in 44.3% of their decisions (638 of 1440), while in the expensive treatment they imitated in 31.5% of the decisions (321 of 1020).[8] This shows that in both treatments players imitated to a significant degree.

For the interpretation of the numbers, it should be noted that in the first periods it was not yet obvious that some players obtained higher payoffs, while in the second half of the experiment several uninformed players had learnt to play the game successfully themselves. The share of imitation decisions in the periods when imitation was sensible is therefore even higher than the numbers show. For example, the share of uninformed players' imitation decisions in periods 6 to 23 is 54% in the cheap treatment and 37% in the expensive treatment. This is also reflected in the distribution of imitation decisions over the 30 periods, which is shown in figure 3.1. and table 3.4 in the appendix. Imitation really starts in the 5th or 6th period, i.e., when it becomes obvious what players are successful. It levels off towards the end when more players understand the rhythm of changes. Note that the distribution of points in the table could only be fully learnt after the 14th period.[9]

— insert figure 3.1. here —

[8] Informed players mainly existed to serve as imitation examples, i.e., to have high average payoffs and give uninformed players a reason to imitate. Their behavior is therefore not included in the analysis of the extent of imitation. For completeness, when I analyze the mechanics of imitation below, the few imitation decisions of informed players (10 in all sessions) are nevertheless included.

[9] The rhythm of changes was such that the first three periods were not connected to the rhythm, and then a sequence over 9 periods was repeated 3 times.

The higher price of imitation in the expensive treatment led to a lower, but still considerable amount of imitation. The difference in the average number of imitations per player is significant at the 1%-level (Wilcoxon-Mann-Whitney test). It should, however, be noted that the size of the reduction is mainly driven by one of the three sessions of this treatment with a very low amount of imitation (14% vs. 46% and 36% in the other two sessions). Since otherwise there were no significant differences between the treatments, for the following analysis the data is aggregated over both treatments.

The observed amount of imitation can be attributed to an *intention* to imitate, since genuine learning as well as randomization strategies can be ruled out as motives for imitating other players in this experiment.
To see this, consider first "genuine" learning strategies, i.e., strategies that increase the player's knowledge about the game. If a player imitates in the experiment, he does not learn the payoff of any cell choice from his own behavior. The cell choices and payoffs of other players are displayed throughout, independently of the players choice. Accordingly, imitation does not increase a player's knowledge of the game. Rather, by reducing the number of data points, it (weakly) decreases the available information.
Consider, second, randomization, i.e., that players are confused or bored and randomly choose actions between which they are indifferent. For this strategy to include imitation in the experiment, imitation had to be costless. But even in the cheap treatment, imitation cost uninformed players an equivalent of 70 EUR-Cent. In the expensive treatment this increased to an equivalent of 2.1 EUR. With an expected payoff from the entire experiment of about 10 to 15 EUR, few participants can be expected to ignore these costs. Any cell choice that a player wishes to play, including the random choice, can be achieved more cheaply if he chooses it himself.
Hence, players who imitate can be expected to do so intentionally.

3.3.2 Imitation of players or repetition of actions

The second question the experiment addresses is whether participants repeat successful past actions (cell choices) or imitate successful players.

The extent to what players imitate other players was outlined above. Do they also "imitate" by repeating past actions? On average, 25,7% of the players that chose their own cell chose the cell that was most successful last period. However, the pattern of cells was such that doing so yielded one of the 2 highest payoffs in 14 out of the 29 possible imitation periods (48%), a fact that was fairly easy to observe. In addition, choosing always cell "c" was a good strategy for a "simple" learner, since it yielded high payoffs in 23 periods. This "c" strategy would have resulted in 12 periods of repetition of the highest payoff from last period, or 41%. Since the observed 25% are clearly below these figures, the data does not suggest an intention to repeat past actions beyond what random and simple learning imply.

Is this result particular to the experimental design? The answer is yes and no. It is particular, because it will not hold in very stable settings, where what is the best action in one period is most likely also the best action in all other periods. Optimal play in such settings can be learned by observing it once, and imitation is not necessary. The result could not be obtained in the settings of previous experiments either, because there players did not have the opportunity to imitate other players' current behavior.

However, the result is also fairly general, because the crucial features of the design are frequently encountered in real life. What makes imitation of successful players attractive is that one can obtain high profits without actually understanding the game. In an economic context, this applies, e.g., to situations where markets change, but not all players know when or how, or how to react to a given change. For example, in oligopolistic markets, some firms may be good in predicting their competitors behavior, while others lack the experience to do so. If given the chance, the latter may then prefer to imitate the former, instead of repeating the formers' behavior of another period.

It is also interesting to see whether participants, in addition to imitating successful other players, learnt how to play the game themselves. Figure 3.2 and table 3.4 show the average number of points obtained from own cell choices, conditional on players actually choosing their own cell. It shows an increase in the second half of the experiment, which indicates that some people have learnt the pattern, and choose their own actions. The increase is significant at $p < 0.01$ (Wilcoxon-Mann-Whitney test). The players who do not learn the pattern mostly continue to imitate.

— insert figure 3.2. here —

3.3.3 Choice of imitation examples

The third, and somewhat more complex question is, which players are actually imitated. In particular, do people use multi-period memory to identify imitation examples? Or is it a good approximation to consider only last period's outcomes when explaining players' imitation decisions?

Approaching this question, I make the assumption that only such players become imitation examples, i.e., are imitated, that have maximum payoffs according to some measure.[10] I call this the *maximization rule*. In general, this measure, which I denote S

[10] This means that I do not follow the approach of Schlag (1998).

for *score*, is of the form

$$S_t^i = \sum_{n=t-s}^{t-1} \delta^{t-1-n} P_n^i \tag{3.1}$$

where δ denotes the discount factor applied to the results of previous periods, s denotes the number of periods taken into consideration, and P_n^i denotes the number of points player i received in period n. t is the period when the imitation decision is made, such that only the results of periods prior to t can be taken into account. Equation (3.1) gives one S per period per player, which can then be compared to this period's scores of all other players of the same session. $\Delta S_t^i = S_t^i - S_t^{max}$ denotes the difference between player i's score and the maximum score S^{max} in period t in this session. I then assess which parameters s and δ in (3.1) minimize the number of imitation decisions that deviate from the maximization rule, i.e., all decisions where for the imitated player i $\Delta S_t^i < 0$.

This approach yields three results. I first set $\delta = 1$ and then choose s optimally for each imitation decision. With this approach, 98% of all imitation decisions (946 of 969) imitate a player for whom $S_t^i = S_t^{max}$. If one allows for $\Delta S = -1$ as error margin, this goes up to more than 99% (963 of 969). However, to obtain these figures, one has to allow for different values of s across imitating players and periods.

Second, I leave $\delta = 1$ but apply the same s to all imitation decisions. I then consider all periods when $t - s > 0$, i.e., periods where s previous periods could have been taken into account. For example, to assess the share of decisions that are consistent with considering the payoffs of the previous 13 periods, I consider all decisions in periods 14 to 30. This approach shows that for $s = 1$, in 85% of the imitation decisions in relevant periods players imitated a player i for whom $S_t^i = S_t^{max}$, i.e., the decisions were in line with the maximization rule. This seems strong support for the assumption that players consider only one past period when making their decision. However, for considering up to 17 periods, i.e., $2 \leq s \leq 17$, in more than 80% of the decisions the imitating players also follow a player for whom $S^i = S^{max}$. Up to $s = 26$, still more than half of the decisions are consistent with imitating a player with the maximum average payoff. This result is summarized in figure 3.3. and table 3.4 in the appendix.

— insert figure 3.3. here —

Third, I set $s = t - 1$, i.e., I consider the results of all available periods, and optimize the discount factor such that it minimizes the number of deviations from the maximization rule. Since players' scores may now differ in only very small amounts, I repeat the analysis for different error margins. Figure 3.4 shows the number of deviations from the maximization rule depending on δ. An error margin of, e.g., 0.1 means that all imitation decisions for which the score of the imitated player i was $S^i \geq S^{max} - 0.1$ are accepted as being consistent with the maximization rule.

— insert figure 3.4. here —

The data show that for error margins of 0.01 and 0.1 points a discount factor of 0 - considering only last period - yields the least number of deviations from the maximization rule (see table 3.4 in the appendix). However, the mathematical abilities and calculation efforts that are implied in these low error margins are high. If one allows for some limitations in the subjects' cognitive abilities or efforts, and accepts error margins of 0.5 to < 1 point from the maximum, higher discount factors, i.e., a higher weight on previous periods, yield less deviations from maximization. Indeed, even a discount factor of $\delta = 1$, i.e., attaching equal weight to all previous periods, leads to higher compliance with the maximization rule than considering only the last period.

Finally, a regression shows that the points of up to four periods prior to the imitation decision have a significant impact on the decision. I run a panel regression with fixed effects, where the independent variables are a player's achieved points per period. The dependent variable is the share of others that imitated a certain player in a given period, relative to all players that imitated in this period. Hence, if in period 12 six players imitated, and four of them imitated player 8, this variable for player 8 in period 12 is $\frac{2}{3}$. The results are displayed in table 3.4 in the appendix. They show that the points of the two periods prior to the decision are significant at $p < 0.01$, while the points of three and four periods prior to the decision are significant at $p < 0.05$.[11]

3.3.4 Test of other learning rules

As the last step of the analysis, I consider some of the rules suggested in the literature, and test whether the data is consistent with their implications. Dixon (2000) and Oechssler (2002) propose a win-stay lose-shift learning rule, according to which individuals' reference level for gains/losses is the average population payoff. Players stick with an action as long as it returns payoffs above their reference level, and experiment with other actions if payoffs fall below the reference level. Applied to this experiment, one would expect that if individuals who chose their own action last period and obtained below-average payoffs would switch to imitating other players, while those who obtained above-average payoffs keep on choosing their own action. The results show, however, that only 30% of the

[11]It should be noted that the independent variables suffer from multicollinearity, with average correlation coefficients between the significant variables of between 0.40 and 0.60. This is to some extent also reflected in the correlation between the dummy variables and the independent variables. This problem does, however, not seem so severe that it compromises the results of the regression.

players with below-average payoffs turned to imitation in the next period. Nevertheless, among those that started imitating, 70% had below-average payoffs in the previous period.

Alos-Ferrer (2004) suggests a rule according to which individuals "imitate the strategy that yielded the highest payoff in memory". This rule is almost fully consistent with the players' imitation decisions if one assumes that strategies equal players.
With all the results that are displayed on the screen being in the "memory" of players, only 9 imitation decisions ($< 1\%$) are such that the imitated player does NOT have the maximum payoff of 15 points in some period prior to the decision. Unfortunately, this result is not particularly instructive. The reason is that the design of the experiment, i.e., its stochastic component, is such that after only few periods, even pure randomization would most likely lead to the majority of players having obtained the maximum payoff in some period. Then, no matter what player is imitated, the decision is in line with the rule.
The same result obtains for the approach suggested by Bergin and Bernhardt (1999). In their model, players imitate one of the actions that performed best relative to other actions in a particular period, and consider all periods within their memory. If one again assumes that players' memory in the experiment is comprised of all past periods, only 6 imitation decisions are not consistent with this rule ($< 1\%$). However, the interpretation of this result suffers from the same problems as the rule of Alos-Ferrer.

3.4 Conclusion

The experiment shows that players have an intention to imitate other players they perceive as being more successful. When choosing imitation examples, they consider more than only last period's payoffs, but focus on players that are successful in the long run. The results of the experiment apply to market situations that are characterized by a changing environment and players that have different levels of information or experience. More generally, they show that imitation exists as a strategy to increase payoffs, even if no information or skill is acquired. This lends support to the interpretation of the results of earlier experiments. It suggests that the behavior observed there is not only consistent with imitation, but that players indeed had the intention to imitate.

Some limitations of the experimental design should finally be mentioned. First, it is not able to determine with ultimate certainty the number of periods that individual players consider when choosing imitation examples. To obtain this, the design would have to guarantee that different lengths of memory in most cases imply different decisions, which is too often not the case here. Second, and for similar reasons, the results do not yield clear implications for which of the behavioral rules that are suggested in the literature players comply with. To explicitly test for their validity, the stochastic structure of the game would have to be amended such that not almost any behavior is in line with these rules.

Appendix A

Period	Average number of imitations, treatment "cheap"	Average number of imitations, treatment "expensive"	Average number of points obtained from own cell choice
1	0.00	0.00	8.0
2	0.75	2.11	10.0
3	1.75	2.17	9.1
4	2.25	2.78	7.3
5	5.00	3.61	9.4
6	6.50	5.44	11.3
7	3.50	2.61	9.4
8	6.75	5.50	10.8
9	5.00	3.22	8.7
10	6.50	6.06	9.5
11	6.75	5.28	10.9
12	6.50	4.67	9.5
13	8.00	4.61	9.2
14	8.00	6.17	12.8
15	7.50	5.11	11.7
16	7.75	4.67	10.1
17	8.25	6.44	10.7
18	7.00	5.72	12.0
19	7.00	4.33	12.7
20	5.00	4.00	13.5
21	6.25	4.72	12.5
22	6.75	4.33	12.2
23	5.75	5.00	13.4
24	4.75	4.67	13.9
25	4.50	3.94	12.1
26	6.00	5.11	12.8
27	5.50	3.22	13.2
28	4.75	3.56	13.9
29	3.75	3.28	14.1
30	4.00	4.00	14.1

Table 3.1: Average imitations and points per period. (Fig. 1 and 2)

s	share		s	share		s	share
1	0.856		11	0.832		21	0.760
2	0.830		12	0.836		22	0.734
3	0.830		13	0.825		23	0.718
4	0.834		14	0.810		24	0.681
5	0.834		15	0.806		25	0.629
6	0.826		16	0.803		26	0.561
7	0.828		17	0.800		27	0.488
8	0.839		18	0.786		28	0.449
9	0.823		19	0.775		29	0.400
10	0.830		20	0.758			

Table 3.2: Share of imitation decisions in the relevant periods that are consistent with imitating the player with the maximum average points over s periods. (Fig. 3)

discount factors	error margins 0.01	0.1	0.5	0.999	1.0
0.00	*0.924*	*0.924*	0.924	0.924	*0.973*
0.05	0.895	0.922	0.928	0.965	0.966
0.10	0.900	0.917	0.933	0.965	0.966
0.15	0.896	0.911	0.933	0.966	0.966
0.20	0.897	0.911	0.934	0.967	0.967
0.25	0.897	0.911	0.936	0.968	0.968
0.30	0.898	0.910	0.937	0.968	0.968
0.35	0.891	0.908	*0.939*	0.972	0.972
0.40	0.891	0.906	*0.939*	0.972	0.972
0.45	0.891	0.912	0.937	*0.973*	*0.973*
0.50	0.896	0.906	*0.939*	0.970	0.970
0.55	0.893	0.905	*0.939*	0.970	0.970
0.60	0.894	0.904	0.938	0.970	0.970
0.65	0.895	0.905	0.935	0.972	0.972
0.70	0.895	0.903	0.935	0.971	0.971
0.75	0.895	0.901	0.934	0.971	0.971
0.80	0.895	0.902	0.935	0.972	*0.973*
0.85	0.895	0.904	0.932	0.971	0.971
0.90	0.895	0.905	0.935	0.971	0.971
0.95	0.894	0.904	0.935	0.969	0.969
1.00	0.895	0.906	0.937	0.968	0.968

Table 3.3: Share of imitation decisions that are consistent with imitating the player with the maximum average points over all periods, depending on discount factors and error margins allowed for the maximum rule. (Fig. 4) Optimum shares are emphasized.

Fixed-effects (within) regression — Number of obs =2880

R-sq: within = 0.0404 — Group variable (i): m

between = 0.4132 — Number of groups = 96

overall = 0.2032

corr(u_i, Xb) = 0.3844 F(10,2774)= 11.67 Prob > F = 0.0000

Variable	**Coefficient** (std. error)	**t**	$P > \|t\|$	95% Conf. Intervall
points(-1)	.0012688 (.0004766)	2.66	0.008	.0003342 .0022033
points(-2)	.0020659 (.0004977)	4.15	0.000	.0010901 .0030417
points(-3)	.0011261 (.0005032)	2.24	0.025	.0001393 .0021128
points(-4)	.0010724 (.0005099)	2.10	0.036	.0000725 .0020723
points(-5)	.0005475 (.0005169)	1.06	0.290	-.000466 .001561
points(-6)	.0003268 (.0005279)	0.62	0.536	-.0007084 .0013619
points(-7)	.0006156 (.0005385)	1.14	0.253	-.0004403 .0016716
points(-8)	.0000704 (.0005449)	0.13	0.897	-.0009981 .0011388
points(-9)	.0007364 (.0005543)	1.33	0.184	-.0003504 .0018232
points(-10)	-.0000528 (.0005273)	-0.10	0.920	-.0010868 .0009812
const	.0191582 (.0051821)	3.70	0.000	.0089969 .0293194

σ_u = .1650182 σ_e = .13010806

ρ = .61665651 (fraction of variance due to u_i)

F test that all $u_i = 0$: F(95, 2774) = 40.79 Prob > F = 0.0000

Table 3.4: Panel Regression. The variable $points(-s)$ denotes the number of points the individual received s periods prior to the imitation decision.

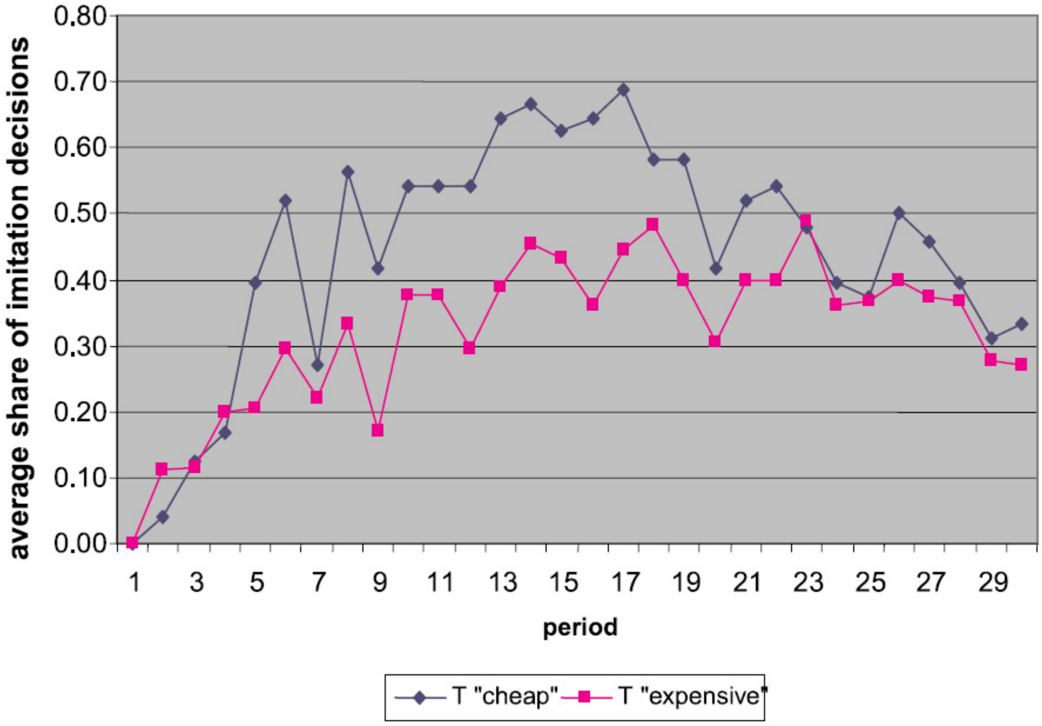

Figure 3.3: Average share of imitation decisions of uninformed players per period per treatment.[1]

[1]The parallel ups and downs for both treatments in periods 4–10 are probably due to a special feature of the change of the point distribution in the table. In some periods at the beginning, it may have seemed as if the rhythm of changes was simpler than it actually was, which may have caused some people in each treatment to give up imitation and choose a cell themselves. After finding out that they were mistaken, they started imitating again. Due to the limited number of players, however, it is impossible to reach a conclusive interpretation of this effect.

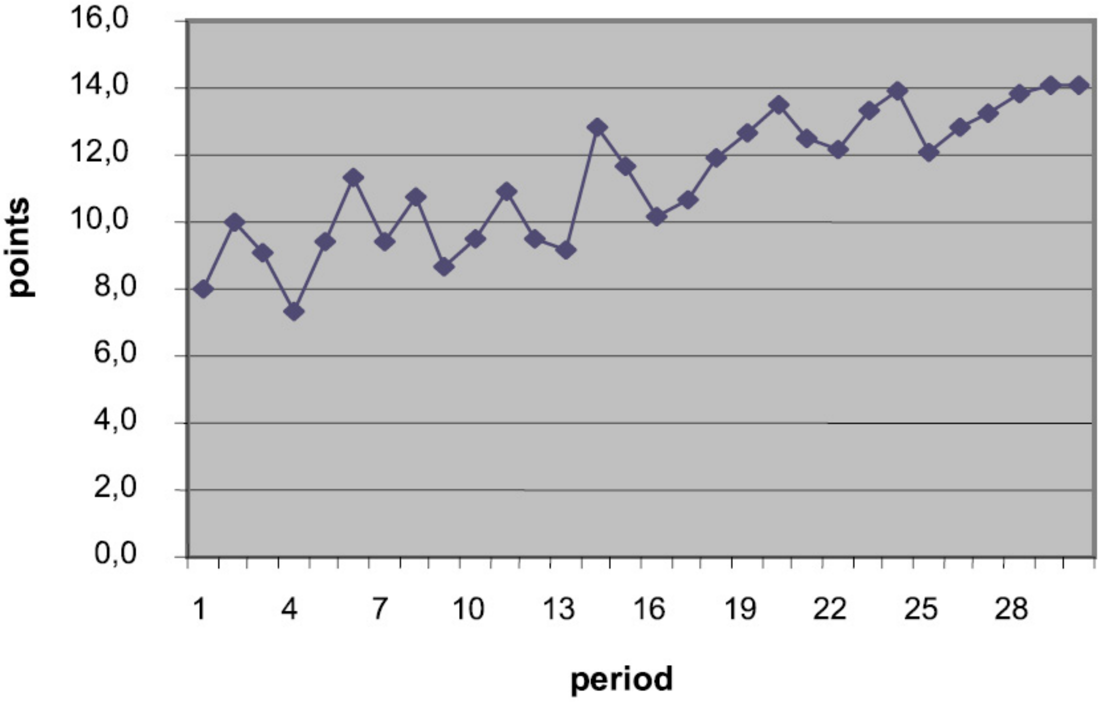

Figure 3.4: Average number of points obtained from choosing a cell oneself, contingent on players actually choosing their own cell.

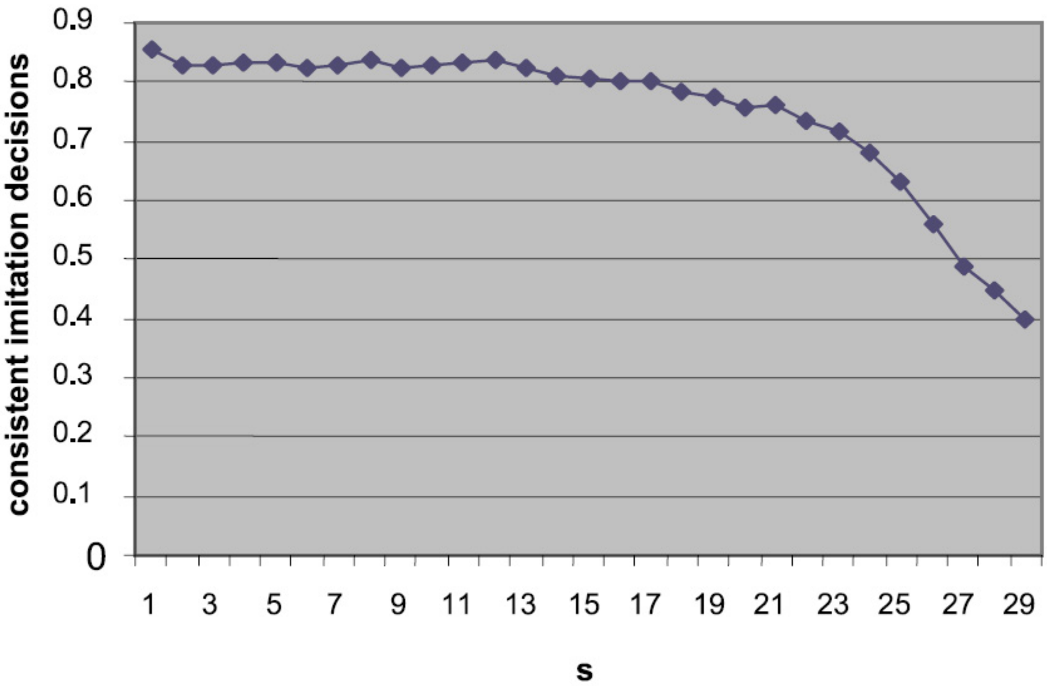

Figure 3.5: Share of imitation decisions that imitate the player with the maximum score if considering s periods, relative to all decisions in periods $t > s$. $\delta = 1$.

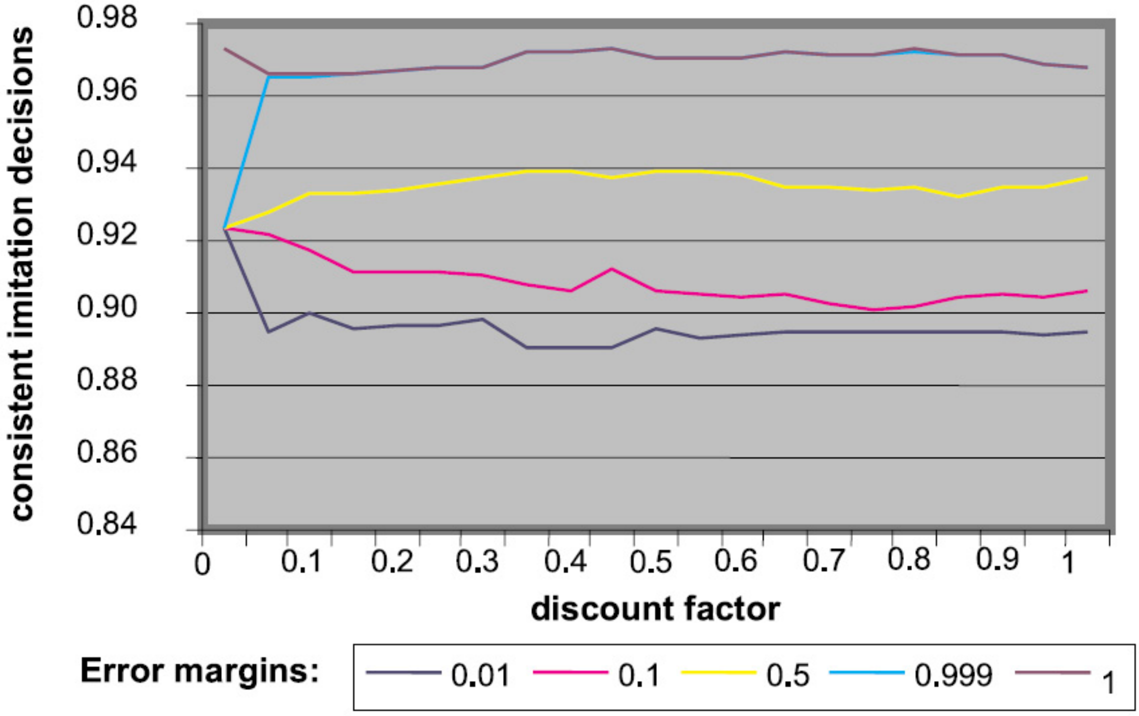

Figure 3.6: Share of imitated players with maximum score in the period previous to the imitation decision, depending on the discount factor applied to past periods and the error margins allowed.

Appendix B

This is the translated version of the instructions informed players received:

Instructions

The experiment you will now participate in is part of a research project financed by the Deutsche Forschungsgemeinschaft (DFG). It aims at analyzing economic decision making.

In the experiment you can earn a considerable amount of money, which depends on your decisions. Accordingly, it is important that you read the instructions carefully.

Please note that these instructions are only for your use. You are not allowed to pass on any information to other participants. Similarly, during the entire experiment it is not permitted to talk to other participants. If you have a question, please raise your hand. We will than come to you and answer your question. Please do not ask your question allowed. If you do not comply with these rules, we have to stop the experiment.

General information

The experiment consists of 30 periods. In each period you make a decision. According to these decisions you receive points. At the end of the experiment, these points are exchanged into Euro and paid out in cash. The course of the experiment, your decisions and the payoffs are explained in detail in what follows.

Decision

In each period a table with nine cells appears on your screen, which may look like this:

a1	b2	c3
b3	a3	a2
c2	c1	b1

Cells are always labelled with a1-a3, b1-b3, c1-c3. In each period you can choose one of those cells and receive points for your decision.

Cells that are labelled with the letter *a* return the highest number of points. cells that are labelled with the letter *b* return medium numbers of points. cells that are labelled with the letter *c* return the lowest number of points. However, cells that are labelled with the same letter can still return different numbers of points. For example, cell a1 can return a higher or a lower number of points than cells a2 and a3 etc.

The pattern of the table, i.e., the points returned by the cells, changes. This change occurs according to a certain rhythm. The logic of the letters - a for the highest points, b for medium points, c for the lowest points - is preserved by the change, but the relation of points within the letters can change. Hence, in one period cell b1 may return a higher number of points than cell b2, while in the next period b2 returns a higher number of points than b1.

When you have decided for a cell, please type it in under "Your cell choice" and press "OK". The number of points returned by the cells then appears on your screen and is added to your account. The choice of cells happens just once every round, it cannot be amended later on.

In each period, the distribution of points over the different cells is the same for all participants. It is, however, possible, that other participants have more or less information about this distribution than you do yourself.

Instead of choosing a cell yourself, you can also choose one of the other participants, whose cell choice is then also valid for you. To do so, you simply type in the number of the player you have chosen under "Your chosen player". You then receive the same number of points as this player, less one point as "fee". You only learn the cell choice of this player after you have chosen him. Note that you can only choose a cell OR a player, not both. Just as for the cell choice, the choice of a player is binding. Even if you are not happy with the cell choice of this player afterwards, you cannot change your number of points anymore.

If you want to choose a cell yourself, you simply leave the box "Your chosen player" empty.

Example:

You see the following table:

a1	b2	c3
b3	a3	a2
c2	c1	b1

Case 1: You choose cell a1 and receive 12 points. (This number of points is fictitious. It has no relation to the true number of points in the experiment.)

Case 2: You decide to follow the choice of player 3. In the box "Your chosen player" you type "3". This player chooses cell a1. He receives 12 points, and you receive 12-1=11 points.

After each period cell choices and points of all players are displayed. The first row shows the number of the period, the following rows show the cell choices and points of all players. F1 denotes the cell choice of player 1, P1 his number of points. F2 denotes the cell choice of player 2, P2 his number of points etc. The cell choices and points of all players are visible on screen throughout the entire experiment. If a player chose another player rather than a cell, the box of is cell choice stays empty.

After all 30 periods are played one period is drawn randomly, which is then relevant for payoffs. The number of points that you received in this period is then exchanged into EUR according to the following rate:

$$2 \text{ points} = 1 \text{ EUR}$$

In addition you receive 3 EUR for your participation. The sum of show up fee and the payoff from your decisions is then paid out in cash immediately.

Questions

Please answer the following questions to ensure that you have understood the instructions.

1. Player 1 chooses a cell with the label b, player 2 chooses a cell with the label c. Who receives the higher number of points?

2. Player 1 chooses a cell with the label a. Player 2 also chooses a cell with the label a. Who receives the higher number of points?

3. Player 1 decides to choose the cell that player 2 chooses. Player 2 receives 10 points. How many points does player 1 receive?

4. In period 4, player 1 chooses an a cell in the upper right corner of the table. In period 5, he again chooses the cell in the upper right corner. Does he receive the same number of points?

This is the translated version of the instructions uninformed players received:

Instructions

The experiment you will now participate in is part of a research project financed by the Deutsche Forschungsgemeinschaft (DFG). It aims at analyzing economic decision making.

In the experiment you can earn a considerable amount of money, which depends on your decisions. Accordingly, it is important that you read the instructions carefully.

Please note that these instructions are only for your use. You are not allowed to pass on any information to other participants. Similarly, during the entire experiment it is not permitted to talk to other participants. If you have a question, please raise your hand. We will than come to you and answer your question. Please do not ask your question allowed. If you do not comply with these rules, we have to stop the experiment.

General information

The experiment consists of 30 periods. In each period you make a decision. According to these decisions you receive points. At the end of the experiment, these points are exchanged into Euro and paid out in cash. The course of the experiment, your decisions and the payoffs are explained in detail in what follows.

Decision

In each period the following table appears on your screen:

a	b	c
d	e	f
g	h	i

cells are labelled with the letters *a* to *i*. In each period you can choose one of those cells and receive points for this. The letters do not contain any information about the numbers of points that are assigned to these cells. They simply serve to improve clarity. The points that are assigned to these cells change. However, this change occurs according to a regular pattern.

When you have decided for a cell, please type it in under "Your cell choice" and press "OK". The number of points returned by the cells then appears on your screen and is added to your account. The choice of cells happens just once every round, it cannot me amended later on.

In each period, the distribution of points over the different cells is the same for all participants. It is, however, possible, that other participants have more or less information about this distribution than you do yourself.

Instead of choosing a cell yourself, you can also choose one of the other participants, whose cell choice is then also valid for you. To do so, you simply type in the number of the player you have chosen under "Your chosen player". You then receive the same number of points as this player, less one point as "fee". You only learn the cell choice of this player after you have chosen him. Note that you can only choose a cell OR a player, not both. Just as for the cell choice, the choice of a player is binding. Even if you are not happy with the cell choice of this player afterwards, you cannot change your number of points anymore.

If you want to choose a cell yourself, you simply leave the box "Your chosen player" empty.

Example:

You see this table:

a	b	c
d	e	f
g	h	i

Case 1: You choose cell a and receive 12 points. (This number of points is fictitious. It has no relation to the true number of points in the experiment.)
Case 2: You decide to follow the choice of player 2. In the box "Your chosen player" you type "3". This player chooses cell a. He receives 12 points, and you receive 12-1=11 points.

After each period cell choices and points of all players are displayed. The first row shows the number of the period, the following rows show the cell choices and points of all players. F1 denotes the cell choice of player 1, P1 his number of points. F2 denotes the cell choice of player 2, P2 his number of points etc. The cell choices and points of all players are visible on screen throughout the entire experiment. If a player chose another player rather than a cell, the box of is cell choice stays empty.

After all 30 periods are played one period is drawn randomly, which is then relevant for payoffs. The number of points that you received in this period is then exchanged into EUR according to the following rate:

$$1.5 \text{ points} = 1 \text{ EUR}$$

In addition you receive 3 EUR for your participation. The sum of show up fee and the payoff from your decisions is then paid out in cash immediately.

Questions

Please answer the following questions to ensure that you have understood the instructions.

1. Player 1 chooses a cell with the label e, player 2 chooses a cell with the label g. Who receives the higher number of points?

3. Player 1 decides to choose the cell that player 2 chooses. Player 2 receives 10 points. How many points does player 1 receive?

4. In period 4, player 1 chooses cell c. In period 5, he again chooses cell c. Does he receive the same number of points?

Bibliography

C. Alos-Ferrer. Cournot versus walras in dynamic oligopolies with memory. *International Journal of Industrial Organization*, 22:193–217, 2004.

C. Altavilla, L. L., and P. Sbriglia. Social learning in market games. *Journal of Economic Behavior and Organization (forthcoming)*, 2006.

J. Apesteguia, S. Huck, and J. Oechssler. Imitation - theory and experimental evidence. *Journal of Economic Theory (forthcoming)*, 2006.

J. Bergin and D. Bernhardt. Comparative dynamics. *Working Papers 981, Queen's University, Department of Economics*, 1999.

H. Dixon. Keeping up with the joneses: Competition and the evolution of collusion. *Journal of Economic Behavior and Organization*, 43:223–238, 2000.

U. Fischbacher. z-tree: A toolbox for readymade economic experiments. *Working Paper No. 21, University of Zurich*, 1999.

D. Fudenberg and L. Imhof. Imitation processes with small mutations. *Journal of Economic Theory (forthcoming)*, 2006.

S. Huck, H. Normann, and J. Oechssler. Learning in cournot oligopoly: An experiment. *Economic Journal*, 109:C80– C95, 1999.

J. Josephson and A. Matros. Stochastic imitation in finite games. *Games and Economic Behavior*, 49:244–259, 2004.

D. Levine and W. Pesendorfer. The evolution of cooperation through imitation. *Games and Economic Behavior (forthcoming)*, 2006.

J. Oechssler. Cooperation as a result of learning with aspiration levels. *Journal of Economic Behavior and Organization*, 49(3):405–409, 2002.

K. Schlag. Why imitate, and if so, how? a boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78:130–56, 1998.

K. Schlag. Which one should i imitate? *Journal of Mathematical Economics*, 31:493–522, 1999.

R. Selten and J. Apesteguia. Experimentally observed imitation and cooperation in price competition on the circle. *Bonn Econ Discussion Paper 19/2002*, 2002.

R. Selten and A. Ostmann. Imitation equilibrium. *Homo Oeconomicus*, 43:111–149, 2001.

F. Vega-Redondo. The evolution of walrasian behavior. *Econometrica*, 65:375–384, 1997.

Chapter 4

Do Public Banks have a Competitive Advantage?

4.1 Introduction

Public banks that hold state guarantees on their deposits often enjoy lower funding costs than private banks in the same market. Wherever this is the case, this cost difference tends to be blamed to distort competition in favor of the public bank.

In this paper, I show that the support of the state may turn out to be a disadvantage for the public bank. The reason is that public (i.e., state-owned) and private banks do not only differ in funding costs, but also in their perceived objective function. To what extent the true objective functions differ is subject to debate. But public banks usually have the *mandate* to support the economy, which they cannot publicly breach. Accordingly, most borrowers will assume that if they take out a loan from the public bank their firm will not get liquidated at the first sign of financial difficulties. Independently of the public bank's true objective function this perception may suffice to allow private banks to enter the market, separate the borrower pool, and obtain profits in equilibrium. What was meant as a policy to support the economy may turn out to increase the interest rates for all borrowers.

Consider a loan market with safe and risky firms, where risky firms are those that have a higher probability of experiencing financial distress. The incumbent public bank initially serves the whole market at a uniform pooling rate. Now the market opens for competition. If private banks have higher funding costs but are otherwise equal to the public bank, their loan offers are not competitive, and the public bank continues to serve the whole market. This is the stylized case that may be cited to show the adverse effect of state guarantees.

However, since private banks are not restricted in setting their policy, they can offer a contract that includes the liquidation of all borrowers in financial distress but an interest rate lower than that of the public bank. Due to its mandate, the public bank cannot compete by offering a similar contract. The private banks' loan is more attractive for safe than for risky firms, since risky firms have a higher probability of being in financial distress, that is, a higher probability of inefficient liquidation under the private banks' loan. Hence, private banks can induce firms to separate by self-selection and lend only to safe firms. Safe firms produce higher expected returns, such that the private banks can offer lower interest rates than the public bank and still overcome their cost disadvantage. The public bank is left with the risky firms.

A key component of the private banks' strategy to separate firm types is the credibility of their threat to liquidate all firms in financial distress. Since their lower interest rate attracts all firms, only the liquidation threat keeps risky firms from applying for their loans. Accordingly, private banks have to develop a long-term strategy, and ensure that the profits from repeatedly lending to safe firms exceed the one-time profits from extending the loans of distressed but viable firms instead of liquidating them. This eliminates the incentive to deviate from the announced strategy and makes the liquidation threat self-enforcing. It implies that even under perfect competition private banks can sustain profits in equilibrium.

Two of the model's assumptions merit additional explanation. First, by focusing on the competition between the two types of banks, I implicitly assume that firms do not have other sources of external funding. In particular, they cannot issue stocks or bonds. This seems a realistic assumption for the sector of small and medium-sized enterprises, which form a large part of the economy, for example, in Germany. In addition, it applies to almost all firms in countries with less developed financial markets. However, ruling out other sources of funding also means that I ignore any implications that the competition between banks and financial markets may have on the behavior of banks. (For literature on how different firms choose between bank finance and, e.g., bonds, see Diamond (1984), Houston and Christopher (1996), Johnson (1997), Bolton and Freixas (2000). For example, Chemmanur and Fulghieri (1994) have a model where banks compete with financial markets. There, establishing a reputation for auditing firms in distress creates a competitive advantage for the bank over bond holders. Obviously, their result contrasts sharply with mine, and a combination of both strategies could hardly be pursued. Which strategy a bank prefers may ultimately depend on the relative importance of the different groups of customers.

The second important assumption is the perceived restriction of the public bank's objective function. The evidence on what objective public banks actually follow is mixed. Bichsel and Spielmann (2004) do not find that public banks in Switzerland set lower rates

than private ones, suggesting that they maximize profits just as their private competitors. In contrast, Sapienza (2002) finds that public banks set lower rates than private banks in Italy, though she cannot clearly determine the motive for this behavior. More generally, Levy-Yeyati et al. (2004) find no strong evidence for public banks either promoting or hindering economic development.
However, public banks have a *mandate* to support the economy (see their statutes, e.g., in Germany and Switzerland). This means that, whatever their true objective function is, they are not able to announce profit maximization as their goal, but at least have to *claim* to follow policies in support of economic development. This justifies the assumption that public banks are *perceived* as granting firms in financial difficulties a chance to complete their projects if they turn out to be viable. Apart from this perception, however, I do not assume that public banks differ in their policies from private banks. Rather, I analyze their behavior as profit maximizers. Hence, in contrast with the literature on soft budget constraints (SBC; see, e.g., Dewatripont and Maskin (1995), Maskin (1996)) I do not assume that the state asks the public bank to make unprofitable loans, but to assess the viability of all firms in distress. Projects that are not viable are liquidated. Hence, the interference of the state does not soften the budget constraints of the public bank's borrowers, but restricts the bank's strategy space within the scope of profitable business strategies. "Risky loans" in this context are those with a positive, though lower NPV. Bad projects in the sense of the SBC literature, i.e., projects with an ex-ante negative NPV, are assumed to be detected through a pre-loan audit. This audit is able to distinguish negative NPV from positive NPV projects, i.e., to determine the rough quality of a project, but not to distinguish between different types of profitable projects. This assumption can be interpreted, e.g., as banks being able to assess the technological quality of a project, but not the quality of the management.

Contrary to some recent work that found state-owned banks to be less efficient than their privately-owned competitors (Caprio et al., 2005), (La Porta et al., 2002), I do not assume the public bank to be inefficient per se. On the one hand, such an assumption would obscure the focus of the paper, which is on the implications of the competition between private and public banks. On the other hand, in some countries, public banks were found to be as efficient or even more efficient than private financial institutions (see Altunbas et al. (2001), for evidence from Germany), such that no final conclusion can be drawn on the prevailing situation.

The argument I develop is based on a simple self-selection mechanism as introduced by Rothschild and Stiglitz (1976). But rather than themselves offering two different contracts, the private banks use the restriction of the public bank's strategy space and offer a contract that this bank is unable to offer. There emerges a separation of agents and principals, where the private banks separate intentionally, while the public bank separates unintentionally.

In sections 4.2 and 4.3 I develop the model and describe the pooling equilibrium. Separation is analyzed in section 4.4, while section 4.5 derives the credibility condition. Section 4.6 concludes.

4.2 Model

There are two types of firms, safe and risky, indexed s and r, respectively. Firm types are private information, i.e., only the firm itself knows its type. Lending occurs recurrently in each period. In this and the next two sections, I analyze only one loan cycle, $t = 0$ to $t = 1$. In section 4.5 I extent the analysis to an infinite horizon. Each firm has a single investment project that requires external finance in $t = 0$ and produces output x_j in $t = 1$ (see figure 4.1 for the sequence of events). Investment is normalized to 1. All projects have non-negative net present value. As mentioned above, this could be the result of a pre-contract audit, which assesses the hard facts of a project, but not the soft facts like the risk attitude of the management, its skill, etc.

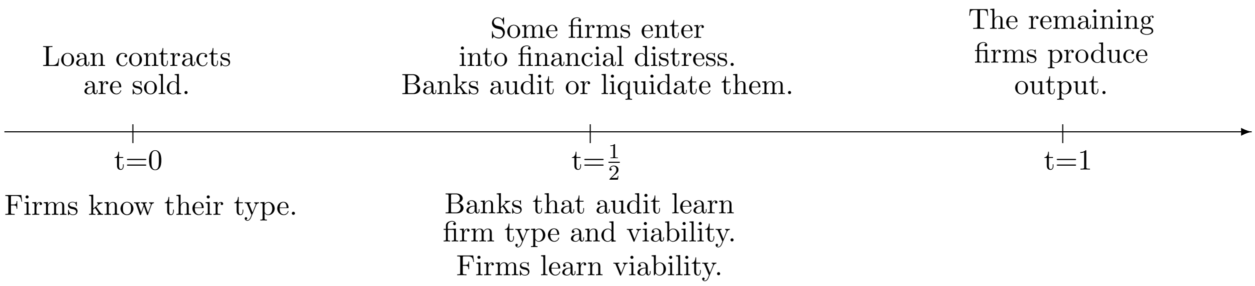

Figure 4.1: Sequence of events in one loan cycle

In $t = \frac{1}{2}$, some firms enter into financial distress. In this case, with probability p_v the firm is viable and able to produce output x_j in $t = 1$. With probability $1 - p_v$ it produces nothing. The a priori quality of a firm (safe or risky) is not related to the probability that a distressed firm is viable. If liquidated in $t = \frac{1}{2}$, firms of either type return a liquidation value of y.

Safe firms have a lower probability of distress than risky firms, $p_s < p_r$, and a higher expected output, $(1 - p_s)x_s > (1 - p_r)x_r$. But their output in case of success is lower than that of risky firms, $x_s < x_r$. This assumption captures the fact that risky firms are not bad per se but might well have a higher possible output than safe firms, e.g., because their managers take higher risks. Relaxing this assumption strengthens rather than weakens the argument.

There are two types of banks in the market: an incumbent public bank, indexed pu, and a number of private banks, indexed pr, that attempt to enter the market. In the model, the sector of private banks is treated as one entity. Neither type is assumed to be budget-constrained, i.e., both could serve the whole market. However, banks are not allowed to incur expected losses in any period. This means that I do not consider price wars as in the entry games of Benoit (1983), Benoit (1984) or Fulghieri and Nagarajan (1996), where banks can fight competitors at the price of making losses in these periods. The public bank holds a state guarantee on its liabilities and enjoys lower funding costs than the private banks, $r_{pu} < r_{pr}$, where r_i includes face value and interest. R_i denotes the loan repayment that bank i charges its borrowers, with $R_i \geq r_i \geq 1$.

Financial distress in $t = \frac{1}{2}$ gives banks the right to foreclose loans, even though the loans are due only in $t = 1$. This may be due, e.g., to the firm defaulting on coupon payments (which I do not explicitly model).[1] Distress can be caused by two events. First, temporary liquidity shortages can force economically viable firms into distress. These firms produce output x_j in $t = 1$ if their loan is extended. Second, firms can be unprofitable due to, e.g., strategic mistakes, unfavorable market development etc. that occurred between $t = 0$ and $t = \frac{1}{2}$. These firms will not produce output in $t = 1$ even if their loan is extended. Their value in $t = 1$ is zero. Hence, in $t = \frac{1}{2}$ it is efficient for a bank to extend the loan in the first case but liquidate the firm in the second.

To distinguish the two cases, banks can audit firms. The audit reveals unprofitable firms with certainty and viable firms with an error margin. q denotes the quality of the audit, i.e., the conditional probability that an economically viable firm in financial distress is identified as such.[2] If the firm is identified as viable, the bank extends its loan but renegotiates the loan contract to obtain a share k_j of the output. The firm's output x_j is known in $t = \frac{1}{2}$, i.e., the firm type is revealed through the audit. For the renegotiation, I assume a form of Nash bargaining where both players receive equal shares of the surplus over the original repayment, $\frac{x_j - R_i}{2}$, and the bank additionally receives its repayment R_i. Hence, the bank receives a share $k_j = \frac{x_j + R_i}{2x_j}$ of x_j, while the firm receives $1 - k_j = \frac{x_j - R_i}{2x_j}$ of

[1] An alternative interpretation is that the bank does not grant an essential follow up loan, which leads to the insolvency of the firm due to liquidity problems.

[2] This means that I assume the audit to produce Type II error of $1 - q$ but no Type I error (see Chemmanur and Fulghieri (1994) for a similar renegotiation outcome). A possible explanation for this assumption is that banks only renegotiate a loan if they are certain that the expected output will be produced. If there are doubts, they prefer to liquidate. I do not explicitly include costs of the audit. However, they would strengthen rather than weaken the argument.

x_j. The audit is assumed to have a positive expected value, i.e., it is efficient to audit all firms in distress. Since it does not affect the argument, I set the discount factor between $t = 0$ and $t = 1$ to one.

4.3 Pooling

Consider the static game where banks interact only once. The two types of banks simultaneously enter the market for loans, and compete for borrowers. Since auditing firms in distress is ex-ante efficient, it is a dominant strategy in the static game. Accordingly, private banks cannot credibly commit to liquidating firms in distress. This means that they cannot separate borrower types, since all firms would be attracted if they offered a lower repayment. The quality of the audit is assumed to be the same for all banks, e.g., as the result of a commonly available audit technique. Banks compete only in loan repayments. Firms then obtain profits under pooling of

$$\begin{aligned} P_r^{pool} &= (x_r - R_i^{pool})(1 - p_r + \frac{1}{2}p_r p_v q) \\ P_s^{pool} &= (x_s - R_i^{pool})(1 - p_s + \frac{1}{2}p_s p_v q) \end{aligned}$$

where R_i^{pool} denotes the repayment that bank i charges a pooled firm population. Given the firms' participation constraints, $P_j^{pool} \geq 0$, banks can charge maximum repayments of $R_r^{max} = x_r$ for risky firms and $R_s^{max} = x_s$ for safe firms. Since $x_s < x_r$ by assumption, x_s is binding for the repayment if a bank wants to lend to all firms.

However, depending on the parameters, it may be more profitable for a bank to set $R_i^{pool} = x_r$ and lend to risky firms only, instead of serving the whole market at a repayment of x_s. Since credit rationing is not the focus of this paper, I assume that the share ϕ of safe firms in the population is large enough such that credit rationing is not profitable for the banks, i.e., the profit from lending to safe firms is not overcompensated by charging all risky firms a higher rate:

$$\textbf{A1:} \qquad \frac{\phi}{1-\phi} > \frac{(x_r - x_s)(1 - p_r + \frac{1}{2}p_r p_v q)}{(1-p_s)x_s + p_s(y + p_v q(x_s - y))} \qquad .$$

Banks then obtain profits

$$\begin{aligned} \Pi_{pu}^{pool} &= \phi[(1-p_s)R_{pu}^{pool} + p_s(y + p_v q(\frac{1}{2}(x_s + R_{pu}^{pool}) - y))] \\ &+ (1-\phi)[(1-p_r)R_{pu}^{pool} + p_r(y + p_v q(\frac{1}{2}(x_r + R_{pu}^{pool}) - y))] - r_{pu} \\ \Pi_{pr}^{pool} &= \phi[(1-p_s)R_{pr}^{pool} + p_s(y + p_v q(\frac{1}{2}(x_s + R_{pr}^{pool}) - y))] \\ &+ (1-\phi)[(1-p_r)R_{pr}^{pool} + p_r(y + p_v q(\frac{1}{2}(x_r + R_{pr}^{pool}) - y))] - r_{pr}. \end{aligned}$$

Neither type of bank is allowed to incur losses in expectation, i.e., banks' participation constraints are given by $\Pi_i^{pool} \geq 0$. Banks therefore have to charge minimum repayments of

$$\begin{aligned} R_{pu}^{pool,min} &= \frac{2(r_{pu} - \bar{p}y) + p_v q(2\bar{p}y - \bar{px})}{2(1-\bar{p}) + p_v q\bar{p}} \\ R_{pr}^{pool,min} &= \frac{2(r_{pr} - \bar{p}y) + p_v q(2\bar{p}y - \bar{px})}{2(1-\bar{p}) + p_v q\bar{p}} \end{aligned} \tag{4.1}$$

where $\bar{p} = \phi p_s + (1-\phi)p_r$ and $\bar{px} = \phi p_s x_s + (1-\phi)p_r x_r$.

With $r_{pu} < r_{pr}$, the public bank can charge a lower repayment than the private banks. Here as in the rest of the paper, I assume that if firms are indifferent between the two banks, they stay with their status quo bank. Since the public bank is the incumbent, initially it is the status quo bank for all firms. Accordingly, if banks compete, the public bank can offer $R_{pu}^{pool} = R_{pr}^{pool,min}$, the lowest profitable rate of the private banks, and lend to all firms in the market.

Proposition 4 .
i) Consider the situation where A1 is fulfilled and the public bank is the only lender in the market. In equilibrium the public bank with audit policy q charges all firms $R_{pu}^{pool*} = x_s$ *and serves the whole market.*
ii) Consider the situation where A1 is fulfilled and the public bank competes with private banks in the market for loans. In equilibrium the public bank with audit policy q charges all firms $R_{pu}^{pool*} = min\{x_s, R_{pr}^{pool,min}\}$ *and serves the whole market.*

The proof is in the appendix.

4.4 Separation

Consider now the case when it is the private banks' policy never to audit any firm in distress. In this section, I will simply assume that this policy is credible in the stage game, and that private banks cannot credibly change their policy from one period to the next. The conditions under which this is actually the case are derived in section 4.5.
As before, in each period all banks enter the market simultaneously. If loans were granted in the previous period, outcomes are realized and payments made before new loans are made.

Risky firms obtain profits under separation of

$$\begin{aligned} P_r^{sep,pu} &= (1-p_r)(x_r - R_{pu}^{sep}) + \frac{1}{2}p_r p_v q(x_r - R_{pu}^{sep}) \\ P_r^{sep,pr} &= (1-p_r)(x_r - R_{pr}^{sep}) \end{aligned}$$

from the public and private bank loan, respectively, with R_{pu}^{sep} and R_{pr}^{sep} denoting the banks' repayments under separation. Similarly, safe firms obtain

$$\begin{aligned} P_s^{sep,pu} &= (1-p_s)(x_s - R_{pu}^{sep}) + \frac{1}{2}p_s p_v q(x_s - R_{pu}^{sep}) \\ P_s^{sep,pr} &= (1-p_s)(x_s - R_{pr}^{sep}). \end{aligned}$$

As before, maximum repayments from firms' participation constraints are $R_r^{max} = x_r$ and $R_s^{max} = x_s$. Separation is obtained if firms self-select, i.e., if risky firms prefer the loan of the public bank, while safe firms prefer the loan of the private banks. This yields the incentive constraints

$$\begin{aligned} P_r^{sep,pu} &> P_r^{sep,pr} \quad \text{for risky firms and} \\ P_s^{sep,pu} &< P_s^{sep,pr} \quad \text{for safe firms.} \end{aligned}$$

Given the public bank's repayment R_{pu}, private banks have to charge a minimum repayment to deter risky firms from choosing its loan of

$$R_{pr}^{min} = R_{pu} - \frac{p_r}{2(1-p_r)} p_v q(x_r - R_{pu}).$$

To attract safe firms they can charge a maximum repayment just below

$$R_{pr}^{max} = R_{pu} - \frac{p_s}{2(1-p_s)} p_v q(x_s - R_{pu}) \quad . \tag{4.2}$$

Banks obtain profits per contract of

$$\begin{aligned} \Pi_{pu}^{sep} &= (1-p_r)R_{pu}^{sep} + p_r(y + p_v q(\frac{1}{2}(x_r + R_{pu}^{sep}) - y)) - r_{pu} \\ \Pi_{pr}^{sep} &= (1-p_s)R_{pr}^{sep} + p_s y - r_{pr} \quad . \end{aligned}$$

From their participation constraints result the banks' minimum feasible repayments under separation as

$$\begin{aligned} R_{pu}^{sep,min} &= \frac{r_{pu} - p_r y - p_r p_v q(\frac{1}{2}x_r - y)}{1 - p_r + \frac{1}{2}p_r p_v q} \\ R_{pr}^{sep,min} &= \frac{r_{pr} - p_s y}{1 - p_s} \quad . \end{aligned} \tag{4.3}$$

Separation is feasible if the minimum repayment that private banks have to set under separation in order not to make losses is lower than the maximum repayment they can set in order to attract safe firms. If the public bank tries to avoid separation, i.e., market entry of private banks, the lowest repayment it can charge is $R_{pu}^{pool,min}$ from (4.1), which depends on r_{pu}. This defines the critical disadvantage in funding costs r_{pr}^*, such that for funding costs below r_{pr}^* the private banks are able to enter the market:

$$r_{pr}^* = (1-p_s)R_{pu}^{pool,min} - p_s(p_v q(\frac{1}{2}(x_s - R_{pu}^{pool,min}) - y).$$

The details are in the appendix.

Assuming the credibility of the liquidation threat, the competitive equilibrium of the stage game is then defined as follows:
In equilibrium, private banks set the minimum feasible separation repayment from (4.3), while the public bank sets the maximum repayment that ensures separation given the private banks' equilibrium repayment:

$$\begin{aligned} R^*_{pr} &= \frac{r_{pr} - p_s y}{1 - p_s} \\ R^*_{pu} &= \frac{R^*_{pr} + \frac{p_s}{2(1-p_s)} p_v q x_s}{1 + \frac{p_s}{2(1-p_s)} p_v q}. \end{aligned}$$

4.5 Credibility

I now drop the assumption that the private banks' audit policy is credible and derive instead the conditions under which this is the case. For this I consider the game in which the lending process of the stage game is repeated an infinite number of times. The private banks' discount factor, reflecting their time preference and continuation probability (or time horizon), is δ.

Since the audit is assumed to have a higher expected return than liquidation, the public bank's policy to audit distressed firms is self-enforcing. For the private banks, credibility has to be achieved through reputation building. However, sticking to the threat to liquidate distressed firms without an audit leaves both players, bank and viable firm, worse off in the short run. Although the bank can execute the liquidation threat because distress gives it power over the firm, short-term rational behavior would induce it to audit all firms, as the expected return from the audit exceeds the liquidation value. Given rational expectations, risky firms would anticipate the bank's deviation from its liquidation policy and free-ride on the cheaper loan.

In order to make foregoing short-term profits from the audit profitable, and thus the liquidation threat credible and self-enforcing, the discounted profits from lending to safe firms in the future have to exceed the expected profits from extending the loan of distressed but economically viable firms today. Given the private banks' repayment under separation, R^{sep}_{pr}, credibility requires that

$$\frac{\Pi(R^{sep}_{pr})}{1 - \delta} \geq p_s p_v q (\frac{1}{2}(x_s + R^{sep}_{pr}) - y) \tag{4.4}$$

where $\Pi(R^{sep}_{pr})$ is the banks' profit per contract if charging repayment R^{sep}_{pr}. The details are in the appendix. If a private bank deviates from its liquidation strategy, the credibility of the liquidation threat is lost. Risky firms are then attracted by the lower repayment and separation fails. Accordingly, private banks cannot change their audit policy from one period to the next without compromising the policy's credibility. This justifies the assumption of section 4.4.

In the competitive equilibrium of section 4.4, the profits of the private banks are zero, such that (4.4) would fail. In order for credibility to be achieved, the profit of the private banks must be positive. This means that for a separating equilibrium that ensures credibility to exist, the following condition has to be satisfied: The minimum repayment from (4.4) that private banks have to set in order to credibly commit to liquidation and separate borrowers is lower than the maximum repayment from (4.2) they can set in order to attract safe firms, given that the public bank sets its lowest feasible pooling rate from (4.1). This yields the critical rate r_{pr}^{**} for the private banks' funding costs as

$$r_{pr}^{**} = (1-p_s)R_{pr} + p_s y - (1-\delta)p_s p_v q(\frac{x_s + R_{pr}}{2} - y) \tag{4.5}$$

where $R_{pr} = \frac{2(r_{pu}-\bar{p}y)+p_v q(2\bar{p}y-\bar{p}x)}{2(1-\bar{p})+p_v q\bar{p}} - \frac{p_s}{2(1-p_s)} p_v q(x_s - \frac{2(r_{pu}-\bar{p}y)+p_v q(2\bar{p}y-\bar{p}x)}{2(1-\bar{p})+p_v q\bar{p}})$.

In the competitive equilibrium of the repeated game, the private banks then set repayments such that (4.4) is fulfilled with equality, while the public bank sets the maximum separating repayment, given the equilibrium repayment of the private banks.

Proposition 5 *Assume that A1 is satisfied and $r_{pr} < r_{pr}^{**}$. In the competitive equilibrium of the infinitely repeated game, private banks liquidate all firms in distress, lend only to safe firms and set*

$$R_{pr}^{**} = \frac{(1-\delta)p_s p_v q(\frac{1}{2}x_s - y) - p_s y + r_{pr}}{1 - p_s - \frac{1-\delta}{2}p_s p_v q}.$$

The public bank with audit policy q lends only to risky firms and sets

$$R_{pu}^{**} = \frac{R_{pr}^{**} + \frac{p_s}{2(1-p_s)} p_v q x_s}{1 + \frac{p_s}{2(1-p_s)} p_v q}.$$

Setting R_{pr}^{**}, private banks make positive profits in equilibrium. Competition does not drive profits to zero, because this would render the liquidation threat non-credible, separation would fail and the private banks would be forced to leave the market.

What does separation mean for borrowers? As usual in self-selection models, the riskier firms are worse off under separation than under pooling. However, since private banks can sustain positive profits in equilibrium, safe firms may also be charged higher rates under separation than under pooling. In particular, if

$$1 - \delta > \frac{(r_{pr} - p_s y)a - (1-p_s)b}{\frac{1}{2}p_s p_v q((2y - x_s)a - b}$$

with $a = (2(1-\bar{p}) + p_v q\bar{p})$ and $b = (2(r_{pr} - \bar{p}y) + p_v q(2\bar{p}y - \bar{p}x))$, then $R_{pu}^{pool*} < R_{pr}^{**}$. That is, if the private banks' discount factor is sufficiently low (but still high enough to achieve separation), safe firms pay higher rates under separation than under pooling.

4.6 Conclusion

Most public banks directly or indirectly hold state guarantees on their deposits and enjoy lower funding costs than their privately owned competitors. This is often blamed to distort competition in favor of the public banks. The model shows that if public banks also have the mandate to support economic development, private banks can enter the market despite their cost disadvantage. With the public bank being *perceived* as supporting borrowers in financial distress, private banks have an instrument to separate firms and lend only to the safe types. This does not depend on the public bank's true objective function.
Interestingly, since the private banks' liquidation strategy is not credible if they earn zero profits, positive profits are sustainable in equilibrium even if there is perfect competition between private banks. The entry of private banks into the market may then lead to a deterioration of lending conditions for all firms, relative to the situation when private banks can only threaten to enter into the pooled market. The restriction of the public bank's policy space, which is meant to support economic development, may therefore result in the opposite effect.

Economic policy makers seem to react to the argument of competition distortion. For example, to comply with EU standards, the state's guarantee for its banks' liabilities was abandoned in Germany in July 2005. This was advertised as a step towards a level playing field for public and private banks. The results of this paper, however, cast doubt on this conclusion. Rather, they imply that as long as public banks are - at least in the public opinion - restricted in setting their policy, while private banks are not, competition is potentially distorted, with all the adverse effects this may induce. If economic policy makers want public banks to continue supporting the economic development, but otherwise want to foster free competition, the conditions on the loan market may not respond as desired.

Appendix

Proof of Proposition 4

Part i)

$x_r > x_s$ per assumption. If the public bank sets $R_{pu}^{pool} > x_r$, it violates the participation constraints of all firms, does not lend and obtains zero profit. If it sets $R_{pu}^{pool} < x_s$ it lends to all firms, but obtains lower profits than with $R_{pu}^{pool} = x_s$. If it sets $x_s < R_{pu}^{pool} \leq x_r$, it lends only to risky firms. Per assumption, this is less profitable for all $R_{pu}^{pool} \leq x_r$ than lending to the whole population at x_s. $\sqcap$

Part ii)

First, for $x_s < R_{pu}^{pool,min}$ the proof of part i) applies. Second, $x_s > R_{pu}^{pool,min}$. If the private banks offer a repayment below $R_{pr}^{pool,min}$ they incur losses. If the public bank charges $R_{pu}^{pool} < R_{pu}^{pool,min}$ it serves the whole market as before, but makes lower profits. If it charges $R_{pu}^{pool} > R_{pu}^{pool,min}$, the private banks offer $R_{pu}^{pool,min}$, serve the whole market, and the public bank makes zero profits. $\sqcap$

Derivation of r_{pr}^*

The maximum repayment the private banks can set in order to attract safe firms, dependent on the public bank's repayment, is given by safe firms' incentive constraint as

$$R_{pr}^{max} = R_{pu} - \frac{p_s}{1-p_s} p_v q \frac{1}{2}(x_s - R_{pu}).$$

The lowest feasible repayment the public bank can set under separation is $R_{pu}^{sep,min}$ as determined by its participation constraint from equation (4.3). Hence, one obtains

$$R_{pr}^{max} = R_{pu}^{sep,min} - \frac{p_s}{1-p_s} p_v q \frac{1}{2}(x_s - R_{pu}^{sep,min})$$

as the highest incentive compatible repayment the private banks can charge if the public bank competes with its lowest feasible separation repayment. In order for the private banks' participation constraint to be met, R_{pr}^{max} has to exceed their minimum repayment under separation, $R_{pr}^{sep,min}$ in (4.3). This gives the condition

$$R_{pu}^{sep,min} - \frac{p_s}{1-p_s} p_v q \frac{1}{2}(x_s - R_{pu}^{sep,min}) \geq \frac{r_{pr} - p_s y}{1-p_s} \tag{4.6}$$

Solving (4.6) for r_{pr} yields

$$r_{pr}^* = (1-p_s) R_{pu}^{sep,min} - p_s (p_v q \frac{1}{2}(x_s - R_{pu}^{sep,min}) - y)$$

Derivation of the credibility condition (4.4)

The profit that is lost per contract per period if distressed but viable firms are liquidated is given by $p_s p_v q(\frac{1}{2}(x_s + R_{pr}^{sep}) - y)$. The private banks lend to safe firms, of which p_s enter into distress. A share p_v of these is viable, but only q of them would be identified as such in an audit. Compared to liquidating them, the private bank would obtain an increase in profits by extending their loans of $\frac{1}{2}(x_s + R_{pr}^{sep}) - y$.

The profit the private bank forgoes by having to leave the market after having lost the credibility of its liquidation threat is determined as the present value of obtaining the profit $\Pi(R_{pr}^{sep})$ from cooperation in infinitely many periods, given its discount factor δ.

Note that since the number of firms a private bank lends to is the same over time as long as the market does not change, if the credibility condition holds for one contract, it holds for an arbitrary number of contracts a bank sells each period.

Bibliography

Y. Altunbas, L. Evans, and P. Molyneux. Bank ownership and efficiency. *Journal of Money, Credit & Banking*, 33(4):926–954, 2001.

J.-P. Benoit. Entry with exit: An extensive form treatment of predation with financial constraints. *IMSSS Technical Report (Stanford University)*, (405), 1983.

J.-P. Benoit. Financially constrained entry in a game with incomplete information. *RAND Journal of Economics*, 15:490–499, 1984.

R. Bichsel and C. Spielmann. State-owned banks as competition enhancers, or the grand illusion. *EFMA 2004 Basel Meetings Paper*, 2004.

P. Bolton and X. Freixas. Equity, bonds, and bank debt: Capital structure and financial market equilibrium under asymmetric information. *Journal of Political Economy*, 108 (2):324–351, 2000.

G. Caprio, J. Fiechter, R. Litan, and M. Pomerleano, editors. *The Future of State-Owned Financial Institutions: Policy and Practice*. Brookings Institution Press, 2005.

T. Chemmanur and P. Fulghieri. Reputation, renegotiation and the choice between bank loans and publicly traded debt. *Review of Financial Studies*, 7(3):475–506, 1994.

M. Dewatripont and E. Maskin. Credit and efficiency in centralized and decentralized economies. *Review of Economic Studies*, 62:541–555, 1995.

D. Diamond. Financial intermediation and delegated monitoring. *Review of Economic Studies*, 51(3):393–414, 1984.

P. Fulghieri and S. Nagarajan. On the strategic role of high leverage in entry deterrence. *Journal of Banking & Finance*, 20:1–23, 1996.

J. Houston and J. Christopher. Bank information monopolies and the mix of private and public debt claims. *Journal of Finance*, 51(5):1863–1889, 1996.

S. Johnson. An empirical analysis of the determinants of corporate debt ownership structure. *Journal of Financial and Quantitative Analysis*, 32(1):47–69, 1997.

R. La Porta, F. L. de Silanes, and A. Shleifer. Government ownership of banks. *Journal of Finance*, 57:265–301, 2002.

E. Levy-Yeyati, A. Micco, and U. Panizza. State-owned banks: Do they promote or depress financial development and economic growth? *Working Paper*, 2004.

E. Maskin. Theories of the soft budget constraint. *Japan and the World Economy*, 8: 125–133, 1996.

M. Rothschild and J. Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information (in symposium: The economics of information). *Quarterly Journal of Economics*, 90(4):629–649, 1976.

P. Sapienza. Lending incentives of state-owned banks. *Kellog School of Management, Finance Department, Working Paper*, (319), 2002.

Danksagung

Ich möchte mich bei meinen Betreuern, Prof. Dorothea Kübler und Prof. Franz Hubert bedanken, die mir während meiner Promotion alle akademische Freiheit liessen, und immer bereit waren, meine - realistischen und nicht so realistischen - Ideen zu diskutieren. Ausserdem gilt mein Dank meinen Kollegen an der Humboldt-Universität zu Berlin, der Technischen Universität Berlin, dem Wissenschaftszentrum Berlin und der University of California, Berkeley, mit denen ich meine Forschung auf Seminaren diskutieren konnte. Schliesslich möchte ich mich noch besonders bei meiner Familie und meinem Freund Ladislav bedanken, dafür, dass sie die geduldigen Pilotobjekte meiner Experimente waren, für unzählige Diskussionen, in denen es fast nie um Ökonomie ging, und für ihre Unterstützung, ohne die mein Leben nicht so schön wäre, wie es ist.

Selbständigkeitserklärung

Ich habe ausser der angeführten keine weitere Literatur bei der Erstellung meiner Dissertation verwendet. Auch wurde mir keine Hilfe zu Teil, die über die Diskussion des Forschungsgegenstandes mit meinen Betreuern sowie mit Kollegen auf Konferenzen und Workshops hinausging. Die Dissertation wurde nicht vorher von einer anderen Fakultät begutachtet.

Ich bezeuge durch meine Unterschrift, dass meine Angaben über die bei der Abfassung der Dissertation benutzten Hilfsmittel, über die mir zu Teil gewordene Hilfe sowie über führere Begutachtungen meiner Dissertation in jeder Hinsicht der Wahrheit entsprechen.