Contracting with Independent Contractors: Theory and Experiment

Ala Avoyan* Mauricio Ribeiro† Andrew Schotter‡

Abstract

While many principal-agent problems deal with incentivizing an employee (agent) to work, others exist where the agent is an independent contractor (IC). With ICs, the contract terms must not only incentivize the IC to accept the principal's contract but also incentivize her to allocate more time to it over the other contracts in her portfolio. We study this dual issue in this paper. We generate sharp predictions using standard theory and then we examine them using experimental data. Both theoretical analysis and experimental results indicate that those contract terms that may lead an agent to accept a project may not be the same terms that will motivate her to allocate as much time to it as a principal would like. Hence, our paper suggests that there is a tradeoff between contract choice and time allocation that may need to be considered by principals.

JEL Classification: J22, C91, D86, D91;

Keywords: Contract Choice, Time Allocation, Independent Contractors, Incentive Structures.

*Department of Economics, Indiana University. E-mail: avoyan.ala@gmail.com

†School of Economics, University of Bristol. E-mail: mauricio.ribeiro@bristol.ac.uk

‡Department of Economics, New York University. E-mail: andrew.schotter@nyu.edu

1 Introduction

Consider the following situation. You are a homeowner (Principal) and want to hire a contractor (Agent) to remodel your house. You have one contractor in mind and think she is a skilled worker, but because of that, others also want to hire her. So, you offer her a contract that sets out a compensation scheme for remodeling your house. Since the contractor is not your employee but an independent contractor (IC), the contract must provide her with incentives to devote time to your project instead of the contracts of other homeowners. But what contracts are successful in doing so?

The answer to this problem has important implications for our economy and its gig-economy sector. For example, Manyika et al. (2016) estimate that 20 to 30 percent of the working-age population in the United States and the EU engage in nontraditional work arrangements. While many of these workers might prefer conventional employment inside the firm, a good portion accept contract work as independent contractors, i.e., self-employed individuals who provide certain services to clients for a limited period.

Contracting with an IC presents two problems. The first problem is that the IC might not even accept your contract because someone else might offer her a contract she finds more attractive. This is the *contract-choice problem*. The second problem you might face is that even if the IC accepts your contract, she must allocate her time between your project and the other projects in her portfolio, and she might not allocate enough time to yours.[1] This is the *contract-time allocation problem*. Therefore, a successful contract must overcome these two problems: the contractor must choose it over competing contracts, and once chosen, it must successfully compete for her time with the other contracts she already accepted.

To establish what types of contracts would successfully overcome these two problems, we must first understand how people choose between and allocate their time across contracts to execute a task. In this paper, we design an experiment to study precisely that. In the first part of the experiment, subjects choose one contract from several pairs of contracts. In the second part, they allocate a time budget between the contracts in these pairs. We then discuss how our experimental results can help to design contracts to overcome the contract-choice and contract time-allocation problems.

Although choosing between contracts is a more familiar exercise, allocating time across contracts has been less studied, at least experimentally. When allocating time across contracts,

[1] Successful contractors often work on several projects at once, and the biggest complaint about them is that one cannot get them to pay attention to a job given their other commitments.

subjects self-impose the time they want to spend on each contract's underlying task, which creates an opportunity cost for the time spent on a contract. By allocating more time to a contract, the probability of successfully executing the contract's task increases, but at the expense of reducing the probability of completing the other contract's task. Going back to our contractor, once she has accepted Contracts A and B, if she decides to allocate more of her time to Contract A, the probability of successfully fulfilling Contract A increases, whereas that of Contract B might decrease.

We base our experimental design on a standard model of choice and time-allocation between contracts that makes sharp predictions about behavior in the experiment given the contracts we use. Our contracts pay either a high or a low payoff depending on the Agent's ability to execute a task by a deadline (which the Agent self-imposes when allocating time). The key feature of these contracts is that the probability of getting the high payoff is higher if the Agent succeeds in the task than if she fails.[2] We interpret the contract's low and high payoffs as *hard* incentives, i.e., the Agent's payment, and the probabilities of getting the high payoff conditional on succeeding and failing the task as *soft* incentives, given they influence the Agent's probabilistic assessment of getting the high payoff as a function of time allocated to a contract. Finally, the difficulty of the task being contracted for is proxied by the success rate of others who attempted the same task under a given deadline in the past. Based on this rate, the Agent can form beliefs about how difficult the task is and, hence, how likely she is to succeed in the task by the deadline.

The model makes two main predictions about subjects' behavior in the experiment. The first prediction is that when allocating time between contracts, the *only* contractual feature that matters for the time allocated to its task is the difference in payoffs between succeeding and failing to complete the task. We call this difference the contract's *spread*. The second prediction is that, when choosing between two contracts, the payoff one gets if one fails to complete the contract's task matters independently of its impact on the contract's spread. We call this payoff the contract's *failure payoff*. These two predictions imply that the characteristics of a contract that make it attractive for choice can make it less effective in attracting time. Therefore, the model predicts that our Principal might be unable to design a contract that simultaneously overcomes the contract-choice and contract-time allocation problems — unless she is willing to

[2] Although real-world contracts do not use randomization, the randomization in our contracts is just an experimental way of emulating that a contractor's ability to execute a task by a deadline is always subject to circumstances the contractor cannot control. Therefore, even if the contractor allocates plenty of time to a task, she cannot be sure that she will complete the task by the deadline. But, by allocating more time to a task, she can influence the probability of completing the task by the deadline.

offer a contract that costs more than the ones she is competing with.

In the experiment, we present subjects with pairs of contacts which they first have to choose between and later allocate time across. We designed these contract pairs to achieve three goals. The first goal is to investigate how subjects react to controlled changes in the different parameters of a contract. For instance, how do subjects react to changes on hard as opposed to soft incentives? The second goal is that the model should make sharp predictions about choices and time allocations in the experiment. These predictions then provide a benchmark to interpret subjects' behavior in the experiment. The third goal is that the contract should have the same expected cost to the Principal. Therefore, we focus on contracts that can cost-effectively overcome the contract-choice and contract-time allocation problems.

We designed the pairs of contracts used in the experiment to achieve three goals. The first goal is to investigate how subjects react to controlled changes in the different parameters of a contract. For instance, how do subjects react to changes on hard as opposed to soft incentives? The second goal is that the model should make sharp predictions about choices and time allocations in the experiment. These predictions then provide a benchmark to interpret subjects' behavior in the experiment. The third goal is that one of the contracts in the pair should not cost more (in expectation) than the contract it is competing against. Therefore, we focus on contracts that can *cost-effectively* overcome the contract-choice and contract-time allocation problems.

In risk contracts, the high payoff is higher than in the baseline, but the low payoff is lower. It is a mean-preserving spread of payoffs relative to the baseline. Therefore, the Agent will face a riskier lottery than in the baseline contract, no matter if she solves the task. Although intuition suggests that a risk-seeking Agent should always prefer these contracts, the model predicts that an Agent's choice will also depend on her confidence in her ability to complete the task. If the Agent is sufficiently risk-averse, she should always choose the baseline. If she is sufficiently risk-seeking, she should always choose the risk contract. But for intermediary values of risk-aversion, the model predicts that the Agent should choose a risk contract over the baseline if (and only if) she is sufficiently confident. Interestingly, the model again predicts that the Agent should always allocate more time to risk contracts than the baseline because they have higher spreads

In dominated contracts, the probability of getting the high prize is smaller than in the baseline regardless of whether the subject can solve the task. Therefore, the choice between the baseline and a dominated contract is obvious. Time allocation is, however, trickier. The model predicts that the Agent should allocate *more* time to the baseline than to one dominated contract but *less* time to the baseline than to the other. As we discuss shortly, the predictions about dominated contracts offer a crucial insight into our subjects' time allocation behavior.

Finally, in ambiguous contracts, the contract only provides partial information about the difficulty of its underlying task. Intuition suggests, and the model predicts, that the Agent should choose the baseline contract over an ambiguous contract if (and only if) she is ambiguity averse. We refrain, however, from making predictions about time allocation in ambiguous contracts because they rely on information we cannot elicit from subjects' behavior. In contrast, the other predictions of the model in our experiment only rely on the terms of the contract and some subjects' characteristics that we can elicit in the lab.[3]

We find that, across all subjects, 62.4 percent of the model's predictions are correct. When we split the predictions between choice and time predictions, we get that 66.2 percent of choice predictions and 59.5 of time predictions are correct. While looking at these averages might suggest that the model performs equally well in describing subjects' choices and time allocations, these averages hide a meaningful difference at the individual level. Only 37 of our 120 subjects adhere to the model's choice predictions more frequently than its time predictions. As we discuss in detail later, this is to be expected, given that the model's time predictions are more robust than its choice predictions.

Our results further suggest that a contract's failure payoff matters when subjects choose between contracts, which corroborates one of the main intuitions of the model. However, they seem to matter much more than what the model predicts when a contract's failure payoff is too low. That is, subjects shy away from contracts that compensate them poorly if they do not complete the task, even if these contracts compensate them generously if they do. Our results also suggest that subjects tend to allocate more time to contracts with higher spreads, which corroborates the model's other main intuition. However, they do not react to spreads as much as the model predicts. The reason seems to be that subjects often allocate more time to the contract they choose, even if it has a lower spread than the other contract in the pair. We call this tendency the "attractiveness bias," because allocating more time to the contract one deems more "attractive" in choice can lead to a sub-optimal time allocation. Subjects' choices and time allocations between the baseline contract and one of the dominated contracts provide the clearest evidence for this type of behavior. Consistent with the model, over 95 percent of subjects choose the baseline over the dominated contract but only 11 percent of subjects allocate more time to the dominated contract as the model predicts.

These results have important implications for effectively designing contracts to overcome the contract-choice and contract-time allocation problems. First, the Principal should avoid of-

[3] We elicit these characteristics in Part III of the experiment.

fering contracts with a failure payoff that is too low. Although decreasing a contract's failure payoff will increase its spread, which our results suggest can help with the contract-time allocation problem, the more we decrease failure payoffs, the worse the contract will perform in the contract-choice problem. But then, by the attractiveness bias, these contracts might not perform as well in the contract-time allocation problem as the model predicts. Therefore, our experimental results suggest that an important tension the Principal faces when designing a contract that is not too costly and can overcome the contract-choice and contract-time allocation problems is to decide how much to lower the contract's failure payoffs to increase its spread, which helps with the contract-time allocation problem, while making sure that the failure payoff is not too low, since this will impair the contract's performance in the contract-choice problem and, by the attractiveness bias, reduce — or even over-rule — its efficacy in the contract-time allocation problem.

2 Literature Review

The paper relates to several strands of literature within economics, such as non-exclusive contracting in contract theory, incentivizing worker's productivity in principal-agent models, and attention/time-allocation.

Time-allocation The paper naturally relates to the literature on time-allocation and multi-tasking. Burmeister-Lamp et al. (2012) examine agents who allocate their time between a wage job (their "day job") and a new enterprise (so called hybrid entrepreneurs). This case relates to the comparison in our paper of a given "safe" contract and a "risky" contract which results in a higher outcome if successfully completed but lower outcome than in "safe" if unsuccessful. Burmeister-Lamp et al. (2012) find that entrepreneurs' risk attitudes do not reliably predict their actual time allocation behavior, which contradicts their theoretical findings. In contrast,in our paper when agents make time-allocation decisions between contracts, we show that risk attitudes should have no influence on time-allocation decisions.[4]

Multi-tasking theory Dewatripont et al. (2000) presents an excellent discussion of multitask agency theory papers and their insights. In scenarios involving multiple principals, the phenomenon of effort substitution prompts principals to vie for the agent's attention. Consequently,

[4] Avoyan and Schotter (2020) investigate how players allocate their limited attention across games and find that attention is attracted to particular features of the games. Avoyan et al. (2024) study the discrepancy between planned and actual attention allocation, demonstrating that individuals often fail to accurately predict their own attention distribution across tasks. This finding is crucial for understanding the dynamics of time allocation between tasks in this paper, as it highlights the potential for misalignment between anticipated and actual behavior in decision tasks as seen in strategic settings.

incentive structures tend to be overly severe leading to exclusivity—where the agent is compelled to engage exclusively with a single principal (see, for example, Martimort (1996), Dixit (1998) and Bernheim and Whinston (1998)). While in our paper we have "effort" substitution, since time can either be spent on one or the other contract, we are able generate non-exclusive outcomes due to agent's overconfidence in their skills compared to a representative agent.

Bar-Isaac and Deb (2014) study a different angle by focusing on agent's reputation concerns when facing common or separate audiences who form beliefs based on joint or separate outcomes. Halac and Prat (2016) theoretically examine the issue of worker performance in the context of potential monitoring by a manager, where manager decides their level of costly attention.

Non-exclusive contract theory Several papers examine non-exclusive contracts in insurance markets and optimal contracting in non-exclusive relationships. Ales and Maziero (2016) study the effect of the presence of non-exclusive contracts on a decentralized market. The paper shows that competition and non-exclusivity of insurance contracts significantly reduce the amount of insurance provided. Attar et al. (2011) consider a "market for lemons" under non-exclusivity. Bisin and Guaitoli (2004) study a static moral hazard and Attar et al. (2014) study adverse selection environments under non-exclusivity assumption on contracts. In contrast, in the current paper we bring to light that incentivizing choice (acceptance of contract) is not the same problem as incentivizing time allocation (attention to a contract within a portfolio), which is of particular importance in environments with independent contractors.

Experiments on principal-agent models Experimental papers have also explored principal-agent relationships, but focusing on exclusive contracts. For example, Hoppe and Schmitz (2013) examine a setting where agents have private information. A key observation from the setting is that when a principal extends a contract to an agent without knowledge of the agent's type, there is a trade-off between optimizing ex post efficiency in unfavorable circumstances and maximizing profits in favorable circumstances. Cabrales and Charness (2011) study optimal contracts in a scenario involving hidden information and team production. A principal presents a team of two agents, whose skill levels are unknown, with one of three potential contract options. The paper finds that frequent refusals of the more imbalanced contract options prompt principals to adjust by offering more advantageous contract options. Corgnet et al. (2018) investigate wage-irrelevant goal-setting in principal-agent model and find that agents' performance is better in the presence of goal setting. In contrast to these papers, we are studying the terms of a simple contract to draw agent's attention away from other projects towards the principal's project.

3 Model and Predictions

In this section, we first present the model we use to predict how people choose between and allocate time across two contracts. We then highlight the key factors that the model predicts should drive choice and time allocation. We then introduce the types of contracts we study and state the model's specific predictions about choice and time allocation in the pairs of contracts we consider.

3.1 Environment

Our contracts specify how an agent will be paid for solving a task by a given *deadline* t. Therefore, our contracts are compensation schemes.[5] If the agent succeeds in completing the task by the deadline, she receives a high prize H with probability p_S or a low prize L with the remaining probability, where $L < H$. If she fails in the task, she receives H with probability p_F and L with the remaining probability, where $p_F < p_S$.

Utility of a Contract To evaluate a contract, the agent must form beliefs about the difficulty of the task. To anchor these beliefs, our contracts inform the agent of the fraction $\alpha \in [0, 1]$ of people that could solve the task by the given deadline. We call such α the *completion rate*, and it proxies the difficulty of a task. Therefore, a contract $\mathcal{L}$ is described by the vector $(H, L, \alpha, p_S, p_F, t)$. We we denote the set of these contracts by $\mathbb{L}$. Figure 1 graphically represents a contract $\mathcal{L} \in \mathbb{L}$.

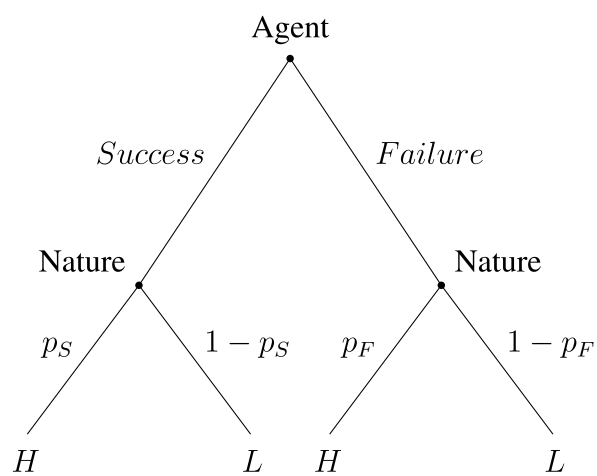

Figure 1: Graphical Representation of a Contract

[5] In our experiment, the task is solving a maze in a given amount of time.

We assume that the agent is a (subjective) expected utility maximizer, which upon learning the difficulty α of the task and the amount of time she has to solve it updates her beliefs about her probability of solving the task and uses this updated belief to calculate the expected utility of the contract.[6] Formally, the agent evaluates a contract $\mathcal{L} \in \mathbb{L}$ by a utility function $U : \mathbb{L} \to \mathbb{R}$ defined as

$$U(\mathcal{L}) \coloneqq \boldsymbol{p}(\alpha, t) \left(p_S u(H) + (1 - p_S) u(L)\right) + \left(1 - \boldsymbol{p}(\alpha, t)\right) \left(p_F u(H) + (1 - p_F) u(L)\right), \quad (1)$$

where $u : \mathbb{R} \to \mathbb{R}$ is a strictly increasing (Bernoulli) utility function and $\boldsymbol{p}(\alpha, t)$ is the agent's belief that she will solve the task in, at most, t units of time given that a fraction α of people have done so.

Choice Between Contracts When choosing between two contracts in a pair, say $\mathcal{L}_1$ and $\mathcal{L}_2$, the agent maximizes U. Hence, for $i \in \{1, 2\}$, she chooses $\mathcal{L}_i$ from the menu $\{\mathcal{L}_1, \mathcal{L}_2\}$ if, and only if, $U(\mathcal{L}_i) \geqslant \max\{U(\mathcal{L}_1), U(\mathcal{L}_2)\}$.

Time Allocation Across Contracts When allocating time across two contracts, the agent allocates a budget of T units of time between the two contracts' underlying tasks so as to maximize the sum of the contracts' expected utilities.

To state the agent's optimization problem, let $\mathcal{L}_{t^\star}$ be the contract with the same characteristics as $\mathcal{L}$, except for the deadline, which is now given by $t^\star$.[7] For each $t \geqslant 0$, define the time-conditional utility function $U(\cdot|t) : \mathbb{L} \to \mathbb{R}$ be $U(\mathcal{L}|t) \coloneqq U(\mathcal{L}_t)$.

Given two contracts $\mathcal{L}_1$ and $\mathcal{L}_2$, suppose that the agent decides to allocate t minutes to $\mathcal{L}_1$, and $T - t$ to $\mathcal{L}_2$. She will then engage in the contracts $(\mathcal{L}_1)_t$ and $(\mathcal{L}_2)_{T-t}$, which she evaluates by $U(\mathcal{L}_1|t)$ and $U(\mathcal{L}_2|T - t)$. We assume that the agent allocates time by solving

$$\max_{t \in [0,T]} \left[U(\mathcal{L}_1|t) + U(\mathcal{L}_2|T - t)\right]. \quad (2)$$

We further assume that, for all $\alpha \in [0, 1]$,

1. $\boldsymbol{p}(\alpha, 0) = 0$;

2. $\boldsymbol{p}(\alpha, \cdot)$ is increasing, concave and continuously differentiable.[8]

[6] For each α, this evaluation formula for contracts is a particular case of the preferences of an agent in the canonical moral hazard problem when contracts only have two outputs, time is replaced by effort, and the cost of effort is zero. See Sinander (2024).

[7] The completion rate α still gives the fraction of people that successfully completed the task when the deadline was t.

[8] If $\boldsymbol{p}$ were linear or convex, the solution of the optimization problem (2) would be to allocate T to the same

The concavity of $\boldsymbol{p}(\alpha, \cdot)$ means that there are decreasing returns to time allocated to a task. That is, as t increases, the cumulative probability of solving the task in at most t seconds increases at decreasing rates.

3.2 Spreads, choice and time allocation

To derive predictions about choice and time allocation, define, for every $\mathcal{L} \in \mathbb{L}$,

$$\begin{aligned} U_S^{\mathcal{L}} &\coloneqq p_S u(H) + (1 - p_S) u(L) \\ U_F^{\mathcal{L}} &\coloneqq p_F u(H) + (1 - p_F) u(L) \\ \Delta_{\mathcal{L}} &\coloneqq U_S^{\mathcal{L}} - U_F^{\mathcal{L}} = (p_S - p_F)(u(H) - u(L)) \end{aligned} \quad .$$

$U_S^{\mathcal{L}}$ is the expected value of the lottery the agent gets if she succeeds in solving the task, $U_F^{\mathcal{L}}$ is the expected value of the lottery the agent gets if she fails in solving the task. We sometimes refer to $U_F^{\mathcal{L}}$ as the agent's *failure payoff*. Finally, we say that $\Delta_{\mathcal{L}}$ is the *spread* of $\mathcal{L}$. Given the agent's utility over contracts, $\Delta_{\mathcal{L}}$ is the product of $(p_S - p_F)$, the *spread in probabilities* of $\mathcal{L}$, and $(u(H) - u(L))$, the *spread in payoffs* of $\mathcal{L}$.

As we now show, although the spreads are important both for choice and time allocation, the model predicts that the difference in failure payoffs matter for choice independently of its impact in spreads whereas it matters for time allocation only through its effect on the spread. This can create a wedge between choice and time allocation predictions in the sense that the model can predict that the agent should choose one contract over the other, but allocate more time to the contract that is not chosen.

Choice Prediction Recall that given two contracts $\mathcal{L}_1$ and $\mathcal{L}_2$, $\mathcal{L}_1$ should be chosen from $\{\mathcal{L}_1, \mathcal{L}_2\}$ if, and only if, $U(\mathcal{L}_i) \geqslant \max\{U(\mathcal{L}_1), U(\mathcal{L}_2)\}$. Substituting the evaluation formula (1) in this expression and manipulating, we get that $\mathcal{L}_1$ is chosen from $\{\mathcal{L}_1, \mathcal{L}_2\}$ if, and only if,

$$\boldsymbol{p}(\alpha_1, t_1)\Delta_{\mathcal{L}_1} - \boldsymbol{p}(\alpha_2, t_2)\Delta_{\mathcal{L}_2} \geqslant U_F^{\mathcal{L}_2} - U_F^{\mathcal{L}_1}.$$

Assuming, as is the case for most choices in our experiment, that $\alpha_1 = \alpha_2 = \alpha$ and that $t_1 = t_2 = t$, $\mathcal{L}_1$ is chosen from $\{\mathcal{L}_1, \mathcal{L}_2\}$ if, and only if,

$$\boldsymbol{p}(\alpha, t)(\Delta_{\mathcal{L}_1} - \Delta_{\mathcal{L}_2}) + \left(U_F^{\mathcal{L}_1} - U_F^{\mathcal{L}_2}\right) \geqslant 0 \tag{3}$$

The inequality (3) implies that choice depends on the agent's beliefs, and in two contractual

contract that would receive more time if $\boldsymbol{p}$ was concave.

features: the difference in the contracts' spreads and in their failure payoffs. Moreover, the difference in failure payoffs matter even after we control for the difference in spreads.

Observation 1 *In two contracts with equally difficult tasks, the difference in failure payoffs can influence choice even after we control for the difference in spreads.*

Time Prediction Solving 2 implies that the time $t_{\mathcal{L}_1}$ allocated to contract $\mathcal{L}_1$ must satisfy:

$$\frac{\boldsymbol{p}'(\alpha_2, T - t_{\mathcal{L}_1})}{\boldsymbol{p}'(\alpha_1, t_{\mathcal{L}_1})} = \frac{\Delta_{\mathcal{L}_1}}{\Delta_{\mathcal{L}_2}},$$

where $\boldsymbol{p}'(\alpha_i, \cdot)$ is the derivative of $\boldsymbol{p}$ with respect to time. Given our assumption that $\boldsymbol{p}$ is concave, the mapping $t \mapsto \frac{p'(\alpha_2, T-t)}{p'(\alpha_1, t)}$ is non-decreasing. Therefore, $t_{\mathcal{L}_1}$ is a non-decreasing function of the ratio of spreads of $\mathcal{L}_1$ and $\mathcal{L}_2$, i.e., the higher $\frac{\Delta_{\mathcal{L}_1}}{\Delta_{\mathcal{L}_2}}$ is, the higher $t_{\mathcal{L}_1}$ will be. In particular, if both tasks are equally difficult (that is, if $\alpha_1 = \alpha_2 = \alpha$), then the contract with higher ratio of spreads will receive more than $\frac{T}{2}$, hence more time than the other contract.

Observation 2 *In two contracts with equally difficult tasks, the one with higher spread will attract more time, and the amount of time it attracts is increasing on its spread. Therefore, failure payoffs should only matter for time allocation through their effect on a contract's spread.*

3.3 Contract Types

We now describe the types of pairs of contracts we study in the experiment and state the model's choice and time allocation predictions in these types. For an explicit derivation and detailed discussion of these predictions, see Appendix A.

The pairs of contracts in the experiment consist of a *baseline* contract, such as the one in Table 1, paired with eight other contracts derived from the baseline by making controlled changes to its characteristics. We consider four different types of controlled changes, which induce four different types of contracts: confidence (C), risk (R), dominated (D), and ambiguous (A).

Table 1: Characteristics of $\mathcal{L}_B$

Characteristics	Value
High Payoff	H_B
Low Payoff	L_B
Completion rate	α_B
Probability High Payoff — Solve	p_S^B
Probability High Payoff — Not Solve	p_F^B

3.3.1 Confidence Contracts

A *confidence* contract $\mathcal{L}_C$ has the same payoffs, completion rate and deadline as $\mathcal{L}_B$, but the probability of getting the high payoff if the agent succeeds in the task is greater than the corresponding probability in $\mathcal{L}_B$ whereas the probability of winning the high payoff if the agent fails in the task is lower than the corresponding probability in $\mathcal{L}_B$.

Formally, $\mathcal{L}_C$ is a confidence contract, if $H_C = H_B$, $L_C = L_B$, $\alpha_C = \alpha_B$ and

$$p_F^C < p_F^B < p_S^B < p_S^C \text{ and } p_F^B + p_S^B = p_S^C + p_F^C.$$

Intuitively, a confident agent, that is, one that believes that she is very likely to succeed in the task, will choose this contract over the baseline contract because the probability of getting H if she succeeds in the task is higher than in the baseline. But given the increase in the spreads of probabilities, the model predicts that confidence contracts should always attract more time.

3.3.2 Risk Contracts

A *risk* contract $\mathcal{L}_R$ has the same probabilities, completion rate and deadline as the baseline contract $\mathcal{L}_B$, but its high payoff is greater than the high payoff of $\mathcal{L}_B$ whereas its low payoff is smaller than the low payoff of $\mathcal{L}_B$. Moreover, the average of of the high and low prizes for $\mathcal{L}_B$ and $\mathcal{L}_R$ are the same. Formally, $\mathcal{L}_R$ is a risk contract if $p_S^R = p_S^B$, $p_F^R = p_F^B$, $\alpha_R = \alpha_B$, and

$$L_R < L_B < H_B < H_R \text{ and } L_R + H_R = L_B + H_B.$$

Intuitively, risk contracts are appealing to risk-seeking agents. However, as we show in Section 3.4.2), how confident the agent is (i.e., the value of $\boldsymbol{p}(\alpha, t)$) also influences the choice between a risk contract and the baseline. But given the increase in the spreads of probabilities, the model predicts that risk contracts should always attract more time.

3.3.3 Dominated Contracts

A *dominated* contract $\mathcal{L}_D$ has the same payoffs, completion rate and deadline as $\mathcal{L}_B$, but both the probabilities of winning the high prize are lower than in $\mathcal{L}_B$. Formally, $H_D = H_B$, $L_D = L_B$, $\alpha_D = \alpha_B$, but

$$p_S^D < p_S^B \text{ and } p_F^D < p_F^B.$$

Clearly, an agent should always choose the baseline contract over a dominated contract. Time allocation, however, depends on the ratio of spreads. Therefore, dominated contracts allow us to test whether the agent understands the relevance of spreads when allocating time. In particular, do subjects understand that an attractive contract in choice should, in some cases, receive less time than a more attractive one?

3.3.4 Ambiguous Contracts

An ambiguous contract has the same payoffs, probabilities of getting the high prize, and deadline as $\mathcal{L}_B$, but the completion rate is ambiguous in the sense that the agent does not know the exact value of the completion rate. Formally, $H_A = H_B$, $L_A = L_B$, $p_S^A = p_S^B$, $p_F^A = p_F^B$, but, for some $\varepsilon > 0$,

$$\alpha_A \in [\alpha_B - \varepsilon, \alpha_B + \varepsilon].$$

Intuitively, choice between an ambiguous contract and the baseline depends on the agent's ambiguity attitude. That is, ambiguity averse agents should choose the baseline over an ambiguous contract, whereas ambiguity seeking agents should to the opposite.

3.4 Predictions

We now state the model's predictions for choice and time allocation between the types of pairs we use in the experiment.

3.4.1 Confidence Contracts

Prediction 1 (Choice C) *An agent should choose the confidence contract over the baseline one if, and only if, $\boldsymbol{p}(\alpha, t) \geqslant \alpha$, i.e., if she believes that her probability of solving the task is higher than the probability of a randomly selected person in the population solving the task.*[9]

Prediction 2 (Time C) *An agent should allocate strictly more time to a confidence contract than to the baseline one.*

Prediction 3 (Time C$_{\boldsymbol{R}}$) *The amount of time allocated to a confidence contract increases as $p_S^C - p_F^C$ increases.*

3.4.2 Risk Contracts

Prediction 4 (Choice R)

1. *If an agent is sufficiently risk-loving, she will choose the risk contract.*

[9] We interpret the condition $\boldsymbol{p}(\alpha, t) \geqslant \alpha$ as saying that the agent is over-confident at completion rate α, where we use over-confidence in the sense of over-placement.

2. *If an agent is sufficiently risk-averse, she will choose the baseline one.*

3. *For intermediate levels of risk aversion, the choice is determined by the agent's confidence attitudes. Specifically, only sufficiently confident subjects will choose the risky contract over the baseline one.*[10]

Prediction 5 (Time R) *An agent should allocate strictly more time to a risk contract than to the baseline one.*

Prediction 6 (Time $\mathbf{R}_{\boldsymbol{R}}$) *The amount of time allocated to the risk contract increases as $H_R - L_R$ increases.*

3.4.3 Dominated Contracts

Prediction 7 (Choice D) *The agent should always choose the baseline contract over a dominated one.*

Prediction 8 (Time D) *When $p_S^D - p_F^D > p_S^B - p_F^B$, more time should be allocated to $\mathcal{L}_D$; and when $p_S^D - p_F^D < p_S^B - p_F^B$, more time should be allocated to $\mathcal{L}_B$.*

Prediction 9 (Time $\mathbf{D}_{\boldsymbol{R}}$) *The amount of time allocated to a dominated contract increases as $p_S^D - p_F^D$ increases.*

3.4.4 Ambiguous Contracts

To make predictions about ambiguous contracts, we need to make assumptions about how the agent resolves the uncertainty about α_A. We assume that the agent reduces ambiguity to risk by taking the completion rate of an ambiguous contract as the convex combination of the end-points of the interval specified in the ambiguous contract, where the weights she assigns to the endpoints measure her ambiguity attitude.[11]

Prediction 10 (Choice A) *Assuming that $\boldsymbol{p}(\cdot, t)$ is increasing given the deadline t, the agent should choose the baseline contract over the ambiguous one if, and only if, she is ambiguity averse.*

Since time predictions about ambiguous contracts require we make assumptions about the cross-derivative of the belief function with respect to time and the completion rate, we refrain from making time predictions for ambiguous contracts ,and let the experimental results inform us about subjects' time allocation when facing such contracts.

[10] See Appendix A for a derivation of the model's formal prediction.

[11] See Appendix A for the formal model the agent uses to resolve the uncertainty about the completion rate.

4 Experimental Design

The experiment consisted of three parts. In Part I, subjects chose a contract from pairs of contracts, where each contract in a pair had a 60-seconds deadline. In Part II, subjects allocated 120 seconds across the contracts in the same pairs. In Part III, we elicited the subjects' characteristics required to test the model. In the end of the experiment, subjects had to solve three mazes to get to their payoffs. We discuss how we selected the three mazes when discussing subject's Payoffs below.

The task associated with each contract is to solve a maze with completion rate α.[12] An important aspect of the design is that subjects only saw the mazes associated with the contracts after Parts I, II, and III. Therefore, the contract's completion rate was the only information they had about the difficulty of a maze when choosing and allocating time.

Table 2 displays the contracts we used in the experiment. $\mathcal{L}_B$ acted as the baseline contract. We then generated two contracts for each type of contract described in Section 3.3: two confidence contracts ($\mathcal{L}_{C_1}$ and $\mathcal{L}_{C_2}$), two risk contracts ($\mathcal{L}_{R_1}$ and $\mathcal{L}_{R_2}$), two dominated contracts ($\mathcal{L}_{D_1}$ and $\mathcal{L}_{D_2}$), and two ambiguous contracts ($\mathcal{L}_{A_1}$ and $\mathcal{L}_{A_2}$). Therefore, subjects had to choose and allocate time in 8 pairs of contracts composed of $\mathcal{L}_B$, and each of these eight derived contracts.

A key feature of the contracts in Table 2 is that their expected cost for a risk and ambiguity neutral principal is, at most, that of the baseline contract. Without this restriction, the principal could easily compete against the baseline by offering a contract that has higher prizes, or higher probabilities of getting the high prize, than the baseline.

Contract	$\mathcal{L}_B$	$\mathcal{L}_{C_1}$	$\mathcal{L}_{C_2}$	$\mathcal{L}_{R_1}$	$\mathcal{L}_{R_2}$	$\mathcal{L}_{D_1}$	$\mathcal{L}_{D_2}$	$\mathcal{L}_{A_1}$	$\mathcal{L}_{A_2}$
High Prize	\$8	\$8	\$8	\$10	\$12	\$8	\$8	\$8	\$8
Low Prize	\$4	\$4	\$4	\$2	\$0	\$4	\$4	\$4	\$4
Completion Rate	.5	.5	.5	.5	.5	.5	.5	[.4, .6]	[0,1]
High Prize | Complete	.6	.8	1	.6	.6	.1	.3	.6	.6
High Prize | Not	.4	.2	0	.4	.4	0	0	.4	.4
Expected Value	6	6	6	6	6	4.4	4.6	—	—

Table 2: List of Contracts in the Experiment

[12] More precisely, the completion rate was the fraction of a pool of undergraduate students that could solve the maze in, at most, 60 seconds.See Appendix B.1 for details on the procedure to generate and calibrate the mazes.

4.1 Part I: Choosing Between Contracts

In Part I, we informed subjects that each contract's deadline was 60 seconds. Therefore, if they choose a contract from a pair and the pair is selected for payment, they would have 60 seconds to solve $\mathcal{L}$'s associated maze. For each pair of contracts, subjects then had to either choose one contract in the pair or declare they were indifferent between them. Therefore, subjects were effectively choosing between (i) the baseline contract, (ii) the alternative contract, or (iii) declaring they were indifferent between these two options.

Figure 2 displays an example of the sample screen in Part I. For each pair of contracts, the subjects stated their choices by clicking the button which had the label of the contract they preferred, or by clicking the 'Either V or W' button if they were indifferent. We randomized the order in which the 8 pairs of contracts were presented to a subject and the order that the contracts in a pair were displayed in the screen (i.e., either on the left or on the right).

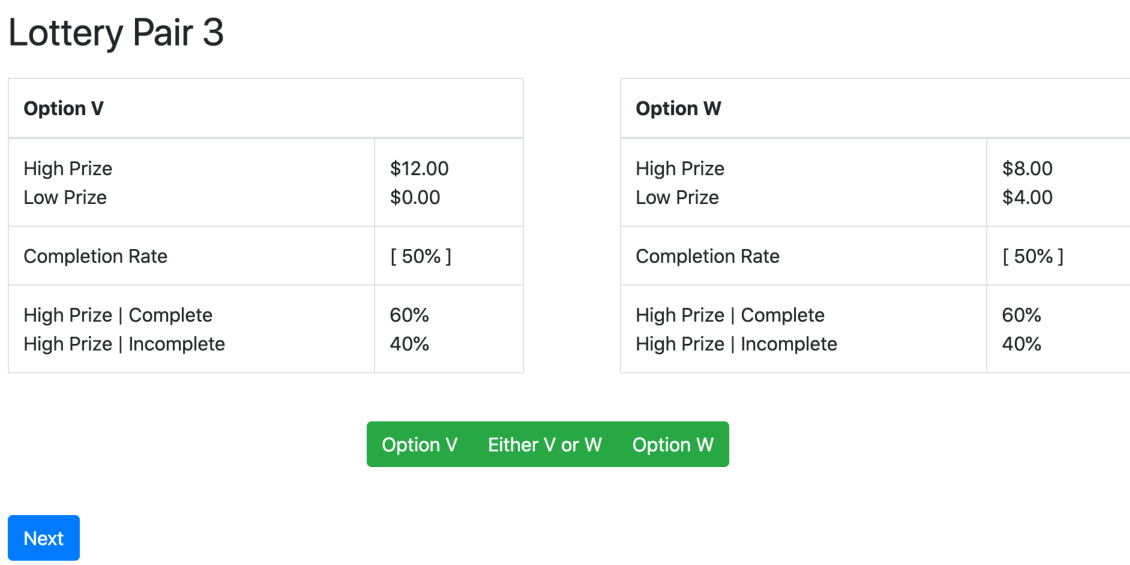

Figure 2: Sample Screen from Part I

The subjects did not receive any feedback between choices, nor did they see the mazes associated with each contract. They were told that at the end of the experiment, i.e., after completing Parts I, II, and III, one of the pairs in Part I would be randomly selected for payment, and they would then get to solve the maze associated with the contract they chose from the pair in at most 60 seconds. If a subject declared indifference in the selected pair, that is clicked the button 'Either V or W,' we randomly selected a contract, and, hence, the maze that the subject had to solve.

4.2 Part II: Allocating Time Between Contracts

After completing Part I, subjects proceeded to Part II. In Part II, we presented subjects with the same 8 pairs of contracts and asked them to state how they would allocate 120 seconds between the contracts (the mazes associated to the contracts). Figure 2 displays a sample screen presented to subjects in Part II.

Subjects had to state how many seconds they wanted to allocate to each option by inputting an integer in the corresponding field. If a subject stated that she wanted to allocate 40 seconds to Option V, then the software automatically assigned 80 seconds to Option W. In this way, we ensured that subjects exhausted their time time budget.

Similarly to Part I, we randomized the order in which the 8 pairs of contracts were presented to a subject and the order that the contracts in a pair were displayed in the screen (i.e., either on the left or on the right).

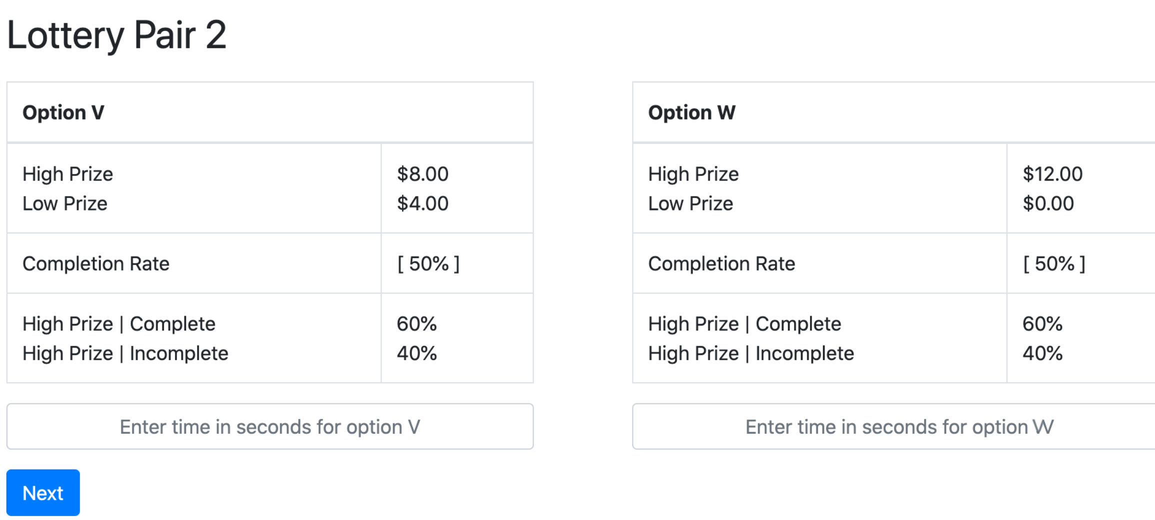

Figure 3: Sample Screen from Part 2

Subjects did not receive any feedback between time allocation decisions, nor did they see the mazes associated with each contract. They were told that at the end of the experiment, i.e., after completing Parts I, II, and III, one of the pairs in Part II would be randomly selected for payment. They would then get to solve the two mazes associated with the selected pair of contracts under the time allocation they stated for that pair.

4.3 Part III: Eliciting Subjects' Characteristics

In Part III, subjects faced 8 different elicitation tasks. The tasks elicited the characteristics we need to test the model's predictions, namely confidence, risk, and ambiguity attitudes. Importantly, we elicited these characteristics using elicitation tasks that we can show are incentive compatible conditional on people behaving consistently with the model. We also elicited other characteristics, e.g., attitude towards the reduction of compound lotteries and alternative measures of confidence and risk attitudes. See Appendix B.2 for further details.

4.4 Payoffs

A subject's payoff in the experiment was the sum of her payoffs in three tasks, each randomly drawn from one of the Parts of the experiment. In the randomly drawn task from Part I, the subject had 60 seconds to solve the maze associated with the contract she chose. The subjects' payoff in this task was the outcome of playing the lottery the contract specifies given their success or failure in solving the maze.

In the randomly drawn task from Part II, the subject had to solve the mazes associated with both contracts in the pair under the time allocation she stated. The subjects' payoff in this task was the sum of the outcomes of the two lotteries the contracts in the pair specify given their success or failure in solving the mazes.

In the randomly drawn task from Part III, the subject's payoff depended on the randomly chosen elicitation task. See Appendix B.2 for an explanation of the payoffs a subject would receive from each elicitation task.

4.5 Implementation

The experiment was conducted at the Center for Experimental Social Science (CESS) laboratory at New York University, using oTree (Chen et al. (2016)) during February and March of 2020. We conducted five sessions, with a total of 120 participants recruited from the general population of NYU students using *hroot* (Bock et al. (2014)). The experiment lasted approximately 60 minutes, and average earnings, including $10 show-up fee, were $31, and ranged from $18 to $46.

5 Results

We now proceed to the results of the experiment. We first examine how successful the model is in predicting subjects' behavior. We are particularly interested in why the model fails and what can we learn from its failures. We then present what changes to the characteristics of the baseline contract lead to a contract that can effectively compete against the baseline, both when subjects choose between and allocate time across contracts.

5.1 Testing the Model: What Can we Learn From its Successes and Failures?

To assess the model's performance in the experiment, we introduce the notion of a prediction matrix.

Definition 1 *Given $N \in \mathbb{N}$ subjects and $M \in \mathbb{N}$ predictions, a **prediction matrix** $(P_{ij})_{i=1,\ldots,N}^{j=1,\ldots,M}$ is a $N \times M$ matrix where*

$$P_{ij} \coloneqq \begin{cases} 1, & j \text{ is true of } i \\ 0, & j \text{ is false of } i \\ NA, & j \text{ is not applicable to } i \end{cases} .$$

A prediction matrix summarizes the model's performance for each subject and each prediction. For each subject i and each prediction j, P_{ij} can take one of three values. If prediction j is correct for subject i, $P_{ij} = 1$. If prediction j is incorrect for subject i, $P_{ij} = 0$. In either of these cases, we say that prediction j is *valid* for subject i. If, however, prediction j cannot be tested for subject i, either because an assumption that is needed to test it fails or because of missing data, $P_{ij} = NA$, we say it is *invalid* (for subject i in prediction j). For the assumptions that each prediction requires to be declared valid, see Appendix A , specifically Table **??**.

Given a prediction matrix, we can calculate how well the model describes the behavior of each subject and its degree of success for each prediction, which leads to the following definitions.

Definition 2 *The **model's degree of success for subject** $i^\star \in \{1, \ldots, N\}$ is defined as*

$$C_{i^\star} \coloneqq \frac{\sum_{\{j: P_{i^\star j} \neq NA\}} P_{i^\star j}}{|\{j : P_{i^\star j} \neq NA\}|}$$

Definition 3 *The **model's degree of success in prediction** $j^\star \in \{1, \ldots, M\}$ is defined as*

$$S_{j^\star} \coloneqq \frac{\sum_{\{i: P_{ij^\star} \neq NA\}} P_{ij^\star}}{|\{j^\star : P_{ij^\star} \neq NA\}|}$$

The model's degree of success for subject i measures how well the model describes subject

i's behavior in the experiment when we restrict attention to valid predictions. The model's degree of success in prediction j is the percentage of subjects that conform to prediction j when we restrict attention to valid predictions.

Table 3 summarizes how well the model accounts for the behavior of subjects in the experiment. On average, we find that 62.4% of all valid predictions are correct. Once we split the predictions between choice and time allocation predictions, we find that 66.2% of valid choice predictions are correct whereas 59.5% of time predictions are correct. However, choice predictions about dominated contracts follow from a simple dominance argument, because the baseline contract awards lotteries that first order stochastically dominate the lotteries awarded by dominated contracts no matter whether a subject succeeds or fails in the task. Therefore, any reasonable model of choice between contracts should predict that the baseline is chosen over a dominated contract, which implies that getting these predictions right only provides weak evidence for the model's predictive success. Once we exclude the choice predictions about dominated contracts, we see that 52.4% of choice predictions are correct. This suggests that the model fares better when predicting time allocations than choice, and, in fact, 83 out of 120 subjects are more consistent with the model in their time allocations than in their choices.

Table 3: Summary of Valid Predictions

Prediction Category	All	Choice	Time	Choice ND
Degree of success	62.4	66.2	59.5	52.4

There is considerable heterogeneity in the degree of success of the model across subjects , as shown in Figure 4. In this figure, for each individual, we present their degree of conformity to choice, dots, and to time, triangles. What we do is rank each subject according to their degree of conformity to choice prediction from lowest to highest and then for each subject present their degree of conformity to our theory's time prediction. This figure, then allows us to see, on an individual level, how subjects differ in their degree of conformity to the time and choice prediction of our theory.

While Figure 4 presents the model's degree of success for each subject, Figure 5 presents the model's degree of success in each prediction. We can then check whether certain predictions are more successful than others, which can help us identify when, and why, the model fails.

First, the predictions about dominated contracts are the most successful ones. This shows that our subjects minimally understand the contracts and react to the incentives they provide. Once we eliminate dominated contracts, we again see that the model fares better in predicting time allocations than in predicting choices. Among the time predictions, the predictions

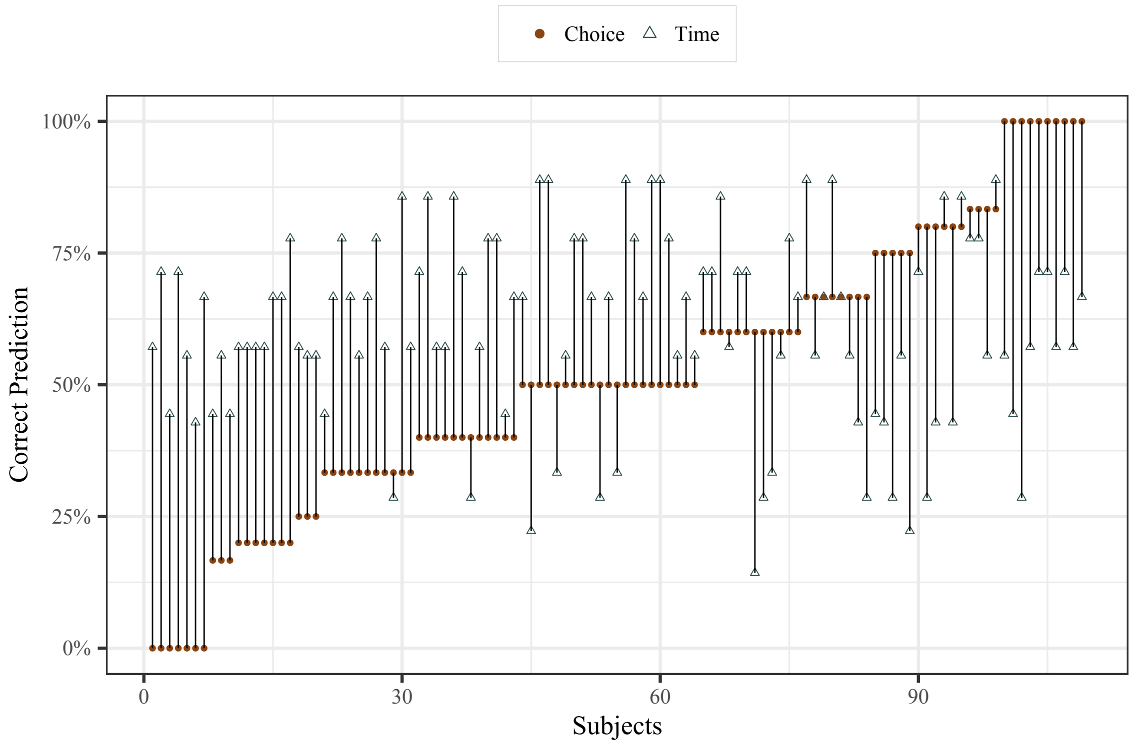

Figure 4: Model's Success by Subject (Dominance Excluded)

Time C_R, Time R_R, and Time D_R are particularly successful. Recall that they test if subjects allocate relatively more time to a contract of a given type if we increase its spread. For instance, prediction Time C_R tests whether contract C_2 receives more time than contract C_1, given that it has a higher spread. The success of these predictions suggest that subjects indeed react to spreads when allocating time, which corroborates one of the model's main intuitions.

But, if this is so, what can explain the failure of predictions Time R_2 and, specially, Time D_2? Moreover, as we show below, predictions about risk contracts and, especially, about confidence contracts are very robust to the specification of preferences over contracts, but they still fail for a significant we assume and it still fails to one third of our subjects.

The previous discussion suggests that two questions require closer scrutiny. First, why are choice predictions less successful than time predictions? Second, do subjects understand the relevance of spreads when allocating time?

Why Are Time Predictions More Successful Than Choice Predictions? One reason why time predictions fare better, on average, than choice predictions is that some time predictions are more robust than choice predictions, in the following sense. Time predictions would still hold

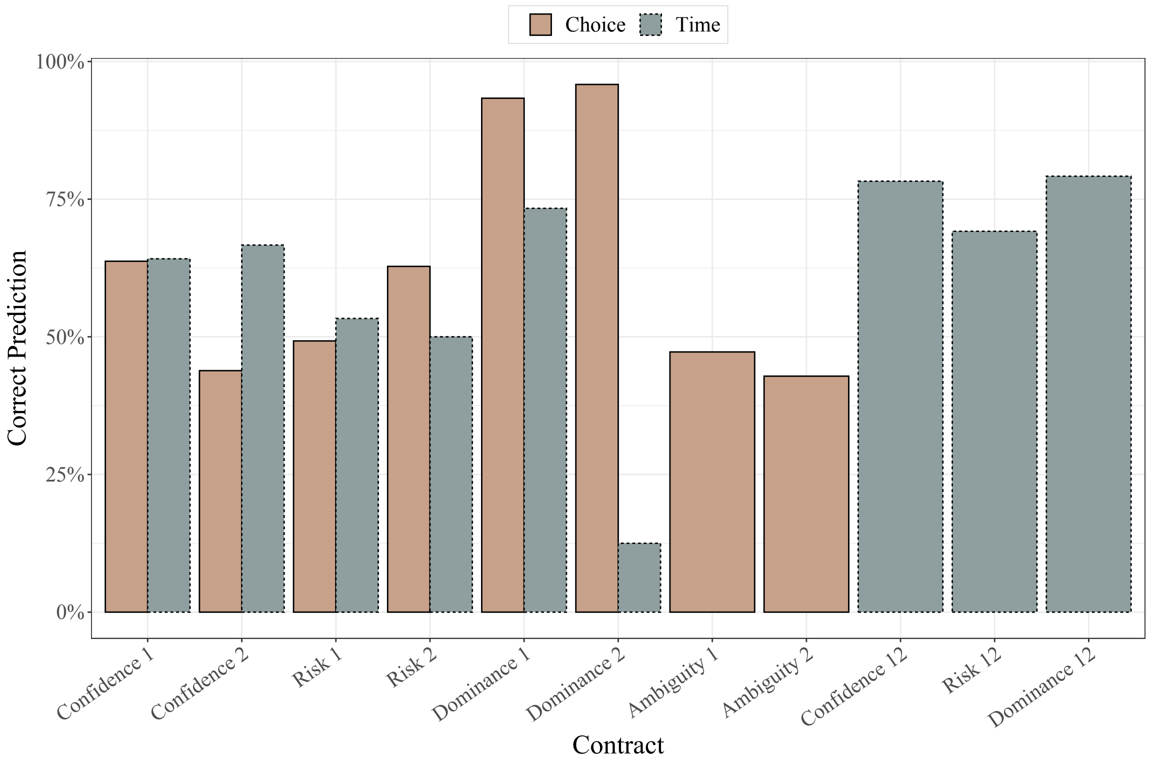

Figure 5: Degree of Success by Prediction

for different types of preferences over contracts, whereas choice predictions are more sensitive to the specific preferences over contracts we assume. In fact, 6 out of the 9 time predictions of the model would hold for much more general preferences over contracts than the ones we use.

To illustrate, suppose that preferences over contracts are represented by the following utility over contracts:

$$U(\mathcal{L}|t) = \rho(\alpha, t)V(L_{\text{Solve}}) + (1 - \rho(\alpha, t))V(L_{\text{NotSolve}}),$$

where $\rho(\alpha, t) : [0, 1] \times \mathbb{R}_+ \to [0, 1]$ should be interpreted as a decision weight, V is a functional over lotteries, and L_{Success} and L_{Failure} are the lotteries one gets if one succeeds and fails in solving the contract's task. We show in Appendix A that if $\rho(\alpha, \cdot)$ is increasing, concave, and continuously differentiable, the model's time predictions about confidence and risk contracts would still hold for several well-known functionals V, including functionals that allow for different types of probability weighting (see Starmer (2000)). We also show in Appendix A that, under the same assumptions on $\rho(\alpha, \cdot)$, the model's time predictions about confidence contracts would still hold if we only assume that V is strictly increasing with respect to first order stochastic dominance.[13]

The same is not true for the model's choice predictions: departures from expected utility over lotteries can lead to very different predictions about choice behavior. Moreover, the choice predictions about ambiguity contracts — Choice_A_1 and Choice_A_2 — rely on the auxiliary model we postulate to resolve the ambiguity in the completion rate of the task.[14] They also rely on the auxiliary assumption that the confidence function is increasing in the completion rate when $t = 60$, an assumption that we can only imperfectly verify. Therefore, it is difficult to know whether the model's relative poor performance in these predictions should be attributed to the model itself, or to these auxiliary assumptions.

Moreover, the model's choice predictions also depend on the subjects' characteristics that we elicit in Part III. Therefore, these characteristics are subject to measurement error (cf., Gillen et al. (2019)). The most important among these characteristics is the value of the confidence function when $\alpha = 0.5$ and $t = 60$, that is, $\boldsymbol{p}(0.5, 60)$. Given the way we elicit $\boldsymbol{p}(0.5, 60)$, we can be over-estimating subjects confidence (Benoît et al. (2022)), which might explain why the prediction Choice_C_2 fails. However, measurement error provides — at best — a partial explanation for the failures of the model. In fact, our model predicts that, if we ignore indifference, either B should be chosen both over C_1 and C_2, or C_1 and C_2 should both be chosen over B. Although this prediction unaffected by measurement error, only the choices of 49% of subjects conform to it.

As we show in detail below, one important reason for the failure of the model's choice predictions is that failure payoffs matter much more than the model predicts when they are too low. Subjects then avoid choosing such contracts, even if they offer higher spreads that, according to our model, should compensate some types of subjects for the low failure payoffs.

The Attractiveness Bias and the Importance of Spreads When Allocating Time As you recall, our time predictions rely almost exclusively on what we called *the spread* of a contract. To investigate whether subjects do indeed pay attention to this spread, consider Figure 5 once more. As we have seen, predictions Time C_R, Time R_R, and Time D_R are particularly successful. Given that the spreads of contracts C_2, R_2, and D_2 are higher than the spreads of C_1, R_1, and D_1, this suggests that subjects react to the increase in the spreads of a contract, as predicted by the model.

[13] The preferences over contracts induced by U includes the broad class of biseparable preferences studied by Ghirardato and Marinacci (2001).

[14] See Appendix A, where we explain the model for the resolution of ambiguity we use. This model is also necessary for the validity of our elicitation of subject's attitude towards ambiguity, as discussed in Appendix B.2

To check if this is so, we ran the following regression:

$$t^i_{(\mathcal{L}_B,\mathcal{L}_X)} = \beta_0 + \beta_1 \left(\frac{\Delta^i_{\mathcal{L}_B}}{\Delta^i_{\mathcal{L}_X}} - 1 \right)$$

where for each subject i and $X \in \{C_1, C_2, R_1, R_2, D_1, D_2\}$,[15] $t^i_{(\mathcal{L}_B,L_X)}$ is the amount of time subject i allocates to contract $\mathcal{L}_B$ in the pair $(\mathcal{L}_B, \mathcal{L}_X)$, and $\left(\frac{\Delta_{\mathcal{L}^i_B}}{\Delta_{\mathcal{L}^i_X}} - 1 \right)$ is the *normalized* ratio of spreads for subject i.

For each pair $(\mathcal{L}_B, \mathcal{L}_X)$, when the normalized ratio of spreads is 0 for subject i, that is, when $\Delta^i_{\mathcal{L}_B} = \Delta^i_{\mathcal{L}_X}$, the model predicts that the subject should allocate the same amount of time to $\mathcal{L}_B$ and $\mathcal{L}_X$. Therefore, if the model is correct on average, we expect $\beta_0 = 60$. We also expect $\beta_1 > 0$, because the model predicts that as the (normalized) ratio of spreads increases, subjects should allocate more time to $\mathcal{L}_B$ relative to $\mathcal{L}_X$.

Table 4 displays the results of the regression. The results show that, consistent with the model, $\beta_1 = 10.77 > 0$. This means that if the spread of $\mathcal{L}_B$ is twice that of $\mathcal{L}_X$, then subjects allocate roughly 11 more seconds to $\mathcal{L}_B$, which suggests that subject understand the importance of spreads for time allocation. At the same time, $\beta_0 = 65.64$, which is statistically greater than 60. The magnitude of β_0 is puzzling given that the model predicts that the baseline contract should receive less than 60 seconds in 5 out of the 6 pairs of contracts included in the regression. Therefore, what is behind it?

Table 4: The Attractiveness Bias

	Coefficient	Robust Std. Err.	p-value
Spread Ratio	10.77	2.47	0.00
Constant	65.64	1.37	0.00

The success of prediction Choice D_2 and failure of prediction Time D_2 suggest that, instead of basing their time allocation exclusively on the ratio of spreads, subjects are allocating more time to the contract they would choose in pair and, hence, find more attractive. We call this pattern the *attractiveness bias*, and we conjecture that it is a particular case of a process Kahneman et al. (2002) call *attribute substitution*. When asked to allocate time between contracts, a complex question, subjects substitute it by a simpler question, namely what contract they find more attractive (and, hence, would choose). Their first instinct is then to allocate more time to

[15] These are the pairs for which the model makes predictions about time allocation.

the more attractive contract, but they adjust their initial instinct to take into account the spreads of the two contracts.

To test whether the data supports the attractiveness bias, we ran a modified version of our first regression, where we condition time allocation on the contract subjects choose in each pair:

$$t^i_{(\mathcal{L}_B,\mathcal{L}_X)} = \sum_{k\in\{B,X,\text{Indiff}\}} \mathbf{1}^i_{\{Choice_{(\mathcal{L}_B,\mathcal{L}_X)}=\mathcal{L}_k\}} \left[\beta^k_0 + \beta^k_1 \left(\frac{\Delta^i_{\mathcal{L}_B}}{\Delta^i_{\mathcal{L}_X}} - 1\right)\right],$$

where for each subject i and $X \in \{C_1, C_2,\ R_1, R_2,\ D_1, D_2\}$, $t^i_{(\mathcal{L}_B,\mathcal{L}_X)}$ is the amount of time subject i allocates to contract $\mathcal{L}_B$ in the pair $(\mathcal{L}_B, \mathcal{L}_X)$, $\left(\frac{\Delta^i_{\mathcal{L}_B}}{\Delta_{\mathcal{L}^i_X}} - 1\right)$ is the *normalized* ratio of spreads, and $\mathbf{1}^i_{\{Choice_{(\mathcal{L}_B,\mathcal{L}_X)}=\mathcal{L}_k\}}$ is an indicator variable that, for each $k \in \{B, X, \text{Indiff}\}$, takes value 1 when subject i chooses $\mathcal{L}_k$, and 0 otherwise.

This regression amounts to running three separate regressions. One regression for the observations in which a subject chose the baseline contract. One regression for the observations where a subject chose the other contract; and, finally, one regression for the observations in which a subject declared to be indifferent between the two contracts. For each of these cases, we estimate a simple linear regression between the time allocated to the baseline and the normalized ratio of spreads. [16] For each pair $(\mathcal{L}_B, \mathcal{L}_X)$ and for each $k \in \{B, X, \text{Indiff}\}$, the model again implies that $\beta^k_0 = 60$ and $\beta^k_1 > 0$.

Table5 displays the results of the regression. Subjects that choose $\mathcal{L}_B$ allocate, on average, more time to $\mathcal{L}_B$ than what the model predicts they should when the contracts have the same spreads ($\beta^B_0 = 70.57$). Similarly, subjects that choose $\mathcal{L}_X$ allocate, on average, more time to $\mathcal{L}_X$ than what the model predicts when the contracts have the same spreads ($\beta^X_0 = 50.6$). Finally, subjects that declare indifference allocate roughly the same amount of time to both lotteries ($\beta^{\text{Indiff}}_0 = 57.73$). Therefore, subjects allocate time in a way that seems to be consistent with how they choose between - and, hence, value - the contracts.

When choosing the baseline contract or declaring indifference, subjects react to the increase in the ratio of spreads as predicted by the model ($\beta^B_1 = 12.35$ and $\beta^{\text{Indiff}}_1 = 5.31$ are statistically significant) but not when choosing the other contract in the pair (β^X_1 is not statistically significant). Interestingly, when subjects choose the baseline contract but the model predicts that they should allocate more time to the other contract, the ratio $\frac{\Delta_{\mathcal{L}_B}}{\Delta_{\mathcal{L}_X}}$ must be sufficiently

[16] In Appendix ???, we run separate specifications of this regression to take into account heterogeneity across subjects and the existence of a non-linear relation between time allocated to the baseline and the ratio of spreads. The qualitative results remain the same and, hence, we decided to include only this regression.

Table 5: The Attractiveness Bias

	Coefficient	Robust Std. Err.	p-value
Time			
Choose B	70.57	1.62	0.00
Indifference	57.73	3.83	0.00
Choose Other	50.60	1.91	0.00
Choice			
Choose B	12.35	1.70	0.00
Indifference	5.31	2.58	0.04
Choose Other	-1.26	1.38	0.36

small - namely, smaller than 1/6 - so that, on average, subjects allocate more time to the other contract than to the baseline. No pair of contracts in the experiment has such a small ratio $\frac{\Delta_{\mathcal{L}_B}}{\Delta_{\mathcal{L}_X}}$, which suggests that although subjects react to the ratio of spreads, they do not, on average, react enough relative to the attractiveness bias.

5.2 Designing Cost-Effective Contracts to Compete for Time

We now explore what can Principal learn about designing contracts that can cost effectively compete for an independent contractor's time from our experimental results. Table 6 summarize subjects' choices and time allocations in the experiment. For each pair of contracts, the left panel displays the percentages of subjects who chose the baseline contract ($\mathcal{L}_B$), who are indifferent ('Indifferent'), and who choose the other contract (Not $\mathcal{L}_B$). The right panel displays the average amount of time subjects allocate to the baseline ($\mathcal{L}_B$) and to the other contract in the pair (Not $\mathcal{L}_B$).

The contracts that we compare to the baseline contract are the result of varying some of the baseline's features so as not to increase the expected cost of implementing the contract, where the expected cost of a contract $\mathcal{L}$ is defined as

$$\text{EC}(\mathcal{L}) := \alpha(p_S \times H + (1-p_S) \times L) + (1-\alpha)(p_F \times H + (1-p_F) \times L). \quad (4)$$

How effective are contracts that spreads probabilities while keeping expected costs constant? That is, how effective are confidence contracts? Table 6 shows that confidence contracts are effective when people must choose between contracts, but their effectiveness decreases if we the spread in probabilities too much by reducing their failure payoff. Whereas 3 out of 4 subjects choose $\mathcal{L}_{C_1}$ over $\mathcal{L}_B$, less than 2 out of 4 subjects choose $\mathcal{L}_{C_2}$ over $\mathcal{L}_B$. This suggests that

Table 6: Choice and time allocation to contract pairs

Pairs	*Choice (percentage)*			*Time (seconds)*	
	$\mathcal{L}_B$	Indifferent	Not $\mathcal{L}_B$	$\mathcal{L}_B$	Not $\mathcal{L}_B$
B vs C1	17.5	10.0	72.5	51.9	68.1
B vs C2	43.5	11.6	44.9	48.6	71.4
B vs R1	36.7	11.7	51.7	54.1	65.9
B vs R2	52.5	5.83	41.7	54.6	65.4
B vs D1	93.3	3.33	3.3	78.9	41.1
B vs D2	95.8	0.83	3.3	73.0	47.0
B vs A1	33.3	31.7	35.0	60.1	59.9
B vs A2	53.3	20.8	25.8	62.6	57.4

although subjects are willing to trade-off a decrease in p_F for an increase in p_S, there are limits to this willingness. In fact, a robust finding in our experiment is that, when choosing between contracts, subjects shy away from contracts that offer too unfavourable outcomes if they fail the task, i.e., with a low failure payoff , even if symmetrically compensated by a favorable outcome if they succeed. That is, subjects find contracts that offer a poor safety net, unattractive when choosing between contracts.

In time allocation, both $\mathcal{L}_{C_1}$ and $\mathcal{L}_{C_2}$ attract significantly more time, on average, than $\mathcal{L}_B$, namely 68.1 and 71.4 seconds. This suggests, consistent with the model, that increasing the spread of probabilities, i.e., increasing $p_S - p_F$, is effective in attracting people's time. However, the average amount of time subjects allocate to $\mathcal{L}_{C_1}$ is not statistically different from the one they allocate to $\mathcal{L}_{C_2}$, which suggests that perhaps subjects do not react to spreads as much as the model predicts. To put matters into perspective, $\mathcal{L}_B$ has a spread that is only 2 times than that of the dominated contract $\mathcal{L}_{D_1}$, and subjects allocate, on average, 79 seconds to $\mathcal{L}_B$. $\mathcal{L}_{C_2}$ has a spread that is 5 times higher than that of $\mathcal{L}_B$, and subjects allocate, on average, 71 seconds to it.

The attractiveness bias can rationalize these patterns. Given that 72.5% of subjects choose $\mathcal{L}_{C_1}$ over $\mathcal{L}_B$ and 44.9% choose $\mathcal{L}_{C_2}$ over $\mathcal{L}_B$, the attractiveness bias then implies that subjects allocate more time, on average, to $\mathcal{L}_{C_1}$ relative to $\mathcal{L}_B$ than the model predicts given their ratio of spreads, whereas they allocate less time to $\mathcal{L}_{C_2}$ relative to $\mathcal{L}_B$ than the model predicts. This makes the average amount of time allocated to $\mathcal{L}_{C_1}$ and $\mathcal{L}_{C_2}$ closer than what we expect given the model. Moreover, given that 93.3% of subjects choose $\mathcal{L}_B$ over $\mathcal{L}_{D_1}$, the attractiveness bias can also explain why $\mathcal{L}_B$ attracts more time over $\mathcal{L}_{D_1}$ than $\mathcal{L}_{C_2}$ over $\mathcal{L}_B$.

Interestingly, we reach very similar conclusions about the effectiveness of spreading payoffs while keeping the expected cost of the contract the same, that is, about the effectiveness of risk contracts. 51.7% of subjects choose $\mathcal{L}_{R_1}$, whereas 36.7% choose $\mathcal{L}_B$, but only 41.7% of subjects choose $\mathcal{L}_{R_2}$, whereas 51.7% choose $\mathcal{L}_B$. This again suggests that subjects shy away from contracts that can result in outcomes that they deem too unfavorable if they fail even if symmetrically compensated by a favorable outcome if they succeed. Moreover, although both risky contracts attract more time than the baseline, $\mathcal{L}_{R_2}$ attracts the same average amount of time as $\mathcal{L}_{R_1}$. The attractiveness bias can again rationalize this pattern.

The spread of a contract in our model is the result of multiplying the spread in probabilities with the spread in payoffs. We have seen that increases in either of these spreads that keep expected costs constant are effective at attracting people's choices and time, provided that we do not increase the spreads too much. A natural question is then what would happen if we increased both spreads?

If the principal is interested in maximizing the amount of time a contract attracts against a baseline contract $\mathcal{L}_B = (H_B, L_B, \alpha, p_S^B, p_F^B)$ while keeping constant expected costs, then our model suggests that the principal should find a contract $\mathcal{L} = (H, L, \alpha, p_S, p_F)$ that solves

$$
\begin{gathered}
\max_{\mathcal{L}} (H-L)(p_S-p_F) \\
\text{subject to} \\
\alpha = \alpha_B \\
H \geqslant L \geqslant 0 \\
p_S \geqslant p_F \geqslant 0 \\
\text{EC}(\mathcal{L}) = \text{EC}(\mathcal{L}_B)
\end{gathered}
\tag{5}
$$

Given the baseline contract we use in the experiment, any contract that solves this problem is of the form $H = x$, $L = 0$, $p_S = 12/x$, and $p_F = 0$, for some $x \geqslant 12$. To fix ideas, set $x = 1$. Our results suggest that such a contract will fair poorly in choice, and, due to the attractiveness bias, will not attract as much time as we would expect. We should instead offer a less extreme contract, say, one in which $p_S = 0.8$, $p_F = 0.2$, $H = 10$ and $L = 2$.

Our model suggests that the Principal can compete against the baseline contract for people's time while reducing the expected costs of the baseline contract. One way to do so is to increase the spread in probabilities ($p_S - p_F$) while decreasing the sum of these probabilities ($p_S + p_F$) relative to the baseline. Since we increased the spread in probabilities, the model predicts that subjects should allocate more time to the resulting contract than to the baseline. This is precisely what dominated contracts do. As expected, dominated contracts $\mathcal{L}_{D_1}$ and $\mathcal{L}_{D_2}$ do much worse than $\mathcal{L}_B$ in choice: 93.3% of subjects choose the baseline contract over either of these two contracts. What is surprising, at first, is that the baseline fares better in time allocation than $\mathcal{L}_{D_2}$, although $\mathcal{L}_{D_2}$ has a higher probability spread, namely $(0.3 - 0)$, than that of $\mathcal{L}_B$, namely $(0.6 - 0.4)$. But, again, the attractiveness bias rationalizes this pattern: $\mathcal{L}_{D_2}$ is (much) less attractive than $\mathcal{L}_B$, and, hence, subjects allocate (much) less time to it than the model predicts.

Finally, we arrive to ambiguous contracts. $\mathcal{L}_{A_1}$ is chosen as often as $\mathcal{L}_B$, which indicates that subjects are not, in general, averse to a small amount of (symmetric) ambiguity in the completion rate. Nevertheless, when the ambiguity increases and we reach $\mathcal{L}_{A_2}$, subjects choose $\mathcal{L}_B$ twice as often than $\mathcal{L}_{A_2}$, which suggests again that people are averse to the possibility of extreme unfavorable outcomes, such as the completion rate of the ambiguous contract being 0. Interestingly, both ambiguous contracts attract roughly the same time, on average, than the baseline contract. For time allocation, therefore, ambiguity aversion seems not to affect a subject's decisions as much.

Therefore, our experimental result suggest three lessons for designing contracts to overcome the contract-choice and contract-time allocation problems:

1. When allocating time, people consider not only a contract's spread, as predicted by the model, but also how attractive the contract is. The attractiveness of a contract can even over-rule the logic of spreads. Therefore, to overcome the contract-time allocation problem, a contract should not only have high spreads but it should also be attractive;

2. When choosing, subjects avoid contracts with extreme unfavorable outcomes, e.g., that pay very low prizes, or pay low prizes with certainty, or in which the completion rate can be very low. Provided that when increasing the spread of probabilities, the spread of payoffs, or ambiguity, we do not introduce these extreme unfavorable outcomes, the resulting contract can overcome the contract-choice problem;

3. Putting (2) and (3) together, our results suggest that increasing the spreads in probabilities or payoffs is a cost-effective way of competing against the baseline contract both for the choice and the time of subjects provided that we do not introduce extremely unfavorable outcomes when increasing spreads. Therefore, contracts that moderately increase either

one or both spreads can overcome both the contract-choice and contract-time allocation problems.

6 Conclusion

This paper addresses an understudied problem: how principals write contracts for independent contractors. Such contractors differ from regular employees in the sense that they are not entirely reliant on the firm for their income and, unless incentivized, may divert their time to other contracts that they have agreed to work on. The challenge for the principal is how to write a contract that will not only be attractive enough to be chosen but, equally importantly, lead the contractor to allocate time to it as a part of her contract portfolio.

Our findings highlight several key insights for designing effective contracts. First, the failure payoffs play a critical role in the contract-choice problem. Contracts with too low a failure payoff are less likely to be selected, even if they offer higher spreads. This indicates a potential trade-off for principals in balancing between incentivizing time allocation and ensuring the contract is chosen in the first place. Secondly, while higher spreads generally lead to more time allocated to a contract, the "attractiveness bias" we observed suggests that people tend to allocate more time to the contract they find more appealing at the choice stage. This behavior can result in sub-optimal time allocation, which the theory does not capture.

Overall, our results suggest that successful contract design must consider both hard and soft incentives and the behavioral tendencies of ICs. Principals must carefully calibrate failure payoffs and spreads to address both the contract-choice and contract-time allocation problems effectively. Moreover, understanding the attractiveness bias can help in designing contracts that are not only chosen but also receive adequate time commitment.

References

Agranov, Marina and Pietro Ortoleva, "Stochastic choice and preferences for randomization," *Journal of Political Economy*, 2017, *125* (1), 40–68.

Ales, Laurence and Pricila Maziero, "Non-exclusive dynamic contracts, competition, and the limits of insurance," *Journal of Economic Theory*, 2016, *166*, 362–395.

Attar, Andrea, Thomas Mariotti, and François Salanié, "Nonexclusive competition in the market for lemons," *Econometrica*, 2011, *79* (6), 1869–1918.

_ , _ , and _ , "Nonexclusive competition under adverse selection," *Theoretical Economics*, 2014, *9* (1), 1–40.

Avoyan, Ala and Andrew Schotter, "Attention in games: An experimental study," *European Economic Review*, 2020, *124*, 103410.

_ and Giorgia Romagnoli, "Paying for Inattention," *Economics Letters*, 2023, *226*, 111114.

_ , Mauricio Ribeiro, Andrew Schotter, Elizabeth R Schotter, Mehrdad Vaziri, and Minghao Zou, "Planned vs. Actual Attention," *Management Science*, 2024, *70* (5), 2912–2933.

Bar-Isaac, Heski and Joyee Deb, "(Good and bad) reputation for a servant of two masters," *American Economic Journal: Microeconomics*, 2014, *6* (4), 293–325.

Benoît, Jean-Pierre, Juan Dubra, and Giorgia Romagnoli, "Belief elicitation when more than money matters: controlling for "control"," *American Economic Journal: Microeconomics*, 2022, *14* (3), 837–888.

Bernheim, B Douglas and Michael D Whinston, "Exclusive dealing," *Journal of political Economy*, 1998, *106* (1), 64–103.

Bisin, Alberto and Danilo Guaitoli, "Moral hazard and nonexclusive contracts," *RAND Journal of Economics*, 2004, pp. 306–328.

Bock, Olaf, Ingmar Baetge, and Andreas Nicklisch, "hroot: Hamburg registration and organization online tool," *European Economic Review*, 2014, *71*, 117–120.

Burmeister-Lamp, Katrin, Moren Lévesque, and Christian Schade, "Are entrepreneurs influenced by risk attitude, regulatory focus or both? An experiment on entrepreneurs' time allocation," *Journal of Business Venturing*, 2012, *27* (4), 456–476.

Cabrales, Antonio and Gary Charness, "Optimal contracts with team production and hidden information: An experiment," *Journal of Economic Behavior & Organization*, 2011, *77* (2), 163–176.

Charness, Gary and Uri Gneezy, "Portfolio choice and risk attitudes: An experiment," *Economic Inquiry*, 2010, *48* (1), 133–146.

Chen, Daniel L, Martin Schonger, and Chris Wickens, "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Corgnet, Brice, Joaquín Gómez-Miñambres, and Roberto Hernán-Gonzalez, "Goal setting in the principal–agent model: Weak incentives for strong performance," *Games and Economic Behavior*, 2018, *109*, 311–326.

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole, "Multitask agency problems: Focus and task clustering," *European economic review*, 2000, *44* (4-6), 869–877.

Dixit, Avinash K, *The making of economic policy: A transaction-cost politics perspective*, MIT press, 1998.

Ghirardato, Paolo and Massimo Marinacci, "Risk, ambiguity, and the separation of utility and beliefs," *Mathematics of operations research*, 2001, *26* (4), 864–890.

Gillen, Ben, Erik Snowberg, and Leeat Yariv, "Experimenting with measurement error: Techniques with applications to the caltech cohort study," *Journal of Political Economy*, 2019, *127* (4), 1826–1863.

Gneezy, Uri and Jan Potters, "An experiment on risk taking and evaluation periods," *The Quarterly Journal of Economics*, 1997, *112* (2), 631–645.

Halac, Marina and Andrea Prat, "Managerial attention and worker performance," *American Economic Review*, 2016, *106* (10), 3104–3132.

Holt, Charles A and Susan K Laury, "Risk aversion and incentive effects," *American economic review*, 2002, *92* (5), 1644–1655.

Hoppe, Eva I and Patrick W Schmitz, "Contracting under incomplete information and social preferences: An experimental study," *Review of Economic Studies*, 2013, *80* (4), 1516–1544.

Kahneman, Daniel, Shane Frederick et al., "Representativeness revisited: Attribute substitution in intuitive judgment," *Heuristics and biases: The psychology of intuitive judgment*, 2002, *49* (49-81), 74.

Manyika, James, Susan Lund, Jacques Bughin, Kelsey Robinson, Jan Mischke, and Deepa Mahajan, "Independent Work: Choice," *Necessity and the Gig Economy, San Francisco, CA: McKinsey Global Institute*, 2016.

Martimort, David, "Exclusive dealing, common agency, and multiprincipals incentive theory," *The RAND journal of economics*, 1996, pp. 1–31.

Moore, Don A and Paul J Healy, "The trouble with overconfidence.," *Psychological review*, 2008, *115* (2), 502.

Sinander, Ludvig, "Optimism, overconfidence, and moral hazard," *arXiv preprint arXiv:2304.08343*, 2024.

Starmer, Chris, "Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk," *Journal of economic literature*, 2000, *38* (2), 332–382.

Appendices

A Derivation of the Predictions

In this section, we derive the formally derive the predictions of the model we used to structure the experiment. Recall that, given two contracts $\mathcal{L}_1$ and $\mathcal{L}_2$, $\mathcal{L}_1$ is chosen from the menu $\{\mathcal{L}_1, \mathcal{L}_2\}$ if, and only if,

$$\boldsymbol{p}(\alpha_1, t_1)\Delta_{\mathcal{L}_1} - \boldsymbol{p}(\alpha_2, t_2)\Delta_{\mathcal{L}_2} \geqslant N_{\mathcal{L}_2} - N_{\mathcal{L}_1},$$

and that the amount of time allocated to $\mathcal{L}_\infty$ satisfies

$$\frac{\boldsymbol{p}'(\alpha_2, T - t_{\mathcal{L}_1})}{\boldsymbol{p}'(\alpha_1, t_{\mathcal{L}_1})} = \frac{\Delta_{\mathcal{L}_1}}{\Delta_{\mathcal{L}_2}}.$$

We refer to the first expression and the *Choice Formula* and to the second expression as the *Time Allocation Formula* in what follows.

A.1 Confidence Contracts

A confidence contract $\mathcal{L}_C$ relative to $\mathcal{L}_B$ is one in which $H_C = H_B$, $L_C = L_B$, $\alpha_C = \alpha_B$ and

$$p_F^C < p_F^B < p_S^B < p_S^C.$$

Manipulating the Choice Formula, we get that $\mathcal{L}_C$ is chosen from the menu $\{\mathcal{L}_B, \mathcal{L}_C\}$ if, and only if,

$$\frac{\boldsymbol{p}(\alpha_B, t)}{1 - \boldsymbol{p}(\alpha_B, t)} \geqslant \frac{p_F^B - p_F^C}{p_S^C - p_S^B}.$$

Since $\frac{\boldsymbol{p}(\alpha_B, t)}{1 - \boldsymbol{p}(\alpha_B, t)}$ is an increasing function of $\boldsymbol{p}(\alpha_B, t)$, a sufficiently confident agent, i.e. one that assigns a sufficiently high probability to solving the task, should choose $\mathcal{L}_C$ over $\mathcal{L}_B$.

Manipulating the Time Allocation Formula, we get that the amount of time allocated to $\mathcal{L}_C$ satisfies

$$\frac{\boldsymbol{p}'(T - t_{\mathcal{L}_C}, \alpha_B)}{\boldsymbol{p}'(t_{\mathcal{L}_C}, \alpha_B)} = \frac{p_S^C - p_F^C}{p_S^B - p_F^B}$$

Since the left hand side is increasing on $t_{\mathcal{L}_C}$ and the right hand side is greater than one, we conclude that $t_{\mathcal{L}_C} > \frac{T}{2}$. That is, the agent should always allocate more time to a confidence contract. Moreover, the amount of time allocated to $\mathcal{L}_C$ is increasing in $p_S^C - p_F^C$.

In the experiment, we use two confidence contracts , namely $\mathcal{L}_{C_1}$ and $\mathcal{L}_{C_2}$ (see Table 2). For both contracts, the model predicts that a subject chooses $\mathcal{L}_{C_i}$ from the menu $\{\mathcal{L}_B, \mathcal{L}_{C_i}\}$ if, and only if, $\boldsymbol{p}\left(\frac{1}{2}, 1\right) \geqslant \frac{1}{2}$. The model also predicts that subjects should always allocate more time to $\mathcal{L}_{C_i}$ than to $\mathcal{L}_B$ for each $i \in \{1, 2\}$, i.e., $t_{\mathcal{L}_{C_i}} > 60$. Finally, given that $p_S^{C_2} - p_F^{C_2} > p_S^{C_1} - p_F^{C_1}$, our theory predicts that agents should allocate more time to $\mathcal{L}_{C_2}$ (relative to $\mathcal{L}_B$) than to $\mathcal{L}_{C_1}$ (relative to $\mathcal{L}_B$), i.e., $t_{\mathcal{L}_{C_2}} > t_{\mathcal{L}_{C_1}}$.

A.2 Risk Contracts

A risk contract $\mathcal{L}_R$ relative to $\mathcal{L}_B$ is on in which $p_S^R = p_S^B$, $p_F^R = p_F^B$, $\alpha_R = \alpha_B$,

$$L_R < L_B < H_B < H_R,$$

and $H_R - H_B = L_B - L_R$.

Manipulating the Choice Formula, we get that $\mathcal{L}_R$ is chosen from the menu $\{\mathcal{L}_R, \mathcal{L}_B\}$ if, and only if,

$$K_{(H_B,L_B)}^{(H_R,L_R)}(u) \geqslant \frac{1 - \left(\boldsymbol{p}(\alpha_B, t)p_S^B + (1 - \boldsymbol{p}(\alpha_B, t))p_F^B\right)}{(\boldsymbol{p}(\alpha_B, t)p_S^B + (1 - \boldsymbol{p}(\alpha_B, t))p_F^B)},$$

where $K_{(H_B,L_B)}^{(H_R,L_R)}(u) = \frac{u(H_R)-u(H_B)}{u(L_B)-u(L_R)}$ measures the agent's risk attitude (see Appendix ?). Define

$$\boldsymbol{q}(\alpha_B, t) \coloneqq \boldsymbol{p}(\alpha_B, t)p_S^B + (1 - \boldsymbol{p}(\alpha_B, t))p_F^B,$$

and $\mathcal{L}_R$ is chosen from $\{\mathcal{L}_R, \mathcal{L}_B\}$ if, and only if,

$$K_{(H_B,L_B)}^{(H_R,L_R)}(u) \geqslant \frac{1 - \boldsymbol{q}(\alpha_B, t)}{\boldsymbol{q}(\alpha_B, t)}.$$

Therefore, a sufficiently risk seeking agent will always choose $\mathcal{L}_R$.[17] Similarly, a sufficiently risk averse subject will always choose $\mathcal{L}_B$.[18] Since the right hand side is a decreasing function $\boldsymbol{p}(\alpha_B, t)$, for intermediate values of $K_{(H_B,L_B)}^{(H_R,L_R)}(u)$, the more confident a subject is, i.e., the higher $\boldsymbol{p}(\alpha_B, t)$, the more prone she is to choose $\mathcal{L}_R$ from $\{\mathcal{L}_R, \mathcal{L}_B\}$.

Manipulating the Time Allocation Formula, we get that the amount of time allocated to $\mathcal{L}_R$ satisfies

$$\frac{\boldsymbol{p}'(T - t_{\mathcal{L}_R}, \alpha_B)}{\boldsymbol{p}'(t_{\mathcal{L}_R}, \alpha_B)} = \frac{u(H_R) - u(L_R)}{u(H_B) - u(L_B)}.$$

Since the left hand side is increasing on $t_{\mathcal{L}_R}$ and the right hand side is greater than one if we

[17] This will happen whenever $K_{(H_B,L_B)}^{(H_R,L_R)}(u) \geqslant p_S^B / p_F^B$

[18] In fact, this happens whenever $K_{(H_B,L_B)}^{(H_R,L_R)}(u) < \frac{p_{N,B}}{p_{S,B}}$.

assume that u is strictly increasing, we conclude that $t_{\mathcal{L}_C} > \frac{T}{2}$. That is, the agent should always allocate more time to $\mathcal{L}_R$. Moreover, the amount of time allocated to $\mathcal{L}_R$ is increasing in $u(H_R) - u(L_R)$.

In the experiment, we use two risk contracts, namely $\mathcal{L}_{R_1}$ and $\mathcal{L}_{R_2}$ (see Table 2). The model predicts that a subject chooses $\mathcal{L}_{C_1}$ from the menu $\{\mathcal{L}_B, \mathcal{L}_{C_1}\}$ if, and only if,

$$K_{(8,4)}^{(10,2)}(u) \geqslant \frac{1 - \boldsymbol{q}(0.5, 60)}{\boldsymbol{q}(0.5, 60)},$$

and chooses $\mathcal{L}_{C_2}$ from the menu $\{\mathcal{L}_B, \mathcal{L}_{C_2}\}$ if, and only if,

$$K_{(8,4)}^{(12,0)}(u) \geqslant \frac{1 - \boldsymbol{q}(0.5, 60)}{\boldsymbol{q}(0.5, 60)}.$$

The model also predicts that subjects should always allocate more time to $\mathcal{L}_{R_i}$ than to $\mathcal{L}_B$ for each $i \in \{1, 2\}$, i.e., $t_{\mathcal{L}_{R_i}} > 60$. Finally, provided that $u(12) - u(0) > u(10) - u(2)$, our theory predicts that agents should allocate more time to $\mathcal{L}_{R_2}$ (relative to $\mathcal{L}_B$) than to $\mathcal{L}_{R_1}$ (relative to $\mathcal{L}_B$).

A.3 Dominated Contracts

A dominated contract $\mathcal{L}_D$ relative to $\mathcal{L}_B$ is one in which $H_D = H_B$, $L_D = L_B$, $\alpha_D = \alpha_B$,

$$p_S^T < p_S^B \text{ and } p_F^T < p_F^B.$$

Manipulating the Choice Formula, we get that $\mathcal{L}_D$ is chosen from the menu $\{\mathcal{L}_B, \mathcal{L}_D\}$ if, and only if,

$$\frac{\boldsymbol{p}(\alpha_B, t)}{1 - \boldsymbol{p}(\alpha_B, t)} \leqslant \frac{p_F^B - p_F^T}{p_S^T - p_S^B}.$$

Since the left hand side is always positive $\frac{\boldsymbol{p}(\alpha_B,t)}{1-\boldsymbol{p}(\alpha_B,t)}$ and the left hand side is always negative, then agents should never choose $\mathcal{L}_D$ from the menu $\{\mathcal{L}_B, \mathcal{L}_D\}$.

Manipulating the Time Allocation Formula, we get that the amount of time allocated to $\mathcal{L}_D$ satisfies

$$\frac{\boldsymbol{p}'(T - t_{\mathcal{L}_D}, \alpha_B)}{\boldsymbol{p}'(t_{\mathcal{L}_D}, \alpha_B)} = \frac{p_S^T - p_F^T}{p_S^B - p_F^B}$$

Since the left hand side is increasing on $t_{\mathcal{L}_D}$, if $\frac{p_S^T - p_F^T}{p_S^B - p_F^B} > 1$, then $t_{\mathcal{L}_D} > \frac{D}{2}$. If $\frac{p_S^T - p_F^T}{p_S^B - p_F^B} < 1$, then $t_{\mathcal{L}_D} < \frac{D}{2}$. Moreover, the amount of time allocated to $\mathcal{L}_D$ is increasing in $p_S^T - p_F^T$.

In the experiment, we use two dominated contracts, namely $\mathcal{L}_{T_1}$ and $\mathcal{L}_{T_2}$ (see Table 2).

For both contracts, the model predicts that a subject should always choose $\mathcal{L}_{T_i}$ from the menu $\{\mathcal{L}_B, \mathcal{L}_{T_i}\}$. The model also predicts that subjects should allocate more time to $\mathcal{L}_{T_2}$ than to $\mathcal{L}_B$, i.e., $t_{\mathcal{L}_{T_2}} > 60$, but should allocate more time to to $\mathcal{L}_B$ than to $\mathcal{L}_{T_1}$, i.e., $t_{\mathcal{L}_{T_1}} < 60$. Finally, given that $p_S^{T_2} - p_F^{T_2} > p_S^{T_1} - p_F^{T_1}$, the model predicts that agents should allocate more time to $\mathcal{L}_{T_2}$ (relative to $\mathcal{L}_B$) than to $\mathcal{L}_{T_1}$ (relative to $\mathcal{L}_B$), i.e., $t_{\mathcal{L}_{T_2}} > t_{\mathcal{L}_{T_1}}$.

A.4 Ambiguous Contracts

An ambiguous contract $\mathcal{L}_A$ relative to $\mathcal{L}_B$ is one in which $H_A = H_S$, $L_A = L_S$, $p_S^A = p_S^B$, $p_F^A = p_F^B$, but, for some $\varepsilon > 0$,

$$\alpha_A \in [\alpha_B - \varepsilon, \alpha_B + \varepsilon]$$

We postulate that the agent evaluates $\mathcal{L}_A$ in two steps. She first resolves the ambiguity with respect to α_A, i.e. decides which one of the contracts in $\{\mathcal{L}_x : x \in [\alpha - \varepsilon, \alpha + \varepsilon]\}$ she is facing, and then evaluates the resulting contract using (1).

She resolves the ambiguity with respect to α_A by substituting it by

$$\alpha_\varepsilon := \gamma_\varepsilon(\alpha_B + \varepsilon) + (1 - \gamma_\varepsilon)(\alpha_B - \varepsilon),$$

for some weight γ_ε, which measures the agent's ambiguity attitude.

Manipulating the Choice Formula, we get that $\mathcal{L}_A$ is chosen from the menu $\{\mathcal{L}_B, \mathcal{L}_A\}$ if, and only if, $\boldsymbol{p}(\alpha_\varepsilon, t) \geqslant \boldsymbol{p}(\alpha_B, t)$. If we assume that $p(\cdot, t)$ is strictly increasing, we have that $\mathcal{L}_A$ is chosen from $\{\mathcal{L}_A, \mathcal{L}_B\}$ if, and only if, $\alpha_\varepsilon \geqslant \alpha_B$. Or, equivalently, if

$$\gamma_\varepsilon \geqslant 1/2.$$

Therefore, $\mathcal{L}_A$ is chosen from $\{\mathcal{L}_A, \mathcal{L}_B\}$ if, and only if, the agent is ambiguity seeking (see Appendix ?).

Manipulating the Time Allocation Formula, we get that the amount of time allocated to $\mathcal{L}_A$ satisfies

$$\frac{\boldsymbol{p}'(T - t_{\mathcal{L}_A}, \alpha_B)}{\boldsymbol{p}'(t_{\mathcal{L}_A}, \alpha_\varepsilon)} = 1$$

Therefore, predictions about time allocation rely the assumptions one willing to make about the ratio of derivatives with respect to time evaluated at different completion rates. Since we refrain in making such assumptions, we also refrain from making predictions about time allocation between the baseline and the ambiguous lotteries.

In the experiment, we use two ambiguous lotteries, namely $\mathcal{L}_{A_1}$ and $\mathcal{L}_{A_2}$ (see Table 2). The model predicts that a subject should choose $\mathcal{L}_{A_i}$ from the menu $\{\mathcal{L}_B, \mathcal{L}_{A_i}\}$ if, and only if, $\alpha_{\varepsilon_i} \geqslant 0.5$, i.e. if the subject is ambiguity seeking. The models makes no predictions about time allocation.

A.5 The Model's Predictions in the Experiment

Tables ? to ?? summarize the predictions that the model makes about the choice and time allocation of each subject for each type of contract. For confidence, risk, and dominated contracts, we have 5 predictions per subject. For ambiguous contracts, we have only 2 choice predictions per subject. Therefore, in total, we have 17 predictions per subject. Choice predictions, however, depend on the successful elicitation of subject's characteristics or on additional assumptions being satisfied, which are also included in Tables ? to ??.

Table 7: Confidence contracts - Predictions

	$(\mathcal{L}_B, \mathcal{L}_{C_1})$	$(\mathcal{L}_B, \mathcal{L}_{C_2})$
Choice_C$_i^*$	$p(0.5,1) > 0.5$ iff $\{\mathcal{L}_{C_1}\}$ $p(0.5,1) = 0.5$ iff $\{\mathcal{L}_B,\mathcal{L}_{C_1}\}$ $p(0.5,1) < 0.5$ iff $\{\mathcal{L}_B\}$	$p(0.5,1) > 0.5$ iff $\{\mathcal{L}_{C_2}\}$ $p(0.5,1) = 0.5$ iff $\{\mathcal{L}_B,\mathcal{L}_{C_2}\}$ $p(0.5,1) < 0.5$ iff $\{\mathcal{L}_B\}$
Time_C$_i$	$t_{\mathcal{L}_{C_1}} > t_{\mathcal{L}_B}$	$t_{\mathcal{L}_{C_2}} > t_{\mathcal{L}_B}$
	$t_{\mathcal{L}_{C_2}} > t_{\mathcal{L}_{C_1}}$	

*Assumption(s):

(i) Successful elicitation of $p(0.5, 1)$

Table 8: Risk contracts - Predictions

	$(\mathcal{L}_B, \mathcal{L}_{R_1})$	$(\mathcal{L}_B, \mathcal{L}_{R_2})$
Choic_R$_i^*$	$K_{(6,4)}^{(8,2)}(u) > \frac{1-\boldsymbol{q}(0.5,1)}{\boldsymbol{q}(0.5,1)}$, iff $\{\mathcal{L}_{R_1}\}$ $K_{(6,4)}^{(8,2)}(u) = \frac{1-\boldsymbol{q}(0.5,1)}{\boldsymbol{q}(0.5,1)}$ iff $\{\mathcal{L}_B,\mathcal{L}_{R_1}\}$ $K_{(6,4)}^{(8,2)}(u) < \frac{1-\boldsymbol{q}(0.5,1)}{\boldsymbol{q}(0.5,1)}$ iff $\{\mathcal{L}_B\}$	$K_{(6,4)}^{(10,0)}(u) > \frac{1-\boldsymbol{q}(0.5,1)}{\boldsymbol{q}(0.5,1)}$, iff $\{\mathcal{L}_{R_2}\}$ $K_{(6,4)}^{(10,0)}(u) = \frac{1-\boldsymbol{q}(0.5,1)}{\boldsymbol{q}(0.5,1)}$ iff $\{\mathcal{L}_B,\mathcal{L}_{R_2}\}$ $K_{(6,4)}^{(10,0)}(u) < \frac{1-\boldsymbol{q}(0.5,1)}{\boldsymbol{q}(0.5,1)}$ iff $\{\mathcal{L}_B\}$
Time_R$_i$	$t_{\mathcal{L}_{R_1}} > t_{\mathcal{L}_B}$	$t_{\mathcal{L}_{R_2}} > t_{\mathcal{L}_B}$
	$t_{\mathcal{L}_{R_2}} > t_{\mathcal{L}_{R_1}}$	

*Assumption(s):

(i) Successful elicitation of $p(0.5, 1)$

(ii) For $(\mathcal{L}_B, \mathcal{L}_{R_1})$, the elicited u must satisfy $u(10) \geqslant u(8) \geqslant u(4) > u(2)$

(iii) For $(\mathcal{L}_B, \mathcal{L}_{R_2})$, the elicited u must satisfy $1 \geqslant u(8) \geqslant u(4) > 0$

Table 9: Ambiguity contracts - Predictions

	$(\mathcal{L}_B, \mathcal{L}_{A_1})$	$(\mathcal{L}_B, \mathcal{L}_{A_2})$
Choice_A$_i^*$	$\alpha_{[0.4,0.6]} > 0.5$, iff $\{\mathcal{L}_{A_1}\}$ $\alpha_{[0.4,0.6]} = 0.5$ iff $\{\mathcal{L}_B,\mathcal{L}_{A_1}\}$ $\alpha_{[0.4,0.6]} < 0.5$ iff $\{\mathcal{L}_B\}$	$\alpha_{[0,1]} > 0.5$, iff $\{\mathcal{L}_{A_2}\}$ $\alpha_{[0,1]} = 0.5$ iff $\{\mathcal{L}_B,\mathcal{L}_{A_2}\}$ $\alpha_{[0,1]} < 0.5$ iff $\{\mathcal{L}_B\}$
Time_A$_i$	—	—
	—	

*Assumption(s):

(i) Successful elicitation of $\boldsymbol{p}(x, 1)$, for $x \in \{0.2, 0.5, 0.8\}$

(ii) $\boldsymbol{p}(0.8, 1) > \boldsymbol{p}(0.5, 1) > \boldsymbol{p}(0.2, 1)$

Table 10: Dominated contracts - Predictions

	$(\mathcal{L}_B, \mathcal{L}_{R_1})$	$(\mathcal{L}_B, \mathcal{L}_{R_2})$
Choice_D$_i$	$\{\mathcal{L}_B\}$	$\{\mathcal{L}_B\}$
Time_D$_i$	$t_{\mathcal{L}_{R_1}} < t_{\mathcal{L}_B}$	$t_{\mathcal{L}_{R_2}} > t_{\mathcal{L}_B}$
	$t_{\mathcal{L}_{R_2}} > t_{\mathcal{L}_{R_1}}$	

A.6 Robustness of Time Predictions

The model's time predictions about confidence and risk contracts would still hold for different preferences over contracts. In fact, assume that subjects evaluate contracts using the formula

$$U(\mathcal{L}|t) = \rho(\alpha, t)V(L_{\text{Solve}}) + (1 - \rho(\alpha, t))V(L_{\text{NotSolve}}),$$

where $\rho(\alpha, t) : [0, 1] \times \mathbb{R}_+ \to [0, 1]$ should be interpreted as decision weight, V is a functional over lotteries, and L_{Solve} and L_{NotSolve} are the lotteries one gets if one solves or does not solve the task associated to $\mathcal{L}$. Assume that $\rho(\alpha, \cdot)$ is increasing, concave, and continuously differentiable. Taking first-order conditions of the time allocation problem

$$\max_{t \in [0,T]} \left[U(\mathcal{L}_B|t) + U(\mathcal{L}_X|T - t) \right]. \tag{6}$$

and assuming that $\alpha_1 = \alpha_2 =: \alpha$, we get to

$$\frac{\rho'(\alpha, T - t_{\mathcal{L}_B})}{\rho'(\alpha, t_{\mathcal{L}_B})} = \frac{V(L^B_{\text{Solve}}) - V(L^B_{\text{NotSolve}})}{V(L^X_{\text{Solve}}) - V(L^X_{\text{NotSolve}})}.$$

Assume first that, for some super-modular function $f : [0, 1] \times \mathbb{R} \to \mathbb{R}$, we have, for any lottery $(p, H; 1 - p, L)$,

$$V((p, H; 1 - p, L)) = f(p, H) + f(1 - p, L).$$

Recall that in the baseline, confidence, and risk contracts, $p_F = 1 - p_S$. Therefore, for confidence contracts,

$$\frac{V(L^B_{\text{Solve}}) - V(L^B_{\text{NotSolve}})}{V(L^C_{\text{Solve}}) - V(L^C_{\text{NotSolve}})} = \frac{[f(p^B_S, H) - f(p^B_S, L)] - [f(1 - p^B_S, H) - f(1 - p^B_S, L)]}{[f(p^C_S, H) - f(p^C_S, L)] - [f(1 - p^C_S, H) - f(1 - p^C_S, L)]},$$

and, for risk contracts,

$$\frac{V(L^B_{\text{Solve}}) - V(L^B_{\text{NotSolve}})}{V(L^R_{\text{Solve}}) - V(L^R_{\text{NotSolve}})} = \frac{[f(p_S, H_R) - f(1 - p_S, H_R)] - [f(p_S, L_R) - f(1 - p_S, L_R)]}{[f(p_S, H_B) - f(1 - p_S, H_B)] - [f(p_S, L_B) - f(1 - p_S, L_B)]}.$$

If f is super-modular, i.e., for every $p_1, p_2 \in [0, 1]$ and $x_1, x_2 \in \mathbb{R}$,

$$f(\max\{p_1, p_2\}, \max\{x_1, x_2\}) - f(p_1, x_1) \geqslant f(p_2, x_2) - f(\min\{p_1, p_2\}, \min\{x_1, x_2\}),$$

and this inequality is strict whenever (p_1, x_1) and (p_2, x_2) are not ranked by the component-wise ordering of $\mathbb{R}^2$, then

$$\max\left\{\frac{V(L^B_{\text{Solve}}) - V(L^B_{\text{NotSolve}})}{V(L^C_{\text{Solve}}) - V(L^C_{\text{NotSolve}})}, \frac{V(L^B_{\text{Solve}}) - V(L^B_{\text{NotSolve}})}{V(L^R_{\text{Solve}}) - V(L^R_{\text{NotSolve}})}\right\} < 1,$$

and, hence, $t_{\mathcal{L}_B} < T/2$.

As an illustration, assume that $f(p, x) = w(p)u(x)$, where $w : [0, 1] \to [0, 1]$ and $u : \mathbb{R} \to \mathbb{R}$ are strictly increasing. We can think of w as a probability weighting function. Then, f is super-modular. Therefore, time predictions about risk and confidence contracts are robust to probability weighting.

Time predictions about confidence contracts also continue to hold provided that V is (strictly) increasing with respect to first order stochastic dominance, because if $\succsim_{FOSD}$ is the ranking of lotteries according to first order stochastic dominance, then

$$L^C_{\text{Solve}} \succ_{FOSD} L^B_{\text{Solve}} \succ_{FOSD} L^B_{\text{NotSolve}} \succ_{FOSD} L^c_{\text{NotSolve}}$$

Therefore, if V is strictly increasing with respect to first order stochastic dominance,

$$\frac{V(L^B_{\text{Solve}}) - V(L^B_{\text{NotSolve}})}{V(L^C_{\text{Solve}}) - V(L^C_{\text{NotSolve}})} < 1,$$

and, hence, $t_{\mathcal{L}_B} < T/2$.

B Further details on the experimental design

B.1 Maze selection

We employed an algorithm to generate mazes, focusing primarily what are referred to as "perfect" mazes—those characterized by having precisely one solution. The key parameters governing maze generation include width, height, Compactness Factor (CF), and Dead End Index (DEI).

The width and height parameters dictate the dimensions of the maze, specifying the number of cells it comprises. Meanwhile, the CF, serves as a metric for assessing maze compactness. A maze with high CF value is indicative of a compact structure where the solution path is relatively short in relation to the overall maze size. In contrast, a low CF value signifies a less compact arrangement, with the solution path traversing a larger portion of the maze.

The DEI quantifies the distribution of dead ends within a maze. A maze with a high DEI features a dispersed arrangement of lengthy dead ends, resulting in a higher DEI value. Conversely, a maze with a low DEI value exhibits a concentrated distribution of shorter dead ends.

All the mazes used in the paper where of size 20 by 20. We then varied the CF and DEI to find mazes for which the competition rates required for our lotteries. A separate group of subjects played number of different mazes with varying CF and DEI, and they were simply paid for all the mazes they solved in a given time frame. We used these data to associate mazes to our contracts. That is, if contract states a completion rate of .5, that contract would have a maze for which 50% of the subjects had solved the maze by 60 seconds.

Figure 6 presents a sample screen for completing a maze. The subjects had to navigate the blue square to the green dot by using keyboards up, down, left and right arrow keys.

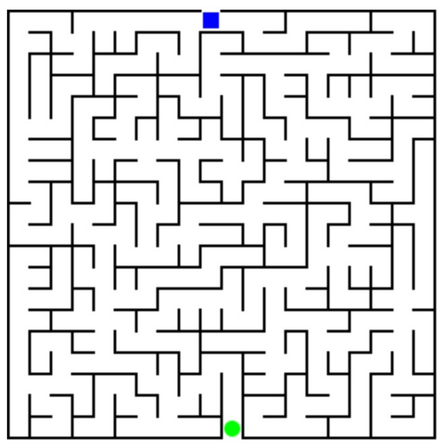

Figure 6: Sample screen of a maze

B.2 Details of Part III: The Characteristics and their Measurement

In Part III of the experiment, we elicit the characteristics of subjects we need to test the predictions of our model: confidence, risk, and ambiguity attitudes. We also elicited other characteristics of our subjects.

B.3 Task 1: Confidence Attitude

The agent's belief in his probability of success in a task associated with a contract, i.e., the function $\boldsymbol{p}$, plays a key role in our analysis. We now show that it can be used to define a notion of (comparative) over-confidence.

Suppose that a subject is told that 50% of people that tried were able to solve the task in at most t minutes. An overconfident subject would then think: "Since I am usually better than others on this task, if the average person solves the task with 50% probability, my probability of solving the task is more than 50%." An under-confident subject, on the other hand, would think: "Since I am usually worse than others on this task, if the average person solves the task with 50% probability, my probability of solving the task is less than 50%." The next definition formalizes this intuition.

Definition 4 *Given $t \geqslant 0$ and a completion rate α, we say that a agent is **over-confident** at (α, t) if $\boldsymbol{p}(\alpha, t) > \alpha$; **confident-neutral** if $\boldsymbol{p}(\alpha, t) = \alpha$; and **under-confident** if $\boldsymbol{p}(\alpha, t) < \alpha$.*

Given $(\alpha, t) \in [0, 1] \times \mathbb{R}_+$, we can elicit $\boldsymbol{p}(\alpha, t)$ as follows. Consider the personal contract $\mathcal{L} = (H, L, 1, 0, \alpha, t)$. The value of this contract is

$$U(\mathcal{L}) = \boldsymbol{p}(\alpha, t)u(H) + (1 - \boldsymbol{p}(\alpha, t))\, u(L).$$

For each $p \in [0, 1]$, define the contract $\mathcal{L}_p$ that pays the lottery $L_p := (p, H; 1 - p, L)$ for sure and and note that

$$U(\mathcal{L}_p) = pu(H) + (1 - p)u(L).$$

Assuming that u is strictly increasing, there exists a unique value of $p^\star \in [0, 1]$ such that $U(\mathcal{L}) = U(\mathcal{L}_{p^\star})$. We then have that $U(\mathcal{L}_p) \geqslant U(\mathcal{L})$ if, and only if, $p \geqslant p^\star$.

We can thus bound $\boldsymbol{p}(\alpha, t)$ through a multiple price list [INSERT REFERENCE]. More specifically, we use a multiple price list with 11 lines, in which on every line $\ell \in \{0, \ldots, 10\}$, we ask the subject to choose between the $\mathcal{L}$ and the contract $\mathcal{L}_{0.1\ell}$. If at line $\ell^\star \in \{0, \ldots, 10\}$, the subject switches from the choice of $\mathcal{L}$ to the choice of $\mathcal{L}_{0.1\ell^\star}$, then we know that

$$p(\alpha, t) \in (0.1(\ell^\star - 1), 0.1\ell^\star].$$

B.4 Task 2: Ambiguity Attitude

Recall that given an ambiguous contract $\mathcal{L}_A$ relative to $\mathcal{L}_B$, where $\alpha_A \in [\alpha_B - \varepsilon, \alpha_B + \varepsilon]$, we assume that the agent first resolves the ambiguity with respect to α_A by substituting it by

$$\alpha_\varepsilon := \gamma_\varepsilon(\alpha_A + \varepsilon) + (1 - \gamma_\varepsilon)(\alpha_A - \varepsilon),$$

for some weight γ_ε. Note that we can interpret γ_ε as capturing the agent's attitude to ambiguity. We say that the agent is *ambiguity-averse* if $\gamma_\varepsilon < \frac{1}{2}$; *ambiguity-neutral* if $\gamma_\varepsilon = \frac{1}{2}$; and *ambiguity-loving* if $\gamma_\varepsilon > \frac{1}{2}$.

In the experiment, we use a Multiple Price List to elicit bounds for α_ε . We use a straightforward modification of the procedure described above to elicit the confidence function $\boldsymbol{p}(\alpha, t)$.

B.5 Task 3: Risk Attitude

We define the risk attitude by the agent's preferences chooses between contracts and their expected values. Formally, given any contract $F \in \Delta(\mathbb{R})$, its expected value is given by $\mathbb{E}_F(x)$. We say that the agent is:[19]

1. *Risk averse* if, for every $F \in \Delta(\mathbb{R})$, $U(\mathbb{E}_F(x)) \geqslant U(F)$;
2. *Risk neutral* if, for every $F \in \Delta(\mathbb{R})$, $U(\mathbb{E}_F(x)) = U(F)$;
3. *Risk seeking* if, for every $F \in \Delta(\mathbb{R})$, $U(\mathbb{E}_F(x)) \leqslant U(F)$.

In our model, the decision-maker's risk attitude is captured by the curvature of her (Bernoulli) utility function u. Therefore, the decision-maker is *risk averse* if, and only if, u is concave; *risk neutral* if, and only if, u is linear; and *risk loving* if, and only if, u is convex.

From now on we focus on a risk averse agent, but clearly the results are easily adapted to the case of a risk neutral or risk seeking agent. Fix $H_1, H_2, L_1, L_2 \in \mathbb{R}$ with $H_2 > H_1 > L_1 > L_2$ and $H_2 - H_1 = L_1 - L_2$. Define

$$K_{(H_1,L_1)}^{(H_2,L_2)}(u) \coloneqq \frac{u(H_2) - u(H_1)}{u(L_1) - u(L_2)}.$$

Since u is strictly increasing, this is well-defined, and we have that $K_{(H_1,L_1)}^{(H_2,L_2)}(u) \leqslant 1$ whenever u is concave. Moreover, given any strictly increasing concave function φ, we have that[20]

$$K_{(H_1,L_1)}^{(H_2,L_2)}(\varphi \circ u) = \frac{\varphi(u(H_2)) - \varphi(u(H_1))}{\varphi(u(L_1)) - \varphi(u(L_2))} \leqslant \frac{u(H_2) - u(H_1)}{u(L_1) - u(L_2)} = K_{(H_1,L_1)}^{(H_2,L_2)}(u).$$

Hence, the more concave u is,[21] the smaller $K_{(H_1,L_1)}^{(H_2,L_2)}(u)$ will be. Hence, $K_{(H_1,L_1)}^{(H_2,L_2)}(u)$ captures

[19] We abuse notation and let $\mathbb{E}_F(x)$ be the personal contract that pays $\mathbb{E}_F(x)$ in all contingencies.

[20] By the concavity of φ, whenever $H_2 > H_1 > L_1 > L_2$, we have that

$$\frac{\varphi(u(H_2)) - \varphi(u(H_1))}{u(H_2) - u(H_1)} \leqslant \frac{\varphi(u(H_1)) - \varphi(u(L_1))}{u(H_1) - u(L_1)} \leqslant \frac{\varphi(u(L_1)) - \varphi(u(L_2))}{u(L_1) - u(L_2)}.$$

[21] Given two strictly increasing and concave functions $u_1, u_2 : \mathbb{R}^n \to \mathbb{R}$, we say that u_1 is *more concave than* u_2 if there exists a concave $\varphi : \mathbb{R} \to \mathbb{R}$ such that $u_1 = \varphi \circ u_2$.

the degree of risk aversion of a risk averse agent. We can then test for risk aversion - and even elicit its intensity - by eliciting u and calculating $K_{(H_1,L_1)}^{(H_2,L_2)}(u)$.

To elicit u, we follow the procedure in.Fix two payoffs $\bar{L}, \bar{H} \in \mathbb{R}$ with $\bar{L} < \bar{H}$, and set $u(\bar{L}) = 0$ and $u(\bar{H}) = 1$. To elicit $x \in (\bar{L}, \bar{H})$, notice that there exists a unique $p_x \in (0,1)$ such that

$$u(x) = U(x) = U(L_x) = p_x u(\bar{H}) + (1 - p_x)u(\bar{L}) = p_x,$$

where $L_x = (p_x, \bar{H}; 1 - p_x, \bar{L})$ Therefore, to elicit $u(x)$, we need to elicit the probability the probability p_x that would make her indifferent between the contract $(p_x, \bar{H}; 1 - p_x, \bar{L})$ and receiving x for sure.

To make the revelation of p_x incentive compatible, we use the BDM mechanism [INSERT CITATION]. We ask subjects to state the value of p_x that would make them indifferent between the contract $(p_x, \bar{H}; 1 - p_x, \bar{L})$ and receiving x for sure. Suppose a subject states p_s. We draw a random number p uniformly from $[0, 1]$. If $p \leqslant p_s$, then the subject wins x for sure. If $p > p_s$, the subject gets to play the contract $(p, \bar{H}; (1 - p), \bar{L})$.

There are two pivotal cases. Suppose first that $p_s < p \leqslant p_x$. The subject then plays the contract $(p, \bar{H}; (1-p), \bar{L})$ which is weakly worse than what he would get if he stated p_x, namely x. Suppose now that $p_x \leqslant p < p_s$. Then, the subject receives x for sure, which is weakly worse what he would get by stating p_x, namely the contract $(p, \bar{H}; (1 - p), \bar{L})$. Therefore, stating p_x is a weakly dominant strategy for the subject.

B.6 Tasks 4 to 8: curvature, risk, over-placement

Task 4 was included to provide us with an insight into the shape of our subjects' "probability of success function" which describes the probability of solving a maze with difficulty α as a function of time. Such a function is an important part of our theory and we used Task 4 to gain some insight into the shape (concavity or convexity) of such a function. This task is a modified version of the procedure introduced in Avoyan and Romagnoli (2023).

Subjects were told that they would have 3 minutes to solve a maze but had to decide how they would be paid for solving it. There are three options to choose from: Option A, Option B and Option C. In Option A the subjects could earn \$10 if they successfully solved the maze in 2 minutes and \$0 if they did not. In Option B, a fair coin will be flipped and if it lands on Heads the subject will earn \$10 if they have successfully solved the maze in 1 minute and \$0 otherwise. If the coin lands Tails the subject will earn \$10 if they solve the maze in 3 minutes and \$0 otherwise. They could choose Option C if they were indifferent between Option A and

Option B.

In choosing their payment scheme our subjects reveal what they feel is the relationship between their ability to solve a given maze and given the time. They are asked to choose between a payment of \$10 if they solve the maze in 2 minutes and a payment of \$10 if they solve the maze in a convex combination of times (1 minute and 3 minutes). This is comparable to eliciting risk preference by offering a sure payment versus a convex combination of payments with a fixed probability. If a subject's probability of success function is concave then she would prefer A to B and if convex the preference would be reversed. Linear function would elicit an indifferent response.

Task 5 was a risk aversion elicitation task using a price-list procedure from Holt and Laury (2002). *Task 6* is a measure of subjects over precision (Moore and Healy (2008)), which elicits a subject's belief about how sure she is about the truth about a given objective fact such as the distance of the moon from the earth. Subjects were asked how far the moon was from the earth and were paid by how close their answer was to the truth. They were also asked what percentage of subjects their answer was closer to the truth than. In other words, if they said 75% then they believed their answer was closer to the true distance than 75% of other subjects. They were rewarded for the accuracy of this guess.

Task 7 uses the task developed by Gneezy and Potters (1997) and Charness and Gneezy (2010) to elicit a subject's risk aversion by asking them to allocate 100 tokens between a safe and risky option. The risky project has a 50 percent chance of success and returns two and a half times the investment if successful, and nothing otherwise.

We included a second risk elicitation task for two reasons. First, in certain environments it better correlated to subject behavior than the Holt and Laury (2002) price-list task. Second, and more importantly, this task is necessary when combined with Task 8, to allow us to test for compound lottery aversion. *Task 8* repeats Task 7 but this time making the risky contract a two-stage compound contract in an attempt to get an understanding for our subjects' attitude toward compound lotteries (this task is introduced in Agranov and Ortoleva (2017)). Subjects with neither aversion nor attraction to compound lotteries should invest the same number of tokens in both tasks. The subjects who dislike (like) compound lotteries will invest less (more) in Task 8.