Failures in Contingent Reasoning: The Role of Uncertainty

January 8, 2019

Alejandro Martínez-Marquina
(Stanford University)

Muriel Niederle
(Stanford University, NBER and SIEPR)

Emanuel Vespa
(UC Santa Barbara)

Abstract

We propose a new channel to account for the difficulties of individuals with contingent reasoning: the presence of uncertainty. When moving from an environment with one state of known value to one with multiple possible values, two changes occur. First, the number of values to consider increases. Second, the value of the state is uncertain. We show in an experiment that this lack of certainty, or the loss of the Power of Certainty, impedes payoff maximization and that it accounts for a substantial portion of the difficulties with contingent reasoning.

1 Introduction

Economic agents regularly face environments with uncertainty, where an underlying state can have multiple possible payoff-relevant values, and optimal behavior requires agents to engage in contingent reasoning. However, individuals routinely fail to make profit-maximizing choices in many such environments, indicating pervasive difficulties with contingent reasoning.[1] In this paper we propose a new source for these difficulties, the presence of uncertainty. Specifically, when moving from an environment with one state of known value to one where the value of the state is uncertain and can have several possible realizations, two changes occur. First, the number of values of the state the agent needs to consider increases. Second, in addition, an environment with more than one value introduces uncertainty. In this paper we propose that the lack of certainty impedes payoff maximization and show that it can account for a substantial portion of the difficulties with contingent reasoning.

Consider a generic decision problem with uncertainty, where each state ω_i occurs with probability p_i, let's call this the *probabilistic* problem. To compute the payoff of an action, the agent has to consider the payoff for each contingency. Compare the rate with which agents select the optimal action in the probabilistic problem to the rate when there is no need for contingent reasoning because the state of the world is realized and known. We refer to the latter as the *one-value* problem. Let the *Complexity of Contingent Reasoning* be defined as the difference between these two maximization rates. We construct the new *deterministic* problem in which the agent faces all n one-value problems at the same time with the constraint that the action she selects must be the same for all n one-value problems. The agent's payoff is the sum of the payoffs in each contingency appropriately weighted by p_i so that the payoff for each action is equal to the expected payoff in the *probabilistic* problem. There is no uncertainty in the deterministic problem, but the agent still has to consider all n contingencies ω_i before making a choice.

Using the deterministic problem we decompose the *Complexity of Contingent Reasoning* into two parts. First is *Computational Complexity,* which is the decrease in the maximization rate between the one-value and the deterministic problem, a difference due to having to consider all n states ω_i as opposed to a single state.[2] Second is the loss of the *Power of Certainty*: the decrease in the maximization rate between the *deterministic* and the *probabilistic* problem, due to the introduction of uncertainty. Specifically, instead of receiving the expected payoff with certainty, the agent receives a lottery where each state-contingent payoff occurs with the appropriate probability. Our main contribution is to propose and show the existence of the *Power of Certainty*, as well as its relative importance in accounting for the *Complexity of Contingent Reasoning.* We also provide some suggestive evidence as to a possible mechanism.

[1] For evidence see for example Shafir and Tversky (1992), Friedman (1998), Charness and Levin (2009), Rabin and Weizsäcker (2009), Esponda and Vespa (2014, 2018b,a), Cason and Plott (2014), Louis (2015), Eyster and Weizsäcker (2016), Enke (2017), Araujo et al. (2018), Moser (2018), and Ngangoué and Weizsäcker (2018).

[2] Previous work in psychology (e.g. Hamilton, 1878, Lennie, 2003) suggests that dealing with two values is more difficult than dealing with one value.

The probabilistic problem we consider is the two-value version of the Acquiring-a-Company problem (Samuelson and Bazerman, 1985). An agent decides whether or not to purchase a firm of value v. The agent knows that with equal chance the firm's value is either v_L or v_H, where $0 < v_L < v_H$. Without knowing the realization of v, the agent submits a price p for the firm. The agent acquires the firm when $p \geq v$ and her payoff is $1.5 \times v - p$, otherwise she does not buy the firm and has a payoff of zero. Charness and Levin (2009) show that even in such a simple version of the problem with only two contingencies to think of, a large proportion of participants fail to select the profit-maximizing price. To evaluate the extent of these difficulties, we also consider the *one-value* problem where the agent knows the value of the company before submitting a price.[3]

To provide evidence for the *Power of Certainty* we construct the following *deterministic* version of the two-value Acquiring-a-Company problem. There are two firms of known value: one of value v_L and another of value v_H. The agent submits a price p that is sent to each firm separately: if $p < v_L$, the agent buys none of the firms, if $v_L \leq p < v_H$, the agent only buys the firm of value v_L at a price p, and if $p \geq v_H$, the agent buys both firms, each at a price p. While there is no uncertainty in the *deterministic* problem, the agent still has to consider both values v_L and v_H to compute payoff maximizing prices, just as in the *probabilistic* problem.[4]

We conduct an experiment using the probabilistic and deterministic problems described above. In part 1, the first 20 rounds, subjects encounter the classic case where $2v_L < v_H$. This ensures that $p = v_L$ is the dominant action in the *deterministic* problems, as well as for all agents who are not very risk seeking in the *probabilistic* problems. Subjects then submit prices in part 2 which consists of 5 rounds where $1.5v_L > v_H$. In part 2, $p = v_H$ is the dominant action in not only the *deterministic* but also the *probabilistic* problems, since the lottery resulting from $p = v_H$ first-order stochastically dominates the lottery resulting from $p = v_L$. When we consider only subjects who both submit $p = v_L$ when $2v_L < v_H$ and $p = v_H$ when $1.5v_L > v_H$ as payoff maximizing, we find that there are twice as many payoff-maximizing subjects in the *deterministic* than in the *probabilistic* problems.

Before attributing this difference to the *Power of Certainty*, note that we may have underestimated the fraction of payoff maximizing subjects in the *probabilistic* problems. A subject who submits $p = v_H$ when $2v_L < v_H$ might still be payoff maximizing if they are very risk seeking. Note, however, that experimental subjects are in general risk averse, and often even excessively so given the low stakes, see Rabin (2000). Nonetheless, we have two strategies to address the role of risk-seeking preferences. First, the difference in the maximization rate between the probabilistic and the deterministic problems is also present, and almost similar in size, when we only consider part 2, where $1.5v_L > v_H$, and risk preferences do not affect the payoff-maximizing choice. Second, we directly measure whether subjects are risk seeking in environments that mirror the outcomes of the probabilistic problem. Specifically, in three questions subjects choose between two lotteries that are constructed the following way: For a specific v_L and v_H realization with $2v_L < v_H$, the

[3]Evidence from other settings (Esponda and Vespa, 2014, Fragiadakis et al., 2017, Ngangoué and Weizsäcker, 2018) suggests that not all subjects behave optimally even when the contingency is known.

[4]Note that payoffs in the probabilistic problem are multiplied by two to make payoffs across treatments comparable.

safe and risky lottery correspond to the lotteries subjects in the probabilistic problem receive when submitting $p = v_L$ and $p = v_H$, respectively. We find that the difference in the payoff maximization rate between the probabilistic and the deterministic problem is basically unchanged once we control for risk attitudes, confirming the existence of the *Power of Certainty*.

To evaluate the quantitative significance of the *Power of Certainty*, we benchmark the extent to which agents are profit maximizing in a simple version of the problem, with only one firm of known value. We show that a substantial part of the welfare loss subjects incur in the *Complexity of Contingent Reasoning* is due to the loss of the *Power of Certainty*.

Finally, we aim to shed some light as to the reason behind the *Power of Certainty*, why the problem with two firms of known value is so much simpler than the problem of one firm with two possible values. While these two scenarios are clearly different, in both do subjects have to take into account four outcomes in order to compute payoff-maximizing prices: for each value v_L and v_H the payoff from submitting $p = v_L$ or $p = v_H$. To gain some insight into the subjects' thought processes, and specifically whether they think of those four outcomes, we had subjects provide incentivized advice to a new individual for the case $v_L = 20$ and $v_H = 120$. In the deterministic problems subjects are much more likely to mention all four outcomes corresponding to $v, p \in \{v_L, v_H\}$. Furthermore, controlling for whether subjects' advice mentions those four outcomes reduces the difference in the subjects' propensity to maximize payoffs by more than half, and the remaining difference fails to be significant.

In a final section we provide evidence of the *Power of Certainty* in a second well-known and different environment, proving its robustness. In the Probability-Matching problem, described in detail in Section 5, we explore problems with probabilities different from fifty fifty, and where risk preferences play no role.

This paper is part of a larger literature showing that, and trying to understand why, agents fail to make profit-maximizing choices. Results in strategic settings, even in environments where fairness concerns play no role, generated a literature accounting for agents' mistakes.[5] There is also a literature in decision making that has documented behavioral biases that can explain failures to maximize payoffs.[6] In addition, active work on heuristics (Tversky and Kahneman, 1974) often studies difficulties related to prediction or updating exercises.[7] Here we focus on a simple decision problem where the state can have multiple possible values, the probability of each value being realized is known, and where evaluating the payoffs for an action only requires computing payoffs

[5]For empirical evidence see for example Kagel and Levin (1986), Nagel (1995), Costa-Gomes and Crawford (2006), Costa-Gomes and Weizsäcker (2008), Ivanov et al. (2010), Eyster et al. (2018), Esponda and Vespa (2018b) and Fragiadakis et al. (2017). For theoretical literature see e.g. Stahl and Wilson (1995), Camerer et al. (2004), Eyster and Rabin (2005), Jehiel (2005), Crawford and Iriberri (2007), Esponda (2008), Crawford et al. (2013). Difficulties with profit-maximization may also depend on how the problem is presented; see e.g. Glazer and Rubinstein (1996).

[6]Failure for profit maximization can be due to various behavioral biases such a hyperbolic discounting (see e.g. Laibson, 1997 and O'Donoghue and Rabin, 1999) or probability weighting (Kahneman and Tversky, 1979).

[7]For recent empirical work see for example Mobius et al. (2014), Levin et al. (2016), Vespa and Wilson (2016), Ambuehl and Li (2017), Enke and Zimmermann (2018); and for models see e.g. Rabin and Schrag (1999), Caplin and Leahy (2001), Bénabou and Tirole (2002) and Brunnermeier and Parker (2005).

specific to each possible value of the state. We propose that aggregating over multiple possible values of the state is especially difficult when there is uncertainty. While there are several recent models aiming to understand difficulties in such basic environments (for example Sims, 2003, Gennaioli and Shleifer, 2010, Bordalo et al., 2012, Kőszegi and Szeidl, 2012, Gabaix, 2014, Caplin and Dean, 2015 and Caplin et al., 2018), none directly captures that the difficulties can arise from the presence of uncertainty itself.

In the next section we describe the experimental design as well as some hypotheses and predictions. The main results are in section 3. In section 4 we provide some insight as to the underlying cause of the *Power of Certainty*. In section 5 we provide evidence of the *Power of Certainty* in a second environment. We then discuss the related literature and possible applications and connect the *Power of Certainty* hypothesis to concepts in psychology. Finally, we summarize and conclude.

2 Experimental Design

2.1 A Conceptual Framework

Consider a decision problem where the state of the world can have n possible values, $\Omega = \{\omega_1, \omega_2, ..., \omega_n\}$, and each contingency ω_i has a probability p_i to be realized, with $\sum p_i = 1$. The agent selects an action a of her action set A, without knowing which state has realized. Her payoff $u(a, \omega_i)$ depends on the action and the realized value of the state, and her expected payoff is given by $\sum p_i u(a, \omega_i)$. Define this as the *probabilistic* problem. Let the *one-value* problem be one where the agent has the same action set A and the same utility function u where, however, the value of the state is realized and known, some $\omega \in \Omega$. Let the *Maximization Rate* be a measure that allows us to compare the rate with which agents maximize their payoffs across problems. In this paper our measure is the fraction of individuals who make profit-maximizing choices. However, other measures could be used, for example, the improvement in payoffs between some benchmark payoff (perhaps resulting from random choice) and the payoff from optimal choice. Which measure of how much agents maximize payoffs is most appropriate may depend on the setting.

The Complexity of Contingent Reasoning Optimal behavior in the *probabilistic* problem with $n > 1$ requires the agent to engage in contingent reasoning. In order to compute the expected payoff of a certain action the agent needs to think of the payoff consequences of her action for each contingency, and then aggregate by appropriately weighting the probability of receiving each payoff. Meanwhile, optimal behavior in the *one-value* problem does not require contingent reasoning as the value of the state is known. The difference in the *Maximization Rate* between the *one-value* and the *probabilistic* problem is defined as being due to the *Complexity of Contingent Reasoning*.

In this paper we propose a new concept, which we show to be a significant component of the *Complexity of Contingent Reasoning*. To do so, we construct a new problem: the *deterministic* problem. The agent selects an action a from the same action set A which is applied to all n *one-value* problems defined by $\omega_1, \omega_2, ..., \omega_n$, respectively. The payoff of the agent is the sum of the payoffs in the n *one-value* problems $\omega_1, \omega_2, ..., \omega_n$ weighted by $p_1, p_2, ..., p_n$, that is, $\sum p_i u(a, \omega_i)$. The *deterministic* and the *probabilistic* problem differ in the payoff the agent receives for a given action. In the *deterministic* problem the agent receives the weighted sum of the payoffs of the *one-value* problems. In the *probabilistic* problem, the agent receives the payoff of one of the *one-value* problems, where each has a p_i chance to be the relevant *one-value* problem, respectively. That is, in both cases, payoffs of all n *one-value* problems defined by $\omega_1, \omega_2, ..., \omega_n$ are relevant. The main difference is that while there is certainty in the *deterministic* problem, there is uncertainty in the *probabilistic* problem. Note, hence, that the *deterministic* and *probabilistic* problem are truly different, and not just different frames of the same problem.

The Power of Certainty Hypothesis *(PoC)* This hypothesis claims that individuals are better at maximizing payoffs in the *deterministic* problem, where there is no uncertainty, rather than in the related *probabilistic* problem with uncertainty. The increase in the *Maximization Rate* when moving from the *probabilistic* to the *deterministic* problem is due to *PoC*.

We can now decompose the *Complexity of Contingent Reasoning* (the difference in the *Maximization Rate* between the *probabilistic* problem and a given *one-value* problem) into two components. The first is the difference between the *probabilistic* and the *deterministic* problem, which measures *PoC*. The second is the difference between the *deterministic* and the *one-value* problem, which we define as *Computational Complexity*. This latter difference captures difficulties arising from the fact that agents need to consider n contingencies in the *deterministic* problem relative to one in the *one-value* problem, holding certainty constant.

Table 1 summarizes the insights from the conceptual framework. Moving from the first to the second row increases the number of contingencies or values of the state the agent has to consider; changes in maximization rates are attributed to *Computational Complexity*. Moving from the first to the second column introduces uncertainty, and hence changes in maximization rates measure the impact of the loss of *PoC*.

			Power of Certainty Uncertainty	
			No	Yes
Computational	Number of	1	One-Value	-
Complexity	Values	n	Deterministic	Probabilistic

Table 1: Conceptual Framework: Summary

Next we present an environment to ascertain the existence of *PoC*, to evaluate its relative importance in the *Complexity of Contingent Reasoning,* and to explore a potential mechanism.

2.2 Basic Environment

The experimental design is built around a simplification of the Acquiring-a-Company game (Samuelson and Bazerman, 1985). A firm of value v is for sale. The action set of the agent is to submit a price $p \in \{0, 1, 2, ..., 150\}$, and the agent's profits are $(1.5v - p)$ if $p \geq v$ and 0 if $p < v$. This is the *one-value* problem. In the related *probabilistic* problem, the value of the firm is unknown to the agent, who knows that v is equally likely to be v_L or v_H, with $0 < v_L < v_H < \infty$.

When deciding about a price, note that any price p not in $\{v_L, v_H\}$ is dominated. For instance, making an offer higher than v_H is dominated by $p = v_H$, which also leads to buying the firm regardless of v but at a lower price. Likewise, making an offer $p = v_L$ dominates any price strictly between v_L and v_H. Finally, offering a price below v_L guarantees that a firm will not be purchased, but is dominated by $p = v_L$, which leads to either zero or strictly positive profits.

Depending on the agent's price $p \in \{v_L, v_H\}$, the expected profit is:

$$\pi(p) = \begin{cases} \frac{1}{2}(1.5v_L - v_L) & \text{if } p = v_L \\ \frac{1}{2}(1.5v_L - v_H) + \frac{1}{2}(1.5v_H - v_H) & \text{if } p = v_H \end{cases}. \tag{1}$$

Note that when $1.5v_L < v_H$, the agent receives a negative profit if $v = v_L$ and $p = v_H$. When $2v_L$ is strictly lower than v_H, both a risk-neutral and a risk-averse agent prefer $p = v_L$. A risk-seeking agent may still prefer to submit a price $p = v_H$ due to the positive returns when $v = v_H$. In contrast, when $1.5v_L \geq v_H$, all agents prefer a price $p = v_H$ to a price $p = v_L$, since the lottery resulting from $p = v_H$ first order stochastically dominates the lottery resulting from $p = v_L$.

In the experiments we present all three problems, the *one-value*, the *probabilistic* and the *deterministic* problem. To ease presentation, we multiply payoffs in the probabilistic and the deterministic problem by two.

2.3 Main Treatments

We first describe the two main treatments of our between-subjects design, the Probabilistic (PROB) and the Deterministic (DET) treatment. In both treatments subjects faced the same five parts, differing only as described below.

Probabilistic Treatment (PROB)

In PROB the subject submits a price p for a single firm that could have either of two values $\{v_L, v_H\}$, each with 50 percent chance. The subject does not know the value of the firm when submitting a price. If $p < v_L$, the subject buys no firm and the profit is zero. If $v_L \leq p < v_H$, she only buys the firm if the firm is of value v_L at a price p. Finally, if $p \geq v_H$, she always buys the firm. PROB corresponds to the *probabilistic* problem where payoffs are multiplied by two, so that expected profits are:

$$\pi^{Prob}(p) = \begin{cases} 0 & \text{if } p < v_L \\ (1.5v_L - p) & \text{if } v_L \leq p < v_H \\ (1.5v_L - p) + (1.5v_H - p) & \text{if } v_H \leq p \end{cases} . \tag{2}$$

Deterministic Treatment (DET)

In DET there are two firms for sale, one of value v_L and another of value v_H. The subject submits a unique price p that is sent to each firm separately. The price p determines whether the agent buys none, one, or two firms, each at price p. DET corresponds to the *deterministic* problem where payoffs are multiplied by two, so that profits are:

$$\pi^{Det}(p) = \begin{cases} 0 & \text{if } p < v_L \\ (1.5v_L - p) & \text{if } v_L \leq p < v_H \\ (1.5v_L - p) + (1.5v_H - p) & \text{if } v_H \leq p \end{cases} . \tag{3}$$

Parts 1-5 of the experiment

In part 1 of each treatment subjects faced 20 rounds with values $\{v_L, v_H\}$ such that $2v_L < v_H$. The first question always uses $v_L = 20$ and $v_H = 120$. For the other 19 questions we chose the two values $\{v_L, v_H\}$ randomly under the following constraints: Both v_L and v_H had to be even numbers and furthermore $10 \leq v_L \leq 30$ and $80 \leq v_H \leq 140$.[8] We decided to select those values randomly in order to not bias our results by unknowingly selecting values (or sequences of values) that would favor specific results (for a discussion on the value of random games see also Fragiadakis et al., 2017). While all subjects saw the same 19 sets of values $\{v_L, v_H\}$, the order was randomized at the subject level. Because subjects could submit any price $p \in \{0, 1, 2, ..., 150\}$, underbidding is always possible and overbidding (which can lead to losses) is bounded. Since in part 1 all values $\{v_L, v_H\}$ satisfied $2v_L < v_H$, any subject in PROB who is not very risk-seeking would optimally submit a price of $p = v_L$ in all rounds. In DET, submitting a price of $p = v_L$ is the dominant strategy.

Part 2 consisted of 5 rounds and values $\{v_L, v_H\}$ are selected such that $1.5v_L > v_H$.[9] This means $p = v_H$ is optimal regardless of risk attitudes, in both PROB and DET. We chose random combinations

[8] We constrained the selected values to be even numbers so that computations such as $1.5v$ result in integers. Subjects are not told the domains from which the values are drawn. They are simply shown the realizations for a specific round and asked to submit a price. The 19 pairs we implemented in rounds 2-20 were: {10, 86}, {10, 96}, {12, 112}, {12, 136}, {12, 138}, {14, 122}, {14, 140}, {16, 98}, {16, 106}, {18, 80}, {18, 86}, {18, 90}, {18, 122}, {20, 130}, {22, 106}, {24, 126}, {24, 128}, {24, 134}, {28, 126}.

[9] In the experiment there was a clear "break" between part 1 and part 2, see the Instructions Appendix for details. Given that the problems require computations, we were concerned that mixing up problems where (for not very risk-seeking agents) $p = v_L$ and $p = v_H$ was a dominant strategy would increase the fraction of subjects who would not submit payoff-maximizing prices.

of values with the constraint that both v_L and v_H had to be even numbers and furthermore $48 \leq v_L \leq 54$, $54 \leq v_H \leq 64$ and $v_L \neq v_H$.[10] While once more all subjects saw the same set of five values $\{v_L, v_H\}$, the order was randomized at the subject level.

In part 3, subjects provided incentivized advice for the case $v_L = 20$ and $v_H = 120$. They recommended what price to submit and why to a future participant which we refer to as the advisee. We told subjects that the advisee will be presented with advice from 5 different participants and that she will select which of the 5 pieces of advice was the most helpful. We told them that they would receive the profits the advisee made in this problem provided their advice is the one selected by the advisee. By forcing $v_L = 20$ and $v_H = 120$ to be in the first round of part 1, we ensure that all subjects have the same amount of experience (none) when they encountered this problem and 24 rounds passed before they provided advice.

In part 4, subjects are presented with 10 rounds that are similar to those in part 1. However, now, subjects faced the treatment of the problem they did not face so far. This means that subjects who faced 20 rounds of probabilistic problems in part 1 faced 10 rounds of deterministic problems in part 4 and vice versa. In round 1 of part 4 we fix $v_L = 20$ and $v_H = 120$. For the remaining 9 rounds we pre-selected 9 sets of values with the same criterion as described in part 1, though all were different from those in part 1.[11] Different subjects faced these 9 pairs of values in a different random order.

Part 5 consisted of three questions in which subjects selected one of two lotteries. In all lotteries there was a 50-50 chance of obtaining a low (π^L) or a high (π^H) payoff. For simplicity we describe a lottery as $\mathcal{L}\left(\pi^L, \pi^H\right)$. In the three questions agents chose between $\mathcal{L}(0, 10)$ and $\mathcal{L}(-250, 140)$, $\mathcal{L}(0, 20)$ and $\mathcal{L}(-180, 120)$, and $\mathcal{L}(0, 30)$ and $\mathcal{L}(-70, 80)$, respectively. These lotteries correspond to PROB cases $(v_L, v_H) = \{\{10, 140\}, \{20, 120\}, \{30, 80\}\}$, where the agent either submits $p = v_L$ or $p = v_H$, respectively.[12] All subjects faced these three questions in that same order, though they were not told how those lotteries were constructed.

Throughout part 1, and in fact throughout all parts of PROB and DET, and throughout all other treatments, subjects did not receive any feedback. That is, after each question they answered (e.g. submitted a price, in part 1), subjects simply received the next question. We did this because Charness and Levin (2009) document that feedback and experience do not remove overbidding ($p > v_L$) in the probabilistic case in a problem akin to those of part 1. In addition, while it may be interesting to address how learning changes the answers, we were mostly concerned that learning would be much more rapid in DET than in PROB. The reason is that feedback may be more informative in DET than in PROB. In PROB, feedback is the outcome of a lottery: Sometimes submitting a high price $p = v_H$, despite leading to losses in expectation, may result in large gains. In DET, feedback consists of the profit, a number. In fact, in all our part 1 problems, $3v_L < v_H$,

[10] The five pairs we implemented were: {48, 54}, {50, 54}, {52, 58}, {52, 64}, {54, 58}.

[11] The 9 pairs we implemented in rounds 2-10 of part 4 were: {10, 82}, {10, 110}, {14, 98}, {16, 80}, {24, 122}, {26, 96}, {28, 92}, {28, 96}, {30, 88}.

[12] The $\{10, 140\}$ and $\{30, 80\}$ combinations correspond to the most extreme values subjects could have encountered in part 1 and part 4.

which implies that submitting a price $p = v_H$ leads to a negative profit in DET. Indeed, experiments have shown that learning may be slower when feedback is the result of a lottery rather than the expected outcome, for a specific example with such a direct comparison see Bereby-Meyer and Roth (2006). While the impact of feedback may be an additional channel leading to better performances in *deterministic* than *probabilistic* problems, we study here the effects absent this potentially important channel.

2.4 Advisee Treatments

We have a Probabilistic (ADVPROB) and a Deterministic (ADVDET) Advisee treatment. The subject (advisee) goes through the same instructions and understanding tests as subjects in the main treatments. A subject then receives the first question of part 1 from the main treatment, where the values are $v_L = 20$ and $v_H = 120$. Before being given a chance to submit a price for that question, the subject receives five pieces of advice, one from each of five subjects from the corresponding main treatment. The subject sees one piece of advice at the time and answers whether the advice is useful (selecting from: very useful, somewhat useful or not useful at all). Subsequently, the subject sees all 5 pieces of advice at once and indicates which advice is most helpful. Finally, the subject submits a price for the $\{20, 120\}$ question and then faces the same 19 questions of part 1 of the main treatment, in random order. The subject then faces part 2 of the main treatment.

2.5 One-Value Treatments

In the One-Value treatments subjects are in part 1 and 2 confronted with a total of 25 rounds of a simplification of part 1 and part 2 of the main treatments. Specifically, in the first 25 rounds each subject submits a price for only a single firm of which they know the value. These 25 rounds correspond to part 1 and part 2 of the main treatments where, in each round, we randomize at the individual level which value (which firm) the subject can buy.

We implement two One-Value treatments. In the Probabilistic One-Value treatment, subjects are presented with the instructions that correspond to PROB. After they have read the instructions, but before they start part 1 and 2 we tell them that they will actually know the value of the firm before they submit a price. In the Deterministic One-Value treatment, subjects read the instructions that correspond to DET. Once they finish reading those instructions we tell them that in part 1 and 2 there will actually be only one of the two firms available for sale and that they will know which one of the two they can buy. That is, in each treatment we reduce the two-contingency problem to a one-contingency problem, but retain the initial description of the two-contingency problem. The remaining parts, part 3, part 4 and part 5 of the Probabilistic (Deterministic) One-Value treatment (which we refer to as ONEVALUE$_{\text{PROB}}$ and ONEVALUE$_{\text{DET}}$) mirror the essential parts of the main treatments, namely part 1, part 2 and part 5 of PROB (DET).

Table 2 summarizes our experimental design.

	Main Treatments PROB and DET	One-Value Treatments ONEVALUE$_{\text{PROB}}$ and ONEVALUE$_{\text{DET}}$	Advisee Treatments ADVPROB and ADVDET
Part 0	-	-	Reads Advice
Part 1	20 Rds $2v_L < v_H$	20 Rds	20 Rds $2v_L < v_H$
Part 2	5 Rds $1.5v_L > v_H$	5 Rds	5 Rds $1.5v_L > v_H$
Part 3	Advice	20 Rds $2v_L < v_H$	-
Part 4	10 Rds $2v_L < v_H$ switched	5 Rds $1.5v_L > v_H$	-
Part 5	Risk Lotteries	Risk Lotteries	-

Table 2: Summary of Experimental Design

2.6 Empirical Roadmap for the Main Results

Samuelson and Bazerman (1985) and Charness and Levin (2009) documented that many subjects fail to maximize payoffs in a setting comparable to PROB, see also Ball et al. (1991) and Selten et al. (2005). Given these findings, we expect a substantial *Complexity of Contingent Reasoning* effect. Therefore, we expect to be able to measure PoC in this environment.

By definition, PoC is the change in the fraction of payoff-maximizing subjects between the *probabilistic* and the *deterministic* treatment, or, in our case PROB and DET. The payoff maximizing strategy in DET is $p = v_L$ and $p = v_H$ in part 2. In PROB this strategy is payoff maximizing for subjects who are not very risk-seeking. At a first pass, we confine attention to those strategies. Note that by requiring subjects to change their submitted price between part 1 and part 2, we rule out subjects who use rules of thumb that do not require an adjustment to the firm values at hand.[13]

We use two strategies to check that we did not underestimate payoff-maximization in PROB, which in turn would lead to overestimation of PoC. First, we use the lottery choices in part 5 to provide a measure of risk-seeking preferences in an environment that closely mirrors the choices of submitting a price in part 1 (see Niederle, 2016 for a discussion of the virtues of such an approach). That is we have a direct test of whether subjects prefer the risky lottery that corresponds to $p = v_H$ over the safer lottery that corresponds to $p = v_L$.[14] We can then control for the participants lottery choices when assessing their propensity to submit $p = v_L$ in part 1. The second strategy consists of only considering actions in part 2, where a price of $p = v_H$ leads to a lottery that first-order stochastically

[13]For example, if the experiment consistent only of part 1 and not of part 2, a rule of thumb could be to submit $p = v_L$ i independent of the values of v_L and v_H. Such a rule of thumb might be more prevalent in DET, which could lead us to overestimating PoC. This rule of thumb could capture that it feels "cheaper" to submit $p = v_L$ than $p = v_H$. Alternatively, it could be that subjects prefer to think about one firm only. While in PROB the subject can buy at most one firm, she can buy two firms in DET if she submits a price $p \geq v_H$. A subject in DET who wants to avoid buying two firms could submit a price below v_H. Such a "constraint" would vastly simplify the problem, and could lead to fewer prices that can lead to losses and more selections of $p = v_L$.

[14]In addition, we can use the lottery choices to address losses due to presenting the complicated probabilistic problem rather than the simple lottery choices that result from submitting either $p = v_L$ or $p = v_H$ (see also Ambuehl et al., 2017).

dominates the lottery resulting from any other price p, and as such is a dominant strategy not only in DET but also in PROB. The disadvantage of this second approach is that individuals who always submit $p = v_H$ independent of the parametrization will be miscounted as payoff-maximizing.[15]

After establishing the existence of *PoC*, we want to assess its relative importance compared to *Computational Complexity* in accounting for the *Complexity of Contingent Reasoning.* This requires knowing the *Maximization Rate* in the *one-value* problem, and hence analyzing choices in the ONE-VALUE treatment. While only very few papers directly compare choices when subjects know the contingency to when they do not, the evidence from e.g. Esponda and Vespa (2014), Ngangoué and Weizsäcker (2018), Fragiadakis et al. (2017) suggests that the *Maximization Rate* in the one-value problem will not be 100% and hence making this assumption would lead to overestimation of *Computational Complexity.*[16,17]

2.7 Understanding the *Power of Certainty*

The evidence from parts 1, 2 and 5 of PROB, DET, and the One-Value treatments allows us to evaluate *PoC* and its importance relative to *Computational Complexities.* This introduces new questions, in particular, what is the mechanism behind *PoC*? This is where we will use subjects' answers in part 3, where, recall, subjects provide incentivized advise to a future participant for the $v_L = 20$, $v_H = 120$ case. By analyzing what subjects choose to write down in their advice and correlating this with the subjects' own behavior, we assess possible channels for *PoC.*

2.8 Procedures

Our subjects are Amazon Turk workers located in the US with a rating of 90 percent or higher. Subjects received a link to a qualtrics survey (see the Instructions Appendix which contains all the surveys). Participants knew that if they made more than two mistakes in the instructions for part 1 of any treatment, they were not allowed to continue the experiment. Upon finishing the experiment,

[15] For the main comparisons across treatments we assume that preferences satisfy first-order stochastic dominance, which is akin to subjects preferring more money to less. We use this assumption directly in part 2 to justify that $p = v_H$ is the dominant price in PROB. In part 1 we use the assumption indirectly. First, note that a part 1 probabilistic problem requires a choice between 151 lotteries given that we allow subjects in part 1 to submit any price in $\{0, ..., 150\}$. However, 149 of these lotteries are first-order stochastically dominated by either the lottery corresponding to $p = v_L$ or the one to $p = v_H$. First-order stochastic dominance delivers that subjects would only select a lottery corresponding to either $p = v_L$ or $p = v_H$. Hence, part 5 choices capture subjects' actual preferences for the corresponding probabilistic problem, that is, first-order stochastic dominance allows us to use part 5 choices as a control for preferences in part 1 problems.

[16] Esponda and Vespa (2014) consists of an extreme version of such a result where, while there are multiple contingencies, the agent has a dominant action in one contingency and is indifferent between her actions in all other contingencies. They show that many subjects who understand what to do in each contingency still fail to behave optimally when there is uncertainty over which contingency is realized. In a sequential treatment where subjects are shown the contingency, approximately 70 percent of subjects eventually use the dominant strategy. If subsequently subjects are asked for an action without knowing the the state, only about 20 percent of subjects submit the dominant strategy.

[17] In Section 3.3, where we present our findings on *Computational Complexity,* we discuss some possible sources of mistakes in the one-value problem.

we asked survey questions pertaining to the sex, age, ethnicity, education and state of residence of the participant. We have a total of 880 Amazon Turk workers with unique IP addresses, of which 44.8 percent are female, 74.8 percent are white, 49 percent are at or below the median age of 32 and 38.9 percent have low schooling ('Some college,' 'High School,' or lower education level).[18]

Probabilistic and Deterministic Treatments We aimed to recruit 200 participants per treatment. In total 425 Amazon Turk workers with a unique IP address started the experiment, 211 in PROB and 213 in DET, of which 23 and 31 made more than 2 mistakes in the instructions for part 1, respectively, a nonsignificant difference in drop rate. Eventually, we were left with 188 and 183 subjects in PROB and DET, respectively.

Payments in PROB and DET were determined as follows. A participant received \$4 for finishing the survey, as well as another \$4 at the beginning of part 1 to which they can add or subtract depending on their choices in the experiment.[19] In each of part 1, part 2 and part 4 we randomly select one round for payment. In part 3 subjects are paid the advisee's payoff in case their advice was selected. In part 5 we randomly select one of the three lottery choices and pay subjects based on their chosen lottery. Payoffs in all questions are expressed in points which are subsequently converted to dollars. In parts 1, 3 and 5 we paid 3 cents per point. In parts 2 and 4 we paid 1 cent per point.[20] On average a participant received \$8.7 (including the \$4 for finishing the survey) and took approximately 40 minutes to complete the survey.

We paid special attention to ensure that the instructions of PROB and DET are as similar as possible. For example, when describing the problem in PROB we use the phrase "transaction of *the* company," while in DET we use "transaction for *each* company." In Appendix A we provide a brief summary of how we explained each problem in the instructions and show screenshots for part 1 round 1 of PROB and DET (see Figure 6). In the Instructions Appendix we provide the full instructions.

Advisee Treatments We recruited 90 Mturkers who had not participated in any earlier treatment, half of which were assigned to each treatment. We drop 4 subjects in the probabilistic and 5 subjects in the deterministic version who make more than two mistakes in the instructions. Eventually we are left with 41 subjects in ADVPROB and 40 in ADVDET.[21]

[18]The regression tables presented in the results section and the online appendices control for these demographic variables. In particular, we construct a gender dummy, an ethnicity dummy that takes value 1 if the responder selected 'White' and 0 otherwise, an age dummy that takes value 1 if the responder is at or below the median age, and a low schooling dummy that takes value 1 if the responder selected 'Some college,' 'High School' or lower education level.

[19]The \$4 payment virtually ensures that at the end of part 1, no agent has negative profits, given that the highest possible loss comes from a price of 150 and a value of the firm of 10 (which just results in a loss of half a cent).

[20]Note that in part 2, the average payment in points from submitting the dominant price is higher than in part 1, and furthermore, there are fewer problems in part 2. We therefore used a lower exchange rate in part 2. However, just to ensure subjects do not react too strongly to only 1 cent payments, we also used that payment in part 4. We find no indication that the reduced payments have an effect.

[21]A participant received \$4 for finishing the survey. Participants are then endowed with another \$4 at the beginning of part 1 and can add/subtract the payment from one randomly selected round for part 1 and one randomly selected round for part 2. Payoffs are expressed in points and transformed to dollars at 3 cents per point in part 1 and 1 cent per point in part 2. On average participants received \$8.2.

One-Value Treatments We recruited 468 Mturkers (233 to the probabilistic and 235 to the deterministic version) who had not participated in any earlier treatment. However, 40 subjects made more than two mistakes in the instructions for part 1 which leaves us with 216 subjects in the Probabilistic and 212 in the Deterministic One-Value treatment.[22]

3 The *Power of Certainty (PoC)*

We assess *PoC* by comparing *Maximization Rates* between PROB and DET. We first consider only strategies with $p = v_L$ in the last five rounds of part 1 and $p = v_H$ in all rounds of part 2 as maximizing strategies. We find that more than twice as many subjects in DET are classified as using this strategy than subjects in PROB. We then show two ways, using part 5 lottery choices as well as part 2 choices only, that our results are not due to underestimating *Maximization Rates* in PROB. We use the *One-Value* treatments to assess the *Complexity of Contingent Reasoning*, and show that *PoC* accounts for about half of it, with the remainder being due to *Computational Complexity*. Finally, we show the robustness of our results by replicating them in a second sample in a slightly different set-up. We present all our results using classification of subjects into types. This reduces the chance that the use of simple rules of thumb lead to an overestimate of the *Maximization Rate*. We show in Appendix B that our findings are robust to instead considering the distribution of prices.[23]

3.1 Evidence for The Power of Certainty

In this subsection we describe types using part 1 & part 2 behavior and relegate tests of significance to the next subsection. We classify subjects based on their choices in part 1 and part 2. In part 2, the price $p = v_H$ is the dominant strategy in both PROB and DET. In part 1, $p = v_L$ is a dominant strategy for subjects in DET and for subjects in PROB who are not (very) risk seeking. To classify subjects, we consider the last five rounds in part 1 and all five rounds of part 2. We classify a subject as V_LV_H if she submitted a price $p = v_L$ in – the last five rounds of – part 1 (V_L in part 1) and $p = v_H$ in – all five rounds of – part 2 (V_H in part 2).[24] The fraction of V_LV_H subjects in DET (41.5

[22] Of the 40 subjects, 17 and 23 correspond to the probabilistic and deterministic version, respectively. The difference in the drop rate is not significant across treatments. Participants received $4 for finishing the survey. Participants are then endowed with another $4 at the beginning of part 1 and can add/subtract the payment from one randomly selected round for parts 1, 2, 3 and 4, and from one randomly selected lottery in part 5. In parts 1, 3 and 5 we paid 3 cents per point. In parts 2 and 4 we paid 1 cent per point. On average a participant received $8.6. While it is possible for subjects to have total earnings below the $4 they were endowed with to potentially lose during the experiment, of the 880 subjects that are in our final sample, only 24 did so, for those subjects we paid only the $4 they were guaranteed, and we did not implement their full losses above the $4 they were endowed with initially.

[23] The approach we follow in the text, which classifies subjects into types, demands consistent behavior from subjects, in contrast to the analysis on submitted prices in Appendix B. This means that there will be a difference in levels. For example, the aggregate frequency of prices $p = v_L$ is higher than the frequency of subjects who submit $p = v_L$ in all periods. While there is a difference in levels, our conclusions comparing outcomes across treatments are not affected. In the appendix we also present classifications of subjects into types allowing for small deviations, and again, we reach similar conclusions.

[24] The classification of subjects into types using exclusively the last five periods of part 1 is presented in Table 17 of Online Appendix B. Later in this section we present the classification of types using only part 2 (see Table 18).

percent) is more than twice the corresponding fraction in PROB (19.5 percent), confirming the PoC hypothesis.

The vast majority of subjects in both treatments can be classified as using either the payoff-maximizing V_LV_H strategy, or one of three alternatives. The second strategy includes behavior that is potentially profit-maximizing in PROB. It requires that subjects are classified as V_H in part 2 and, in addition, either strictly mix between v_L and v_H in the last five rounds of part 1 ($MixV_H$), or select $p = v_H$ in all the last five rounds of part 1 (V_HV_H). We group subjects who exhibit either behavior into the type $\{MixV_H, V_HV_H\}$. Note that while Mix could correspond to a profit-maximizing type in part 1, it would be quite knife edged.

In PROB, the $\{MixV_H, V_HV_H\}$ type contains 24.5 percent of all subjects, more than the percent classified as V_LV_H. One interpretation is that many of those subjects are risk seeking and profit maximizing. While we relegate a more careful analysis of this hypothesis to the next section, note that another hint that those may not all be payoff-maximizing subjects is that we also have many such subjects in DET (15.9 percent). In DET, these strategies are dominated and as such clear "mistakes," and we do not, ceteris paribus, expect subjects in PROB to be less error-prone than in DET.

The third type we consider consists of strategies where subjects exclusively submit prices that are either v_L or v_H (focal prices), which are not payoff maximizing in either treatment: These consist of part 1/part 2 strategies: V_LV_L, V_LMix, $MixV_L$, $MixMix$, V_HV_L, and V_HMix, which we group as a third strategy type labeled 'Focal.' In total, 18.7 percent of subjects in PROB are classified as using Focal strategies that cannot be payoff maximizing, compared to 8.3 percent in DET. The prevalence of Focal strategies in PROB also suggests that the higher incidence of V_HV_H and $MixV_H$ types in PROB may be due to a general similar incidence of types in PROB who use only v_L and v_H prices than in DET, though in PROB they are less likely to be used in a profit-maximizing way. In addition, among the Focal strategies, there are two that show an insensitivity towards the fact that values of the firms are changing between part 1 and part 2, namely V_LV_L and *MixMix*. Such subjects are almost exclusively present in PROB, where 8.0 and 3.2 percent of participants are classified in these categories compared to 1.1 and 1.6 percent in DET, respectively.

Finally, we also find slightly more subjects in PROB systematically submitting dominated prices that are neither v_L or v_H. This fourth strategy type, which we label 'Dominated' or 'Dom,' consists of subjects submitting $p \notin \{v_L, v_H\}$ in rounds 16-25. There are 21.8 percent of subjects classified as this type in PROB compared to 15.9 percent in DET.

Table 3 summarizes the classification of subjects into the four strategy types using the last five rounds of part 1 and all five rounds of part 2.[25] Figure 1 shows the classification if we also include the first fifteen rounds of part 1. For each of the four strategies, each figure reproduces the proportion of types using data from rounds 16-25. At round 1, we show the proportion of types if we were to

[25] Subjects classified as 'Residual' submit $p \notin \{v_L, v_H\}$ at least once and either $p = v_L$ or $p = v_H$ at least once in rounds 16-25.

Types	$V_L V_H$	$\{MixV_H, V_H V_H\}$	Focal	Dom	Residual	Participants
PROB	19.7	24.5	18.7	21.8	15.4	188
DET	41.5	15.9	8.2	15.9	18.5	183

Table 3: Part 1 and Part 2 Type Classification [as % of participants]

Notes: Types are defined based on the prices p_1 submitted in the last five rounds of part 1 and p_2 submitted in all five rounds of part 2. Type $V_L V_H$: $p_1 = v_L$ and $p_2 = v_H$. Type $\{MixV_H, V_H V_H\}$: $p_1 \in \{v_L, v_H\}$ and at least one $p_1 = v_H$ and $p_2 = v_H$. Type 'Focal': $p_1, p_2 \in \{v_L, v_H\}$ and at least one $p_2 = v_L$ (corresponds to $V_L V_L$, $V_L Mix$, $MixV_L$, $MixMix$, $V_H V_L$, or $V_H Mix$). Type 'Dom': $p_1, p_2 \notin \{v_L, v_H\}$. Residual: All remaining subjects.

demand subjects to follow the corresponding part 1 portion of the strategy for all 20, and not only the last 5 rounds. Finally, we also show how the proportions would change if the part 1 portion of the strategy would have to be followed from any round n onwards, for $1 \leq n \leq 15$.

In particular, Figure 1a follows the proportion of subjects who are eventually classified as $V_L V_H$. The fraction of payoff-maximizing subjects in DET is almost twice that of PROB for any given n, another strong indication of a significant PoC effect. Despite there being no feedback, there seems to be some learning in both treatments: the number of subjects classified as $V_L V_H$ increases by roughly 50 percent when we move from considering subjects who submit v_L from round 1 onwards compared to those who submit v_L in the last 5 rounds (and always submit v_H in part 2).

In Online Appendix B, we relax the definition of types in two ways and show that we would reach similar conclusions. First, we classify subjects as a given type even if their prices conform to the type in only four of the last five rounds in part 1 and four of the five rounds in part 2, see Table 16 and Figure 3. Second, we allow subjects to submit a price that is slightly higher than the price describing the type of the subject. Although we explicitly asked a question in the instructions, it is possible that some subjects think that they have to bid slightly above v in order to buy the firm. In the second robustness exercise a subject is classified as $V_L V_H$ if $p \in [v_L, v_L + 2]$ in the last 5 rounds of part 1 and if $p \in [v_H, v_H + 2]$ in all rounds of part 2. See Table 20 for details. Finally, to alleviate concerns that Amazon Turk workers are different from laboratory subjects, we compare round 1 choices where the values were selected to match the parametrization of Charness and Levin (2009), to the choices of their subjects. We find no large differences.[26]

3.2 Risk-Seeking Preferences

While the previous section already presented evidence suggesting that the prevalence of $p = v_H$ prices in rounds 1-20 in PROB is not due to subjects being risk seeking, we now provide two direct

[26] In the first round of Charness and Levin's "shifted two-value treatment" (which however is not the first round of play fir subjects), where the value is selected with equal chance from {20,119}, 48.7 percent of subjects submit a price equal to the low value, 27.9 percent submit a dominated price and 23.4 percent submit a price equal to the high value. In the first round of PROB, 42.0 percent of subjects submit a price equal to the low value, 40.4 percent submit a dominated price, and 18.6 percent submit a price equal to the high value. We have therefore no indication that the Amazon Turk subjects behave differently to the undergraduate students in Charness and Levin (2009).

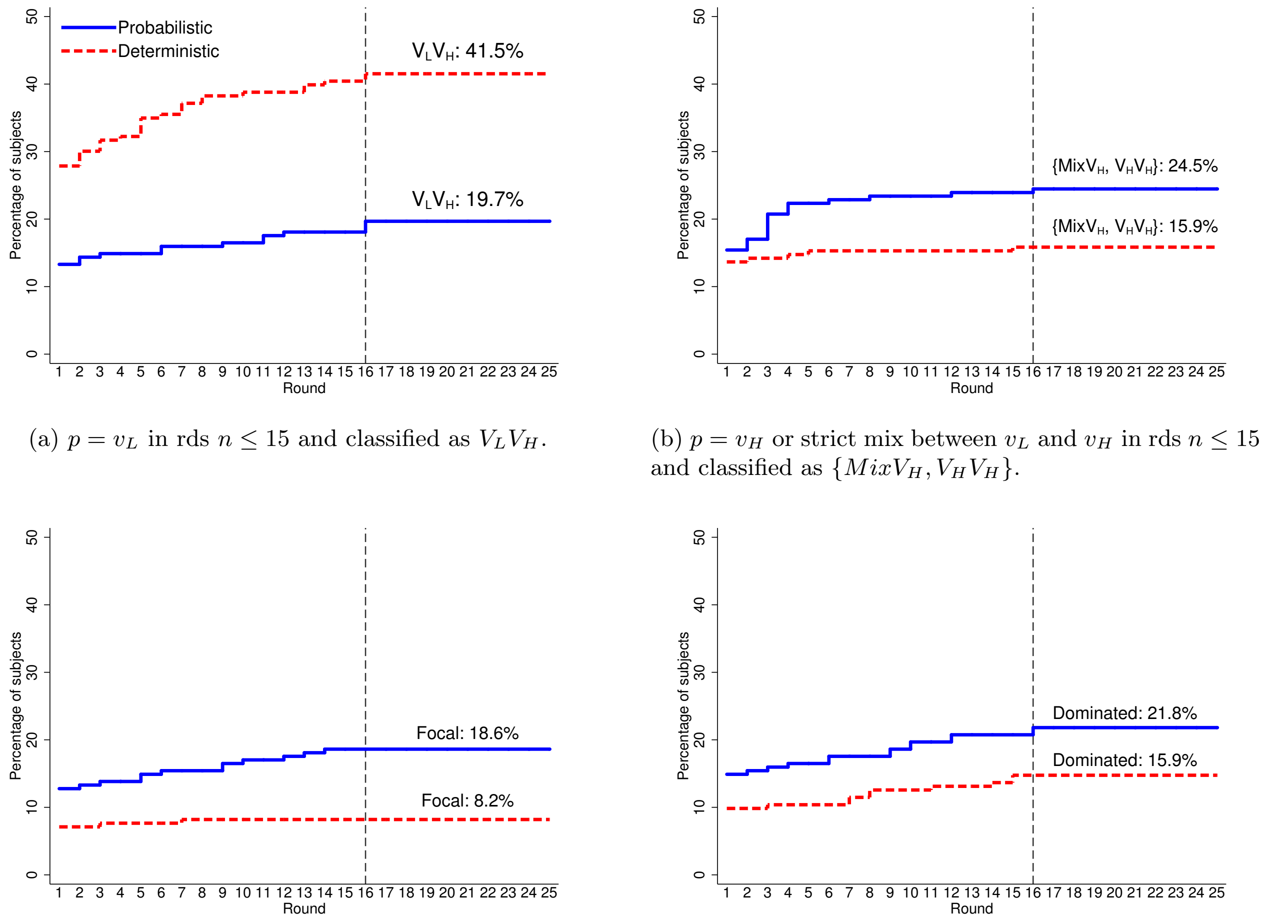

(a) $p = v_L$ in rds $n \leq 15$ and classified as $V_L V_H$.

(b) $p = v_H$ or strict mix between v_L and v_H in rds $n \leq 15$ and classified as $\{MixV_H, V_H V_H\}$.

(c) $p \in \{v_L, v_H\}$ in rds $n \leq 15$ and classified as Focal.

(d) $p \notin \{v_L, v_H\}$ in rds $n \leq 15$ and classified as Dominated.

Figure 1: Evolution of Types (part 1 & part 2).

tests of this hypothesis. This will show that the *PoC* effect is not due to underestimating payoff maximizing types in PROB.

First, recall that in part 5 subjects chose between three sets of lotteries, where there is a 50:50 chance of obtaining a low (π^L) or a high (π^H) payoff. For simplicity, we describe a lottery as $\mathcal{L}\left(\pi^L, \pi^H\right)$. Participants choose between $\mathcal{L}(0, 10)$ and $\mathcal{L}(-250, 140)$, $\mathcal{L}(0, 20)$ and $\mathcal{L}(-180, 120)$, and $\mathcal{L}(0, 30)$ and $\mathcal{L}(-70, 80)$, which correspond to the probabilistic problems $\{v_L, v_H\}$ of $\{10, 140\}$, $\{20, 120\}$, and $\{30, 80\}$, if subjects were to submit $p = v_L$ and $p = v_H$, respectively. The cases of $\{10, 140\}$ and $\{30, 80\}$ present the extreme values subjects subjects could have encountered in part 1, and represent the case where the difference in expected returns between $p = v_L$ and $p = v_H$ is maximized or minimized, respectively. Overall, 70.0 percent of subjects in DET and 68.6 percent in PROB never took any of the lotteries that correspond to submitting $p = v_H$.[27]

[27] The risky alternative was selected only once by 15.4 percent (15.9 percent), twice by 4.8 percent (4.9 percent), and in all three lotteries by 11.2 percent (9.3 percent) of subjects in PROB (DET), respectively.

	(1) $V_L V_H$	(2) $\{MixV_H, V_H V_H\}$	(3) Focal	(4) Dom	(5) Dom or Res
Det	0.244***	−0.102*	−0.053	−0.074	−0.089
	(0.060)	(0.056)	(0.048)	(0.053)	(0.063)
Num Risky	−0.062**	−0.008	−0.003	0.073***	0.073**
	(0.031)	(0.029)	(0.025)	(0.027)	(0.033)
Num Risky × Det	0.043	0.040	−0.011	−0.052	−0.071
	(0.046)	(0.043)	(0.036)	(0.040)	(0.048)
Constant	0.286***	0.223***	0.203***	0.147**	0.287***
	(0.069)	(0.065)	(0.055)	(0.061)	(0.073)
Observations	371	371	371	371	371

Table 4: Main Treatments: Estimation output using last 5 rounds of part 1 and part 2 for the classification of types

Notes: Results from a linear regression. The dependent variable takes value 1 if the subject is classified as (1) $V_L V_H$, (2) $\{MixV_H, V_H V_H\}$ (3) Focal, (4) Dom - Dominated, (5) Dom or Res - Dominated or Residual type. Det is a treatment dummy that equals 1 if the subject participated in DET. Num Risky is the number of risky lotteries the subject chose in part 5 (from 0 to 3). The regression also includes demographic controls (Gender, Ethnicity, Age and Schooling) and a control for the number of errors in the instructions. The full output is presented in Table 15 of Online Appendix B.

The difference in *Maximization Rates* between PROB and DET is robust to controls for risk preferences: A linear regression on the classification of subjects (Table 4) shows that there are significantly more subjects classified as $V_L V_H$ in DET compared to PROB, when controlling for subjects' lottery choices in part 5, confirming *PoC*. Specifically, we use the number of risky lotteries a subject chose, from 0 to 3.[28]

Table 4 column (2) shows that the incidence of $V_H V_H$ and $MixV_H$, while higher in PROB than in DET, is not driven by subjects who are very risk seeking. In fact, subjects who select more risky lotteries in part 5 seem to be less likely to maximize payoffs. While taking more lotteries significantly reduces the chance for subjects to be classified as $V_L V_H$, it increases the chance to be classified as the *Dominated* type. Furthermore, the effect of risk is never significantly different in DET compared to PROB, though risk preferences have no bearing in DET.

For a second test that *PoC* is not due to underestimating payoff maximization in PROB, we turn to part 2 of PROB and DET. Any subject who prefers a lottery that first-order stochastically dominates another maximizes payoffs by submitting $p = v_H$. However, when we focus on V_H types (subjects who submit $p = v_H$ in all five rounds of part 2), we may also capture subjects who are selecting $p = v_H$ for a different reason, such as, for example, subjects who use the dominated strategy $V_H V_H$ or $MixV_H$ in DET. This could affect the levels of estimated *Maximization Rates*. Nonetheless, the *PoC* effect is roughly similar to the one we found when focusing on part 1 and part 2: 65 percent of subjects in DET are classified as V_H compared to 47.9 percent in PROB.[29]

[28] Our results are robust to instead using a dummy indicating whether a subject chose at least one risky lottery, see Table 22 in Online Appendix C.

[29] To test for the treatment effect we run a linear regression with the same right-hand side specification as regressions reported in Table 4, but where the left-hand side variable takes value 1 if the subject is classified as type V_H. The treatment effect estimate is 0.21 and the p-value 0.001. For a complete description of types in part 2 see Table 18 in Online Appendix B.

To summarize, as expected, we do not find evidence that *PoC*, the difference in *Maximization Rates* between PROB and DET as measured by V_LV_H types is driven by subjects being risk seeking and hence an underestimation of the *Maximization Rate* in PROB.

3.3 The Role of the *Power of Certainty* (*PoC*)

After confirming the existence of *PoC*, we now assess its importance in accounting for the *Complexity of Contingent Reasoning.* To assess either of these two costs, we need to assess the *Maximization Rate* in the *one-value* problem. Let $MR(X)$ be the maximization rate in treatment X, then the importance of *PoC* in accounting for the *Complexity of Contingent Reasoning* is given by

$$\frac{MR(\text{DET}) - MR(\text{PROB})}{MR(one - value) - MR(\text{PROB})}.$$

While it may be tempting to simply assume that the *Maximization Rate* in the One-Value treatment is 100%, other experiments suggest that this may not be the case. In order not to overestimate the *Complexity of Contingent Reasoning,* we first assess the role of the "background complexity" which captures how much the *Maximization Rate* is below 100% in the One-Value treatment.

We therefore analyze the 25 rounds in the One-Value treatment where the problem consists of a single firm of known value. Note that we find no statistical difference depending on whether subjects were introduced to the One-Value environment using probabilistic or deterministic instructions.[30] Consequently, we pool all data from both treatments and we label the joint treatment as ONEVALUE. When we classify subjects based on the prices they submit in the last 10 rounds, 62.6 percent submit only $p = v$, a strategy type that we refer to as V. A total of 13.6 percent of subjects submit dominated prices: 8.6 percent submit only $p > v$, 1.2 percent submit only $p < v$, and 3.8 percent submit either $p < v$ or $p > v$ in the last 10 rounds. Subjects who submit each, a dominated price and $p = v$ in at least one of the last 10 rounds are classified as the Residual type, which covers 23.8 percent, see Table 23 in Online Appendix D.[31]

If, at a first pass, we put aside risk-seeking subjects, and hence estimate $MR(\text{PROB}) = |V_LV_H| = 19.7$ (see Table 3), and we have $MR(\text{ONE-VALUE}) = 62.6$, then the *Complexity of Contingent Reasoning* involves a reduction in the *Maximization Rate* of 42.9 percentage points. One interpretation is that amongst subjects who consistently submit payoff-maximizing prices in ONEVALUE, only 31.5 percent are able to do so when the firm has two possible values each with 50 percent chance.[32] Given that

[30] We run a panel regression in which the left-hand side is a dummy variable (1 if $p = v$) and the right-hand side includes a treatment dummy (1 if the subject was introduced to the one-firm problem after reading deterministic instructions), demographic controls and a constant. The estimated treatment dummy is quite small and not significant in the whole sample (-0.012, p-value = 0.717) or if we constrain the sample to prices in the last 10 rounds (0.005, p-value = 0.883).

[31] If we relax the classification constraints, more subjects are classified as V. For example if we allow for two rounds (out of the last 10) in which the subject did not submit $p = v$, the proportion of subjects classified as V increases from 62.6 to 67.1 percent. See Online Appendix B for other robustness checks. Furthermore, roughly three quarter of prices are equal to v. See Table 10 of Appendix B for the distribution of submitted prices.

[32] For a qualitative comparison, we summarize the corresponding findings of Esponda and Vespa (2014). In their

MR(DET)= 41.5 percent, the lack of *PoC* accounts for more than half of the total effect of the *Complexity of Contingent Reasoning.* If we restrict attention to subjects who never took a risky lottery, both in the ONEVALUE as well as the main treatments, we find that *PoC* accounts for 44.2 percent of the *Complexity of Contingent Reasoning*, see Table 23 in Online Appendix D.

Instead of using the fraction of subjects who submit payoff-maximizing strategies, we can use two measures in the payoff space. Measure 1 is based on the expected profits subjects receive in each treatment. Measure 2 is based on the fraction of gains a subject made when comparing their payoffs to, on the one hand, random behavior, and, on the other hand, optimal behavior. To reduce noise, we use, for both measures, all 25 rounds. We find that for the median subject both measures indicate that about 70 percent of the payoff loss incurred by the *Complexity of Contingent Reasoning* is due to the loss of *PoC,* see Online Appendix D for details.

To summarize, we find that in both the strategy space and the payoff space, there is robust evidence that a substantial part of the *Complexity of Contingent Reasoning* can be attributed to the lack of *PoC*. More concretely, there is an important improvement, both in the action and payoff space, when we compare the two contingency environment of DET relative to the two contingency environment of PROB, where the major difference is whether the two contingencies are certainties.

3.4 Robustness and Replication of the Relative Importance of *PoC*

In this section we provide a robustness test and replication of our main result on *PoC*.[33] In the ONEVALUE treatments subjects are first confronted with 25 rounds of a simplification of part 1 and part 2 of the main treatments. Then, in part 3, part 4 and part 5 of the ONEVALUE treatments subjects face part 1, part 2 and part 5 of either PROB or DET, we refer to the treatments as $\text{ONEVALUE}_{\text{PROB}}$ and $\text{ONEVALUE}_{\text{DET}}$, respectively. Therefore, apart from a simple replication, subjects in the ONEVALUE treatments are trained for 25 periods to think about the value of the company when submitting a price before they encounter part 1 and 2 of the main treatments.

Using the same classification of types in part 3 and part 4 of the ONEVALUE treatments we used in part 1 and part 2 of the main treatments, focusing first on one firm significantly increases the chance with which subjects are classified as the dominant type $V_L V_H$ by eight percentage points and decreases the chance to be classified as the Dominated type or being unclassified (Residual type) in PROB. While changes are similar in DET, they are not significant, see Tables 25 and 26 in Online Appendix D. It seems that focusing on one firm of known value helps subjects to subsequently only

treatment where the contingency is known, about 76 percent of subjects consistently make the optimal decision. However, when subjects subsequently make a decision in which the state is unknown approximately 22 percent of subjects make the optimal choice. Even though the environment in Esponda and Vespa (2014) is quite different from ours, and they use NYU students rather than Amazon Turk workers, the finding on "Background Complexity" and the *Complexity of Contingent Reasoning* is qualitatively similar. Background Complexity could be due to many sources and our experiment was not designed to identify them. The findings suggest, however, that it may be relevant for future research to study when and why subjects have difficulties with problems in which the contingency is known.

[33]For the value of replications, see e.g. Coffman and Niederle (2015), Camerer et al. (2016), Coffman et al. (2017).

submit prices $p \in \{v_L, v_H\}$.[34]

We confirm that there is a large difference in the $V_L V_H$ classification between the probabilistic and deterministic problems, namely 27.8 versus 47.2 percent when conditioning on all 212 subjects in ONEVALUE$_{\text{PROB}}$ and 216 in ONEVALUE$_{\text{DET}}$. When we condition on the 139 subjects in ONEVALUE$_{\text{PROB}}$ and the 129 in ONEVALUE$_{\text{DET}}$ who are classified as V, then 41.7 and 65.1 are classified as $V_L V_H$, respectively, see Table 24 in Online Appendix D for details. The linear regression using all 428 subjects in Table 5 confirms that subjects in ONEVALUE$_{\text{DET}}$ are significantly more likely to be classified as $V_L V_H$ than subjects in ONEVALUE$_{\text{PROB}}$, and less likely to be classified as using a Focal (and hence a dominated) strategy.

The coefficients on V in regressions (4) and (5) of Table 5, where V is a dummy controlling for whether the subject submitted $p = v$ in the last 10 rounds of parts 1 and 2 in which subjects submit prices for one firm of known value, show that being able to solve the one-firm problem significantly reduces instances of prices $p \notin \{v_L, v_H\}$. However, subjects classified as V are not only more likely to be classified as $V_L V_H$, but also as using Focal strategies, which are strategies that are dominated (though only using prices $p \in \{v_L, v_H\}$).

	(1) $V_L V_H$	(2) $\{MixV_H, V_H V_H\}$	(3) Focal	(4) Dom	(5) Dom or Res
Det	0.255***	−0.045	−0.178***	−0.048	−0.032
	(0.055)	(0.054)	(0.047)	(0.036)	(0.041)
Num Risky	−0.091***	0.102***	−0.015	−0.007	0.004
	(0.029)	(0.029)	(0.025)	(0.019)	(0.022)
Num Risky × Det	0.018	−0.082**	0.023	0.012	0.041
	(0.040)	(0.039)	(0.034)	(0.026)	(0.030)
V	0.367***	0.103**	0.104***	−0.258***	−0.574***
	(0.044)	(0.043)	(0.037)	(0.029)	(0.033)
Constant	0.179***	0.191***	0.144**	0.223***	0.487***
	(0.067)	(0.066)	(0.057)	(0.044)	(0.051)
Observations	428	428	428	428	428

Table 5: ONEVALUE Treatments: Estimation output using last 5 rounds of part 3 and part 4 for the classification of types.

Notes: Results from a linear regression. The dependent variable takes value 1 if the subject is classified as (1) $V_L V_H$, (2) $\{MixV_H, V_H V_H\}$ (3) Focal, (4) Dom, (5) Dom or Res. Det=1 is a dummy variable that takes value 1 if the observation corresponds to ONEVALUE$_{\text{DET}}$ and 0 if it corresponds to ONEVALUE$_{\text{PROB}}$. Num Risky is the number of risky lotteries the subject chose in part 5 (from 0 to 3). The variable V is a dummy variable that equals 1 if the subject selected $p = v$ in the last 10 periods of problems with one firm (last 5 periods of part 1 and 5 periods of part 2). The regression also includes demographic controls (Gender, Ethnicity, Age and Schooling) and a control for the number of errors in the instructions. The full output is presented in Table 27 of Online Appendix D.

Finally, we use subjects in the ONEVALUE treatments to reassess the role of *PoC* on the *Complexity of Contingent Reasoning*. Using only $V_L V_H$ as the payoff-maximizing strategy, of the 62.6 percent of subjects are classified as V in parts 1 and 2 in the ONEVALUE treatments, 41.7 were classified as

[34] This is consistent with Charness and Levin (2009), who find that giving feedback after every period in the (probabilistic) acquiring-a-company problem reduces the submission of strictly dominated prices.

$V_L V_H$ in ONEVALUE$_{\text{PROB}}$ and 65.1 in ONEVALUE$_{\text{DET}}$ (see Table 24 in Online Appendix D for details). Hence, *PoC* accounts for 40.1 percent of the *Complexity of Contingent Reasoning.* If in addition we restrict attention to subjects who never took a risky lottery in part 5, the *PoC* effect accounts for 42.7 percent.[35]

Using the profit measures described in Section 3.3, we replicate that *PoC* captures a substantial portion of possible gains in the profit space.[36]

4 Understanding the Power of Certainty

In this section we shed light on the underlying cause of *PoC*. We first analyze the advice subjects provided to another participant. Using the advice as a window into the subjects' thought process, we ask whether subjects mention the four possible outcomes associated with submitting a price $p = v_L$ and $p = v_H$. We then explore to what extent differences in advice are correlated with differences in strategies between PROB and DET.

4.1 Advice

Subjects provide advice to another participant for the $v_L = 20$, $v_H = 120$ problem. In both PROB and DET the vast majority of subjects provide a numerical advice (90.4 and 92.4 percent) and/or some explanation (77.7 and 79.2 percent, respectively). Furthermore, the length of the written advice is, while similar between PROB and DET, if anything a little longer in PROB.[37]

The pattern of numerical recommendations matches our earlier findings. While 63.9 percent of subjects in DET recommend submitting $p = v_L$, 43.1 percent provide the same recommendation in PROB (Table 29 of Online Appendix E). Meanwhile, 32.5 percent of subjects in PROB recommend to either submit $p = v_H$ or mix between v_L and v_H, compared to 17.5 percent in DET.

[35] Approximately two-third of subjects who participate in ONEVALUE treatments never take a risk in part 5, and 67.8 percent of these subjects are classified as V in parts 1 and 2. Out of subjects who are classified as V and never take a risk in part 5, 51.1 percent are classified as $V_L V_H$ in ONEVALUE$_{\text{PROB}}$, compared to 72 percent in ONEVALUE$_{\text{DET}}$, so that *PoC* accounts for 42.7 percent of the *Complexity of Contingent Reasoning.*

[36] According to *Measure 1,* of the total payoff gain for the median subject comparing the first 25 rounds of ONEVALUE to ONEVALUE$_{\text{PROB}}$, 56.4 (75.4) percent is achieved by the median subject in the DET (if restricted to subjects who took no risky lottery in part 5). *Measure 2* is based on a measure of relative payoffs, and in this case 54.6 (91.1) percent of the total gain is achieved by the median subject in the DET (if restricted to subjects who took no risky lottery in part 5).

[37] The median (mean) number of words used in the advice equals 42 (54.7) in PROB and 34 (45.1) in DET. Using a quantile regression on the median (linear regression) with a treatment dummy on the right-hand side and the number of words on the left-hand side we find that the difference in the median (mean) is significant at the 10 (5) percent level (p-values of 0.06 and 0.03, respectively).

		PROB	DET
All Outcomes	Both Explicit	1.6	13.7
	One Explicit	0.0	8.7
	Qualitative	14.4	30.1
	Mention Outcomes Only	10.6	0.0
	Total	26.6	52.5

Table 6: Mentioning All Outcomes in the Part 3 Advice [as % of participants]

Notes: 188 participants in PROB and 183 in DET. There are four outcomes (v, p) where $v, p \in \{v_L, v_H\}$. The four rows divide subjects who mention all four outcomes dependent on whether subjects explicitly report payoffs associated with $p = v_L$ and (or) $p = v_H$: Both (One) Explicit; report payoffs qualitatively (Qualitative) or only mention all outcomes without addressing payoffs (Mention Outcomes Only).

A subject who understood that submitting prices $p \notin \{v_L, v_H\}$ is dominated needs to be aware of and take into account the four outcomes (v, p), where $v, p \in \{v_L, v_H\}$, in order to compute the payoff-maximizing price.[38] Table 6 shows that 26.6 percent of subjects in PROB mention all four outcomes in at least some way compared to 52.5 percent in DET, a significant difference.[39] Only 3 subjects (1.6 percent) in PROB explicitly compute payoffs for both $p = v_L$ and $p = v_H$, compared to 13.7 percent of subjects in DET. An additional 8.7 percent of subjects in DET compute the payoff associated with either $p = v_L$ or $p = v_H$.[40] While one reason subjects do not mention all four outcomes is because they do not take all of them into account, another could be that advisors expect advisees to prefer simple advice. Since the written advice may only be used to justify a recommendation and hence may leave out some outcomes, we do not only analyze differences in the written advice. Rather, the main analysis is on how the advice correlates with the strategies of advisors in parts 1 and 2, though the advice may only be a noisy signal of what outcomes individuals thought about and focused on when submitting prices.

Table 7 shows that in both treatments, mentioning all four outcomes is positively correlated with being classified as $V_L V_H$ and negatively correlated with being classified as using a dominated strategy that is not focal ($p \notin \{v_L, v_H\}$). Furthermore, controlling for whether subjects mention all four outcomes, reduces the difference in the $V_L V_H$ classification between PROB and DET by more than half, and it is no longer significant.[41] Furthermore, among subjects who mention all four outcomes, the modal classification is $V_L V_H$: 52.0 and 64.6 percent in PROB and DET, respectively, compared

[38] The advice was classified following the protocol described in Online Appendix H. The classification was verified by a research assistant. Examples of each class of advice are provided in Online Appendix G.

[39] To test for significance, we run a linear regression in which the left-hand side variable is a dummy that takes value 1 if the subject submitted advice that mentions all four outcomes and on the right-hand side we include a constant and a treatment dummy (1=DET). The coefficient on the treatment dummy is positive and significant at the 1 percent level (p-value <0.01).

[40] Some subjects in PROB mention all four outcomes without, however, providing any guidance as to how to use these outcomes to compute payoffs.

[41] We also test the null hypothesis that the treatment effect coefficient is equal to .244, the estimate reported in Table 4; the p-value equals .085.

	(1) $V_L V_H$	(2) $\{MixV_H, V_H V_H\}$	(3) Focal	(4) Dom	(5) Dom or Res
Det	0.117	0.021	0.004	−0.084	−0.142*
	(0.073)	(0.074)	(0.064)	(0.070)	(0.082)
Advice mentions all outcomes	0.395***	−0.088	0.062	−0.184***	−0.369***
	(0.068)	(0.069)	(0.059)	(0.065)	(0.076)
Advice mentions all outcomes × Det	0.025	−0.148	−0.112	0.090	0.234**
	(0.092)	(0.093)	(0.080)	(0.088)	(0.103)
Num Risky	−0.029	−0.015	0.003	0.058**	0.042
	(0.029)	(0.029)	(0.025)	(0.028)	(0.032)
Num Risky × Det	0.004	0.050	−0.016	−0.036	−0.039
	(0.042)	(0.043)	(0.037)	(0.040)	(0.047)
Constant	0.132*	0.257***	0.178***	0.220***	0.433***
	(0.069)	(0.070)	(0.060)	(0.066)	(0.077)
Observations	371	371	371	371	371

Table 7: Main Treatments: Estimation output using last 5 rounds of part 1 and part 2 for the classification of types

Notes: Results from a linear regression. The dependent variable takes value 1 if the subject is classified as (1) $V_L V_H$, (2) $\{MixV_H, V_H V_H\}$ (3) Focal, (4) Dom, (5) Dom or Res. Det is a treatment dummy that equals 1 if the subject participated in DET. Advice mentions all outcomes is a dummy that equals 1 if the advice of the subject mentions all four outcomes (v, p) with $v, p \in \{v_L, v_H\}$. Num Risky is the number of risky lotteries the subject chooses in part 5 (from 0 to 3). The regression also includes demographic controls (Gender, Ethnicity, Age and Schooling) and a control for the number of errors in the instructions. The full output is presented in Table 35 of Online Appendix E.

to only 8.0 and 16.1 percent among all other subjects ($p < 0.01$ and $p < 0.01$, respectively).

To further our understanding of what subjects focus on and how this relates to the strategies they use, we exclude from our sample subjects who mention all four outcomes (as we already have a pretty good understanding of their behavior). A large fraction of subjects mention no outcome (40.4 and 29.5 percent in PROB and DET, respectively). For a complete classification see Table 31 in Online Appendix E. The largest category of subjects who mention some (but not all) outcomes are subjects who mention large gains, that is they highlight the gains that can be obtained for a firm of value $v = v_H$ when submitting a price of $p = v_H$.[42] We have 12.8 and 6 percent of such subjects in PROB and DET, respectively. The modal classification of these subjects is *Mix* or V_H in part 1: 62.5 and 54.6 percent in PROB and DET, respectively, compared to only 29.8 and 27.6 percent among the other subjects ($p < 0.01$ and $p = 0.08$, respectively). However, they are probably not using *Mix* or V_H in part 1 because they are risk-seeking. Indeed, 62.5 and 90.9 percent of subjects who mention large gains in PROB and DET, respectively, select the safe alternative (corresponding to $p = v_L$) when they face the lotteries that correspond to the problem $(v_L, v_H) = (20, 120)$ in part 5.

The final substantive group consists of subjects who only mention large losses. These are subjects who mention or highlight the losses that can result for a firm of value $v = v_L$ when submitting a

[42] These are subjects who mention only $\{(v_L, v_L)$ and $(v_H, v_H)\}$, only $\{(v_H, v_H)\}$, only $\{(v_L, v_H)$ and explicitly $(v_H, v_H)\}$ or only $\{(v_L, v_L)$, (v_H, v_L) and $(v_H, v_H)\}$. We consider an outcome to be mentioned if it is mentioned explicitly or implicitly. An example in which (v_H, v_H) is only implicitly mentioned would be the following: "submitting 120 can lead to a gain, but it's not worth the risk." In this case we consider that the subject is implicitly mentioning that if $v = v_H$, there would be a gain.

price of $p = v_H$.[43] We have 13.3 and 3.8 percent of such subjects in PROB and DET, respectively. These subjects are quite often classified as V_L in part 1: 40 and 100 percent in PROB and DET, respectively, compared to only 14.2 and 15.0 percent among the other subjects ($p < 0.01$ and $p < 0.01$, respectively). While we do not have sufficiently many subjects in DET for a detailed analysis, subjects in PROB who only mention large losses are more likely to be classified as a Focal type with V_L in part 1 (such as $V_L V_L$ or $V_L Mix$), that is a type who is definitely not profit maximizing, namely 21.9 percent compared to only 5.7 percent among other subjects ($p < 0.01$).[44]

A further 3.2 and 4.4 percent of subjects in PROB and DET, respectively, only mention outcomes (v, p) associated with $p = v_L$, that is either only $\{(v_L, v_L)\}$ or only $\{(v_L, v_L)$ and $(v_H, v_L)\}$, and a final 3.7 and 3.8 percent of subjects, respectively, make a mistake, that is, compute at least one of the payoffs wrongly.

To summarize, the results point towards subjects in DET being much more likely to consider all four outcomes (v, p) where $v, p \in \{v_L, v_H\}$. Subjects who mention all four outcomes in their advice are significantly more likely to be classified as $V_L V_H$. Furthermore, controlling for whether the advice mentions all four outcomes significantly reduces and in fact accounts for the differences between PROB and DET. It seems that the fact that the two companies exist in DET, rather than being possible states in PROB, makes it easier for subjects to think about the outcomes they receive for each firm for a given price. This, in turn, seems to be responsible for the *Power of Certainty.*

4.2 Across Treatment Adaptation and Mentioning All Outcomes

We provide a second piece of evidence as to the relevance of thinking about all four outcomes, or more precisely, of subjects mentioning all four outcomes in the advice. In part 4, subjects encounter 10 problems that have the same characteristics as those of part 1, including that $2v_L < v_H$. Subjects now, however, play the opposite treatment than in part 1, that is, subjects in DET encounter probabilistic problems and subjects in PROB face deterministic problems in part 4. Note that part 4 comes after part 3, where subjects provided advice. So, instead of connecting the part 3 advice to the strategies subjects were classified as before giving advice, we now use the advice as a control for the future part-4 behavior.

We classify subjects in part 4 based on the prices submitted in the last 5 rounds. Note that in the deterministic problems (encountered in part 4 of PROB), $p = v_L$ is the dominant strategy, while in probabilistic problems very risk-seeking subjects might find that $p = v_H$ maximizes their earnings. All other prices are dominated. Using linear regressions we show that subjects who are classified as $V_L V_H$ in part 1 are significantly more likely to be classified as V_L^4 in part 4 (that is submitting $p = v_L$ in the last 5 rounds of part 4), and less likely to be classified as Dom^4 (submitting dominated

[43]These are subjects who mention only $\{(v_L, v_H)\}$, only $\{(v_L, v_H)$ and implicitly $(v_H, v_H)\}$, only $\{(v_L, v_L)$, (v_H, v_L) and $(v_L, v_H)\}$ or only $\{(v_L, v_L)$, (v_L, v_H) and $(v_H, v_H)\}$.

[44]The written advice correlates in expected ways with the price advisors recommend, see Table 29 of Online Appendix E.

prices $p \notin \{v_L, v_H\}$ in part 4). Table 8 also shows that subjects who mention all four outcomes are more likely to be classified as V_L^4 and less likely to be classified as Dom^4 or V_H^4.

	(1) V_L^4	(2) V_H^4	(3) Mix^4	(4) Dom^4	(5) Dom^4 or Res^4
Det	−0.103	0.045	0.084*	0.010	−0.026
	(0.063)	(0.056)	(0.050)	(0.052)	(0.059)
Num Risky	−0.059*	0.054*	−0.025	0.036	0.030
	(0.032)	(0.029)	(0.026)	(0.027)	(0.030)
Num Risky × Det	−0.061	0.016	0.038	−0.017	0.006
	(0.047)	(0.042)	(0.037)	(0.039)	(0.044)
Advice mentions all outcomes	0.134**	−0.106**	0.030	−0.089*	−0.057
	(0.057)	(0.051)	(0.045)	(0.047)	(0.053)
$V_L V_H$	0.331***	−0.030	−0.049	−0.188***	−0.252***
	(0.059)	(0.053)	(0.047)	(0.049)	(0.055)
Constant	0.312***	0.210***	0.068	0.289***	0.409***
	(0.074)	(0.066)	(0.059)	(0.061)	(0.069)
Observations	371	371	371	371	371

Table 8: Main Treatments: Estimation output using part 4 types

Notes: Results from a linear regression. The dependent variable takes value 1 if the subject is classified as (1) V_L^4, (2) V_H^4, (3) Mix^4, (4) Dom^4, or (5) Dom^4 or Res^4. Det is a treatment dummy that equals 1 if the subject participated in DET, which in this case means that those subjects are facing probabilistic problems. Num Risky is the number of risky lotteries the subject chose in part 5 (from 0 to 3). $V_L V_H$ is a dummy that takes value 1 if the subject was classified as $V_L V_H$ in parts 1 and 2. Advice mentions all outcomes is a dummy that equals 1 if the advice of the subject mentions all four outcomes (v, p) with $v, p \in \{v_L, v_H\}$. The regression also includes demographic controls (Gender, Ethnicity, Age and Schooling) and a control for the number of errors in the instructions. The full output is presented in Table 21 of Online Appendix B.

The results from part 4 confirm that mentioning all four outcomes is also relevant in determining a subject's ability to submit optimal prices in an environment that is slightly different than the one they encountered before.

4.3 Advisee Treatment

In this section, we briefly comment on the effects of receiving advice. Subjects in the Advisee treatments saw a piece of advice from five different subjects about what price to submit in round 1 of part 1, which corresponds to the problem $v_L = 20$ and $v_H = 120$. Subjects in ADVPROB (ADVDET) received advice from subjects in PROB (DET). The subjects then faced part 1 and part 2 from the main treatments.

We point out that the results on the effect of advice are speculative as the advisee treatments were only conducted with the purpose to incentivize the advice (part 3) of the main treatments. Specifically, the advice participants received was not calibrated to be comparable across treatments ex-ante. In particular, subjects in ADVDET are more likely to receive advice that mentions all four outcomes, given that four outcomes are mentioned more frequently by participants in DET.[45]

[45] Table 32 in Online Appendix E shows the classification into types dependent on whether subjects received an advice that mentions all four outcomes, whether they selected it or not. All but one subject in ADVDET received an

However, since advisees received advice from five subjects, 97.5 percent of subjects in ADVDET received information on all four outcomes across the five pieces of advice, which is not very different from the 95.1 percent of subjects in ADVPROB.[46]

Classifying advisees just like subjects in PROB and DET, 24.4 percent of subjects are classified as $V_L V_H$ in ADVPROB, which is similar to the 19.7 percent in PROB. In contrast, almost three quarters of subjects in ADVDET (72.5 percent) are classified as $V_L V_H$, which is almost double the fraction of such subjects in DET (41.5 percent), a significant difference. For statistical tests see Table 33 in Online Appendix E and Table 36 for the type classification. Not only is the advice helpful in deterministic but not in probabilistic problems, the difference is also significant, see Table 34 in Online Appendix E. While previous work has often shown that decisions made with advice are closer to the predictions of economic theory than choices made without advice, this is not borne out in probabilistic problems, see Schotter (2003) for a literature survey. This highlights that explaining and interpreting advice in PROB seems to be relatively difficult, and much more difficult than in DET.

Furthermore, subjects who received advice that mentions all four outcomes and recognized its significance are more likely to be classified as $V_L V_H$ than those who did not, see the regression in Table 34 of Online Appendix E. This provides some additional hint that thinking about the four possible outcomes (v, p) with $v, p \in \{v_L, v_H\}$ is crucial to be able to behave optimally in the Acquiring-a-Company problem.

5 Robustness: *PoC* and Probability Matching

We consider a second environment that is famous in its own right, in which we test the *PoC* hypothesis to evaluate its robustness. In the Probability-Matching problem, we vary the probabilities of events to be 60:40, 70:30 or 80:20 instead of just 50:50. Second, like in part 2 of the Acquiring-a-Company problem, there is a dominant strategy, so risk preferences play no role. The drawback is that the dominant strategy is on the boundary of the action set, which often reduces optimization rates.[47] Moreover, as we describe in the next paragraph, the action set itself involves lotteries. While we would expect these features to make the deterministic problem more difficult and hence reduce the relative importance of *PoC*, the environment does provide a stress test for the *PoC* hypothesis. In other words, the goal is to test if even in this famously hard problem, where teaching subjects not to do probability matching has been difficult, there is still a significant *PoC* effect.

advice that mentioned all four outcomes. This is why we will not be able to use the variable "received advice that mentions all four outcomes" when accounting for type classification in the Advisee treatments.

[46] A linear regression with a treatment dummy on the right-hand side indicates that the treatment effect is small (0.024) and not statistically significant (p-value 0.577).

[47] It has been well documented that when the optimal action is on the boundary, it tends to be selected less often. An example is the case of voluntary contribution games, where the equilibrium is not to contribute. For a detailed discussion see Vesterlund (2016).

In the Probability-Matching problem, there are ten boxes of which, for example, six are green and four are red (orange in the experiment). The subject guesses a color, green or red, and to do so selects how many of 100 tickets are green and how many are red. One ticket is randomly chosen (with equal chance), and the color of that ticket determines the subject's guess.

In probabilistic problems, one of the ten boxes is randomly selected with equal chance. The subject receives \$10 if her guess matches the color of the box and \$0 otherwise. Many similar versions of this problem have been extensively studied, and the typical finding is that subjects fail to behave optimally (i.e. fail to select the dominant strategy of 100 green tickets).[48] Instead, many subjects use a strategy referred to as 'probability matching' (hence the name of the problem), which in our setting involves selecting 60 green and 40 red tickets.[49] In the deterministic problem the subject's guess is applied to each of the 10 boxes and the subject receives \$1 for each box that matches the color of the guess and \$0 for all other boxes.[50]

In terms of the framework introduced in Section 2.1, following the example above, there are two possible values of the state, $\omega \in \{g, r\}$, with the probability of g(reen) and r(ed) being $p_{\text{g}} = 0.6$ and $p_{\text{r}} = 0.4$, respectively. Let a be the agent's action, where with a slight abuse of notation $a \in A = \{\,(a, 100-a) |\, a \in \{0, 1, ..., 100\}\}$. Let $\text{Pr}_{\text{G}}(a) = \frac{a}{100}$ be the probability that the guess of the subject is G(reen) and $\text{Pr}_{\text{R}}(a) = 1 - \frac{a}{100}$ that it is R(ed). The expected payment for a given action a in the probabilistic problem is $\pi^{Prob}(a) = [\text{Pr}_{\text{G}}(a)p_{\text{g}} + \text{Pr}_{\text{R}}(a)p_{\text{r}}]10$, since the payment for matching the color is \$10. In deterministic problems the payment is given by $\pi^{Det}(a) = [\text{Pr}_{\text{G}}(a)p_{\text{g}} + \text{Pr}_{\text{R}}(a)p_{\text{r}}]10$, since all ten boxes are used for payment.[51]

Subjects either faced probabilistic (PROB^{PM}) or deterministic (DET^{PM}) problems. In each case subjects faced 12 problems, where the number of green boxes is 2, 3, 4, 6, 7, or 8, and the payoffs are either as described or half as much. The first problem subjects faced involves six green and four red boxes and payoffs are \$5 and \$0.5, respectively. The remaining 11 problems were randomized on the individual level.

Subjects were Amazon Turk workers recruited as described in Section 2.8, received \$2 for completing the experiment as well as the earnings from a randomly chosen round. Using the same procedures as in the Acquiring-a-Company treatments, we were left with 267 and 266 subjects in PROB^{PM} and DET^{PM}, respectively. On average a participant received \$6.8 (including the \$2 for finishing the survey)

[48]Note that the only requirement for this strategy to be optimal is that when one lottery first order stochastically dominates another, the agent prefers it.

[49]For early evidence see Siegel and Goldstein (1959), and Fantino (1998) and Vulkan (2000) for surveys.

[50]An alternative implementation would have subjects facing the same problem several times with the subject guessing a single color in each repetition. The standard finding in the literature in this case is that subjects switch the color of choice. We decided not to use this implementation because we wanted to face subjects with many problems of varying probabilities and it may be cumbersome that each problem is repeated many times. However, the alternative implementation would completely eliminate the use of any randomization in the deterministic treatment, so we speculate that it may have led to larger treatment effects.

[51]The One-Value problem corresponds to asking subjects to guess the color of a single box of known color, suggesting that "background complexity" may be minimal. For this reason we did not conduct the One-Value treatment that corresponds to this environment.

	(1) OT	(2) Optimal	(3) PMT	(4) PM
Det	0.103**	0.087**	−0.077*	−0.054*
	(0.045)	(0.043)	(0.043)	(0.029)
Constant	0.293***	0.322***	0.216***	0.208***
	(0.064)	(0.057)	(0.061)	(0.045)
Observations	533	6396	533	6396

Table 9: Probability-Matching Treatments: Treatment Effects

Notes: Results from a linear regression. (1) The dependent variable takes value 1 if the subject is classified as Optimal Type (OT). (2) The dependent variable takes value 1 if the choice is optimal. (3) The dependent variable takes value 1 if the subject is classified as the Probability-Matching Type (PMT). (4) The dependent variable takes value 1 if the choice is exactly probability matching. In columns (2) and (4) standard errors are clustered by subject. Det is a dummy variable that takes value 1 if the observation corresponds to the deterministic treatment. All regressions control for: number of errors in instructions, number of errors in instructions interacted with Det, Gender, Ethnicity, Age and Education Level. The full output is presented in Table 38 of Online Appendix I.

and took approximately 15 minutes to complete the survey.[52]

A guess is classified as Optimal if all 100 tickets match the majority color of the 10 boxes and as PM (Probability Matching) if they match the proportion of green and red boxes. A subject is classified as OT, the Optimal Type, or PMT, the Probability-Matching Type, if all her guesses in the last 5 problems are Optimal, or PM, respectively; for detailed classification see Online Appendix I.

Linear regressions in Table 9 include controls similar to those in the Acquiring-a-Company regressions, for full results see Table 38 of Online Appendix I. The output in column (1) shows that there are about one third more subjects classified as the Optimal Type (OT) in DET^{PM} than in PROB^{PM}, this 10 percentage-points treatment effect is significant. In contrast, PROB^{PM} has about one-third more subjects (an 8 percentage-points treatment effect) classified as PMT, as shown in Column (3). The regressions in columns (2) and (4) confirm that results are similar when instead of type classification of subjects we use only the classification of guesses. In conclusion, we find a significant *PoC* effect in the Probability-Matching problem, which is notoriously difficult for subjects.

6 Discussion

We demonstrated a *PoC* effect in the Acquiring-a-Company problem (and replicated it in the very different Probability-Matching problem) and showed that it can have a significant role in accounting for the *Complexity of Contingent Reasoning.* In this section we first provide another way to assess the importance of *PoC*. We then extensively discuss potential underlying mechanisms for *PoC*.

[52]The procedures, including subject restrictions and final survey, were as in the Acquiring-a-Company experiment, see Online Appendix F for a summary of the instructions and the Instructions Appendix for details of all the Qualtrics surveys. Of the 576 Amazon Turk workers with unique IP addresses who started the experiment, 47.8 percent are female, 75.7 percent are White, 52.6 percent are at or below the median age of 34 and 35.6 percent have low schooling (‘Some college,’ ‘High School’ or lower education level). Of the 285 subjects in PROB^{PM} and 291 in DET^{PM}, 18 and 25 made more than 2 mistakes in the instructions for part 1, respectively. The difference in the drop rate is not significant across treatments.

A perhaps more speculative way of assessing the importance of *PoC* is to compare it to other mechanisms that often improve the performance of agents, specifically learning. In the Acquiring-a-Company problem, which received more attention from economists than Probability Matching, we can use data from the two-value problem $v_L = 20$, $v_H = 119$ in Charness and Levin (2009) to address the effect of learning. They had 111 subjects (who first encountered other problems) play 30 rounds, and after each round subjects received feedback in terms of payoffs. We can compare the fraction of $p = v_L$ prices in rounds 1-10 and 11-20 in PROB and DET, as well as the comparable rounds in Charness and Levin (2009). In rounds 1-10 in PROB, 42.3 percent of prices are v_L, compared to 52.3 in DET. That is, *PoC* leads to a roughly twenty-five percent (or ten percentage-points) increase in v_L prices. In rounds 11-20, the percent of v_L prices are 43 in PROB and 55.5 in DET, leaving the relative importance of *PoC* roughly unchanged. In Charness and Levin (2009), where subjects are always confronted with the $v_L = 20$, $v_H = 119$ problem, 50.9 percent of prices are $p = v_L = 20$ in the first 10 rounds, compared to 56 percent in the second 10 rounds (and 54 percent in the last 10 rounds). This suggests a learning effect of only five percentage-points, or a ten percent increase in $p = v_L$ prices in the second compared to the first set of 10 rounds. This indicates that *PoC* is important in increasing payoff maximization even relative to other well-studied mechanisms such as learning.

If future research confirms the significant role of *PoC,* it will be important to develop a satisfying theoretical model. In the remainder of this section, we provide some speculation as to an underlying mechanism of *PoC.*

Our findings suggest that thinking through all contingencies is challenging in probabilistic settings, but individuals are (more) able to do so in deterministic settings. As we showed, subjects in deterministic problems are more likely to mention consequences of all contingencies in their advice compared to subjects in probabilistic problems. Another example that individuals have problems thinking through the state space, which can be alleviated when they receive some help is by Esponda and Vespa (2018a). They consider probabilistic settings in which optimal behavior requires satisfying Savage's sure-thing principle (Savage, 1972).[53] The sure-thing principle asserts that if a person prefers A to B if event X realizes and is indifferent between A and B if X does not realize, then the person prefers A to B prior to learning whether X results or not.[54] Esponda and Vespa (2018a) first reproduce well-known failures of standard theory in decision-theory environments (e.g. the Ellsberg Paradox) and in games (e.g. overbidding in a second-price auction). They then show that failures

[53]Other experiments also show that for specific environments choices may not be consistent with the sure-thing principle and they refer to discrepancies as the disjunction effect (e.g. Shafir and Tversky, 1992, Tversky and Shafir, 1992, Shafir, 1994, and Croson, 1999). The psychology model for the disjunction effect is that individuals do not make choices based on consequences of decisions, but rather based on reasons for making one choice over another. For example, suppose a student would want to take a vacation in Hawaii when they pass a big test to celebrate and when they fail this big test in order to console themselves. Then, if the student has to decide whether to go to Hawaii before knowing the test results, the student has no reason to go, violating the sure thing principle, see also Shafir et al. (1993).

[54]In the Acquiring-A-Company and Probability-Matching problems the sure-thing principle applies in a trivial manner because there is no set of states where the agent is indifferent between available actions. That is, X equals the full state space.

are reduced in a treatment that keeps uncertainty unchanged but makes focusing on event X salient. Their results show that without help many subjects do not partition the state space into X realizing and X not realizing, which suggests that many subjects do not think through the state space. If, as the evidence so far suggests, PoC indeed helps subjects think through all contingencies before submitting a choice, this would be helpful for agents also when the sure-thing principle applies in a non-trivial manner (i.e. when set X is not equal to the full state space).

To understand why the deterministic treatment may help subjects think through all contingencies compared to the probabilistic treatment, note that in the deterministic treatment, to compute the payoff of a certain action requires considering all contingencies. By contrast, in the probabilistic treatment, it is feasible to compute *a* possible payoff of an action without having to compute *all* possible payoffs of that action. For example, an agent who evaluates submitting $p = 120$ in the probabilistic Acquiring-a-Company problem with $v_L = 20$ and $v_H = 120$ can compute the payoff she receives if $v = v_H$ without having to compute the payoff if $v = v_L$. In the deterministic setting this is not possible. A subject who wants to compute the payoff of some action needs to consider all contingencies. For example, the agent who evaluates the payoff of submitting $p = 120$ would buy both firms and hence, in order to compute payoffs, needs to consider both contingencies. The arguments for Probability Matching are similar.

One potential mechanism for why agents do not think through all contingencies in probabilistic settings is the "control heuristic" (Thompson et al., 1998) from the psychology literature. This heuristic is a shortcut that involves two elements: the intention of the subject to achieve an outcome, and the perceived connection between the subject's action and the desired outcome. In the Acquiring-a-Company problem the outcome of desire would be the highest-payoff outcome $(p = v_H, v = v_H)$. The subject may reason that this desired outcome will only occur if she submits $p = v_H$. However, the subject's misperception would come from the agent believing that by selecting $p = v_H$ it will be more likely for the value of the firm to be v_H. The control heuristic can also rationalize some of our findings in the Probability-Matching problem. The desired outcome would be to always match the color of the random box in the probabilistic setting. The subject may reason that this outcome will only occur if the fraction of green tickets matches the fraction of green boxes. The misperception would come from the agent believing that the right ticket will be selected for a given random box. In both the Acquiring-a-Company and the Probability-Matching problem, one reason for the difficulties in the probabilistic treatment may come from the subject's belief that her actions can influence which state of the world realizes.[55]

While we do not have a model, and the recent theoretical literature on biases in decision making cannot directly account for our results (e.g. Sims, 2003, Bordalo et al., 2012, Gabaix, 2014, and Caplin et al., 2018), there is one paper that predicts a difference between probabilistic and deterministic problems in our two environments. Li (2017) can account for the main comparative static result in submitted strategies, though we conjecture perhaps not exactly for the right reasons. The

[55] The psychology literature, through Gigerenzer and Hoffrage (1995), also provides some indication that a deterministic setting might help in problems such as Bayesian updating.

paper asks under what circumstances a strategy-proof mechanism is obviously so, that is such that agents may recognize this fact and hence submit the dominant strategy. Specifically, a dominant strategy is defined as obviously dominant if the best outcome from a deviation is no better than the worst outcome from following the dominant strategy.

In contrast to the motivating environments in Li (2017), we have no strategic interaction but the concept of obvious strategy proofness can still be applied. Submitting $p = v_L$ is a dominant and indeed obviously-dominant strategy in part 1 of DET (where $2v_L < v_H$), since the return from $p = v_L$ is $0.5v_L$, while that from $p = v_H$ is $1.5v_L - 0.5v_H < 0.5v_L$. However, $p = v_L$ is not an obviously-dominant strategy in PROB: the best outcome from $p = v_H$ is $0.5v_H$, which is larger than the worst outcome from $p = v_L$, which is 0. Similarly, for part 2, it is possible to show that $p = v_H$ is obviously dominant only in DET but not in PROB. Likewise in the Probability-Matching problem, the worst outcome from the dominant strategy is 0, while the best outcome from deviation is the maximal payment of \$10. If agents are able to recognize dominant strategies only if these are obviously dominant, then we would expect a comparative static between PROB and DET such as the one observed in the paper.

One reason we believe that obvious dominance may not be exactly the right concept in our environment is that we were able to link the failure to play the dominant strategy to the failure to mention all four outcomes (v, p), where $v, p \in \{v_L, v_H\}$. Furthermore, it is not straightforward to link an agent who can only compare sets, as in the best outcome from deviation to the worst outcome using the dominant strategy, to an agent who seems to actually not take all possible outcomes into account. One question open for future research is to design an experimental test that assesses if certainty helps even if profit-maximizing strategies are already obviously dominant in a probabilistic setting.

7 Conclusion

In economic environments with uncertainty there is well documented evidence that individuals often fail to behave optimally. In this paper we study what can cause these difficulties. To do this, we first define the *Complexity of Contingent Reasoning* as the reduction in the rate with which agents use payoff-maximizing strategies when moving from an environment with one state of known value (a *one-value* problem) to one where the state can have several possible values (a *probabilistic* problem). We propose a new hypothesis to account for this *Complexity of Contingent Reasoning*, the presence of uncertainty. Specifically, there are two changes when moving from a *one-value* to a *probabilistic* problem. First, the number of states an agent needs to consider increases. We attribute the resulting problems in payoff maximization to *Computational Complexity*. Second, the *probabilistic* problem introduces uncertainty or eliminates what we term the *Power of Certainty* (*PoC*).

To show the existence and relevance of *PoC*, we introduce a new environment, the *deterministic* problem, where the agent's action set and state-specific payoffs are identical to the *probabilistic* problem, and the agent's action is similarly applied to all possible values of the state. However, in the *deterministic* problem, the payoff of an action is the sum of payoffs from each state, weighted by the probability with which the state occurs in the *probabilistic* problem. That is in both the *probabilistic* and the *deterministic* problem all possible states are payoff relevant. The main difference is that while there is uncertainty in the *probabilistic* problem, there is certainty in the *deterministic* problem.

To provide evidence for *PoC*, we ran experiments in an Acquiring-a-Company problem and in a Probability-Matching problem. In both cases, subjects are significantly more likely to behave optimally in the *deterministic* than in the *probabilistic* setting. We also use the Acquiring-a-Company problem to show that *PoC* can represent a large proportion of the difficulties captured by the *Complexity of Contingent Reasoning*. We show that about half of this increase in the failure to maximize payoffs is due to the loss of *PoC* between the *one-value* and the *probabilistic* problem. Finally, we inquired into possible mechanisms behind these results. Our findings suggest that the *deterministic* problem may make it easier for subjects to consider all payoff-relevant outcomes when deciding on an action.

While we provided evidence for *PoC* in two different settings, our framework suggests *PoC* is relevant in any problem with uncertainty. For example, our results suggest that firms are more likely to be optimizing relative to consumers. A standard argument is that it is competition among firms that eliminates those firms who fail to maximize profits, but our findings provide a new explanation: The *PoC* hypothesis suggests that any agent may find it easier to optimize when they face an environment in which all events are realized (possibly with different frequencies or weights), relative to when playing a lottery in which one of many possible events will result. *PoC* asserts that individuals who make choices for a firm may find it easier to optimize relative to an individual making a choice as a consumer.

PoC may not only be relevant in decision problems (as the ones we study in this paper), but also in strategic settings. For example, consider a classic independent private-value environment in which agents bid in a second price auction. Experiments have shown that individuals often fail to submit the dominant strategy that consists of bidding their value when bidding against one opponent with several possible private values. Consider instead the following problem that keeps the values of the opponent the agent has to consider constant, but eliminates uncertainty. Suppose the agent faces as many second price auctions as there are possible values for the opponent, where now, however, in each auction the agent faces an opponent of known value. The restriction is that the agent has to submit the same bid in all auctions. It is easy to see that bidding her own value is still a dominant strategy for the agent, though *PoC* may make it easier for the agent to realize this.

Finally, for policy purposes it may be useful to explore to what extent the gains from describing a problem without uncertainty can also be captured when, in fact, there is uncertainty. Note that in the Acquiring-a-Company problem risk preferences were not such a large concern in part 1 because most individuals are risk averse. In part 2, as well as in the Probability-Matching problem, we had a dominant strategy for any subject who, when one lottery first-order stochastically dominates another, would prefer the former. However, there are many problems, for example when deciding whether to purchase an insurance, where a subjects risk preferences are crucial when deciding about the optimal choice. It remains an open question to what extent *PoC* could help subjects in such environments as well.

We provide three possible approaches to extend the impact of *PoC* to environments where risk preferences play a role. For the first approach, note that in the current experiments we designed the deterministic problem so that payoffs correspond to expected utility payoffs in the probabilistic problem for a risk-neutral individual. However, for any given risk parameter, one could generate a deterministic problem that corresponds to the expected utility of each action given that risk parameter. Future work can assess the role of *PoC* for such a risk-parameter adjusted deterministic version of a probabilistic problem.

For the second and third approach we use once more the Acquiring-a-Company problem for illustrative purposes. Second, suppose we present the agent with the *deterministic* problem, with the following twist. After having purchased none, one or two firms, instead of paying the agent for both firms, we afterwards randomize which of the two firms the agent was actually allowed to purchase and hence which counts for payment. We could call this the 'Late Lottery' treatment which is, of course, just a different frame of the probabilistic treatment. However, narrow bracketing (basically thinking of the problem while ignoring the 'late randomization device') may lead the agent to treat this, at first, as the *deterministic* problem, which may hence help the agent to think through the full state space. It remains to be seen whether the agent can then widen the bracket and re-introduce her risk preferences before making a decision. Third, we could tell the agent that one of the two firms has already been chosen, v_L or v_H, the agent simply does not know yet which one, we could call this the 'Unknown Resolved Lottery' treatment. This manipulation could reduce any mistaken belief of the agent that her actions can influence the outcome of the state, and could provide a test for the role of the control heuristic in accounting for losses incurred by the lack of *PoC*. We leave it to future research to determine the usefulness of such framing devices.

A final way in which our design could prove useful concerns not directly *PoC*, but rather *deterministic* problems. Specifically, the *deterministic* problem can provide a clear indication as to the costs due to *Computational Complexity* in terms of failure to maximize payoffs. This could be helpful in providing a bound as to the role of "complex" preferences in accounting for individual choices. Put differently, it has been notoriously difficult to assess whether a subject's choice is due to a 'mistake' or an idiosyncratic preference. The *deterministic* problem can be useful in providing a lower bound for how prone individuals are to make mistakes. This is because the failure to optimize in a *deterministic* setting means leaving money on the table (which, in environments where fairness concerns do not play a large role, is indicative either of a 'mistake' or of a cost that the agent is willing to pay perhaps for not fully paying attention). For example, assume subjects only submit prices $p \in \{v_L, v_H\}$ in part 1 of the Acquiring-a-Company problem. Then one could be tempted to attribute $p = v_H$ prices to risk-seeking preferences. Suppose, however, that among this same set of subjects a certain fraction, say 20 percent, also submit those $p = v_H$ prices in *deterministic* problems (where they hence select less over more money). It seems then more than plausible that $p = v_H$ prices of those 20 percent of participants in *probabilistic* problems are *not* an indication of risk-seeking preferences! We leave it to future research to determine whether a measure of computational complexity can be useful in bounding the role of 'complex' preferences in accounting for idiosyncratic choices and can help uncover 'mistakes.'

Bibliography

Ambuehl, Sandro and Shengwu Li, "Belief updating and the demand for information," *Games and Economic Behavior*, 2017, *109*, 21–39.

_ **, B. Douglas Bernheim, and Annamaria Lusardi**, "A Method for Evaluating the Quality of Financial Decision Making, with an Application to Financial Education," *National Bureau of Economic Research*, 2017.

Araujo, Felipe, Stephanie Wang, and Alistair J. Wilson, "The Times are A-Changing: Dynamic Adverse Selection in the Laboratory," *Working paper*, 2018.

Ball, Sheryl B, Max H Bazerman, and John S Carroll, "An evaluation of learning in the bilateral winner's curse," *Organizational Behavior and human decision Processes*, 1991, *48* (1), 1–22.

Barrett, Garry F and Stephen G Donald, "Consistent tests for stochastic dominance," *Econometrica*, 2003, *71* (1), 71–104.

Bénabou, Roland and Jean Tirole, "Self-confidence and personal motivation," *The Quarterly Journal of Economics*, 2002, *117* (3), 871–915.

Bereby-Meyer, Yoella and Alvin E Roth, "The speed of learning in noisy games: Partial reinforcement and the sustainability of cooperation," *The American Economic Review*, 2006, *96* (4), 1029–1042.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, "Salience theory of choice under risk," *The Quarterly Journal of Economics*, 2012, *127* (3), 1243–1285.

Brunnermeier, Markus K and Jonathan A Parker, "Optimal expectations," *The American Economic Review*, 2005, *95* (4), 1092–1118.

Camerer, Colin F, Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Johan Almenberg, Adam Altmejd, Taizan Chan et al., "Evaluating replicability of laboratory experiments in economics," *Science*, 2016, *351* (6280), 1433–1436.

_ **, Teck-Hua Ho, and Juin-Kuan Chong**, "A cognitive hierarchy model of games," *The Quarterly Journal of Economics*, 2004, *119* (3), 861–898.

Caplin, Andrew and John Leahy, "Psychological expected utility theory and anticipatory feelings," *The Quarterly Journal of Economics*, 2001, *116* (1), 55–79.

_ **and Mark Dean**, "Revealed preference, rational inattention, and costly information acquisition," *The American Economic Review*, 2015, *105* (7), 2183–2203.

_ **,** _ **, and John Leahy**, "Rationally inattentive behavior: Characterizing and generalizing Shannon entropy," *Working Paper*, 2018.

Cason, Timothy N and Charles R Plott, "Misconceptions and game form recognition: Challenges to theories of revealed preference and framing," *Journal of Political Economy*, 2014, *122* (6), 1235–1270.

Charness, Gary and Dan Levin, "The origin of the winner's curse: a laboratory study," *American Economic Journal: Microeconomics*, 2009, *1* (1), 207–236.

Coffman, Lucas C and Muriel Niederle, "Pre-analysis plans have limited upside, especially where replications are feasible," *Journal of Economic Perspectives*, 2015, *29* (3), 81–98.

_ **,** _ **, and Alistair J Wilson**, "A Proposal to Organize and Promote Replications," *American Economic Review (Papers & Proceedings)*, 2017, *107* (5), 41–45.

Costa-Gomes, Miguel A. and Georg Weizsäcker, "Stated beliefs and play in normal-form games," *The Review of Economic Studies*, 2008, *75* (3), 729–762.

Costa-Gomes, Miguel and Vincent P Crawford, "Cognition and behavior in two-person guessing games: An experimental study," *The American Economic Review*, 2006, *96* (5), 1737–1768.

Crawford, Vincent P. and Nagore Iriberri, "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?," *Econometrica*, 2007, *75* (6), 1721–1770.

Crawford, Vincent P, Miguel A Costa-Gomes, and Nagore Iriberri, "Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications," *Journal of Economic Literature*, 2013, *51* (1), 5–62.

Croson, Rachel TA, "The disjunction effect and reason-based choice in games," *Organizational Behavior and Human Decision Processes*, 1999, *80* (2), 118–133.

Enke, Benjamin, "What You See Is All There Is," *Working paper*, 2017.

_ **and Florian Zimmermann**, "Correlation neglect in belief formation," *Review of Economic Studies (forthcoming)*, 2018.

Esponda, Ignacio, "Behavioral equilibrium in economies with adverse selection," *The American Economic Review*, 2008, *98* (4), 1269–1291.

_ **and Emanuel Vespa**, "Hypothetical Thinking and Information Extraction in the Laboratory," *American Economic Journal: Microeconomics*, 2014, *6* (4), 180–202.

_ **and** _ , "Contingent Preferences and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory," *Working Paper*, 2018.

_ **and** _ , "Endogenous sample selection: A laboratory study," *Quantitative Economics*, 2018, *9* (1), 183–216.

Eyster, Erik and Georg Weizsäcker, "Correlation neglect in portfolio choice: lab evidence," *Working Paper*, 2016.

_ **and Matthew Rabin**, "Cursed equilibrium," *Econometrica*, 2005, *73* (5), 1623–1672.

_ **,** _ **, and Georg Weizsäcker**, "An Experiment on Social Mislearning," *Working paper*, 2018.

Fantino, Edmund, "Behavior analysis and decision making," *Journal of the experimental analysis of behavior*, 1998, *69* (3), 355–364.

Fragiadakis, Daniel, Daniel Knoepfle, and Muriel Niederle, "Who is Strategic?," *Working Paper*, 2017.

Friedman, Daniel, "Monty Hall's three doors: Construction and deconstruction of a choice anomaly," *The American Economic Review*, 1998, *88* (4), 933–946.

Gabaix, Xavier, "A sparsity-based model of bounded rationality," *The Quarterly Journal of Economics*, 2014, *129* (4), 1661–1710.

Gennaioli, Nicola and Andrei Shleifer, "What comes to mind," *The Quarterly Journal of Economics*, 2010, *125* (4), 1399–1433.

Gigerenzer, Gerd and Ulrich Hoffrage, "How to improve Bayesian reasoning without instruction: frequency formats.," *Psychological review*, 1995, *102* (4), 684.

Glazer, Jacob and Ariel Rubinstein, "An Extensive Game as a Guide for Solving a Normal Game," *Journal of Economic Theory*, 1996, *1* (70), 32–42.

Hamilton, William, *Lectures on metaphysics*, Sheldon, 1878.

Ivanov, Asen, Dan Levin, and Muriel Niederle, "Can relaxation of beliefs rationalize the winner's curse?: an experimental study," *Econometrica*, 2010, *78* (4), 1435–1452.

Jehiel, Philippe, "Analogy-based expectation equilibrium," *Journal of Economic Theory*, 2005, *123* (2), 81–104.

Kagel, John H. and Dan Levin, "The winner's curse and public information in common value auctions," *The American Economic Review*, 1986, pp. 894–920.

Kahneman, Daniel and Amos Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica*, 1979, pp. 263–291.

Kőszegi, Botond and Adam Szeidl, "A model of focusing in economic choice," *The Quarterly Journal of Economics*, 2012, *128* (1), 53–104.

Laibson, David, "Golden eggs and hyperbolic discounting," *The Quarterly Journal of Economics*, 1997, *112* (2), 443–478.

Lennie, Peter, "The cost of cortical computation," *Current Biology*, 2003, *13* (6), 493–497.

Levin, Dan, James Peck, and Asen Ivanov, "Separating Bayesian updating from non-probabilistic reasoning: An experimental investigation," *American Economic Journal: Microeconomics*, 2016, *8* (2), 39–60.

Li, Shengwu, "Obviously strategy-proof mechanisms," *American Economic Review*, 2017, *107* (11), 3257–87.

Louis, Philippos, "The Barrel of Apples Game: Contingent thinking, inference from others' actions, and strategic heterogeneity," *Working Paper*, 2015.

Mobius, Markus M, Muriel Niederle, Paul Niehaus, and Tanya S Rosenblat, "Managing self-confidence: Theory and experimental evidence," *Working Paper*, 2014.

Moser, Johannes, "Hypothetical thinking and the winner's curse: An experimental investigation," *Working Paper*, 2018.

Nagel, Rosemarie, "Unraveling in guessing games: An experimental study," *The American Economic Review*, 1995, *85* (5), 1313–1326.

Ngangoué, Kathleen and Georg Weizsäcker, "Learning from unrealized versus realized prices," *Working Paper*, 2018.

Niederle, Muriel, "Gender," *in* The Handbook of Experimental Economics, Volume 2*, edited by John H Kagel and Alvin E Roth*, 2016.

O'Donoghue, Ted and Matthew Rabin, "Doing it now or later," *American Economic Review*, 1999, pp. 103–124.

Rabin, Matthew, "Risk aversion and expected-utility theory: A calibration theorem," *Econometrica*, 2000, *68* (5), 1281–1292.

_ **and Georg Weizsäcker**, "Narrow bracketing and dominated choices," *American Economic Review*, 2009, *99* (4), 1508–43.

_ **and Joel L Schrag**, "First impressions matter: A model of confirmatory bias," *The Quarterly Journal of Economics*, 1999, *114* (1), 37–82.

Samuelson, W. F. and M. H. Bazerman, "The winner's curse in bilateral negotiations. In V. Smith (ed.), Research in Experimental Economics," 1985, *3*.

Savage, Leonard J, *The foundations of statistics*, Courier Corporation, 1972.

Schotter, Andrew, "Decision making with naïve advice," *American Economic Review*, 2003, pp. 196–201.

Selten, Reinhard, Klaus Abbink, and Ricarda Cox, "Learning direction theory and the winner?s curse," *Experimental Economics*, 2005, *8* (1), 5–20.

Shafir, Eldar, "Uncertainty and the difficulty of thinking through disjunctions," *Cognition*, 1994, *50* (1), 403–430.

_ **and Amos Tversky**, "Thinking through uncertainty: Nonconsequential reasoning and choice," *Cognitive Psychology*, 1992, *24* (4), 449–474.

_ **, Itamar Simonson, and Amos Tversky**, "Reason-based choice," *Cognition*, 1993, *49* (1), 11–36.

Siegel, Sidney and Donald Aaron Goldstein, "Decision-making behavior in a two-choice uncertain outcome situation.," *Journal of Experimental Psychology*, 1959, *57* (1), 37.

Sims, Christopher A, "Implications of rational inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Stahl, Dale O and Paul W Wilson, "On players? models of other players: Theory and experimental evidence," *Games and Economic Behavior*, 1995, *10* (1), 218–254.

Thompson, Suzanne C, Wade Armstrong, and Craig Thomas, "Illusions of control, underestimations, and accuracy: a control heuristic explanation.," *Psychological Bulletin*, 1998, *123* (2), 143.

Tversky, Amos and Daniel Kahneman, "Judgment under Uncertainty: Heuristics and Biases," *Science*, 1974, *185* (4157), 1124–1131.

_ **and Eldar Shafir**, "The disjunction effect in choice under uncertainty," *Psychological Science*, 1992, *3* (5), 305–310.

Vespa, Emanuel and Alistair J. Wilson, "Communication with multiple senders: An experiment," *Quantitative Economics*, 2016, *7* (1), 1–36.

Vesterlund, Lise, "Using experimental methods to understand why and how we give to charity," *Handbook of Experimental Economics*, 2016, *2*, 91–151.

Vulkan, Nir, "An economist's perspective on probability matching," *Journal of economic surveys*, 2000, *14* (1), 101–118.

Appendices

A Acquiring-a-Company: Summary of Instructions

Below we present a summary of the instructions of the Acquiring-a-Company problem, where we highlight commonalities and differences between PROB and DET. In bold we display sentences that are common to both treatments. In regular font are instructions specific to PROB and in italics those specific to DET. For screenshots of the round 1 part 1 problem with $v_L = 20$ and $v_H = 120$ in PROB and DET, see Figure 6 in Online Appendix F and for full instructions see the Instructions Appendix.

Summary of Instructions:

There is a company for sale. You can buy one company or none. The company's value is either A or B. [*There are two companies for sale. You can buy two companies, one or none. The first company's value is A. The second company's value is B.*] **A and B represent two numbers. The numbers change from round to round.**

In each round you will learn the value A and B that the company may have. With equal chance the company's value is A or B. In each round, the interface is programmed to toss a coin and assign value A if heads comes up and value B if tails comes up. You will not know which of the two values is selected. [*In each round, you will learn the value A of the first company and the value B of the second company.*]

You will submit one price. You can submit any price from 0 to 150. This is the price that you are willing to pay for the company [*a company*].

You do not know if the company for sale is of value A or of value B**.** [*You do know that the two companies for sale are the first company of value A and the second company of value B.*]

Transaction for the [*each*] **company**:

- **If the price you submit is higher than or equal to the value of the company, you buy the company.**
- **If you buy, you increase the company's value by** 50%**. This means that if you buy** the company of value A [*the first company of value A*], **the value to you is** $1.5 \times A$**. If you buy** the company of value B [*the second company of value B*]**, the value to you is** $1.5 \times B$**.**
- If you buy the company, your profit is: 2 times (1.5 x value of the company - price)**.** [*For each company you buy, your profit is: 1.5 x value of the company - price.*]

- [*For each company,*] **If the price you submit is lower than the value of the company, you don't buy the company and your profit is zero.**

[*The price you submit is the same for both transactions. You can make money from both transactions.*]

Your payoff for the round is the profit from the transaction. [*Your payoff for the round is the sum of the profit from the two transactions, the transaction of the first company and the transaction of the second company.*]

B Main Findings Using Submitted Prices: Summary

In the main part of the paper we analyze results and provide evidence for the *Power of Certainty* by classifying subjects into types. In this appendix we show that we reach similar conclusions if we use the aggregate distribution of submitted prices in part 1.

Most of our analysis groups prices into five categories: $p < v_L$, $p = v_L$, $p \in (v_L, v_H)$, $p = v_H$ and $p > v_H$. We first provide a more detailed classification in Table 10, which further disaggregates two categories: $p \in (v_L, v_H)$ and $p > v_H$, to evaluate whether some prices are especially common.[56] We do not, however, observe a large mass of prices at any particular price $p \in (v_L, v_H) \cup (v_H, 150]$. The most common category consists of prices that are one or two units above the value of the firm, which is the focus of one of our robustness checks for the classification of types in Online Appendix B.

	PROB		DET	
	All Rounds	Last 10	All Rounds	Last 10
$p < v_L$	1.12	0.85	4.10	3.39
$p = v_L$	42.69	43.03	53.91	55.52
$p \in \{v_L + 1, v_L + 2\}$	5.53	4.90	5.14	4.86
$p \in \{1.5 \times v_L, \frac{v_L+v_H}{2}\}$	1.84	1.01	0.79	0.76
$p \in (v_L, v_H)$ and not yet classified	15.45	14.57	9.29	8.31
$p = v_H$	21.33	23.94	17.16	16.99
$p \in \{v_H + 1, v_H + 2\}$	3.16	3.41	2.10	2.57
$p \in \{v_L + v_H, 1.5 \times v_H\}$	0.48	0.53	1.66	2.09
$p > v_H$ and not yet classified	8.40	7.77	5.85	5.52
Total	100.0	100.0	100.0	100.0

Table 10: Distribution of Prices by Treatment (in %)

Notes: 188 and 183 participants in PROB and DET, respectively. All rounds include data for 3,760 (188 subjects × 20 rounds) and 3,660 (183 subjects × 10 rounds) in PROB and DET, respectively. Last 10 rounds include data for 1,880 (188 subjects × 10 rounds) and 1,830 (183 subjects × 10 rounds) in PROB and DET, respectively.

[56] An even more detailed disaggregation is provided in Table 11 of Online Appendix A.

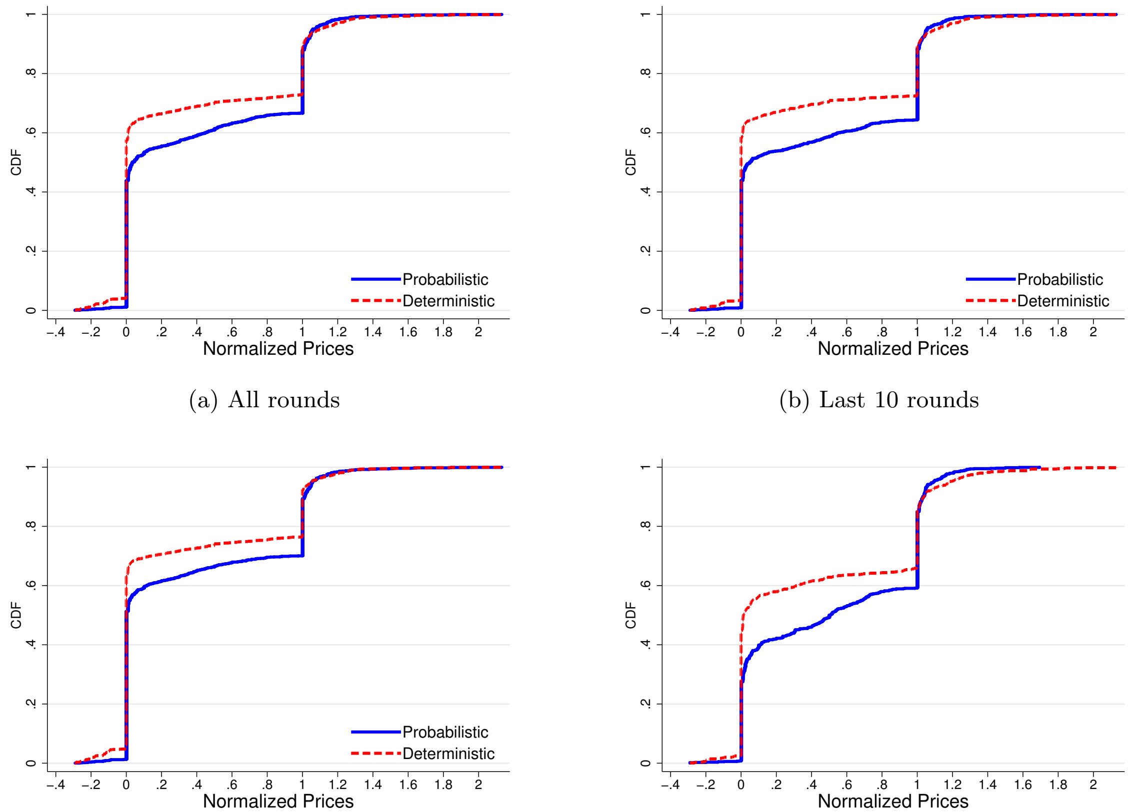

(a) All rounds

(b) Last 10 rounds

(c) Subjects who never took a risk

(d) Subjects who took risk once or more

Figure 2: Main Treatments Part 1-Distribution of Normalized Prices by Treatment

Notes: Normalized Prices (p^N) are computed for each question and subject in the following manner: $p^N = \frac{p-v_L}{v_H-v_L}$, where p is the price submitted by the subject for a question with values $\{v_L, v_H\}$. Notice that if $p = v_H$, then $p^N = 1$ and if $p = v_L$, then $p^N = 0$. Values of $p^N \in (0,1)$ indicate subjects submitting dominated prices in the interval $p \in (v_L, v_H)$. The case $p^N < 0$, corresponds with $p < v_L$ and, finally, the case $p^N > 1$ takes place if $p > v_H$. 'All rounds' uses the answers submitted for all 20 problems and 'Last 10 rounds' for the last 10 problems subjects faced in part 1. 'Subjects who never took a risk' includes subjects who never selected a risky lottery in part 5. 'Subjects who took risk once or more' involves subjects who selected at least one risky lottery in part 5.

Table 12 of Online Appendix A shows that, when considering all 20 rounds, there are substantially more prices $p = v_L$ in DET (53.9 percent) than in PROB (42.7 percent). The numbers are virtually unchanged when we only consider the last 10 rounds. Column (1) of Table 13 of Online Appendix A presents a random-effects estimation that confirms this treatment effect to be significant. Focusing on just the payoff-maximizing price $p = v_L$ we find evidence consistent with the PoC hypothesis.

The differences in submitted prices across treatments can be easily summarized by Figure 2. We show the cumulative distribution of normalized prices. We compute a normalized price $p^N = \frac{p-v_L}{v_H-v_L}$, such that $p^N = 0$ indicates $p = v_L$, and $p^N = 1$ corresponds to $p = v_H$. The figure shows the cumulative distribution of normalized prices in part 1 for all rounds (panel a) and the last 10 rounds (panel b). The distribution of normalized prices in PROB first order stochastically dominates the distribution in the DET and the difference is statistically significant.[57] Panels c and d confirm that the results

[57]We test for first-order stochastic dominance using the test in Barrett and Donald (2003). The test consists of two

hold if we distinguish between subjects who did and did not select at least one risky lottery in part 5.

In Online Appendix A we also show that a treatment effect for part 2 is also present when we use prices instead of the classification of types.

steps. We first test the null hypothesis that the distribution in the deterministic treatment either first order stochastically dominates or is equal to the distribution in the probabilistic treatment. We reject this null hypothesis using all rounds or the last 10 rounds, the corresponding $p-$value is 0 in both cases. We then test the null hypothesis that the distribution in the probabilistic treatment first order stochastically dominates the distribution in the deterministic treatment. We cannot reject the null in this case, with a corresponding $p-$value of 0.817 using all rounds and 0.239 using the last 10 rounds.