

When Humans Aren't Optimal: Robots that Collaborate with Risk-Aware Humans

Minae Kwon*, Erdem Biyik*, Aditi Talati†, Karan Bhasin‡, Dylan P. Losey*, Dorsa Sadigh*

* Stanford University, † American High School, ‡ The Harker School

[mnkwon, ebiyik, dlosey, dorsa]@stanford.edu, talatiaditi02@gmail.com, karanbhasin03@gmail.com

Abstract

In order to collaborate safely and efficiently, robots need to anticipate how their human partners will behave. Some of today's robots model humans as if they were also robots, and assume users are always *optimal*. Other robots account for human limitations, and relax this assumption so that the human is *noisily* rational. Both of these models make sense when the human receives deterministic rewards: i.e., gaining either \$100 or \$130 with certainty. But in real-world scenarios, rewards are rarely deterministic. Instead, we must make choices subject to *risk* and *uncertainty*—and in these settings, humans exhibit a cognitive bias towards *suboptimal* behavior. For example, when deciding between gaining \$100 with certainty or \$130 only 80% of the time, people tend to make the risk-averse choice—even though it leads to a lower expected gain! In this paper, we adopt a well-known *Risk-Aware* human model from behavioral economics called Cumulative Prospect Theory and enable robots to leverage this model during human-robot interaction (HRI). In our user studies, we offer supporting evidence that the Risk-Aware model more accurately predicts suboptimal human behavior. We find that this increased modeling accuracy results in safer and more efficient human-robot collaboration. Overall, we extend existing rational human models so that collaborative robots can anticipate and plan around suboptimal human behavior during HRI.

CCS Concepts

• **Mathematics of computing** → Probabilistic inference problems; • **Computing methodologies** → Cognitive robotics; Theory of mind.

Keywords

cognitive HRI; cumulative prospect theory; human prediction

ACM Reference Format:

Minae Kwon*, Erdem Biyik*, Aditi Talati†, Karan Bhasin‡, Dylan P. Losey*, Dorsa Sadigh*. 2020. When Humans Aren't Optimal: Robots that Collaborate with Risk-Aware Humans. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction (HRI '20)*, March 23–26, 2020, Cambridge, United Kingdom. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3319502.3374832>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.

HRI '20, March 23–26, 2020, Cambridge, United Kingdom

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6746-2/20/03...\$15.00

<https://doi.org/10.1145/3319502.3374832>

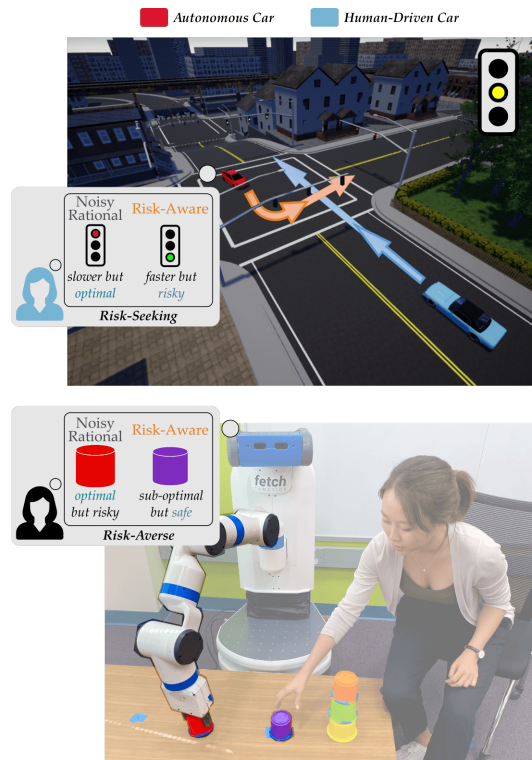


Figure 1: Robots collaborating with humans that must deal with risk. (Top) Autonomous car predicting if the human will try to make the light. (Bottom) Robot arm anticipating which cup the human will grab. In real world scenarios, people exhibit a cognitive bias towards irrational but Risk-Aware behavior.

1 Introduction

When robots collaborate with humans, they must *anticipate* how the human will behave for seamless and safe interaction. Consider a scenario where an autonomous car is waiting at an intersection (see top of Fig. 1). The autonomous car wants to make an unprotected left turn, but a human driven car is approaching in the oncoming lane. The human's traffic light is yellow, and will soon turn red. Should the autonomous car predict that this human will stop—so that the autonomous car can safely turn left—or anticipate that the human will try and make the light—where turning left leads to a collision?

Previous robots anticipated that humans acted like robots, and made *rational* decisions to maximize their reward [1, 18, 25, 36, 44, 45]. However, assuming humans are always rational fails to account for the limited time, computational resources, and noise

that affect human decision making, and so today’s robots anticipate that humans make *noisily* rational choices [13, 17, 32, 37, 47]. Under this model, the human is always most likely to choose the action leading to the highest reward, but the robot also recognizes that the human may behave suboptimally. This makes sense when humans are faced with deterministic rewards: e.g., the light will definitely turn red in 5 seconds. Here, the human knows whether or not they will make the light, and can accelerate or decelerate accordingly. But in real world settings, we usually do not have access to deterministic rewards. Instead, we need to deal with uncertainty and estimate *risk* in every scenario. Returning to our example, imagine that the human has a 95% chance of making the light if they accelerate: success saves some time during their commute, while failure could result in a ticket or even a collision. It is still rational for the human to decelerate; however, a risk-seeking user will attempt to make the light. How the robot models the human affects the *safety* and *efficiency* of this interaction: a Noisy Rational robot believes it should turn left, while a *Risk-Aware* robot realizes that the human is likely to run the light, and waits to prevent a collision.

When robots treat nearby humans as noisily rational, they *miss out* on how risk biases human decisions. Instead, we assert

To ensure safe and efficient interaction, robots must recognize that people behave suboptimally when risk is involved.

Our approach is inspired by behavioral economics, where results indicate that users maintain a nonlinear transformation between *actual* and *perceived* rewards and probabilities [20, 42]. Here, the human over- or under-weights differences between rewards, resulting in a cognitive bias (a systematic error in judgment) that leads to risk-averse or risk-seeking behavior. We equip robots with this cognitive model, enabling them to anticipate risk-affected human behavior and better collaborate with humans in everyday scenarios (see Fig. 1).

Overall, we make the following contributions:

Incorporating Risk in Robot Models of Humans. We propose using Cumulative Prospect Theory as a Risk-Aware model. We formalize a theory-of-mind (ToM) where the robot models the human as reacting to their decisions or environmental conditions. We integrate Cumulative Prospect Theory into this formalism so that the robot can model suboptimal human actions under risk. In a simulated autonomous driving environment, our user studies demonstrate that the Risk-Aware robot more accurately predicts the human’s behavior than a Noisy Rational baseline.

Determining when to Reason about Risk. We identify the types of scenarios where reasoning about risk is important. Our results suggest that scenarios with close expected rewards is the most important in determining whether humans will act suboptimally.

Safe and Efficient Collaboration when Humans face Uncertainty. We develop planning algorithms so that robots can leverage our Risk-Aware human model to improve collaboration. In a collaborative cup stacking task, shown on the bottom in Fig. 1, the Risk-Aware robotic arm anticipated that participants would choose suboptimal but risk-averse actions, and planned trajectories to avoid interfering with the human’s motions. Users completed the task more efficiently with the Risk-Aware robot, and also subjectively preferred working with the Risk-Aware robot over the Noisy Rational baseline.

This work describes a computationally efficient and empirically supported way for robots to model suboptimal human behavior by extending the state-of-the-art to also account for risk. A summary of our paper, including videos of experiments, can be found *here*.

2 Related Work

Previous work has shown that robots that successfully predict humans’ behavior exhibit improved performance in many applications, such as assistive robotics [2, 14, 22, 23], motion planning [27, 48], collaborative games [26], and autonomous driving [3, 38, 39]. One reason behind this success is that human modeling equips robots with a theory of mind (ToM), or the ability to attribute a mind to oneself and others [34, 41]. Devin and Alami [11] showed ToM can improve performance in human-robot collaboration.

For this purpose, researchers have developed various human models. In robotics, the Noisy Rational choice model has remained extremely popular due to its simplicity. Several works in reward learning [5–7, 9, 32, 37], reinforcement learning [17], inverse reinforcement learning [8, 35, 47], inverse planning [4], and human-robot collaboration [33] employed the noisy rational model for human decision-making. Other works developed more complex human models and methods specifically for autonomous driving [18, 21, 38, 44]. Unfortunately, these models either assume humans are rational or do not handle situations with uncertainty and risk. There have been other human models that take a learning-based approach [30, 31, 43]. While this is an interesting direction, these methods are usually not very data efficient.

In cognitive science, psychology and behavioral economics, researchers have developed other decision-making models. For example, Ordonez and Benson III [28] investigated decision making under time constraints; Diederich [12] developed a model based on stochastic processes to model humans’ process of making a selection between two options, again under a time constraint. Ortega and Stocker [29] proposed a rationality model based on concepts from information theory and statistical mechanics to model time-constrained decision making. Mishra [24] studied decision making under risk from the perspectives of biology, psychology and economics. Halpern et al. [19] modeled the humans as a finite automata, and Simon [40] developed *bounded rationality* to incorporate sub-optimality and constraints. Evans et al. [16] investigated different biases humans can have in decision-making. Among all of these works, Cumulative Prospect Theory (CPT) [42] remains prominent as it successfully models suboptimal human decision making under risk. Later works studied how Cumulative Prospect Theory can be employed for time-constrained decision making [15, 46].

In this paper, we adopt Cumulative Prospect Theory as an example of a Risk-Aware model. We show that it not only leads to more accurate predictions of human actions, but also increases the performance of the robot and the human-robot team.

3 Formalism

We assume a setting where a human needs to select from a set of actions \mathcal{A}_H . Each action $a_H \in \mathcal{A}_H$ may have several possible consequences, where, without loss of generality, we denote the number of consequences as K . For a given human action a_H , we express the probabilities of each consequence and their corresponding rewards

as a set of pairs:

$$C(a_H) = \left\{ \left(p^{(1)}, R_H^{(1)}(a_H) \right), \left(p^{(2)}, R_H^{(2)}(a_H) \right), \dots, \left(p^{(K)}, R_H^{(K)}(a_H) \right) \right\}$$

We outline and compare two methods that use $C(a_H)$ to model human actions: Noisy Rational and Cumulative Prospect Theory (CPT) [42]. CPT is a prominent model of human decision-making under risk [15, 46] and we use it as an example of a Risk-Aware model. Finally, we describe how we can integrate them into a partially observable Markov decision process (POMDP) formulation of human-robot interaction.

Noisy Rational Model. According to the noisy rational model, humans are more likely to choose actions with the highest expected reward, and are less likely to choose suboptimal actions (i.e., they are optimal with some noise). The noise comes from constraints such as limited time or computational resources. For instance, in the autonomous driving example, Noisy Rational model would predict the human will most likely choose the optimal action and decelerate. Denoting the expected reward of the human for action a_H as

$$R_H(a_H) = p^{(1)} R_H^{(1)}(a_H) + p^{(2)} R_H^{(2)}(a_H) + \dots + p^{(K)} R_H^{(K)}(a_H),$$

the noisy rational model asserts

$$P(a_H) = \frac{\exp(\theta \cdot R_H(a_H))}{\sum_{a \in \mathcal{A}_H} \exp(\theta \cdot R_H(a))}, \quad (1)$$

where $\theta \in [0, \infty)$ is a temperature parameter, commonly referred to as the *rationality coefficient*, which controls how noisy the human is. While larger θ models the human as a better reward maximizer, setting $\theta=0$ means the human chooses actions uniformly at random.

Hence, the Noisy Rational model is simply a linear transformation of the reward with the rationality coefficient $\theta \geq 0$, which makes the transformation monotonically non-decreasing. As the model does not transform the probability values, it becomes impossible to model suboptimal humans using this approach. The closest Noisy Rational can get to modeling suboptimal humans is to assign a uniform probability to all actions.

Risk-Aware Model. We adopt Cumulative Prospect Theory (CPT) [42] as an example of a Risk-Aware model. According to this model, humans are not simply Noisy Rational. They may, for example, be suboptimally risk-seeking or risk-averse. For instance, in the autonomous driving example, human drivers can be risk-seeking and try to make the yellow light even though they risk a costly collision. The Risk-Aware model captures suboptimal decision-making by transforming both the probabilities and the rewards. These transformations aim to represent what humans actually perceive. The reward transformation is a pairwise function:

$$v(R) = \begin{cases} R^\alpha & \text{if } R \geq 0 \\ -\lambda(-R)^\beta & \text{if } R < 0 \end{cases}.$$

The parameters $\alpha, \beta \in [0, 1]$ represent how differences among rewards are perceived. For instance, when $\alpha, \beta \in (0, 1)$, the model predicts that humans will perceive differences between large positive (or negative) rewards as relatively lower than the differences between smaller positive (resp. negative) rewards, even though the true differences are equal. $\lambda \in [0, \infty)$ characterizes how much more (or less) important negative rewards are compared to positive rewards. When $\lambda > 1$, humans are modeled as loss-averse, assigning

more importance to losses compared to gains. The reverse is true when $\lambda \in [0, 1)$.

The Risk-Aware model also implements a transformation over the probabilities. The probabilities $(p^{(1)}, p^{(2)}, \dots)$ are divided into two groups based on whether their corresponding true rewards are positive or negative. The probability transformations corresponding to positive and negative rewards (w^+, w^-) are as follows:

$$w^+(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}}, \quad w^-(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{1/\delta}},$$

where $\gamma, \delta \in [0, 1]$.

Without loss of generality, we assume that each of the K rewards are ordered in decreasing order, i.e. $R_H^{(i+1)}(a) \leq R_H^{(i)}(a)$ for all $i \in \{1, 2, \dots, K-1\}$ and $a \in \mathcal{A}_H$. Then, the probability transformation is as follows:

$$\begin{aligned} \pi(C(a_H)) &= (\pi^+(C(a_H)), \pi^-(C(a_H))) \\ \pi^+(C(a_H)) &= \left(w^+(p^{(1)}), w^+(p^{(1)} + p^{(2)}) - w^+(p^{(1)}), \dots \right) \\ \pi^-(C(a_H)) &= \left(\dots, w^-(p^{(K)} + p^{(K-1)}) - w^-(p^{(K)}), w^-(p^{(K)}) \right) \end{aligned}$$

Finally, we normalize probabilities so that $\pi(C(a_H))$ sums to 1:

$$\bar{\pi}_j(C(a_H)) = \frac{\pi_j(C(a_H))}{\sum_{i=1}^K \pi_i(C(a_H))}, \quad \forall j \in \{1, 2, \dots, K\}$$

When $\gamma, \delta \in (0, 1)$, the probability transformations capture biases humans are reported to have [42] by overweighting smaller probabilities and underweighting larger probabilities.

Based on these two transformations, we now extend the human decision making model with the Risk-Aware model:

$$R_H^{\text{CPT}}(a_H) = \bar{\pi}_1(C(a_H)) \cdot v\left(R_H^{(1)}(a_H)\right) + \dots + \bar{\pi}_K(C(a_H)) \cdot v\left(R_H^{(K)}(a_H)\right)$$

$$P(a_H) = \frac{\exp\left(\theta \cdot R_H^{\text{CPT}}(a_H)\right)}{\sum_{a \in \mathcal{A}_H} \exp\left(\theta \cdot R_H^{\text{CPT}}(a)\right)}. \quad (2)$$

In contrast to the Noisy Rational model, the Risk-Aware model's expressiveness allows it to model *both* optimal and suboptimal human decisions by assigning larger likelihoods to those actions.

Formal Model of Interaction. We model the world where both the human and the robot take actions as a POMDP, which we denote with a tuple $\langle \mathcal{S}, \mathcal{O}, \mathcal{O}, \mathcal{A}_H, \mathcal{A}_R, T, r_H, r_R \rangle$. \mathcal{S} is the finite set of states; \mathcal{O} is the set of observations; $\mathcal{O} : \mathcal{S} \rightarrow \mathcal{O}$ defines the shared observation mapping; \mathcal{A}_H and \mathcal{A}_R are the finite action sets for the human and the robot, respectively; $T : \mathcal{S} \times \mathcal{A}_H \times \mathcal{A}_R \times \mathcal{S} \rightarrow [0, 1]$ is the transition distribution. r_H and r_R are the reward functions that depend on the state, the actions and the next state. In this POMDP, we assume the agents act simultaneously. Having a first-order ToM, the human tries to optimize her own cumulative reward given an action distribution for the robot, $P(a_R | s)$. The human value function $V_H(s)$ can then be defined using the following Bellman update:

$$V_H(s) = \max_{a_H \in \mathcal{A}_H} \mathbb{E}_{a_R | s} \mathbb{E}_{s' | s, a_H, a_R} \left[R_H(s, a_H, a_R, s') + \gamma V_H(s') \right].$$

We then use the fact that

$$P(s, s', a_R | o, a_H) = P(s | o) \cdot P(a_R | s) \cdot P(s' | s, a_H, a_R)$$

to construct a set $C(a_H)$ for the current observation o that consists of the pairs $(P(s, s', a_R | o, a_H), V_H(s'))$ for varying s, s' , and a_R .

Table 1: Autonomous Driving. Users were given different amounts of *information* about the likelihood that the light would turn red. Under *risk*, we list two tested probabilities of the light turning red.

<i>Information</i>	<i>Time</i>
None: <i>With some probability the light will turn red.</i>	Timed: 8 s Not Timed: no limit
Explicit: <i>There is a 5% chance the light will turn red.</i>	Risk
Implicit: <i>Of the previous 380 cars that decided to accelerate, the light turned red for 19 cars.</i>	High: 95% Low: 5%

When modeling the human as zeroth-order ToM, $P(a_R | s)$ will simply be a uniform distribution.

Having constructed $C(a_H)$, we can define the human’s utility function for different values of s , and s' . The utility functions for both Noisy Rational and Risk-Aware models are defined as follows: **Noisy Rational:**

$$R_H(o, a_H) = \sum_{s, s' \in \mathcal{S}} \sum_{a_R \in \mathcal{A}_R} P(s, s', a_R | o, a_H) \cdot V_H(s').$$

Risk-Aware:

$$R_H^{CPT}(o, a_H) = \sum_{s, s' \in \mathcal{S}} \sum_{a_R \in \mathcal{A}_R} \tilde{\pi}_i(P(s, s', a_R | o, a_H)) \cdot v(V_H(s')),$$

where the index i corresponds to the event that leads to s' from s with a_R , a_H . An optimal human would always pick the action a_H that maximizes $\mathbb{E}_{s|o} \mathbb{E}_{a_R|s} \mathbb{E}_{s'|s, a_H, a_R} [V_H(s')]$. The robot can obtain $P(a_H|o)$ using Eqn. (1), Eqn. (2) and use it to maximize its own cumulative reward.

Summary. We have outlined two ways in which we can model humans (*Noisy Rational* and *Risk-Aware*), and how we can formalize these models in a human-robot interaction setting. In the following section, we empirically analyze factors that allow the Risk-Aware robot to more accurately model human actions.

4 Autonomous Driving

In our first user study, we focus on the autonomous driving scenario from the bottom of Fig. 1. Here the autonomous car—which wants to make an unprotected left turn—needs to determine whether the human-driven car is going to try to make the light. We asked human drivers whether they would accelerate or stop in this scenario. Specifically, we adjusted the *information* and *time* available for the human driver to make their decision. We also varied the level of *risk* by changing the probability that the light would turn red. Based on the participant’s choices in each of these cases, we learned *Noisy Rational* and *Risk-Aware* human models. Our results demonstrate that autonomous cars that model humans as *Risk-Aware* are better able to explain and anticipate the behavior of human drivers, particularly when drivers make suboptimal choices.

Experimental Setup. We used the driving example shown in Fig. 1. Human drivers were told that they are returning a rental car, and are approaching a light that is currently yellow. If they run the red light, they have to pay a \$500 ticket. But stopping at the light will prevent the human from returning their rental car on time, which also has an associated fine! Accordingly, the human drivers had to

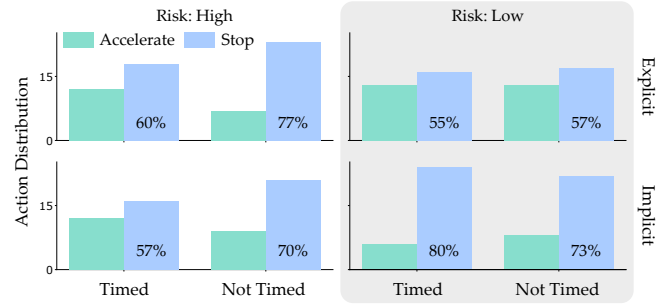


Figure 2: Action distributions for human drivers. Across all surveyed factors (*information*, *time*, and *risk*), more users preferred to stop at the light. Interestingly, stopping was the *suboptimal* choice when the light rarely turns red (*Low*).

decide between *accelerating* (and potentially running the light) or *stopping* (and returning the rental car with a fine).

Independent Variables. We varied the amount of *information* and *time* that the human drivers had to make their decision. We also tested two different *risk* levels: one where accelerating was optimal, and one where stopping was optimal. Our parameters for information, time, and risk are provided in Table 1.

Information. We varied the amount of information that the driver was given on three levels: None, Explicit, and Implicit. Under None, the driver must rely on their own prior to assess the probability that the light will turn red. By contrast, in Explicit we inform the driver of the exact probability. Because probabilities are rarely given to us in practice, we also tested Implicit, where drivers observed other peoples’ experiences to estimate the probability of a red light.

Time. We compared two levels for time: a Timed setting, where drivers had to make their choice in under 8 seconds, and a Not Timed setting, where drivers could deliberate as long as necessary.

Risk. We varied risk along two levels: High and Low. When the risk was High, the light turned red 95% of the time, and when risk was Low, the light turned red only 5% of the time.

Participants and Procedure. We conducted a within-subjects study on Amazon Mechanical Turk and recruited 30 participants. All participants had at least a 95% approval rating and were from the United States. After providing informed consent, participants were first given a high-level description of the autonomous driving task and were shown the example from Fig. 1. In subsequent questions, participants were asked to indicate whether they would accelerate or stop. We presented the Timed questions first and the Not Timed questions second. For each set of Timed and Not Timed questions, we presented questions in the order of their informativeness from None to Explicit. The risk levels were presented in random order¹.

Dependent Measures. We aggregated the user responses into *action distributions*. These action distributions report the percentage of human drivers who chose to accelerate and stop under each treatment level. Next, we learned *Noisy Rational* and *Risk-Aware* models of human drivers for the autonomous car to leverage². To

¹To learn more about our study, please check out our autonomous driving survey link: https://stanfordgsb.qualtrics.com/jfe/form/SV_cUgxaZlvEkdb3ud

²To learn the models, we used the Metropolis-Hastings algorithm [10] and obtained 30 independent samples of the model parameters.

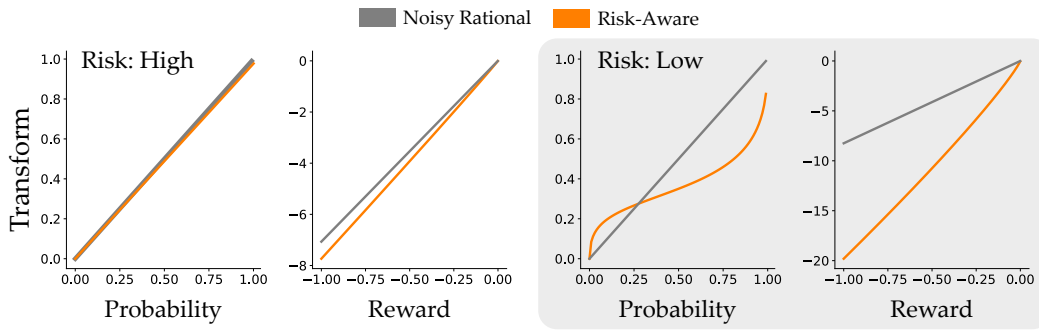


Figure 3: Averaged probability and reward transformations for human drivers that are modeled as *Noisy Rational* or *Risk-Aware*. In scenarios where the light frequently turns red (High), both models produce similar transformations. But when the light rarely turns red (Low), the models diverge: here the *Risk-Aware* autonomous car recognizes that human drivers *overestimate* both the probability that light will turn red and the cost of running the light. This enables *Risk-Aware* autonomous cars to explain why human drivers prefer to stop, even though accelerating is the optimal action when the light rarely turns red.

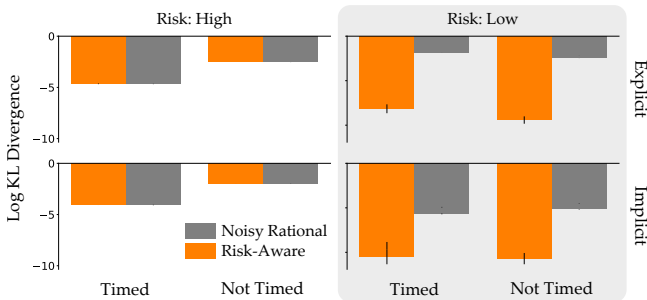


Figure 4: Model accuracy (lower is better). When the light often turns red (High), both models could anticipate the human’s behavior. But when the light rarely turns red (Low), only the *Risk-Aware* autonomous car correctly anticipated that the human would stop.

measure the accuracy of these models, we compared the Kullback-Leibler (KL) divergence between the *true* action distribution and the model’s *predicted* action distribution. We report the *log KL divergence* for both *Noisy Rational* and *Risk-Aware* models.

Hypothesis.

H1. *Autonomous cars which use Risk-Aware models of human drivers will more accurately predict human action distributions than autonomous cars who treat humans as noisily rational agents.*

Baseline. In order to confirm that our users were trying to make optimal choices, we also queried the human drivers for their preferred actions in settings where the expected rewards were *far apart* (e.g., where the expected reward for accelerating was much higher than the expected reward for stopping). In these baseline trials, users overwhelmingly chose the *optimal* action (93% of trials).

Results. The results from our autonomous driving user study are summarized in Figs. 2, 3, and 4. In each of the tested situations, most users elected to stop at the light (see Fig. 2). Although stopping at the light is the optimal action in the High risk case—where the light turns red 95% of the time—stopping was actually *suboptimal* in the Low risk case—where the light only turns red 5% of the time. Because humans chose optimal actions in some cases (High risk) and suboptimal actions in other situations (Low risk), the

autonomous car interacting with these human drivers must be able to anticipate *both* optimal and suboptimal behavior.

In cases where the human was rational, autonomous cars learned similar *Noisy Rational* and *Risk-Aware* models (see Fig. 3). However, the *Risk-Aware* model was noticeably different in situations where the human was suboptimal. Here autonomous cars using our formalism learned that human drivers *overestimated* the likelihood that the light would turn red, and *underestimated* the reward of running the light. Viewed together, the *Risk-Aware* model suggests that human drivers were risk-averse when the light rarely turned red, and risk-neutral when the light frequently turned red.

Autonomous cars using our *Risk-Averse* model of human drivers were better able to predict how humans would behave (see Fig. 4). Across all treatment levels, *Risk-Averse* attained a log KL divergence of -5.7 ± 3.3 , while *Noisy Rational* only reached -3.3 ± 1.3 . This difference was statistically significant ($t(239) = -11.5, p < .001$). Breaking our results down by risk, in the High case both models were similarly accurate, and any differences were insignificant ($t(119) = .42, p = .67$). But in the Low case—where human drivers were suboptimal—the *Risk-Averse* model significantly outperformed the *Noisy Rational* baseline ($t(119) = -17.3, p < .001$).

Overall, the results from our autonomous driving user study support hypothesis H1. Autonomous cars leveraging a *Risk-Aware* model were able to understand and anticipate human drivers both in situations where the human is optimal *or* suboptimal, while the *Noisy Rational* model could not explain why the participants preferred to take a safer (but suboptimal) action.

Follow-up: Disentangling Risk and Suboptimal Decisions. After completing our user study, we performed a simulated experiment within the autonomous driving domain. Within this experiment, we *fixed* the probability that the light would turn red, and then *varied* the human driver’s action distribution. When fixing the probability, we used the High risk scenario where the optimal decision was to *stop*. The purpose of this follow-up experiment was to make sure that our model can also explain suboptimally *aggressive* drivers, and to ensure that our results are not tied to the Low risk scenario. Our simulated results are displayed in Fig. 5. As before, when the human driver chose the optimal action, both *Noisy Rational* and *Risk-Aware* models were equally accurate. But when

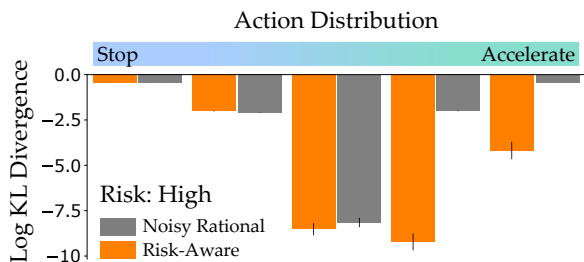


Figure 5: Model accuracy (lower is better) on simulated data. We simulated a spectrum of human action distributions in scenarios where the light often turns red (High risk). The optimal action here is to stop. As the human becomes increasingly optimal, both *Noisy Rational* and *Risk-Aware* provide a similarly accurate prediction. But when the human is suboptimal—accelerating through the light—the *Risk-Aware* autonomous car yields a more accurate prediction.

the human behaved aggressively—and tried to make the light—only the *Risk-Aware* autonomous car could anticipate their suboptimal behavior. These results suggest that the improved accuracy of the *Risk-Aware* model is tied to user *suboptimality*, and not to the particular type of risk (either High or Low).

Summary. We find supporting evidence that *Risk-Aware* is more accurate at modeling human drivers in scenarios that involve decision making under uncertainty. In particular, our results suggest that the reason why *Risk-Aware* is more effective at modeling human drivers is because humans often act *suboptimally* in these scenarios. When humans act rationally, both *Noisy Rational* and *Risk-Aware* autonomous cars can understand and anticipate their actions.

5 Collaborative Cup Stacking

Within the autonomous driving user studies, we demonstrated that our *Risk-Aware* model enables robots to *accurately* anticipate their human partners. Next, we want to explore how our formalism leverages this accuracy to improve *safety* and *efficiency* during HRI. To test the usefulness of our model, we performed two user studies with a 7-DoF robotic arm (Fetch, Fetch Robotics). In an **online** user study, we verify that the *Risk-Aware* model can accurately model humans in a collaborative setting. In an **in-person** user study, the robot leverages *Risk-Aware* and *Noisy Rational* models to anticipate human choices and plan trajectories that avoid interfering with the participant. Both studies share a common experimental setup, where the human and robot collaborate to stack cups into a tower.

Experimental Setup. The collaborative cup stacking task is shown in Fig. 1 (also see the supplemental video). We placed five cups on the table between the person and robot. The robot knew the location and size of the cups *a priori*, and had learned motions to pick up and place these cups into a tower. However, the robot did not know which cups its human partner would pick up.

The human chooses their cups with two potential towers in mind: an *efficient but unstable* tower, which was more likely to fall, or a *inefficient but stable* tower, which required more effort to assemble. Users were awarded 20 points for building the stable tower (which never fell) and 105 for building the unstable tower (which collapsed 80% of the time). Because the expected utility of building the unstable tower was higher, our *Noisy Rational* baseline anticipated that participants would make the unstable tower.

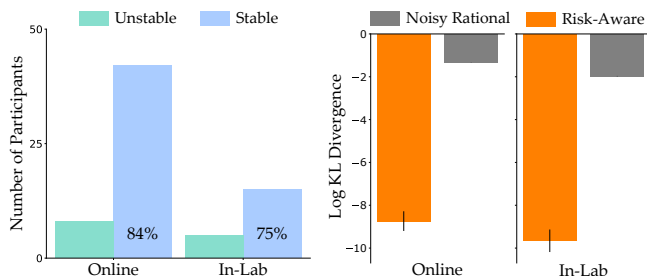


Figure 6: Results from online and in-person user studies during the collaborative cup stacking task. (Left) Although building the unstable tower was *optimal*, more participants selected the stable tower. (Right) Model accuracy, where lower is better. The *Risk-Aware* robot was better able to predict which cups the human would pick up.

Independent Variables. We varied the robot’s model of its human partner with two levels: *Noisy Rational* and *Risk-Aware*. The *Risk-Aware* robot uses our formalism from Section 3 to anticipate how humans make decisions under uncertainty and risk.

5.1 Anticipating Collaborative Human Actions

Our **online** user study extended the results from the autonomous driving domain to this collaborative cup stacking task. We focused on how *accurately* the robot anticipated the participants’ choices.

Participants and Procedure. We recruited 14 Stanford affiliates and 36 Amazon Mechanical Turkers for a total of 50 users (32% Female, median age: 33). Participants from Amazon Mechanical Turk had at least a 95% approval rating and were from the United States. After providing informed consent, each of our users answered survey questions about whether they would collaborate with the robot to build the *efficient but unstable* tower, or the *inefficient but stable* tower. Before users made their choice, we explicitly provided the rewards associated with each tower, and implicitly gave the probability of the tower collapsing. To implicitly convey the probabilities, we showed videos of humans working with the robot to make stable and unstable towers: all five videos with the stable tower showed successful trials, while only one of the five videos with the unstable tower displayed success. After watching these videos and considering the rewards, participants chose their preferred tower type ³.

Dependent Measures. We aggregated the participants’ decisions to find their *action distribution* over stable and unstable towers. We fit *Noisy Rational* and *Risk-Aware* models to this action distribution, and reported the *log KL divergence* between the actual tower choices and the choices predicted by the models.

Hypotheses.

H2. *Risk-Aware robots will better anticipate which tower the collaborative human user is attempting to build.*

Results. Our results from the **online** user study are summarized in Fig. 6. During this scenario—where the human is collaborating with the robot—we observed a bias towards risk-averse behavior. Participants overwhelmingly preferred to build the stable tower (and take the guaranteed reward), even though this choice was suboptimal. Only the *Risk-Aware* robot was able to capture and

³To learn more about our study please check out our cup stacking survey link: https://stanfordgsb.qualtrics.com/jfe/form/SV_0oz1Y04mQ0s3i7P

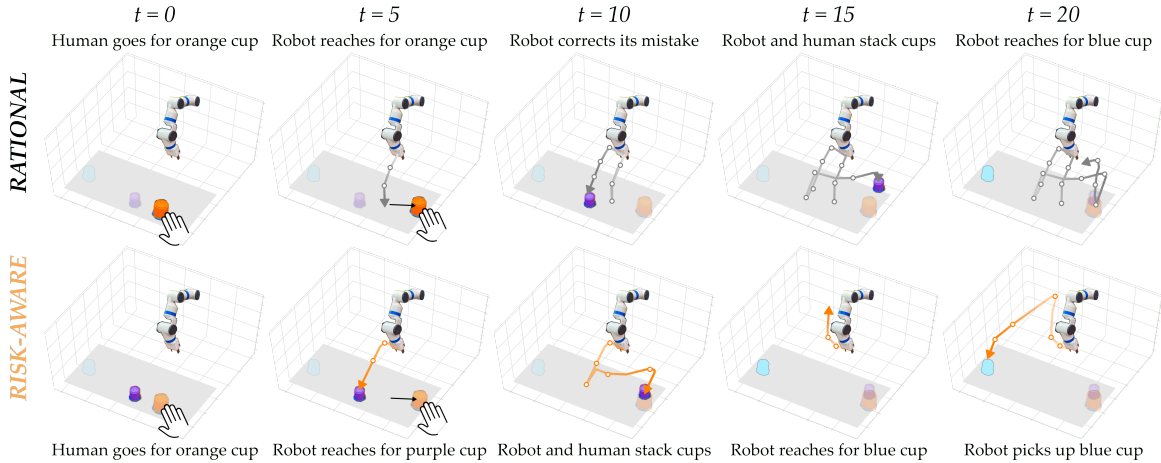


Figure 7: Example robot and human behavior during the collaborative cup stacking user study. At the start of the task, the human reaches for the orange cup (the first step towards a stable tower). When the robot models the human as a *Noisy Rational* partner (top row), it incorrectly anticipates that the human will build the optimal but unstable tower; this leads to interference, replanning, and a delay. The robot leveraging our *Risk-Aware* formalism (bottom row) understands that real decisions are influenced by uncertainty and risk, and correctly predicts that the human wants to build a stable tower. This results in safer and more efficient interaction, leading to faster tower construction.

predict this behavior: inspecting the right side of Fig. 6, we found a statistically significant improvement in model *accuracy* across the board ($t(59) = -21.1, p < .001$). Focusing only on the **online** users, the *log KL divergence* for *Risk-Aware* reached -8.7 ± 3.0 , while *Noisy Rational* remained at -1.3 ± 0.01 ($t(29) = -13.1, p < .001$). Overall, these results match our findings from the autonomous driving domain, and support hypothesis H2.

5.2 Planning with Risk-Aware Human Models

Having established that the *Risk-Aware* robot more accurately models the human’s actions, we next explored whether this difference is *meaningful* in practice. We performed an **in-lab** user study comparing *Noisy Rational* and *Risk-Aware* collaborative robots. We focused on how robots can leverage the *Risk-Aware* human model to improve *safety* and *efficiency* during collaboration.

Participants and Procedure. Ten members of the Stanford University community (2 female, ages 20 – 36) provided informed consent and participated in this study. Six of these ten had prior experience interacting with the Fetch robot. We used the same experimental setup, rewards, and probabilities described at the beginning of the section. Participants were encouraged to build towers to maximize the total number of points that they earned.

Each participant had ten familiarization trials to practice building towers with the robot. During these trials, users learned about the probabilities of each type of tower collapsing from experience. In half of the familiarization trials, the robot modeled the human with the *Noisy Rational* model, and in the rest the robot used the *Risk-Aware* model; we randomly interspersed trials with each model. After the ten familiarization trials, users built the tower once with *Noisy Rational* and once with *Risk-Aware*: we recorded their choices and the robot’s performance during these final trials. The presentation order for these final two trials was counterbalanced.

Dependent Measures. To test *efficiency*, we measured the time taken to build the tower (Completion Time). We also recorded the

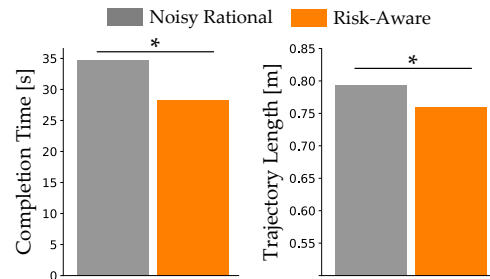


Figure 8: Objective results from our in-lab user study. When participants built the tower with the *Risk-Aware* robot, they completed the task more efficiently (lower Completion Time) and safely (lower Trajectory Length). Asterisks denote significance ($p < .05$).

Cartesian distance that the robot’s end-effector moved during the task (Trajectory Length). Because the robot had to replan longer trajectories when it interfered with the human, Trajectory Length was an indicator of *safety*.

After participants completed the task with each type of robot (*Noisy Rational* and *Risk-Aware*) we administered a 7-point Likert scale survey. Questions on the survey focused on four scales: how enjoyable the interaction was (Enjoy), how well the robot understood human behavior (Understood), how accurately the robot predicted which cups they would stack (Predict), and how efficient users perceived the robot to be (Efficient). We also asked participants which type of robot they would rather work with (Prefer) and which robot better anticipated their behavior (Accurate).

Hypotheses.

H3. *Users interacting with the Risk-Aware robot will complete the task more safely and efficiently.*

H4. *Users will subjectively perceive the Risk-Aware robot as a better partner who accurately predicts their decisions and avoids grabbing their intended cup.*

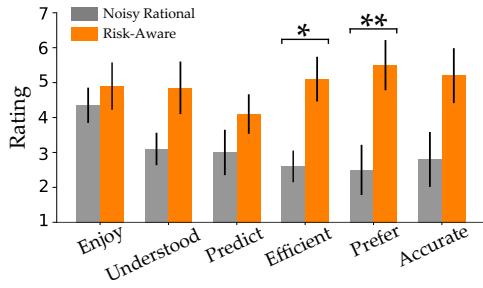


Figure 9: Subjective results from our in-person user study. Higher ratings indicate agreement (i.e., more enjoyable, better understood). Here ** denotes $p < .07$, and * denotes $p < .05$. Participants perceived the *Risk-Aware* robot as a more efficient teammate, and marginally preferred collaborating with the *Risk-Aware* robot.

Results - Objective. We show example human and robot behavior during the *in-lab* collaborative cup stacking task in Fig. 7. When modeling the human as *Noisy Rational*, the robot initially moved to grab the optimal cup and build the unstable tower. But in 75% of trials participants built the suboptimal but *stable* tower! Hence, the *Noisy Rational* robot often interfered with the human’s actions. By contrast, the *Risk-Aware* robot was collaborative: it correctly predicted that the human would choose the stable tower, and reached for the cup that best helped build this tower. This led to improved safety and efficiency during interaction, as shown in Fig. 8. Users interacting with the risk-aware robot completed the task in less time ($t(9) = 2.89, p < .05$), and the robot partner also traveled a shorter distance with less human interference ($t(9) = 2.24, p < .05$). These objective results support hypothesis H3.

Results - Subjective. We plot the user’s responses to our 7-point surveys in Fig. 9. We first confirmed that each of our scales (*Enjoy*, *Understood*, etc.) was consistent, with a Cronbach’s alpha > 0.9 . We found that participants marginally preferred interacting with the *Risk-Aware* robot over the *Noisy Rational* one ($t(9) = 2.09, p < .07$). Participants also indicated that they felt that they completed the task more efficiently with the *Risk-Aware* robot ($t(9) = 3.01, p < .05$). The other scales favored *Risk-Aware*, but were not statistically significantly. Within their comments, participants noticed that the *Noisy Rational* robot clashed with their intention: for instance, “it tried to pick up the cup I wanted to grab”, and “the robot picked the same action as me, which increased time”. Overall, these subjective results partially support hypothesis H4.

Summary. Viewed together, our *online* and *in-lab* user studies *not only* extended our autonomous driving results to a collaborative human-robot domain, but they *also* demonstrated how robots can leverage our formalism to meaningfully adjust their behavior and improve safety and efficiency. Our *in-lab* user study showed that participants interacting with a *Risk-Aware* robot completed the task faster and with less interference. We are excited that robots can actively use their *Risk-Aware* model to improve collaboration.

6 Discussion and Conclusion

Many of today’s robots model human partners as *Noisy Rational* agents. In real-life scenarios, however, humans must make choices subject to uncertainty and risk—and within these realistic settings, humans display a cognitive bias towards *suboptimal* behavior. We

adopted Cumulative Prospect Theory from behavioral economics and formalized a human decision-making model so that robots can now anticipate suboptimal human behavior. Across autonomous driving and collaborative cup stacking environments, we found that our formalism better predicted user decisions under uncertainty. We also leveraged this prediction within the robot’s planning framework to improve *safety* and *efficiency* during collaboration: our *Risk-Aware* robot interfered with the participants less and received higher subjective scores than the *Noisy Rational* baseline. We want to emphasize that this approach is *different from making robots robust* to human mistakes by *always* acting in a risk-averse way. Instead, when humans prefer to take safer but suboptimal actions, robots leveraging our formalism *understand* these conservative humans and increase overall team performance.

Limitations and Future Work. A strength and limitation of our approach is that the *Risk-Aware* model introduces additional parameters to the state-of-the-art *Noisy Rational* human model. With these additional parameters, robots are able to predict and plan around suboptimal human behavior; but if not enough data is available when the robot learns its human model, the robot could overfit. We point out that for all of the user studies we presented, the robots learned *Noisy Rational* and *Risk-Aware* models from the *same amount* of user data.

When learning and leveraging these models, the robot must also have access to real-world information. Specifically, the robot must know the rewards and probabilities associated with the human’s decision. We believe that robots can often obtain this information from experience: for example, in our collaborative cup stacking task, the robot can determine the likelihood of the unstable tower falling based on previous trials. Future work must consider situations where this information is not readily available, so that the robot can identify collaborative actions that are *robust* to errors or uncertainty in the human model.

Finally, we only tested the *Risk-Aware* model in bandit settings where the horizon is 1. Ideally, we would want our robots to be able to model humans over longer horizons. We attempt to address part of this limitation by conducting a series of experiments in a grid world setting with a longer horizon. We found that a *Risk-Aware* robot can more accurately model a sequence of human actions as compared to the *Noisy Rational* robot. Experiment details and results are further explained in the Appendix.

Collaborative robots need algorithms that can predict and plan around human actions in real world scenarios. We proposed an extension of *Noisy Rational* human models that also accounts for suboptimal decisions influenced by risk and uncertainty. While user studies across autonomous driving and collaborate cup stacking suggest that this formalism improves model accuracy and interaction safety, it is only one step towards seamless collaboration.

Acknowledgments

Toyota Research Institute (“TRI”) provided funds to assist the authors with their research but this article solely reflects the opinions and conclusions of its authors and not TRI or any other Toyota entity.

References

- [1] Pieter Abbeel and Andrew Y Ng. 2004. Apprenticeship learning via inverse reinforcement learning. In *International Conference on Machine Learning (ICML)*. ACM.
- [2] Muhammad Awais and Dominik Henrich. 2010. Human-robot collaboration by intention recognition using probabilistic state machines. In *19th International Workshop on Robotics in Alpe-Adria-Danube Region (RAAD 2010)*. IEEE, 75–80.
- [3] Haoyu Bai, Shaojun Cai, Nan Ye, David Hsu, and Wee Sun Lee. 2015. Intention-aware online POMDP planning for autonomous driving in a crowd. In *2015 IEEE international conference on robotics and automation (icra)*. IEEE, 454–460.
- [4] Chris L Baker, Rebecca Saxe, and Joshua B Tenenbaum. 2009. Action understanding as inverse planning. *Cognition* 113, 3 (2009), 329–349.
- [5] Chandrayee Basu, Erdem Biyik, Zhixun He, Mukesh Singhal, and Dorsa Sadigh. 2019. Active Learning of Reward Dynamics from Hierarchical Queries. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [6] Erdem Biyik, Daniel A. Lazar, Dorsa Sadigh, and Ramtin Pedarsani. 2019. The Green Choice: Learning and Influencing Human Decisions on Shared Roads. In *Proceedings of the 58th IEEE Conference on Decision and Control (CDC)*.
- [7] Erdem Biyik, Malayandi Palan, Nicholas C. Landolfi, Dylan P. Losey, and Dorsa Sadigh. 2019. Asking Easy Questions: A User-Friendly Approach to Active Reward Learning. In *Proceedings of the 3rd Conference on Robot Learning (CoRL)*.
- [8] Michael Bloem and Nicholas Bambos. 2014. Infinite time horizon maximum causal entropy inverse reinforcement learning. In *53rd IEEE Conference on Decision and Control*. IEEE, 4911–4916.
- [9] Daniel S Brown, Wonjoon Goo, and Scott Niekum. 2019. Ranking-Based Reward Extrapolation without Rankings. *arXiv preprint arXiv:1907.03976* (2019).
- [10] Siddhartha Chib and Edward Greenberg. 1995. Understanding the metropolis-hastings algorithm. *The american statistician* 49, 4 (1995), 327–335.
- [11] Sandra Devin and Rachid Alami. 2016. An implemented theory of mind to improve human-robot shared plans execution. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 319–326.
- [12] Adele Diederich. 1997. Dynamic stochastic models for decision making under time constraints. *Journal of Mathematical Psychology* 41, 3 (1997), 260–274.
- [13] Anca D Dragan, Kenton CT Lee, and Siddhartha S Srinivasa. 2013. Legibility and predictability of robot motion. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction*. IEEE Press, 301–308.
- [14] Anca D Dragan and Siddhartha S Srinivasa. 2012. *Formalizing assistive teleoperation*. MIT Press, July.
- [15] Espen Moen Eilertsen. 2014. *Cumulative Prospect Theory and Decision Making Under Time Pressure*. Master’s thesis.
- [16] Owain Evans, Andreas Stuhlmüller, and Noah Goodman. 2016. Learning the preferences of ignorant, inconsistent agents. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- [17] Chelsea Finn, Sergey Levine, and Pieter Abbeel. 2016. Guided cost learning: Deep inverse optimal control via policy optimization. In *International Conference on Machine Learning*, 49–58.
- [18] Andrew Gray, Yiqi Gao, J Karl Hedrick, and Francesco Borrelli. 2013. Robust predictive control for semi-autonomous vehicles with an uncertain driver model. In *2013 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 208–213.
- [19] Joseph Y Halpern, Rafael Pass, and Lior Seeman. 2014. Decision Theory with Resource-Bounded Agents. *Topics in cognitive science* 6, 2 (2014), 245–257.
- [20] Daniel Kahneman and Amos Tversky. 2013. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*. World Scientific, 99–127.
- [21] Martin Liebner, Michael Baumann, Felix Klanner, and Christoph Stiller. 2012. Driver intent inference at urban intersections using the intelligent driver model. In *2012 IEEE Intelligent Vehicles Symposium*. IEEE, 1162–1167.
- [22] Dylan P Losey and Marcia K O’Malley. 2019. Enabling Robots to Infer How End-Users Teach and Learn Through Human-Robot Interaction. *IEEE Robotics and Automation Letters* 4, 2 (2019), 1956–1963.
- [23] Dylan P Losey, Krishnan Srinivasan, Ajay Mandekar, Animesh Garg, and Dorsa Sadigh. 2019. Controlling Assistive Robots with Learned Latent Actions. *arXiv preprint arXiv:1909.09674* (2019).
- [24] Sandeep Mishra. 2014. Decision-making under risk: Integrating perspectives from biology, economics, and psychology. *Personality and Social Psychology Review* 18, 3 (2014), 280–307.
- [25] Andrew Y Ng, Stuart J Russell, et al. 2000. Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning (ICML)*.
- [26] Truong-Huy Dinh Nguyen, David Hsu, Wee-Sun Lee, Tze-Yun Leong, Leslie Pack Kaelbling, Tomas Lozano-Perez, and Andrew Haydn Grant. 2011. Capir: Collaborative action planning with intention recognition. In *Seventh Artificial Intelligence and Interactive Digital Entertainment Conference*.
- [27] Stefanos Nikolaidis, David Hsu, and Siddhartha Srinivasa. 2017. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research* 36, 5-7 (2017), 618–634.
- [28] Lisa Ordóñez and Lehman Benson III. 1997. Decisions under time pressure: How time constraint affects risky decision making. *Organizational Behavior and Human Decision Processes* 71, 2 (1997), 121–140.
- [29] Pedro A Ortega and Alan A Stocker. 2016. Human decision-making under limited time. In *Advances in Neural Information Processing Systems*. 100–108.
- [30] Takayuki Osogami and Makoto Otsuka. 2014. Restricted Boltzmann machines modeling human choice. In *Advances in Neural Information Processing Systems*. 73–81.
- [31] Makoto Otsuka and Takayuki Osogami. 2016. A deep choice model. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- [32] Malayandi Palan, Nicholas C Landolfi, Gleb Shevchuk, and Dorsa Sadigh. 2019. Learning Reward Functions by Integrating Human Demonstrations and Preferences. *arXiv preprint arXiv:1906.08928* (2019).
- [33] Stefania Pellegrinelli, Henny Admoni, Shervin Javdani, and Siddhartha Srinivasa. 2016. Human-robot shared workspace collaboration via hindsight optimization. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 831–838.
- [34] David Premack and Guy Woodruff. 1978. Does the chimpanzee have a theory of mind? *Behavioral and brain sciences* 1, 4 (1978), 515–526.
- [35] Deepak Ramachandran and Eyal Amir. 2007. Bayesian Inverse Reinforcement Learning. In *IJCAI*, Vol. 7. 2586–2591.
- [36] Vasumathi Raman, Alexandre Donzé, Dorsa Sadigh, Richard M Murray, and Sanjit A Seshia. 2015. Reactive synthesis from signal temporal logic specifications. In *Proceedings of the 18th international conference on hybrid systems: Computation and control*. ACM, 239–248.
- [37] Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. 2017. Active Preference-Based Learning of Reward Functions. In *Robotics: Science and Systems*.
- [38] Dorsa Sadigh, S Shankar Sastry, Sanjit A Seshia, and Anca Dragan. 2016. Information gathering actions over human internal state. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 66–73.
- [39] Dorsa Sadigh, S. Shankar Sastry, Sanjit A. Seshia, and Anca D. Dragan. 2016. Planning for Autonomous Cars that Leverage Effects on Human Actions. In *Proceedings of Robotics: Science and Systems (RSS)*. <https://doi.org/10.15607/RSS.2016.XII.029>
- [40] Herbert A Simon. 1972. Theories of bounded rationality. *Decision and organization* 1, 1 (1972), 161–176.
- [41] Andrea Thomaz, Guy Hoffman, Maya Cakmak, et al. 2016. Computational human-robot interaction. *Foundations and Trends® in Robotics* 4, 2-3 (2016), 105–223.
- [42] Amos Tversky and Daniel Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty* 5, 4 (1992), 297–323.
- [43] Vaibhav V Unhelkar and Julie A Shah. 2019. Learning Models of Sequential Decision-Making with Partial Specification of Agent Behavior. In *Thirty-Third AAAI Conference on Artificial Intelligence*.
- [44] Ram Vasudevan, Victor Shia, Yiqi Gao, Ricardo Cervera-Navarro, Ruzena Bajcsy, and Francesco Borrelli. 2012. Safe semi-autonomous control with enhanced driver modeling. In *2012 American Control Conference (ACC)*. IEEE, 2896–2903.
- [45] Michael P Vitis and Claire J Tomlin. 2013. A probabilistic approach to planning and control in autonomous urban driving. In *52nd IEEE Conference on Decision and Control*. IEEE, 2459–2464.
- [46] Diana L Young, Adam S Goodie, Daniel B Hall, and Eric Wu. 2012. Decision making under time pressure, modeled in a prospect theory framework. *Organizational behavior and human decision processes* 118, 2 (2012), 179–188.
- [47] Brian D. Ziebart, Andrew Maas, J. Andrew Bagnell, and Anind K. Dey. 2008. Maximum Entropy Inverse Reinforcement Learning. In *Proc. AAAI*. 1433–1438.
- [48] Brian D Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, J Andrew Bagnell, Martial Hebert, Anind K Dey, and Siddhartha Srinivasa. 2009. Planning-based prediction for pedestrians. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 3931–3936.

Appendix

To investigate how well *Risk-Aware* and *Noisy Rational* model humans in more complex POMDP settings, we designed two different maze games. Each game consists of two 17-by-15 grids and these two grids have the exact same structure of walls, which are visible to the player. In each grid, there is one *start* and two *goal* squares. Players start from the same square, and reach either of the goals. Each square in the grids has an associated reward, which the player can also observe. The partial observability comes from the rule that the player does not exactly know which grid she is actually playing at. While she is in the first grid with 95% probability, there is a 5% chance that she might be playing in the second grid. We visualize

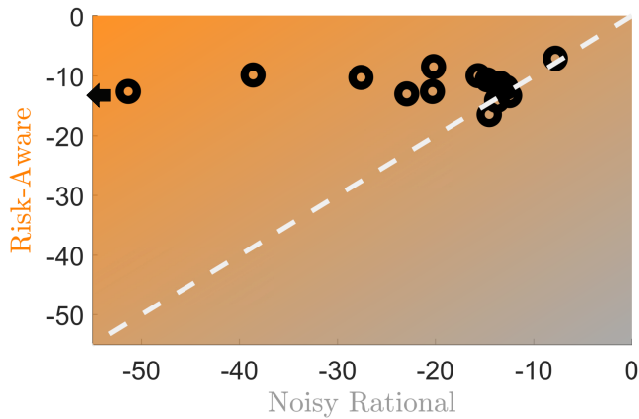


Figure 11: Log-Likelihood values by Risk-Aware and Noisy Rational models. One outlier point is excluded from the plot and is shown with an arrow.

the grids for both games in Fig. 10, and also attach the full mazes in the supplementary material. We restricted the number of moves in each game such that the player has to go to the goals with the minimum possible number of moves. Finally, we enforced a time limit of 2 minutes per game.




Figure 10: Summaries of two games. For each game, we have a maze. The values written on the mazes represent how much reward players can collect by entering those roads. The first numbers in each pair correspond to the 95% grid, and the second one to the 5% grid.

We investigate the effect of both risk and time constraints via this experiment. While it is technically possible for the players to compute the optimal trajectory that leads to the highest expected reward, time limitation makes it very challenging, and humans resort to rough calculations and heuristics. Moreover, we designed the mazes such that humans can get high rewards or penalties if they are in the low-probability (5%) grid. This helps us investigate when humans become risk-seeking or risk-averse.

We recruited 17 users (4 female, 13 male, median age 23), who played both games. We used one game (two grids) to fit the model parameters independently for each user, and the other game (other two grids) to evaluate how well the models can explain the human behavior. As the human actions depend not only the immediate rewards, but also the future rewards, we ran value iteration over the grids and used the values to fit the models as we described in Sec. 3. We again employed Metropolis-Hastings to sample model parameters, and recorded the mean of the samples.

Figure 11 shows the log-likelihoods for each individual user for Risk-Aware and Noisy Rational models. Overall, Risk-Aware explains the test trajectories better. The difference is statistically significant (paired t -test, $p < 0.05$). In many cases, we have seen risk-averse and risk-seeking behavior from people. For example, 12 out of 17 users chose the risk-seeking action in the test maze by trying to get 25 reward with probability 5% instead of getting 2 with 100% probability. Similarly, 15 out of 17 users choose to guarantee 0.9 reward and gain 0.1 more with 5% probability instead of guaranteeing 1.6 reward and losing 10 with 5% probability. This is an example of suboptimal risk-averse action.