

Preferences for Consistency

Armin Falk and Florian Zimmermann

University of Bonn

19th October 2010

Preliminary and incomplete

Abstract

This paper studies how a “preference for consistency” can affect economic decision-making. We propose a three-period model of decision-making where people have a preference for consistency because consistent behavior allows them to signal personal and intellectual strength. This creates a trade-off between choosing according to their updated beliefs and reputational concerns. We then present three experiments that study main predictions and implications of the model. The first experiment is a simple principal-agent game that tests the basic logic of the model. In the main treatment agents solve an estimation task twice and are paid for accuracy. Principals observe the difference between the two estimates and select an agent for a third estimation task. Agents who are selected win a prize and engage in a third estimation task. In the control condition agents only solve the first two estimation tasks. In line with the model, principals select more consistent agents. Anticipating this, agents’ estimates are much more consistent in the main compared to the control treatment. The second experiment studies the role of early commitment. In a simple estimation task we show that subjects disregard valuable information when they have committed to a first statement. Compared to a control condition without early commitment, subjects’ estimates are further away from the correct value, implying lower earnings. The third experiment investigates consistency as a potential source of social influence. The decision context in this experiment is donations to charitable organizations. In the main condition subjects state hypothetically how much they would donate, prior to their actual decision. The hypothetical question is part of a survey. In the control condition subjects complete the same survey except for the hypothetical donation question. When asked hypothetically, participants give biased answers. We show that, in line with the model, this bias carries over to actual donation behavior due to participants’ taste for consistency. The experiment suggests that preferences for consistency can be a powerful channel of social influence.

1 Introduction

The desire to be and appear consistent is a powerful determinant of human behavior. Once we have made a decision or taken a stand, we often experience a strong pressure to live up to that commitment. In this paper we study how preferences for consistency can affect economic decision-making. In a first step we develop a simple model that specifies motives and conditions that lead to consistent behavior.

This allows predicting the role of consistency in different decision contexts. In a second step we provide results from three experiments that test main predictions and implications of the model.

Our model is built on the notion that consistent behavior is associated with “personal and intellectual strength.” (Cialdini, 1984). Evidence from social psychology suggests that while consistency signals positive personality traits, inconsistent behavior is generally viewed as an undesirable personality trait (see Asch, 1956 or Allgeier, Byrne, Brooks and Revnes, 1979). Inconsistent beliefs, words or deeds are often indicative of confusion or even mental illness. In the model decision-makers signal positive traits when acting consistently. They repeatedly face the same choice problem but are uncertain about their preferred outcome. Before taking a decision, they receive signals about their preferred outcome. Decision-makers differ with respect to the strength of signals they receive. The stronger the signals, the better they know their preferred outcome and the higher is their personal and intellectual strength. On the opposite, decision-makers with low personal and intellectual strength receive rather weak signals. Regarding behavior over time, personal and intellectual strength is associated with consistent choices. We assume that decision-makers’ utility consists of two parts, a “standard” part reflecting material concerns and a reputational part capturing decision-makers’ image concerns. They like to be perceived as intellectually strong. This creates a trade-off between choosing according to their updated beliefs and reputational concerns. We show that in a fully separating perfect Bayesian equilibrium decision-makers display preferences for consistency because consistent behavior allows them to signal high personal and intellectual strength. The model allows several interpretations. The value of signaling can be interpreted as hedonic, i.e., people like it if others think positively about them, or as strategic in the sense that people expect benefits from future interactions. Depending on context the signaling of strength is related to ability and knowledge, strong identity and personality or predictability and reliability. Finally, the recipient of the signal can be other people but also the decision-maker himself. In this case the decision-maker learns about himself by observing what he does (e.g., as in Bénabou and Tirole, 2006).

In the second part of the paper we report evidence from three decision experiments. The first experiment tests the basic logic of our model in a simple principal-agent framework. In the main treatment agents solve two estimation tasks and are paid for accuracy. The correct estimation result is the same in both tasks and both, agents and principals, know this. Principals are randomly matched with two agents. They observe the difference between the two estimates for both agents and select one of the two for a third estimation task. Principals are paid according to accuracy of the selected agent in the third estimation task. Therefore, they have an incentive to select the agent they believe is most able. Agents who are selected win a prize and engage in the third estimation task. Selection decisions of the principals inform us whether they value consistent behavior. In order to examine whether agents actually estimate more consistently we need to study a control treatment without principals. Agents in the control treatment simply perform the two estimation tasks, i.e., there is no reputational value of consistency. The results support the model’s predictions: Agents who estimate more consistently have a significantly higher chance of being selected by principals. Anticipating the value of consistency, agents’ estimations in the main treatment are much more consistent than in the control treatment.

The second experiment studies the role of early commitment for the preference for consistency. Intuitively, actively committing to an opinion, belief, intention or action is a precondition for observing consistent or inconsistent behavior. Without commitment, i.e., without taking a stand or an action, observers will not be able to detect possible inconsistencies. Therefore, a decision-maker is not constrained

by reputational concerns and can maximize utility without taking potential reputational costs into account. We test this model prediction in a simple estimation context. In two treatments subjects have to perform an estimation task and receive valuable information regarding the solution of the task. In the main treatment, subjects commit to a first estimate prior to receiving the helpful information and without knowing that they will later receive it. After they have received it they are free to revise their first estimate. In the control treatment, no prior commitment is made. The model predicts that in the main treatment, subjects will want to act consistently with their prior commitment. Therefore they will respond less strongly to the valuable information in comparison to the control treatment. Our findings support this prediction. The deviation of the final estimate from the valuable information is significantly higher in the main treatment than in the control treatment. Thus, the prior commitment makes subjects neglect valuable information leading to lower payoffs.

In the third experiment we examine the role of consistency preferences as a means of social influence. The trick is to “tempt” a person to make a biased statement. In a second step he or she is confronted with a request related to that statement and the pressure to live up to it. This way of influencing behavior is also known as the foot-in-the-door technique. Given that the first statement involves no or only low costs it is relatively easy to provoke biased statements. The preference for consistency will make the person want to live up to the biased statement and act against his or her actual interest. We test this prediction in the context of pro-social decision-making. Subjects decide on how much money they want to donate to some charity organization. We study two treatments: in the main treatment subjects are asked how much they would - hypothetically - donate if they were asked, prior to their donation decision. In the control treatment this question is not asked. The model predicts that due to the preference for consistency, biases in prior statements (how much would I donate?) should carry on to later choices because people feel a pressure to act consistently with the biased statement. Our results confirm this hypothesis. Donations in the main treatment significantly differ from donations in the control treatment, demonstrating the role of consistency as a means of social influence.

Our work is related to a large literature in social psychology. Cialdini, 1984 has put forward the idea that consistency is a signal of personal and intellectual strength. “It is at the heart of logic, rationality, stability and honesty.” (Cialdini, 1984). Cognitive dissonance theory offers another explanation for consistency preferences as consistent behavior helps people to avoid cognitive dissonances. The first theoretical foundation of that concept is developed in Heider, 1946, Newcomb, 1953 and Festinger, 1957. Yet another potential explanation for a taste for consistency can also be found in Cialdini, 1984, i.e., costs of thinking. If thinking is costly, consistent behavior offers a straightforward way to avoid these costs. Several interesting psychological experiments examine the consequences of consistency preferences for behavior. Famous examples are Freedman and Fraser, 1966 and Sherman, 1980 who analyze the effectiveness of the foot-in-the-door technique in different experiments. Cialdini et al., 1978 examine how the taste for consistency can be used to influence behavior via another channel, the so-called low-ball procedure. This technique aims to gain commitment to an action by not fully revealing its true consequences, for example costs. After commitment is reached, true consequences are revealed but the desire to appear consistent makes people stick to their commitment. While all these experiments offer fascinating insights into human behavior, they lack the standards of economic experiments, such as sufficient monetary incentives and no deception of subjects. In the economics literature Eyster, 2002 and Yariv, 2005 have put forward models of consistent choice. They assume that preferences for consistency are rooted in the desire to avoid cognitive dissonance. In Eyster, 2002 people try to rationalize past mistakes by taking

current actions that justify these mistakes, thereby reducing cognitive dissonances. Yariv, 2005 proposes a model where people reduce cognitive dissonance by sticking to beliefs under which past choices were optimal. Different to their approaches our model focuses on the value of signaling and the role of public reputation. Ellingsen and Johannesson, 2004 and Vanberg, 2008 refer to the taste for consistency as a possible reason for why people incur costs of lying. Regarding the modeling approach, our model is based on the model by Prendergast and Stole, 1996.

The remainder of the paper is organized as follows: In the next section we introduce our model. In section 3 we present the three experiments. Section 4 concludes.

2 The Model

In the model we formalize the intuition that “... a high degree of consistency is normally associated with personal and intellectual strength.” (Cialdini, 1984). To do so we modify the model of Prendergast and Stole, 1996. In the model, a decision-maker repeatedly chooses from a choice set but is uncertain about his preferred outcome. He receives a signal about his preferred outcome every period. Decision-makers differ in the strength of the signal they receive. The higher personal and intellectual strength the more precise are the signals. Therefore strength is associated with consistent behavior. Since decision-makers aim at having a reputation for personal and intellectual strength they display preferences for consistency because acting consistently behavior allows them to signal strength.

2.1 Baseline Model

There are three periods, $t = 0, 1, 2$. In periods 1 and 2, a decision-maker (D) chooses x_t from a choice set $X = (-\infty, \infty)$ in public. D has a unique preferred outcome $\mu \in X$, which he is uncertain about. In period 0, D and the public are holding an uninformative prior on μ .¹ In periods 1 and 2, before making his choice, D privately receives a signal about μ , $m_t = \mu + \epsilon_t$, $\epsilon_t \sim N(0, \sigma^2)$.

There are different types of decision-makers who differ in the strength of their signals σ . σ is D’s private knowledge, but is commonly known to be drawn from a uniform distribution over $[\underline{\sigma}, \bar{\sigma}]$. Thus $[\underline{\sigma}, \bar{\sigma}]$ defines the type space in our model. σ represents D’s personal and intellectual strength. It is decreasing in σ .

Upon receiving a signal m_t , D updates his beliefs about μ following Bayes’ rule. Since the period 0 prior is uninformative, D’s updated period 1 belief is completely based on the signal he received in period 1, m_1 . Accordingly, the updated expectation on μ in period 1 is $\hat{\mu}_1 = E(\mu|m_1) = m_1$. The variance of D’s updated period 1 belief is $Var(\mu|m_1) = \sigma^2$. In period 2, D receives an additional signal m_2 and updates again. Since signals in both periods are of equal strength, the updated period 2 expectation about μ is $\hat{\mu}_2 = E(\mu|m_1, m_2) = \frac{1}{2}m_1 + \frac{1}{2}m_2$. The variance of D’s updated period 2 belief is $Var(\mu|m_1, m_2) = \frac{1}{2}\sigma^2$.

In periods 1 and 2, the decision-maker chooses x_t in order to maximize utility. D’s utility function has two components. The first is “standard” outcome-based utility. Decision-makers want to minimize the quadratic distance between x_t and μ

¹Technically we assume that the prior is normally distributed with variance τ^2 , where $\tau \rightarrow \infty$.

$$-(x_t - \mu)^2. \tag{1}$$

Thus, standard utility is always maximized by choosing $x_t = \mu$.

The second component of utility is reputational concerns. Decision-makers like if the public perceives them as receiving strong signals, i.e., as having high personal and intellectual strength. An alternative interpretation of the reputational concern is a desire for a positive self-image (similar as in Bénabou and Tirole, 2006). In this case decision-makers receive a perfect signal about their personal and intellectual strength prior to their decision. Thus, when deciding they know their strength. However, for their self-assessment, the signal is not available, e.g., due to reasons of imperfect recall. Since actions are easier to recall than signals, decision-makers use past actions for their self-assessment. Thus the model is compatible with the intuition that people care about their self-image and construct self-image based on past actions. In the following we describe the model mostly in terms of public reputation but always mean to include a self-signaling interpretation. Reputational concerns are expressed by

$$-\alpha(E(\sigma | x_t, h_{t-1})). \tag{2}$$

$E(\sigma | x_t, h_{t-1})$ denotes the public's (or D's) expectation of D's σ , conditional on D's choice x_t and the history of choices h_{t-1} if such a history exists. The parameter α specifies how much D cares about his reputation. Differences in α might reflect, e.g., the size of the public, the social distance between D and the public, or the strategic value of reputation. In labor relations, e.g., signaling higher abilities may lead to higher wages and promotion. Alternatively, α could reflect a hedonic value of reputation in the sense that people simply enjoy being regarded as intellectually and personally strong.

Putting these two components together, in periods 1 and 2, D solves the following maximization problem, facing a trade-off between maximizing outcome-based utility and gaining reputation:

$$\max_{x_t} U_t = -(x_t - \hat{\mu}_t)^2 - \alpha(E(\sigma | x_t, h_{t-1})) \tag{3}$$

Note that we assume myopic, non-forward looking behavior. Decision-makers are not anticipating period 2 decisions when deciding in period 1. This can be justified in two ways. First, we are interested in situations where future decisions are not anticipated and believe that this is quite common. Second, and more importantly, the model goes through if decision-makers are actually forward looking. The reason is that given a normally distributed signal in period 1, the learning direction is unknown. It is therefore unclear how to improve in terms of consistent behavior even in situations where the decision-maker anticipates future decisions.

2.2 Equilibrium

We now turn to equilibrium behavior. The problem we consider is one of signal-extraction. In line with this class of models we focus on separating equilibria.

In periods 1 and 2, D receives a signal m_t , updates his beliefs $\hat{\mu}_t$ and makes a choice x_t that maximizes (3). The public only observes D's actions x_t and subsequently forms beliefs about the decision-maker's

type σ . Notice from (3) that D's optimal choice x_t only depends on his current belief $\hat{\mu}_t$ and the history of actions h_{t-1} . From this it follows that in equilibrium, we necessarily have pooling of all types of decision-makers (type space $[\underline{\sigma}, \bar{\sigma}]$) who hold the same belief $\hat{\mu}_t$ and share the same history h_{t-1} . We therefore transform the typespace. A type is now defined by the choice history h_{t-1} and the current belief $\hat{\mu}_t$. For the new typespace, we can state the definition of a fully separating Perfect Bayesian Nash Equilibrium (PBE).

Definition: $x_t^*(\hat{\mu}_t, h_{t-1})$, $\hat{\mu}_t^*(x_t, h_{t-1})$ (where $\hat{\mu}_t^*$ is the public's (or according to the self-signaling interpretation, D's own) belief about $\hat{\mu}_t$) is a fully separating Perfect Bayesian Equilibrium iff $\forall h_{t-1}$, $x_t^* \in \operatorname{argmax} [-(x_t - \hat{\mu}_t)^2 - \alpha(E(\sigma | \hat{\mu}_t^*(x_t, h_{t-1}), h_{t-1}))]$ and $\hat{\mu}_t^*$ is the inverse of x_t^* over $\hat{\mu}_t$.

The equilibrium consists of two functions, $x_t^*(\hat{\mu}_t, h_{t-1})$ and $\hat{\mu}_t^*(x_t, h_{t-1})$. $x_t^*(\hat{\mu}_t, h_{t-1})$ describes equilibrium behavior of the decision-maker conditional on the beliefs he holds, and $\hat{\mu}_t^*(x_t, h_{t-1})$ denotes the public's beliefs about $\hat{\mu}_t$, which will be correct in equilibrium. The equilibrium fully separates decision-makers with different beliefs, conditional on the choice history h_{t-1} . Thus, the public can form its expectations of D's σ directly from D's beliefs - $E(\sigma | \hat{\mu}_t, \hat{\mu}_{t-1})$.

We now consider equilibrium behavior in periods 1 and 2 in more detail:

Period 1

In period 1, there exists no choice-history. Thus, the public can form its expectation of σ solely based on $\hat{\mu}_1$, which it infers from the observation of x_1 . However, by only observing x_1 or $\hat{\mu}_1$, respectively, the public gains no additional knowledge regarding the type of the decision-maker σ . The prior, that σ is uniformly distributed on $[\underline{\sigma}, \bar{\sigma}]$, remains unchanged. In other words, we have

$$\frac{\partial E(\sigma | \hat{\mu}_1)}{\partial \hat{\mu}_1} = 0.$$

From the perspective of the decision-maker this means that D cannot affect his reputation. As a consequence, he simply maximizes outcome-based utility by choosing $x_1^* = \hat{\mu}_1$. The public observes D's choice and correctly infers the signal he or she received, i.e., $\hat{\mu}_t^*(x_1) = \hat{\mu}_1$.

Period 2

In period 2, the decision problem is more involving. From observing D's period 2 choice, the public now learns something about D's personal and intellectual strength. Thus when deciding on x_2 D now takes into account possible reputation effects. We will show that in order to signal strength, D chooses x_2 consistently with his period 1 choice. In other words, he picks a period 2 choice which, regarding outcome-based utility, is suboptimally close to his or her period 1 choice.

We start by describing the updating process of the public regarding D's personal and intellectual strength. In a fully separating equilibrium, the public is able to infer the true beliefs $\hat{\mu}_t$ D is holding from his choices. The public's expectation on σ in period 2 therefore depends on $\hat{\mu}_1$ and $\hat{\mu}_2$. For this expectation, we can show that the following Lemma holds.

Lemma 1: For the expectation of σ , the following holds: $\frac{\partial E(\sigma|\hat{\mu}_1, \hat{\mu}_2)}{\partial \hat{\mu}_2} > (<)0$ if $\hat{\mu}_2 > (<)\hat{\mu}_1$.

Lemma 1 states that the public's belief about σ increases in the distance between D's beliefs in periods 1 and 2. Put differently, the closer $\hat{\mu}_1$ and $\hat{\mu}_2$, the stronger is the signal of personal and intellectual strength. In equilibrium, D will therefore signal strength by choosing x_2 close to x_1 , i.e., by acting as if his beliefs in period 1 and 2 were close to each other. This is stated in Proposition 1.

Proposition 1: If α is not "too large", i.e., $\sup_{\hat{\mu}_2} \alpha \frac{\partial^2 E(\sigma|\hat{\mu}_2, \hat{\mu}_1)}{\partial \hat{\mu}_2^2} < \frac{1}{4}$, there exists a unique fully separating Perfect Bayesian Equilibrium in period 2, where $\hat{\mu}_2 > x_2^* > \hat{\mu}_1$ or $\hat{\mu}_2 < x_2^* < \hat{\mu}_1$.

Proposition 1 states the main result of our model, that decision-makers have a preference for consistent behavior. In period 2, decision-makers do not simply maximize standard outcome-based utility. Instead, they take decisions in between their period 2 belief and their period 1 choice. They act more consistently with their period 1 choice in order to signal high personal and intellectual strength. Thus, in equilibrium, they are willing to sacrifice outcome-based utility to increase their reputational utility.

The result of Proposition 1 is mainly driven by D's reputational concerns. The following Lemma states that D's preference for consistency is more pronounced, the larger D's reputational concern is.

Lemma 2: Consider the fully separating PBE for a given reputational concern α_1 and signals $\hat{\mu}_1$ and $\hat{\mu}_2$. Compare this with the fully separating PBE for $\alpha_2 > \alpha_1$ and the same signals $\hat{\mu}_1$ and $\hat{\mu}_2$. Then the following inequality holds: $\hat{\mu}_2 > x_2^*(\alpha_1) > x_2^*(\alpha_2) > \hat{\mu}_1$ or $\hat{\mu}_2 < x_2^*(\alpha_1) < x_2^*(\alpha_2) < \hat{\mu}_1$.

2.3 Reduced Version of the Model

In the following we briefly describe a reduced and simplified version of the model, where we directly assume a taste for consistency. This model can be used to analyse the consequences of a taste for consistency in complex economic environments in a straightforward way. We derive hypotheses for our experiments from this model.

Again there are three periods, $t = 0, 1, 2$. In periods 1 and 2, the decision-maker chooses x_t from a choice set $X = (-\infty, \infty)$. D has a unique preferred outcome $\mu \in X$, which he is uncertain about. In period 0, D and the public are holding an uninformative prior on μ . In periods 1 and 2, before taking a decision, D privately receives a signal on μ , $m_t = \mu + \epsilon_t$, $\epsilon_t \sim N(0, \sigma^2)$. D updates his beliefs following Bayes' rule. Consequently, period 1 belief is $\hat{\mu}_1 = E(\mu|m_1) = m_1$. In period 2, we have $\hat{\mu}_2 = E(\mu|m_1, m_2) = \frac{1}{2}m_1 + \frac{1}{2}m_2$.

We directly assume a taste for consistency in D's utility function. If a choice history h_{t-1} exists, D wants to be consistent with this history. D solves the following maximization problem:

$$\max_{x_t} U_t = -(x_t - \hat{\mu}_t)^2 - \alpha(x_t - h_{t-1})^2.$$

As in the main model α specifies how much D cares about reputation.

If no choice history exists, D simply solves:

$$\max_{x_t} U_t = -(x_t - \hat{\mu}_t)^2.$$

It is straightforward to derive behavior in periods 1 and 2. In period 1 there is no choice history. In the absence of past choices to be consistent with, D simply chooses $x_1^* = \hat{\mu}_1$. In period 2, D faces a trade-off between being consistent with the period 1 choice and following the updated period 2 beliefs. Consequently he will choose $\hat{\mu}_2 > x_2^* > \hat{\mu}_1$ or $\hat{\mu}_2 < x_2^* < \hat{\mu}_1$ respectively.

3 Experimental Evidence

In this section we present evidence from three experiments designed to test main predictions and implications of the model. In particular we address the following issues: First, we study the signal value of consistent behavior in the context of a simple principal agent game. We test whether principals reward consistent behavior and whether agents anticipate this. In the second experiment we investigate the role of commitment. The idea is that without explicit commitment to some point of view, statement or intention, inconsistencies are impossible to detect and consequently consistency is less important. Our third experiment demonstrates the importance of consistency preferences for social influence and manipulation.

3.1 Experiment 1: The Value of Consistency

Design: The first experiment tests the basic logic of our model. The central assumption is that consistent behavior is viewed positively as it signals ability, personal strength or reliability. In the experiment the decision context involves a simple estimation task. We test whether principals infer higher ability from more consistent estimates and whether agents anticipate this.

We study two treatments. In the main treatment, upon arrival, subjects were randomly assigned to the roles of principals and agents.² In each session there were twice as many agents than principals. Subjects were seated in separate rooms according to their roles. The experiment involved two stages. All subjects were informed about both stages at the beginning. In the first stage, agents had to perform two estimation tasks. The task was to estimate how many times the letter “e” appeared in a text with 1966 letters. In both texts the number of e’s was identical and agents and principals were informed about this. The correct number was 233.³ Subjects saw the first text for 60 seconds on their computer screen. Then they had 60 seconds to state their estimate about the number of e’s (first estimate). Without getting any feedback on the first task, agents then saw the second text for 60 seconds. Again they had 60 seconds to provide their estimate for the second text (second estimate). Subjects were paid for accuracy. For both estimates the following rule applied: The maximum profit for each estimate was five Euro, which the agent received if his or her estimate was less than 1 percent above or below the correct number. For every percentage point the estimate deviated from the correct number, 10 Cents were deducted. If the estimate was more than 50 percent above or below the true value, profits were zero. Negative earnings were not possible.

After all agents had completed their two estimates, the second stage began. Two agents were randomly assigned to one principal. The decision of the principal was to select one of the two agents for a third estimation task, which was known to be similar to the first two estimation tasks. The principal was paid

²In the experiment, we used a neutral framing, i.e., subjects were called “participants A”, and “participants B”, respectively.

³The two texts are shown in Appendix A.

according to the precision of the estimate in this third task. Thus principals had an incentive to select the agent who they thought is most able in the estimation task. For their decision, principals were informed about the absolute difference between the first and the second estimate for both agents assigned to them. This information was provided on an answer sheet.⁴ Note that we did not display the four single estimates of the two agents. The reason is that this would have allowed principals to use one agent’s estimates to infer something about the other agent’s strength. Formally, the two estimates of agent 1 are a signal about the true preferences of agent 2, given that both have the same preference (they want to estimate the identical number). If principals observe this signal, agent’s 2 estimates reveal different information about σ than in the absence of that knowledge. Therefore we decided only to provide absolute differences between estimates to the principals, thereby only allowing the kind inference on σ that is captured by our model.

On their answer sheet principals had to select “their” agent. The third task was to estimate how many times the letter “a” appears in a text of again 1966 letters. Principals were paid according to the accuracy of the selected agent’s third estimate. The maximum payoff was ten Euro, which was paid if the agent’s estimate was less than 1 percent above or below the correct number. For every percentage point the estimate deviated from the correct number, 20 Cents were deducted. If the estimate was more than 50 percent above or below the correct value, the principal’s payoff was zero. Agents had an incentive to be selected and to estimate as precise as possible in the third task. For being selected they received a prize of 10 Euro. In addition, they were paid according to accuracy where payment was identical to the payment scheme in the first two estimates.

Principals’ selection decisions inform us about the potential value of estimating consistently. However, to examine whether agents anticipate this and actually behave more consistently we need an additional treatment that eliminates (or reduces) the importance of the reputational concern for consistency. This is what we do in the control treatment. The control treatment is simply the first stage of the main treatment, i.e., agents estimate how many times the letter “e” appeared in the texts used in the main treatment. The payoff scheme for the two tasks was identical to that in the main treatment. Comparing behavior between the two treatments informs us whether agents anticipate the value of consistency and therefore behave more consistently in the main treatment.

Procedural Details: A total of 168 subjects participated in six sessions. In the main treatment, 64 subjects participated as agents and 32 as principals. 72 subjects participated in the control treatment. Subjects were mostly students from various fields at the University of Bonn. No subject participated in more than one session. The experiment was run using the experimental software z-Tree (Fischbacher, 2007). The principals took their choice on an answer sheet. Sessions lasted on average about 60 minutes in the main treatment and 45 minutes in the control treatment. Average earnings were 12.90 Euro for principals and 12.06 Euro for agents, including a show-up fee of eight Euro for principals and four Euro for agents. (One Euro was worth about 1.35 USDollar at the time).

Note that subjects were neither paid nor informed about the correct solution of the estimation task before all sessions had been concluded (main and control treatment). This was done to avoid the possibility that feedback concerning the correct result of the estimation tasks would spread and confound results. In order to receive their payments, subjects were asked to go to the experimenter’s office.

⁴Information was given in an anonymous way, i.e., principals could not link information to actual subjects in the experiment.

Hypotheses: Given that there is a value of consistent behavior, agents who estimate more consistently (low absolute difference between estimates) should have a higher probability of being selected by principals. If agents anticipate this, estimates should be more consistent in the main treatment compared to the control treatment. Formally, the treatment variation of our experiment corresponds to a change in the reputational concern α . In their first estimate, subjects simply answer according to their beliefs, i.e., $x_1 = \hat{\mu}_1$. For their second estimate, subjects face a trade-off between acting according to their updated belief or estimating more consistently with their first estimate. The solution to their maximization problem is simply

$$x_2 = \frac{1}{1 + \alpha} \hat{\mu}_2 + \frac{\alpha}{1 + \alpha} x_1.$$

It immediately follows that an increase in α leads to more consistent behavior, i.e., the absolute deviation between x_2 and x_1 should be smaller in the main treatment, compared to the control treatment. Together with the assumption that consistent behavior signals ability we state our predictions as follows:

Note that we do not model the tournament directly. In the experiment we implemented a tournament incentive structure where agents win a prize if they outperform the other agent. Thus, on an individual level, optimal behavior in the main treatment is likely to depend on an agent's belief about the other agent's choice and ability. For our purposes however it is sufficient to take the tournament incentives as a strategic value of behaving consistently which is not present in the control treatment. In other words the experiment can be thought of as capturing a situation where α represents a strategic value of consistency, for example in a labor market context. In this sense, we test whether consistent behavior signals ability and is rewarded with higher pay or promotion and if this reward is anticipated.

Hypothesis 1 (i) The likelihood that an agent is selected by a principal decreases in the absolute difference between first and second estimate. (ii) The absolute difference between first and second estimate is smaller in the main treatment compared to the control treatment.

Results: The first result concerns the selection decisions of principals. In line with our hypothesis, a higher absolute difference between the two estimates decreases the likelihood of being selected. *Figure 1* shows that the likelihood of being selected is about 70 percent for differences between zero and five and decline for larger differences.⁵ For differences larger than 31, e.g., the likelihood drops to about 22 percent. The decrease in likelihood is significant as a simple Probit regression shows. When we regress the probability of being selected on the absolute differences between the estimates we get a negative and significant coefficient (p-value <0.01). The marginal effect is -0.012, indicating that an increase in the absolute difference of one point decreases the likelihood of being selected by about 1.2 percent. Further evidence comes from the observation that among all principals 75 percent select the agent with the smaller absolute difference. A simple binomial test rejects the null hypothesis that principals randomized with equal probability (p-value <0.01).⁶

We now turn to agents' behavior. The correct answer for both estimations was 233. Using all estimates (main and control treatment), the average first estimate was 220.57 while the average second estimate

⁵All figures are shown in section 6.

⁶In those 25 percent of the cases where agents did not select the agent with smaller difference in estimates, it was usually the case that the smaller difference was extremely small (0,1 or 2). We speculate that very small differences between estimates may have created some conspicuousness among principals.

was 215.64. The variance in estimates was rather high. The standard deviation of all estimates in the first task was 83.29, in the second it was 74.39. At the end of the experiment, we asked agents in the control treatment to briefly describe their estimation strategy for the two estimation tasks.⁷ Almost all decision-makers who answered the question described a similar procedure. First they counted the number of e's for a couple of rows. Then they counted the total amount of rows in the text and projected the total number of e's in the text.

Consistent with our model, in the control treatment, smaller differences between estimates do in fact reflect differences in estimation ability. More consistent agents are more able: The correlation between absolute difference in estimates and precision of estimates (measured as the sum of absolute distance between first estimate and true value and second estimate and true value) is strong and significantly positive (corr. coefficient is 0.429, with p-value <0.001). Thus using the difference between the estimates as a signal of ability is in general justified. Note, however, that differences in estimates do not coincide with better precision in the main treatment. In a separating equilibrium this is what we would expect. One way to interpret this finding is that incentives to signal consistency were so strong that the equilibrium is better characterized as a pooling rather than as a separating equilibrium. It would be interesting to test this idea by gradually reducing the value of α , e.g., by reducing the value of the prize selected agents receive.

Given principals' behavior and the fact that being selected was associated with a relatively large prize (10 Euro plus earnings from a third estimate), the model predicts that agents chose more consistently in the main compared to the control treatment. This is in fact what we find. *Figure 2* shows scatterplots of first and second estimates. The left panel displays observations from the control treatment, the right panel from the main treatment. While estimates are correlated in both treatments, it is obvious that the correlation is much tighter in the main treatment, i.e., that decisions are much more consistent. The correlation coefficients are 0.37 in the control treatment and 0.94 in the main treatment, respectively. *Figure 3* shows a histogram of absolute differences between the two estimates for both treatments. While about 70 percent of agents report an estimation difference below 15 in the main treatment the respective number is only 25 percent in the control treatment. More than 40 percent of agents in the control treatment indicated differences larger than 30 while less than 15 percent do so in the main treatment. The average absolute difference between estimates in the control treatment is 53.8 (std. dev. 70.5) while it is only 17.3 (std. dev. 23.2) in the main treatment. This difference is significant at any conventional level (p-value < 0.0001, using either Wilcoxon rank-sum test or simple OLS regression, regressing the absolute difference on a constant and a treatment dummy).

We summarize our main results as follows:

Result 1 The likelihood of being selected decreases in the absolute difference between the two estimates. Agents anticipate this. As a consequence, the absolute difference between first and second estimate is significantly lower in the main treatment, compared to the control treatment.

⁷The question we asked was the following: "Please briefly describe how you proceeded in the two estimation tasks. How did you get to your estimation results?"

3.2 Experiment 2: The Role of Commitment

Our second experiment studies the role of early commitment. Intuitively, actively committing to an opinion, belief, intention or action is a precondition for observing consistent or inconsistent behaviors. Without commitment, i.e., without taking a stand or an action, observers will not be able to detect possible inconsistencies. Therefore, a decision-maker is not constrained by reputational concerns and can maximize utility without taking a reputational cost into account. In contrast, once we have committed to an opinion or belief, we cannot easily change our mind without revealing some inconsistency. We test this intuition and prediction of the model in the context of an estimation task and show how commitment to an opinion can make people disregard valuable information.

Design: To investigate the importance of commitment we study two treatments, one with commitment (main treatment) and one without (control treatment). The different steps of the experiment are illustrated in *Figure 4*. The main treatment is shown in the upper panel of *Figure 4*. First, subjects were explained the task: Subjects had to estimate the number of peas in a bowl.⁸ Subjects were paid according to the precision of their estimate. If their estimate was less than 5 percentage points above or below the true value of 3000, subjects earned 10 Euro. For every 5 percentage points, the estimate deviated from the true value, we deducted 50 Cents. For example, a subject whose estimate deviated 17 percent from the true value earned 8.50 Euro. Negative earnings from the estimation task were not possible.

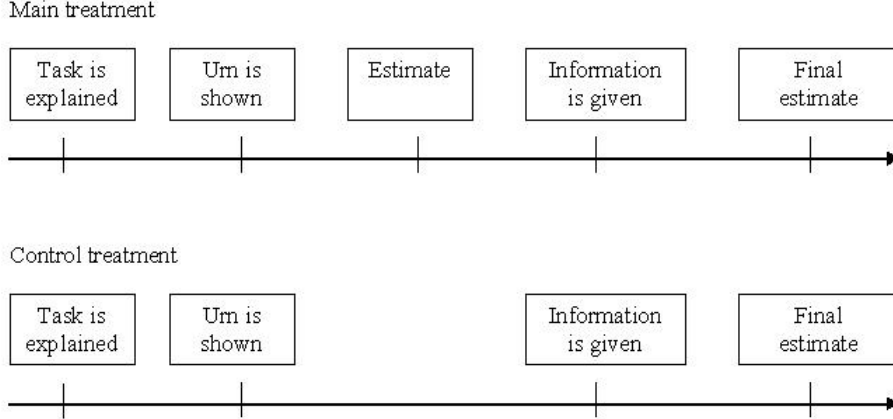
Subjects were seated around a table, which was placed in the middle of the lab.⁹ After subjects knew about the task the bowl was shown. The bowl with peas was placed in the middle of the table. Subjects were told that they had as much time as they wanted to estimate the number of peas. They were asked to write down their estimate on an answer sheet that had been distributed at the beginning of the experiment. As soon as a subject indicated that he or she had written down an estimate, the experimenter went to the subject and copied the subject’s estimate into a table. This means that subjects had written down their first estimate and knew that the experimenter knew this estimate. At this point, subjects had committed to their first estimate. After all subjects had stated their estimates, the experimenter announced that he would now provide subjects with additional and “helpful” information regarding the estimation task. Each subject received an information sheet containing the following sentence. “In the past it has often been the case in various estimation tasks, that the average estimate of all participants is often relatively close to the true value. The estimation task you are facing has also been conducted with a different group of participants. They have also been paid according to precision of their estimates. The average estimate of the number of peas in the bowl of this group was 2615. If you want to, you can now revise your estimate.” After they received the information sheet, subjects had time to revise their information on their answer sheet. After all subjects had indicated that they had specified their final estimate, the experimenter collected their answer sheets and the estimation task ended.

The additional information we provided to subjects was based on a separate experiment we had conducted with 61 different subjects. They faced the same estimation task and were also paid according to the precision of their estimates. The average estimate of that group was 2615. In the results section we show that the additional information was in fact valuable to subjects.

⁸A picture of the bowl is shown in Appendix A.

⁹Subjects were seated sufficiently far away from each other, such that they could not see what other subjects were writing down.

Figure 4: Timing of the experiment



The only difference between main and control treatment was that in the control treatment subjects did not state an estimate prior to receiving the valuable information (see lower panel of *Figure 4*). In the control treatment, subjects saw the bowl with peas for some time prior to receiving the information sheets. The time was approximating the time that subjects in the main treatment needed. During this time subjects could form a belief about the correct number of peas. However, they did not write their estimate down, nor did the experimenter get to know this estimate. In this sense, no commitment was made. After subjects received the information sheets, they stated their estimate on an answer sheet. After all subjects had provided their estimates, the experimenter collected the answer sheets and the estimation task ended.

Procedural Details: A total of 54 subjects participated in four sessions, 28 in the main and 26 in the control treatment. Subjects were mostly undergraduate students from various fields at the University of Bonn. No subject participated in more than one session. The experiment was conducted paper and pencil. Sessions lasted on average about 45 minutes. Subjects earned on average 12.31 Euro, including a show-up fee of 5 Euro.

Subjects were neither paid nor informed about the correct solution of the estimation task before all sessions had been concluded (main and control treatment). This was done to avoid the possibility that feedback concerning the correct result of the estimation tasks would spread and confound results. In order to receive their payments, subjects were asked to go to the experimenter's office.

Hypotheses: In both treatments subjects see the bowl with peas and form a belief about the correct number of peas. In terms of the model, they receive a signal m_1 , and form $\hat{\mu}_1 = m_1$ accordingly. When asked to make a first statement, they will choose $x_1 = \hat{\mu}_1$. Then subjects receive valuable information, an exogenous signal $m_s = \mu + \epsilon_s$, $\epsilon_s \sim N(0, \sigma_s^2)$. They update beliefs following Bayes' rule: $\hat{\mu}_2 = E(\mu|m_1, m_s) = \frac{\sigma_s^2}{\sigma_1^2 + \sigma_s^2} m_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_s^2} m_I$. For their final choice x_2 , subjects now face the trade-off between being consistent with their prior commitment and stating their updated belief. They will choose $\hat{\mu}_2 > x_2^* > \hat{\mu}_1$ or $\hat{\mu}_2 < x_2^* < \hat{\mu}_1$ respectively, thereby partially neglecting valuable information. Compare this with a situation where subjects receive exactly the same signals, but do not commit to a first opinion. Then, when making their final choice x_2 , there is no prior commitment to be consistent with, and subjects will decide based on their updated beliefs, i.e., $x_2 = \hat{\mu}_2 = \frac{\sigma_s^2}{\sigma_1^2 + \sigma_s^2} m_1 + \frac{\sigma_1^2}{\sigma_1^2 + \sigma_s^2} m_I$. Thus, without a prior commitment, subjects are taking the valuable information completely into account when taking their final

decision. Since the distribution of beliefs should be identical between treatments, we hypothesize that on average, final estimates in the main treatment will be further away from the valuable information than final estimates in the control treatment. In as much as information is actually valuable, final estimates will also be closer to the correct solution in the control treatment compared to the main treatment, implying higher earnings of subjects in the control treatment. We summarize our hypotheses as follows:

Hypothesis 2 The absolute difference between the final estimate and the information value of 2615 should be higher in the main treatment, compared to the control treatment. Since subjects in the main treatment disregard valuable information, it follows directly that the quality of estimates and therefore earnings are lower in the main than in the control treatment. Thus, on average, final estimates in the main treatment will be further away from the correct value of 3000 compared to the control treatment.

Results: Pooling data from both treatments the average (final) estimate was 2552.5. The estimation problem is very difficult and answers range from 500 to 5800. Accordingly, the variance was large as indicated by a standard deviation of 1021.865. We chose a difficult task on purpose as it offers an ideal context to provide subjects with helpful information. To show that the information was in fact valuable we simply count the number of subjects in the main treatment whose estimate in the first estimation was closer to 3000 than 2615. It turns out that this holds for only 5 out of 28 subjects. This means that about 82 percent of subjects could improve their (first) estimate by simply choosing 2615 or by moving in this direction.

We now turn to our main variable of interest, the absolute deviation between final estimate and the information value of 2615. *Figure 5* shows a histogram of estimates in intervals of 300 around the information value. In the control treatment about 54 percent of all estimates are in the interval ± 300 around the information value. In contrast only about 28 percent of all estimates in the main treatment lie within this interval. The figure also shows that extreme deviations from the information value are more frequent in the main compared to the control treatment. On average, the deviation in the main treatment is 464.13 points higher than in the control treatment. Related to the information value, this is a difference of about 18 percent. The difference in deviations from the information value is statistically significant (p-value < 0.07 , using Wilcoxon rank-sum test or simple OLS regression, regressing the absolute difference on a constant and a treatment dummy (p-value < 0.02)).

Figure 6 suggests that early commitment in the main treatment affects subjects' final estimate. It shows a scatterplot with subjects' first and final estimates together with a line indicating the information value 2615. The figure reveals that many subjects are either at or close to the 45-degree line indicating a strong resistance to take into account new and valuable information. It also shows that if subjects change, they change in the direction of 2615, as predicted by the model. The correlation between first and final estimate is 0.53 (p-value < 0.004).

The disregard of valuable information is associated with a decrease in the quality of estimates and earnings. On average, estimates in the main treatment are 512.46 points further away from the correct value than estimates in the control treatment. The effect is statistically significant (p-value < 0.03 , using Wilcoxon rank-sum test or simple OLS regression, regressing the deviation from the true value on a constant and a treatment dummy (p-value < 0.01)).

We summarize our results as follows:

Result 2 Commitment is key: it induces subjects to act consistently at the cost of disregarding valuable information and receiving lower payoffs.

3.3 Experiment 3: Social Influence

In our third experiment we address the issue of social influence. The idea is that given a preference for consistency, it is possible to influence and manipulate people. The trick is to “tempt” a person to make a biased statement. In a second step, he or she is confronted with a request related to that statement and the pressure to live up to it. Given that the first statement typically involves no or only low costs, it is relatively easy to provoke statements that do not necessarily reflect a person’s true preferences. As a result the person may end up acting against his or her actual interest. This technique is also known as foot-in-the-door and has been studied, e.g., in Freedman and Fraser, 1966 and Sherman, 1980.

Design: The decision context in this experiment is prosocial decision-making. Subjects’ decision is how much money to donate to some charitable fund. To study the potential of influence we conduct two treatments: in the main treatment subjects are asked how much they would donate if they were asked, prior to their donation decision. In the control treatment this question is not asked. Due to the preference for consistency, biases in prior statements (how much would I donate?) might carry on to later choices because people feel obliged to act consistently with the biased statement.

In both treatments, the experiment started with a survey. Subjects did not know that they were later asked to donate money. The survey consisted of 19 sub-questions and took about 10 minutes. The survey included a short version of the so-called “Big Five” inventory (John et al., 1991) and the cognitive reflection test (Frederick, 2005). The only difference between the main and the control treatment was that in the main treatment, we inserted one additional question, in between the “Big Five” and the cognitive reflection test. We asked subjects how they would hypothetically decide if they were asked to donate money. The question reads as follows: “Imagine, in an experiment you received an amount of 15 Euro in addition to your show-up fee. You had to choose, which part of the 15 Euro you want to donate to a charitable organisation. You could choose between different organisations and could donate any amount from 0 to 15. How much would you donate? Please give an amount between 0 and 15.” In the control treatment, we removed this question from the questionnaire. After all subjects had completed the survey, the donation experiment was announced.

In the donation experiment, subjects received an endowment of 15 Euro and had to decide how much of the endowment to donate to a charity organisation, and how much to keep for themselves. They could donate any amount from 0 to 15. In case they wanted to donate a positive amount, subjects could choose their preferred charitable organizations: they could either choose from a list of 8 organizations, or could name a charity organisation of their own.¹⁰

Note that the only difference between the two treatments was the hypothetical donation question. Moreover, the question was not asked right before the actual experiment but was instead embedded in a survey. The cognitive reflection test involves three questions related to intelligence and impulsiveness and was used to introduce some distraction between prior statements and donation experiment. We also did not remind subjects about their answers to the donation question when presenting the experiment.

¹⁰The list of organizations consists of very well-known and respected charities, Brot für die Welt, Kindernothilfe, German Red Cross, Welthungerhilfe, BUND, Greenpeace, terre des hommes, Aktion Mensch.

We therefore think that we implemented a fairly subtle treatment. It is possible that the effects reported below would become stronger if we had either asked the question right before the experiment or reminded subjects about their answers.

Procedural Details: A total of 64 subjects participated in four sessions, 32 in the main treatment and 32 in the control treatment. Subjects were mostly undergraduate students from various fields at the University of Bonn. No subject participated in more than one session. The experiment was conducted paper and pencil. Sessions lasted on average about 40 minutes and subjects earned on average 13.50 Euro, including a show-up fee of 5 Euro.

Hypotheses: We hypothesized that subjects would give a biased answer when asked hypothetically how much to donate to a charitable organization. In fact we were expecting them to overstate their willingness to donate given that donating to charity is socially desired, and given that it was essentially costless to signal positive characteristics. Note, however, that from the viewpoint of the model it makes no difference whether they over- or understate their true willingness to donate. Any bias should carry over to actual donation behavior. As a consequence, donations will differ between main and control treatment.

To see this suppose a bias in the hypothetical choice in period 1. For example, subjects choose $x_1^* > \hat{\mu}_1$.¹¹ To focus on the role of the preference for consistency as a device for social influence, we assume that subjects use the correct belief $\hat{\mu}_1$ for Bayesian updating.¹² In period 2, the actual donation decision, subjects face the following choice-problem:

$$\max_{x_2} U_2 = -(x_2 - \hat{\mu}_2)^2 - \alpha(x_2 - x_1^*)^2$$

They will choose x_2 , such that the following holds:

$$\hat{\mu}_2 > x_2^* > \hat{\mu}_1^* \text{ or } \hat{\mu}_2 < x_2^* < \hat{\mu}_1^*$$

From this it follows that $E(x_2^*) > E(x^*) = \mu$.

In the control treatment subjects only take one unbiased choice x^* . Subjects receive a signal m and choose $x^* = m$. It immediately follows that:

$$E(x_2^*) > E(x^*) = \mu$$

In the main treatment, subjects face a trade-off between acting consistently with their biased prior statement and acting according to their unbiased beliefs about their true preferences. Subjects solve this trade-off by choosing an actual donation level in between their biased statement and their unbiased beliefs. In the control treatment, subjects don't face such a trade-off and will simply choose according to

¹¹The model does not explain why people may give biased responses in the form of $x_1^* \neq \hat{\mu}_1$, it has to assume them. Possible reasons discussed in the literature are biases due to bounded rationality such as gambler's fallacy, base rate neglect, or hot hand fallacy (see for example Tversky and Kahneman, 1971, Grether, 1980, Charness and Levin, 2005 or Dohmen et al., 2009) or the signaling of pro-social or other motives that do not necessarily correspond to true underlying preferences (see for example Bénabou and Tirole, 2006).

¹²If subjects would use a biased belief this would contribute to an additional bias in subjects' period 2 choice. Thus, if subjects are unaware of the bias the treatment effect would become even larger.

their unbiased beliefs. Therefore, actual donations in the main treatment should be higher or lower than donations in the control treatment, depending on the respective bias. We summarize this prediction as follows:

Hypothesis 3 Subjects' actual donation decision in the main treatment will be biased in the direction of their hypothetical statement. As a consequence, they will donate more or less in the main than in the control treatment, depending on the direction of the bias.

Results: In a first step we check whether subjects gave a biased hypothetical statement about their willingness to donate. To do so we compare hypothetical donations in the main treatment with actual donations in the control treatment. It turns out there is a striking bias: While average hypothetical donations are 4.27 Euro actual donations in the control treatment are 7.63 Euro. Thus actual donations in the control treatment are almost 80 percent higher (!) than hypothetical donations. This difference is statistically significant at any conventional level (p -value < 0.001 , Wilcoxon rank-sum test). The sign of the bias is quite surprising and, at least to us, unexpected. A possible explanation is that in the actual donation decision we presented a list of eight charitable organizations, which may have made donations more concrete and credible, triggering a higher willingness to donate. In the hypothetical question we mentioned a list of possible charities but did not explicitly mention which ones.

Regardless of the sign of the bias, the interesting question is whether the bias carries over to the actual donation decision in the main treatment. Given that the bias is negative, the model predicts lower donations in the main compared to the control treatment. This is in fact the case. *Figure 7* shows a histogram of actual donations in both treatment conditions. 38 percent of donations in the main treatment are 3 Euro or lower, the corresponding fraction in the control treatment is only about 13 percent. Likewise the fraction of donations higher or equal 12 Euro is more than twice as high in the control compared to the main treatment. On average subjects donate 5.38 Euro in the main and 7.63 Euro in the control treatment, i.e., donations are about 42 percent lower in the main than the control condition. This difference is statistically significant using either a Wilcoxon rank-sum test (p -value < 0.03) or simple OLS regression, regressing donations on a constant and a treatment dummy (p -value < 0.02). Suggestive evidence for the importance of the consistency bias comes from the scatterplot shown in *Figure 8*. It shows that many subjects donated exactly the same amount as previously stated. The correlation coefficient between hypothetical and actual donations is 0.77 (p -value < 0.001). Those subjects who deviated typically increased their donation in line with the model's prediction. As a result, actual donations are 1.11 Euro higher than hypothetical donations. We summarize our results as follows:

Result 3 Subjects' donations are influenced by a biased hypothetical first statement. The bias is negative, resulting in lower donations in the main treatment compared to the control treatment.

4 Concluding Remarks

We presented a model that conceptualizes the preference for consistency and allows the analysis of how it affects economic behavior. In the model, people have a preference for consistency, because consistent behavior allows them to signal personal and intellectual strength. We conducted three simple experiments

to test main predictions and implications of our model. In the context of an estimation task, the first experiment shows that consistent behavior is viewed positively as a signal of ability. Agents anticipate this and estimate much more consistently compared to a control treatment without principals, i.e., in a situation where signaling ability has no strategic value. The second experiment shows that explicit commitment is a cornerstone in explaining preferences for consistency as commitment is a prerequisite for detecting inconsistencies. In the experiment such a commitment leads to a neglect of valuable information. The third experiment demonstrates the effectiveness of the preference of consistency as a means of social influence.

In our model we highlight the role of signaling positive traits as a key driver of consistent behavior. Other motives that have been discussed in the literature are reduction of cognitive dissonance and costs of thinking (see Cialdini, 1984). The latter would explain why we apply heuristics or routines in contexts where we face similar decisions over time. In terms of modeling it would imply some fixed cost of thinking or reflection, which would lead to similar behaviors in similar contexts, unless expected costs of consistent behavior become excessively large. Cognitive dissonance reduction as a driver of consistent behavior has been modeled, e.g., by Eyster, 2002 and Yariv, 2005. In Eyster, 2002 people try to rationalize past mistakes by taking current actions that justify these mistakes, thereby reducing cognitive dissonances. Yariv, 2005 proposes a model where people reduce cognitive dissonance by sticking to beliefs under which past choices were optimal. Most likely all three motives are important and often act in concert. The most important difference between a signaling story and fixed costs or cognitive dissonance concerns the role of third parties. We believe that the presence of observers creates a particular strong desire to appear consistent. Moreover, our first experiment demonstrates that cognitive dissonance or cost of thinking alone are not sufficient to understand consistent behavior. While we cannot rule out that these two motives play a role as well, they cannot explain observed treatment differences. This follows simply from the fact that both motives can be relevant in both treatments.

Notice that we consider situations where decision-makers face the same choice problems repeatedly and where the preferred outcome, μ , remains constant over time. In this situation, personal and intellectual strength is reflected by consistent behavior. A good example is decision making related to personality or identity. The latter is not constantly changing but remains relatively stable over long periods of time. Therefore, a person signals a strong personality by acting consistently. In changing environments where the preferred outcome changes, however, personal and intellectual strength could mean to adjust quickly to new environments. In fact sticking to previous points of view would be considered as rigid or even stupid. For example, if subjects in our first or second experiment would have been told that the number of letters or peas, respectively, is (very) different between the two estimations, choosing identical estimates would not signal high ability but the opposite. Put differently, if the circumstances that lead to a first decision have clearly changed before taking a second decision, it is not a sign of strength to stick to the first period's decision. In such a case decision-makers will, due to reasons of consistency, make sure that the public knows that circumstances have in fact changed. This explains why decision-makers (such as politicians) exert much effort to explain the reasons for why they have changed their mind and why acting differently should not be interpreted as inconsistent behavior.

Our results have interesting policy implications. We demonstrated that the preference for consistency can be a powerful channel of social influence. Using this channel, policy-makers can trigger certain desired behaviors, for example cooperativeness or law compliance. Take donations to help disaster victims. A

fundraising strategy could be to tempt people in a first step to commit to a statement like ‘believe supporting these victims is of great importance’. In a second step these people are confronted with an actual donation request. Clearly the direction of influence is not necessarily for good. The preference for consistency might for example also be (and is probably) used by salespersons or marketing experts to lure consumers into purchasing goods they normally would not buy. This points to an additional policy implication, namely consumer protection. Consumers should be protected from purchases under consistency pressure. A possible way to go could be to guarantee by law sufficient time to withdraw from closed sales contracts. This would allow consumers to undo bad consumption choices they later regret. In many countries such laws already exist (at least for certain types of contracts) and our results offer additional justification for that.

The way consistency preferences affect behavior can also be informative for the design of committee procedures and negotiation rules. We have seen that early commitments to a certain belief or opinion can make people disregard valuable information. Therefore, procedures that force committee members to state their opinions on a matter at an early stage risk that these members later on partially neglect new information, in particular if it contradicts their earlier statements. In negotiation contexts, statements of requests at an early point of negotiations can increase the likelihood of negotiation failure because negotiation outcomes below these requests cannot be reached without one party revealing some inconsistency.

We also want to point out the potential role of preferences for consistency in circumventing self-control problems. Consider smoking: A public announcement “I will stop smoking!” creates a huge pressure to live up to that announcement. Continuing to smoke is only possible at the cost of revealing some inconsistency, thereby signalling low personal strength. Thus, explicit public commitments can be very effective in solving problems of self-control. In fact, one reason for the effectiveness of self-help groups could be that they force their members to make public announcements.

As a last point we want to stress consequences of our results for experimental and survey methodology. Our study suggests that there are potentially large interdependencies between behavior in different experiments conducted in one session, between experimental behavior and questionnaire answers or between different survey answers. Depending on the context of the experiment or the survey study, the preference for consistency introduces a bias in behavior that potentially confounds results. This should be kept in mind when designing experimental and survey studies. As a conclusion from our results, variables of major interest should always be elicited first to avoid consistency bias.

5 References

- Allgeier, A.R., Donn Byrne, Barbara Brooks and Diane Revnes (1979). "The Waffle Phenomenon: Negative Evaluations of Those Who Shift Attitudinally", *Journal of Applied Social Psychology*, Vol 9 (2), pp. 170-182.
- Andreoni, James and Douglas Bernheim (2009). "Social image and the 50-50 norm: a theoretical and experimental analysis of audience effects", *Econometrica*, Vol 7 (5), pp. 1607-1636.
- Asch, Solomon (1956). "Studies of independence and conformity: A minority of one against a unanimous majority", *Psychological Monographs*, Vol 70 (Whole no. 416).
- Bénabou, Roland and Jean Tirole (2006). "Incentives and prosocial behavior", *American Economic Review*, Vol 96 (5), pp. 1652-1678.
- Campbell, John and John Cochrane (1999). "By force of habit: a consumption-based explanation of aggregate stock market behavior", *Journal of Political Economy*, Vol 107 (2), pp. 205-51.
- Charness, Gary and Dan Levin (2005). "When Optimal Choices Feel Wrong: A Laboratory Study of Bayesian Updating, Complexity and Affect", *American Economic Review*, Vol 95 (4), pp. 1300-1309.
- Christiano, Lawrence, Martin Eichenbaum and Charles Evans (2005). "Nominal rigidities and the dynamic effects of a shock to monetary policy", *Journal of Political Economy* 113 (1), pp. 1-45.
- Cialdini, Robert, Rodney Bassett, John Cacioppo and John Miller (1978). "Low-Ball Procedure for Producing Compliance: Commitment then Cost", *Journal of Personality and Social Psychology*, Vol. 36 (5), pp. 463-476.
- Cialdini, Robert, (1984). "Influence, the psychology of persuasion", New York: Harper Collins.
- Cialdini, Robert, Melanie Trost and Jason Newsom (1995). "Preference for consistency: the development of a valid measure and the discovery of surprising behavioral implications", *Journal of Personality and Social Psychology*, Vol. 36 (5), pp. 463-476.
- Dohmen, Thomas, Armin Falk, David Huffman, Felix Marklein and Uwe Sunde (2009). "Biased Probability Judgment: Evidence of Incidence and Relationship to Economic Outcomes from a Representative Sample", *Journal of Economic Behavior and Organization*, Vol. 72 (3), pp. 903-915.
- Ellingsen, Tore and Magnus Johannesson (2004). "Promises, threats and fairness", *The Economic Journal*, Vol. 114 (April), pp. 397-420.
- Ellingsen, Tore and Magnus Johannesson (2008). "Pride and prejudice: the human side of incentive theory", *American Economic Review*, Vol. 98 (3), pp. 990-1008.
- Eyster, Erik (2002). "Rationalizing the past: a taste for consistency", working paper.
- Festinger, Leon (1957). "A Theory of Cognitive Dissonance", Stanford, CA: Stanford University Press.
- Fischbacher, Urs (2007). "z-Tree: Zurich toolbox for ready-made economic experiments", *Experimental Economics*, Vol. 10 (2), pp. 171-178.

- Frederick, Shane (2005). "Cognitive Reflection and Decision Making", *Journal of Economic Perspectives*, Vol. 19 (4), pp. 25-42.
- Freedman, Jonathan and Scott Fraser (1966). "Compliance without pressure: the Foot in the Door technique.", *Journal of Personality and Social Psychology*, Vol. 4 (2), pp. 195-202.
- Grether, David (1980). "Bayes Rule as a Descriptive Model: The Representativeness Heuristic", *Quarterly Journal of Economics*, Vol. 95 (3), pp. 537-557.
- Heider, Fritz (1946). "Attitudes and cognitive organization", *Journal of Psychology*, Vol. 21, pp. 107-112.
- John, O. P., Donahue, E. M., Kentle, R. L. (1991) "The Big Five Inventory Versions 4a and 54". Technical report (July), Institute of Personality Assessment and Research, Berkeley, CA.
- Newcomb, Theodore (1953). "An approach to the study of communicative acts", *Psychological Review*, Vol. 60, pp. 393-404.
- Prendergast, Canice and Lars Stole (1996). "Impetuous youngsters and jaded old-timers: acquiring a reputation for learning", *Journal of Political Economy*, Vol. 104 (6), pp. 1105-1134.
- Sherman, Steven (1980). "On the self-erasing nature of errors of prediction", *Journal of Personality and Social Psychology*, Vol. 39 (2), pp. 211-221.
- Tversky, Amos and Daniel Kahneman (1971). "Belief in the law of small numbers", *Psychological Bulletin*, Vol. 2, pp. 105-110.
- Vanberg, Christoph (2008). "Why Do People Keep Their Promises? An Experimental Test of Two Explanations", *Econometrica*, Vol. 76 (6), pp. 1467-1480.
- Yariv, Leeat (2005). "Ill see it when I believe it - a simple model of cognitive consistency", working paper.

6 Figures

Figure 1:

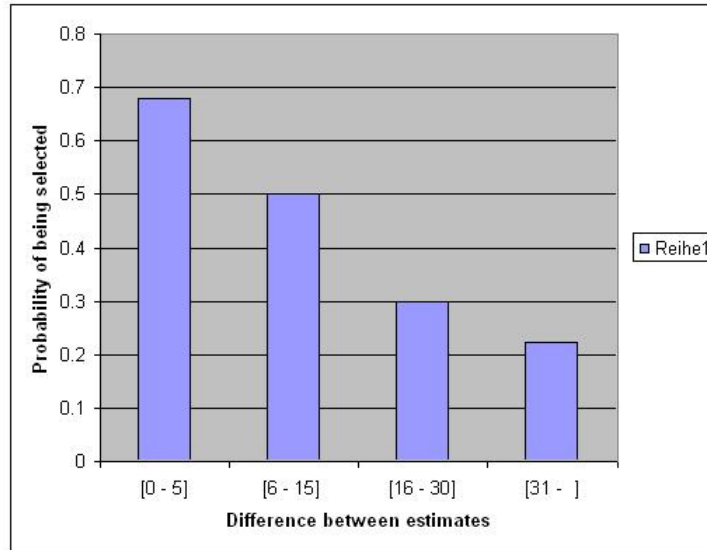


Figure 2:




Figure 3:

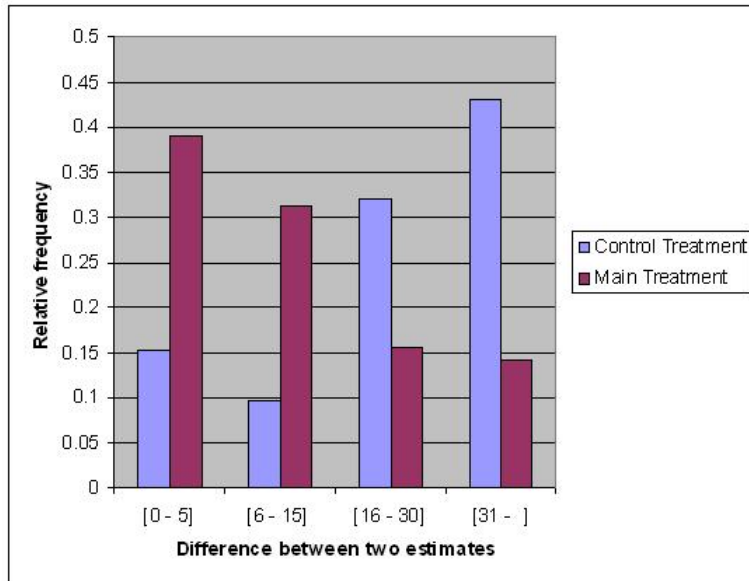


Figure 5:

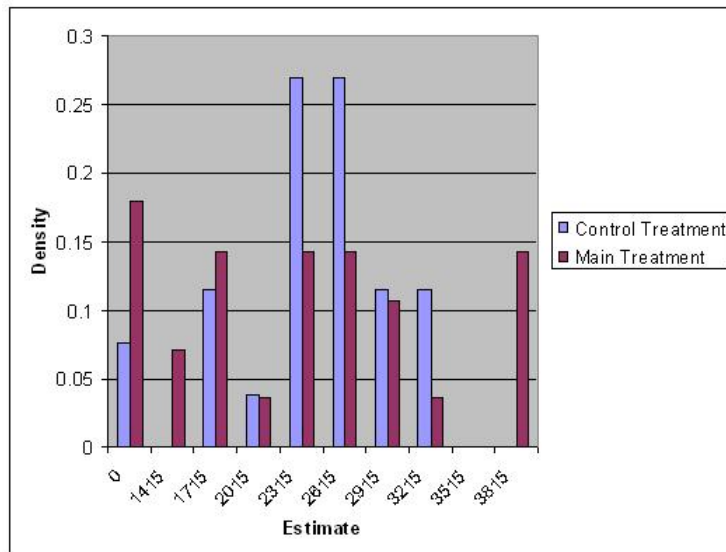


Figure 6:

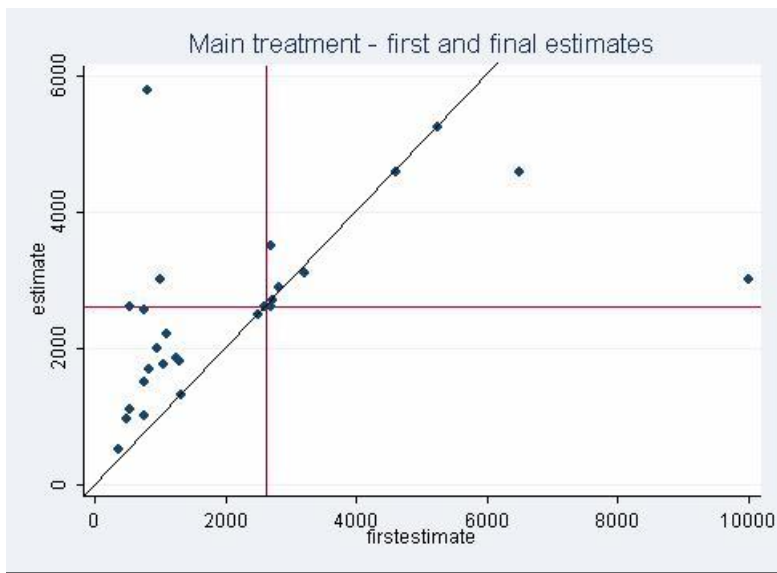


Figure 7:

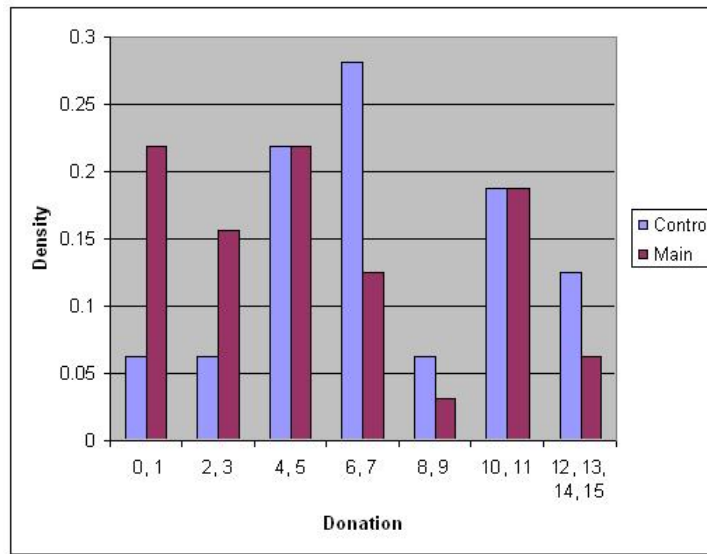
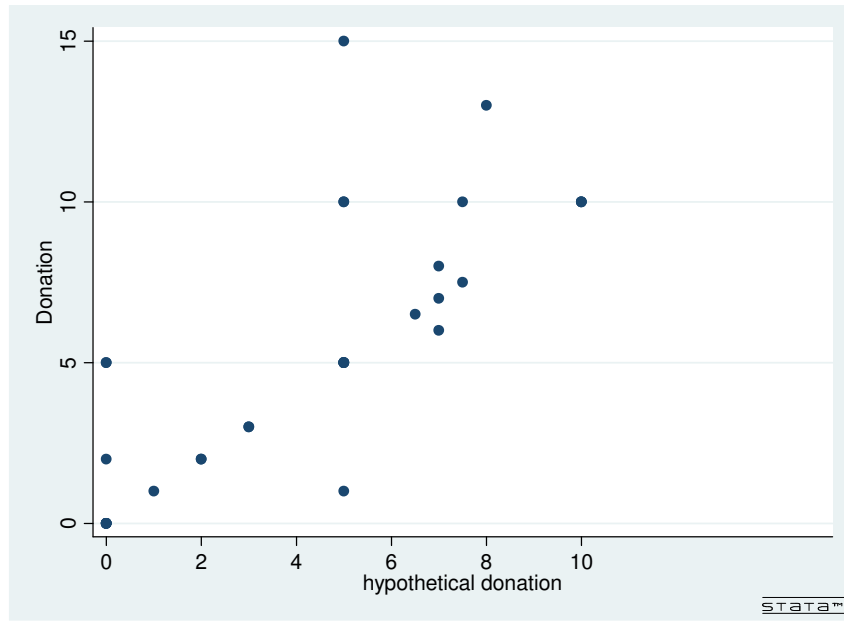


Figure 8:



7 Appendix A

Text1:

Periode	1 von 1	Verbleibende Zeit [sec]: 60
<p>The above theoretical studies all stress the negative consequences of corruption such as low investment and inefficiencies. Leff, 1964 suggests that corruption increases economic growth since it enables firms to circumvent unnecessary regulations and bad laws. In that sense corruption is a sort a device that frees firms from bureaucratic impairments and delay. An example might be a very rigid trade regime that practically forbids firms to exchange goods with foreign partners. Through bribing of the responsible officials, firms can circumvent these strict rules and engage in international trade, thus realizing efficiency gains. Critics of this theory point to the fact that it only applies to useless regulations. However, some researchers proposed positive consequences of corruption. However, corruption might very well be used to circumvent regulations which aim to limit pollution or to combat industrial concentration. It also seems to be difficult, if not impossible, to restrict corruption to areas where it helps getting around bad regulations. Lui, 1985 also stresses a potentially positive role of corruption. He establishes a model where official licences are allocated among firms such that the firms that pay the highest price also get the licence. The most efficient firm will benefit most from the licence and thus will pay the highest bribe and thus obtain the licence. Therefore an allocating system based on bribes would lead to an efficient outcome in the sense that the most efficient firm will obtain the licence. Now what does the empirical evidence say? Does corruption reduce investment and economic growth or is it instead growth-promoting? Regarding the macro evidence, Mauro, 1995 is the first study that attempts to examine the relationship between corruption and growth in a Cross-section of countries, which is why I will also present it here in a little more detail. Mauro uses a data set of the Business International indices on institutional efficiency. The indices reflect the perspectives and assessments of BI's correspondents and analysts of conditions in the regarded countries. In his study, Mauro also uses indices on corruption, red tape, quality of the legal system, political stability and various others. He also analyses links between the corruption index and investment.</p>		

Text2:

Periode	Verbleibende Zeit [sec]: 59
<p>Bureaucratic efficiency might also be endogenous in the regressions, since both indices might not only have an impact on growth and investment but might themselves be influenced by growth and investment. Due to that, Mauro runs regressions using ethnolinguistic fractionalization as an instrument for corruption. Linguistic fractionalization measures the probability that two randomly picked persons of a country will not actually belong to the same ethnolinguistic group. It is also likely to be highly correlated with corruption (and other institutional variables) but can be assumed to have no direct impact on economic growth or investment, which makes it a good instrument. The study shows that corruption and bureaucratic efficiency both have a statistically significant influence on investment, both in OLS and 2SLS estimations. The higher embezzlement (i.e. the lower bureaucratic efficiency) is, the lower is investment. However, contrary to what other researchers claim, the study does not find robust evidence that corruption decreases economic growth. Several of his corruption-growth specifications the coefficient of the corruption index is not significantly different from zero. He also merely finds a statistically significant influence of the bureaucratic efficiency index on economic growth, i.e., less bureaucratic efficiency leads to lower growth. Since the index on bureaucratic efficiency is a broader measure than embezzlement itself, the researcher's study also does not provide robust evidence that embezzlement leads to lower economic growth. Sveenske, 2005 updates Mauro's calculations. With far more observations, he also does not find significant influence of corruption on economic growth. These results are also interesting. The vast majority of the theoretical literature suggests that embezzlement should have a negative impact on economic growth, not in Mauro's (or Sveenske's) between-countries regressions? One of many possible answers is that measuring corruption appears to be difficult for researchers. Many measures are possible and all of them are subject to measurement error with the corresponding estimation problems. Furthermore, growth regressions of the type he conducts are subject to various problems as is discussed in Leevine and Renelet, 1992.</p>	

Picture of bowl with peas:

