Exploring non-adoption of tips for sequential decision-making: views and use of tips, barriers encountered, and problem-solving style

David Lee
UC Santa Cruz
dlee105@ucsc.edu

Wichinpong Park Sinchaisri
UC Berkeley
parksinchaisri@berkeley.edu

Abstract

Prior research on human-AI systems has focused on simple adopt/reject decisions on AI-generated tips, leaving sequential decision-making contexts less explored. In sequential tasks, tips suggest best practices that humans must operationalize across multiple decisions towards a problem-solving objective. In this paper, we explore human non-adoption of such tips in a virtual kitchen management task where disruptions necessitate a change of strategy. A qualitative analysis reveals diverse views of tips: as rules, directional principles, experimental options, or initially ignorable advice. Tips can still benefit those who reject them through creating focal points influencing worker sense-making. The challenge of operationalizing tips can lead to diverse barriers not just related to trust, but also to tip usability and environmental factors. A follow-up quantitative study confirmed three prominent barriers impacting participant intent to use tips. It also found that problem solving style, specifically an "Orientation to Change', may influence one's experience of barriers and tip adoption.

Keywords: human-ai collaboration, algorithmic tips, barriers to adoption of tips, problem-solving style

1. Introduction

AI-powered decision-support systems are increasingly being used in high-stakes settings like healthcare and legal contexts (Kononenko, 2001; Mallari et al., 2020). However, the opaque black-box nature of such systems create significant risks for society, especially in light of human tendencies to over-rely or under-rely on AI when deciding whether to adopt AI tips, phenomenon known as automation bias (Bussone et al., 2015; Jacobs et al., 2021; Lai and Tan, 2019; Mosier et al., 2017) and algorithmic aversion (Castelo et al., 2019; Dietvorst et al., 2015), respectively. A long thread of work has sought to better understand these biases, primarily in binary accept/reject decision-making settings. However, sequential decision-making, where multiple interdependent choices impact both short-term and long-term outcomes towards a larger problem solving objective, remains understudied. In such settings, tips may not correspond to accept/reject decisions for single decisions, but could consist of "best practices" for the overall strategy that a human still needs to figure out how to incorporate into individual decisions.

Bastani et al., 2025 argues that these contexts are particularly interesting not only due to their prevalence, but also due to significant amounts of trace data in organizations that can be used by AI to identify and recommend best practices. They introduce and study an algorithm for generating tips that best bridge the gap between a given person's current strategy and best-performing strategies, finding such tips more effective yet loss adopted compared to intuitive human-provided tips. We expand this inquiry through a mixed-methods exploration of human non-adoption. The sequential decision-making context is particularly interesting because it captures more of the complexities of AI tip adoption beyond just the initial intent to adopt to also figuring out how to adopt a tip.

Our qualitative analysis of survey responses obtained from Bastani et al., 2025 revealed participants' diverse views or use of tips: as rules, directional principles, experimental options, or initially ignorable advice. However, even when rejected outright, tips could facilitate strategy development by creating focal points highlighting potential strategies in the solution space.

Participants encountered diverse barriers when trying to operationalize tips. Some related to not trusting a tip (due to it being counterintuitive or resulting in bad outcomes), but others related to tip usability (being difficult to implement, lacking clarity, being difficult to track whether they were implementing) and to broader environment factors (misaligned incentives).

A subsequent quantitative study reinforced these barriers and highlighted three of the aforementioned barriers (counterintuitive tips, bad outcomes, and difficulty to implement) as particularly prominent in impacting intention to use tips. We also found that those with a stronger "Developer" type in their "Orientation to Change", i.e. those who like structure and rules (using a simplified version of the VIEW problem-solving style assessment Selby et al., 2004), were less likely to express several of these and other barriers, and more likely to comply with the provided tips.

In what follows, after describing related work (**Section 2**) and the experimental setting (**Section 3**), we describe our qualitative findings on views of tips and barriers to adoption (**Section 4**). This is followed by our quantitative findings highlighting three barriers as particularly correlated to future intention to use tips and exploring the role of problem-solving style (**Section 5**). We conclude by discussing implications and directions for AI-assisted decision-making in sequential settings.

2. Related Work

Machine-learning systems are increasingly used to support human decision-making in domains such as healthcare and criminal justice (e.g., Komorowski et al., 2018; Kononenko, 2001; Lee et al., 2021; Mallari et al., 2020). While often accurate, these systems are not always adopted, especially when humans either over-rely on flawed advice (*automation bias*) or reject useful advice from algorithms (*algorithm aversion*) (Dietvorst et al., 2015). Our work examines the latter in a sequential decision-making context, where interpreting and acting on algorithmic advice may be especially complex.

2.1. Factors affecting algorithm aversion

People are more sensitive to algorithm errors than human errors, often withdrawing trust more rapidly when algorithms fail (Bogert et al., 2021; Dietvorst et al., 2015). Aversion also arises when advice feels misaligned with how humans naturally reason. For example, people are less likely to rely on algorithms in tasks that are subjective, moral in nature, or easily comparable to their own judgments (Bigman and Gray, 2018; Castelo et al., 2019; Kawaguchi, 2021). While transparency is often assumed to increase trust, it can sometimes backfire. More interpretable models can overwhelm users with detail, making errors harder to detect unless paired with designs that actively prompt reflection (Buçinca et al., 2021; Poursabzi-Sangdeh et al., 2021). A recent review categorizes factors into four groups: attributes of the algorithm (e.g., opacity, complexity), the individual (e.g., personality, self-efficacy), the task (e.g., subjectivity, stakes), and the context (e.g., organizational purpose, social influence) (Mahmud et al., 2022). These suggest that adoption is not solely a function of performance, but depends on how advice is framed, interpreted, and situated.

2.2. Sequential contexts for decision-making

Most prior research on algorithm aversion examines one-off decisions such as predicting test scores, job fit, or product rankings (Dietvorst and Bharti, 2020; Dietvorst et al., 2015), where the primary question is whether or not to follow a recommendation. Sequential settings introduce new challenges: people must coordinate multiple interdependent actions, manage uncertainty over time, and make trade-offs between short-term and long-term outcomes. An exception is McIlroy-Young et al., 2020's study of human-AI collaboration in chess, which showed that effective interfaces must model human behavior at a granular level to provide useful guidance. Our work continues this direction by focusing on how people interpret and integrate AI-generated advice into evolving strategies.

2.3. Our contribution

Our paper is among the first to explore algorithmic tip adoption in sequential decision-making settings. We identify barriers such as counterintuitive advice, implementation difficulty, and negative outcomes that are less relevant in one-off tasks. These findings extend the algorithm aversion literature by highlighting what it means for a tip to be "optimal" when success depends on interpreting and integrating tips into a longer strategy. Additionally, our investigation into individual problem-solving style (based on a simplified version of the VIEW framework; Selby et al., 2004) introduces a novel behavioral dimension that has not been addressed in prior systematic reviews.

3. Background and Setting

3.1. The task: managing a virtual kitchen

We build on the interactive game environment developed in Bastani et al., 2025, where participants manage a virtual kitchen resembling the popular game *Overcooked* (see **Figure 1**). At each step, they assign cooking-related tasks (e.g., chopping, cooking, plating) to three virtual workers with unknown and heterogeneous skill levels (e.g., how fast to complete each task). The goal is to minimize the number of steps to complete a set of customer orders (e.g., to chop, then cook, then plate four burgers). Participants learn workers' skills only through trial-and-error. Success requires balancing short-term efficiency with long-term task planning. This setting captures the complexity of sequential decision making and offers a rich environment to study how people adapt their strategies, and whether they adopt algorithmically generated best-practice tips to improve performance.

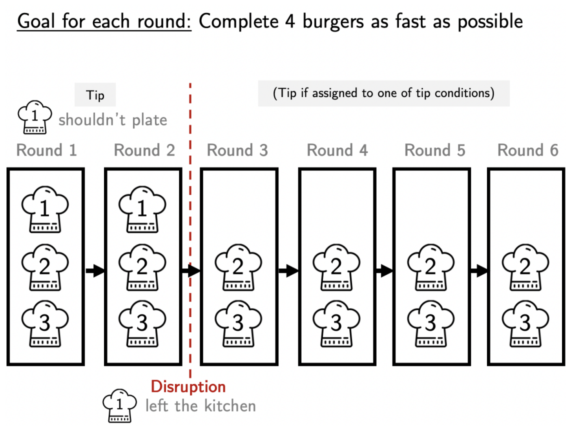

Figure 1: The kitchen management game flow: each participant plays the game with 3 virtual workers for 2 rounds, faces a disruption in the kitchen (i.e., the chef leaves), and continues for 4 rounds. The sequential tasks to complete each burger order: chop → cook → plate.

3.2. The algorithmically generated tips: recommending "optimal" best practices

In Bastani et al., 2025, participants are shown a "best practice" tip before each round, with tips generated using an algorithm that analyzes decision sequences from prior participants (who played without tips) to recommend interpretable, high-impact strategies (e.g., "Server should cook twice [out of four]"). Notably, the server is slow at cooking, so this tip is counterintuitive. However, it is optimal in the disrupted scenario where the most capable worker is unavailable. Since describing a full optimal policy is unrealistic, the algorithm identifies tips that, when incorporated into a human strategy, most improve expected long-term performance. Tips are derived from a Markov Decision Process model of human gameplay, with steps to filter out noisy suggestions. Full details are in Bastani et al., 2025.

3.3. Participants and overall study design

In Bastani et al., 2025, participants were recruited from Amazon Mechanical Turk after a tutorial and comprehension check. They played two training rounds in a simplified kitchen with uniformly skilled workers followed by two rounds in a full-capacity kitchen with three asymmetric workers and then four additional rounds in a disrupted kitchen where the chef, the fastest worker, was removed. Participants had to fulfill the same food orders using only the sous-chef and server.

This setup ensures participants are familiar with the game mechanics before needing to adapt their strategy due to disruption. The task challenges participants to infer worker skills through experience and to balance short-term task completion with longer-term bottleneck avoidance, creating an opportunity to observe whether AI tips aid in re-learning.

3.4. Prior results motivating our study

Bastani et al., 2025's behavioral study with 1,011 participants tested how different types of tips affected performance in the virtual kitchen task. Participants were randomly assigned to one of four conditions: no tip (control), an algorithmic tip, a human-generated tip, or a naive baseline tip. The algorithmic tip, "Server should cook twice," was selected to maximize long-term performance based on past gameplay data. The human-generated tip, "Server should cook once," was the most frequently recommended by prior players when asked what advice would best help future participants. The baseline tip, "Sous-chef should plate twice," was drawn from a simple frequency count of common actions from earlier participants.

Although the algorithmic tip was less intuitive, since the server is typically slow at cooking compared to the more capable chef, it led to better performance. 19% of participants who received it reached the optimal solution in the final round, compared to fewer than 1% in other conditions. However, adoption of the algorithmic tip remained much lower (24–48%) than that of the more intuitive human-generated tip (83–88%). This disconnect between effectiveness and adoption motivates our effort to better understand how users interpret and respond to AI-generated advice in sequential decision-making tasks.

3.5. Discussion on the experimental setting

Our study takes a mixed-methods approach, starting from an analysis of qualitative data obtained from

Bastani et al., 2025 followed by a subsequent quantitative study using the same experimental setup as their study with minor modifications. We see this setting as providing a strong foundation for studying tip adoption in sequential decision-making as an ideal setting should involve a multi-step task where actions interdependently affect long-term outcomes, pose real trade-offs between short- and long-term efficiency, and present algorithmic advice that is both helpful and potentially counterintuitive. It should also reflect real-world conditions in which AI-generated best practices are derived from historical performance traces.

The virtual kitchen meets these criteria. It allows participants to build familiarity before facing disruption, challenges them to infer skill levels and avoid bottlenecks, and features tips that improve performance but are often ignored. While the setting is stylized, it mirrors key managerial challenges in task allocation and scheduling under uncertainty. For instance, project coordinators and plant managers regularly balance performance benchmarks and shifting team capabilities without access to globally optimal policies.

4. Study 1: A rich view of how workers view tips and barriers to adoption

4.1. Method: qualitative analysis

For the first study, we obtained the anonymized survey data from the 247 participants from Bastani et al., 2025 who received the optimal algorithmic tip *"server cooks twice"*. We conducted an inductive thematic analysis of responses to the question, *"What did you think about the tip for these last four rounds and how did you incorporate it in your strategy?"* We first open-coded responses for anything related to how workers engaged with provided tips. These open codes were compared and grouped to identify common themes, leading to focused coding around two research questions: *1) How did workers view or conceptualize the tips? 2) What barriers kept people from using tips, either initially or in later rounds?* **Table 1** summarizes identified themes with illustrative quotes.

In the following theme discussions, our quotes use the syntax ($PID, c_3|c_4|c_5|c_6$, $d_3|d_4|d_5|d_6$) to convey more context, where PID is the participant, c_i refers to the number of times the had the server cook in rounds 3 to 6 (the disrupted rounds) and d_i refers to the duration they achieved in rounds 3 to 6 (with 34 being the minimum duration possible). For example, *(P90,* 2|2|2|2, 36|37|36|37*)* refers to a participant who had the server cook twice in each of the disrupted rounds and achieved a duration of 36, 37, 36, and 37 in those rounds.

4.2. How workers viewed and used tips

Rules: Of the workers who started off with an optimistic view of the tip, a few viewed the tips as *rules* that they *"had to figure out how to incorporate" ($P90, 2|2|2|2, 36|37|36|37$)*. For these workers, tips constrained the solution space they had to explore when sensemaking. For example, workers said: *"I knew that the server took longer to cook but HAD to cook twice so I had to figure out how to incorporate it" ($P90$)* and *"I thought of it as a rule and not a tip, even though it didn't say it was a rule. So, I followed the tip..." ($P108, 2|2|2|2, 34|36|34|34$)*

Directional principles: Another group of workers also had an optimistic view of the tips, but described them as *directional principles* to focus or be more cognizant of. For these workers, they did not feel like they needed to follow it exactly, but the tip guided them in becoming more aware about using the server to cook. Workers said: *"I didn't try to have the serve cook twice, but I was cognizant and more aggressive with having them cook in general- I just didn't track the exact number of times." ($P183, 1|2|1|2, 39|39|38|40$)* and *"It was very helpful. It made me focus on making sure the server cooked more even if that was not his obvious strength." ($P43, 1|2|2|2, 38|34|34|34$)*

Experimental options: Unlike the optimistic view of the previous two groups, others viewed it as an *experimental option* to try. Some viewed it neutrally as something to test before evaluating while others viewed it skeptically, but were willing to give it a shot. They said: *"I tested it out the first round and found out that it worked, so I repeated it during later rounds." ($P214, 2|2|2|2, 36|34|34|34$)*, *"I tried it the first time, but I don't think it was a good tip, so I ignored it the next times." ($P128, 2|1|1|1, 41|38|38|38$)*, and *"To me it didn't make sense. It basically went against everything I was taught, but I tried it." ($P86, 1|2|1|2, 38|35|36|36$)*

Initially ignorable advice: Finally, there were many who were skeptical of the tip and chose not to follow it initially. For some of these workers, tips still played a role by making those options more salient for the workers' own sensemaking and testing processes, or simply as things to try when nothing else worked, like Schelling or focal points (Schelling, 1980). Workers said: *"I thought that it was kind of suspicious at first but as I was figuring the game out myself, I thought that it was correct." ($P195, 2|2|2|2, 35|40|34|34$)*, *"I did not listen to the tip the first two times since he takes more ticks but noticed when I incorporated it, I was more efficient" ($P52, 1|1|2|2, 38|38|36|34$)*, and *"At first I didn't follow it because it seemed counter intuitive since they're slow. But then I had trouble, so I tried it and*

Table 1: Research questions, high-level themes and illustrative quotes.

Category and RQ	Theme	Representative Quote
Views and use of tips: *How did workers view or conceptualize tips?*	Tips are rules "I had to figure out how to incorporate"	"I knew that the server took longer to cook but HAD to cook twice so I had to figure out how to incorporate it"
	Tips are directional principles hinting in the right direction	"I didn't try to have the server cook twice, but I was cognizant and more aggressive with having them cook in general - I just didn't track the exact number of times."
	Tips are experimental options to try out[a]	"I tested it out the first round and found out that it worked, so I repeated it during later rounds."
	Tips are initially ignored as incorrect[b]	"At first I didn't follow it because it seemed counterintuitive since they're slow. But then I had trouble, so I tried it and came out ahead."
Barriers to adoption: *What barriers kept people from using tips?*	The tip felt counterintuitive	"I thought it didn't make sense. The server took 12 ticks to cook, so I had them only cook once because the sous-chef could finish in 8."
	The tip was hard to figure out how to implement	"I had a difficult time incorporating it and using it to my advantage. It always felt like the server took longer than needed when I could have had them doing other tasks."
	Trying to follow the tip resulted in bad outcomes	"I tried it the first time, but I don't think it was a good tip, so I ignored it the next times."
	I wasn't sure what the tip actually meant	"I wasn't sure what it meant. Does chopping count as cooking?"
	I lost track of how many times the server cooked	"It was confusing, I couldn't keep track of if he cooked or not."
	I was worried that exploring the tip would impact short-term performance and payments	"I didn't like it because i believe it took me longer to finish and i didn't receive any bonuses in those weeks."

came out ahead." ($P5, 1|1|2|2, 38|38|34|34$)

4.3. Barriers to adoption

Counterintuitive tips: One of the most common barriers expressed was that the tip was *counterintuitive*, which caused many to not follow the tip initially. As will be seen later, the counterintuitive nature of the tip also compounded some of the later barriers. Workers said: *"The first round, I ignored it because I knew the sous chef would do it quicker." ($P229, 1|2|2|2, 36|34|34|34$)*, *"At first I didn't follow it because it seemed counter intuitive since [servers] are slow." ($P5, 1|1|2|2, 38|38|34|34$)*, and *"I thought it didn't make sense. The server took 12 ticks to cook, so I had them only cook once because the sous-Chef could finish in 8." ($P156, 2|1|1|1, 38|38|38|38$)*

Hard to implement: Other workers talked about how it was *hard to implement*, which related to it being counterintuitive. Because tips need to be incorporated into a broader strategy, workers had to figure out how to apply it, saying: *"I had a difficult time incorporating it and using it to my advantage. It always felt like the server took longer than needed when I could have had them doing other tasks." ($P167, 1|2|2|1, 38|35|40|38$)*, *"I tried to incorporate it into my strategy but somewhere along the way I got lost." ($P96, 2|1|1|2, 37|38|38|40$)*, and *"It wasn't as useful as the tip in the first three rounds. I didn't really know how to implement it into my own strategy, or what it really implied." ($P132, 1|2|1|1, 38|40|36|38$)*

Bad Outcomes: A third challenge was that incorporating the tip (by having the server cook twice) could result in worse outcomes. These *bad outcomes* caused people to abandon the tip, saying: *"I tried it the first time, but I don't think it was a good tip, so I ignored it the next times." ($P128, 2|1|1|1, 41|38|38|38$)*, *"The time I tried to incorporate the tip, I used more ticks than when I ignored it." ($P42, 1|2|1|1, 38|40|36|38$)*, *"I let the server cook twice in the last couple of rounds and it didn't work well. If the game had continued I would have let the server only cook once." ($P101, 1|1|2|2, 36|36|38|39$)*

Lack of clarity: The previously described 3 barriers were the most commonly expressed, but there were also other barriers. For example, a few people felt there was a *lack of clarity* regarding what the tip meant concretely. They said: *"I wasn't sure what it meant. Does chopping count as cooking?" ($P133, 0|1|1|1, 42|36|38|41$)*, *"I thought it was a little too broad, but maybe I'm just stupid because I could not figure out how to finish in less than 40 ticks." ($P27, 1|2|2|2, 40|41|40|41$)*, and *"I* *was really confused about this tip, I wasn't sure what it meant by let the server cook twice. I did this and it did not really help me, but maybe I misinterpreted the tip." ($P135, 1|1|2|1, 38|39|36|39$)*

Hard to track: Another barrier was that it was sometimes *hard to track* how many times the server had cooked so they did not know whether they had implemented the tip or not. This is a variant of the 'hard to implement' barrier, but unlike those quotes where participants focused on the challenge of getting it to work logically, one participant expressed more of a logistical challenge: *"It was confusing, I couldn't keep track of if he cooked or not" ($P88, 1|1|1|1, 38|36|36|35$)*

Misaligned incentives: Finally, one participant touched on *misaligned incentives* in that trying to figure out how to implement the tip could result in lower short-term compensation if the tip is not implemented well: *"i didn't like it because i believe it took me longer to finish and i didnt receive any bonuses in those weeks" ($P225, 1|2|1|2, 38|48|39|39$)*

5. Study 2: Relating barriers, intent, and problem-solving style

5.1. Method: quantitative study design

To investigate patterns observed in participants' earlier open-ended survey responses, we conducted a follow-up study with a new set of participants ($N = 60$, mean age = 37.8, 43.3% female, 76.7% with a 2-year degree, average gameplay duration = 30.5 minutes). Participants completed the same virtual kitchen task described in Section 3 in the *Algorithm* condition (e.g., with the algorithmic tip), and answered new survey questions designed to probe interpretations, barriers, and behavioral responses identified in the prior study. We analyzed relationships among participants' views, intent, compliance, barriers, performance, and problem-solving style using Pearson correlations.

Participants indicated their intent to follow the tip before each of the four disrupted rounds (e.g., *"Do you intend to follow the tip in the coming round?"*[1]). They also completed three matrix-style 5-point Likert-scale questions asking how they initially viewed the tip (*"To what extent do you agree or disagree with the following statements about how you viewed the tip when you first received it going into Round 3?"*[2]), how their perception evolved over time (*"To what extent do you agree or disagree with the following statements about how your thoughts about the tip evolved over the course of Rounds 3-6?"*[3]), and what barriers, if any, hindered adoption (*"To what extent do you agree or disagree with the following statements about barriers that kept you from adopting the tip at any point within Rounds 3-6?"*[4].).

To explore how individual differences in problem-solving approach might influence tip adoption, we included a self-assessment based on the VIEW problem-solving style framework (Selby et al., 2004). VIEW is a validated instrument that captures three dimensions of problem-solving style: Orientation to Change (Explorer vs. Developer), Manner of Processing (External vs. Internal), and Ways of Deciding (People vs. Task preferences). Prior studies have used VIEW to understand creativity, learning styles, and team dynamics (Houtz, 2009; Treffinger et al., 2008).

To keep the study manageable for participants, we implemented a simplified version of VIEW. Participants viewed summary descriptions of the opposing types on each dimension (adapted from the original instrument) and rated themselves on a 7-point scale indicating where they fell between the two ends. For example,

[1]Options were *"Yes, I intend to follow the tip this round"*, *"Maybe, I'll keep it in mind and see"*, *"No, I don't plan to follow the tip this round"*, and *"Other"*.

[2]Rows were: *"I viewed the tip positively as a rule I had to figure out how to follow"*, *"I viewed the tip positively as a hint in the right direction, but not required"*, *"I viewed the tip neutrally as an option to consider trying at some point"*, *"I viewed the tip negatively as likely flawed, but still planned to try it"*, and *"I viewed the tip negatively as likely flawed, so did not intend to follow it"*.

[3]Rows were: *"My view of the tip became more negative by the end"*, *"My view of the tip became more positive by the end"*, *"I went back and forth on whether to use the tip"*, and *"The tip highlighted new ideas I would not have thought of otherwise"*.

[4]Rows were: *"The tip felt counterintuitive"*, *"It was difficult to figure out how to implement the tip"*, *"Trying to follow the tip resulted in bad outcomes"*, *"I wasn't sure what the tip actually meant"*, *"I lost track of how many times the server cooked"*, and *"I was worried that exploring the tip would impact my payment in initial rounds"*

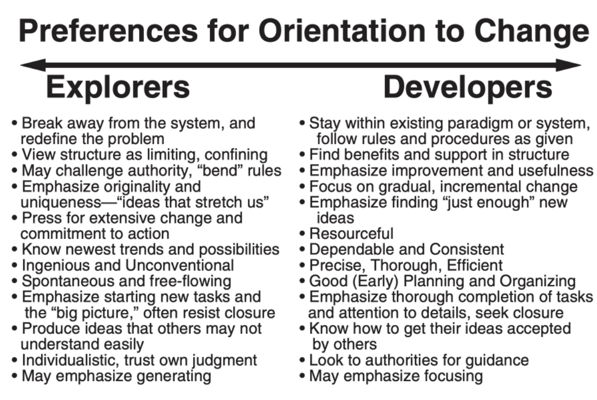

Figure 2: The summarizing figure for "Orientation to Change" drawn from Selby et al., 2004. The other dimensions are similar, with "Manner of Processing" describing people who prefer to process *Externally* vs *Internally* and "Ways of Deciding" describing people who give attention to *People* vs *Task* preferences when making decisions.

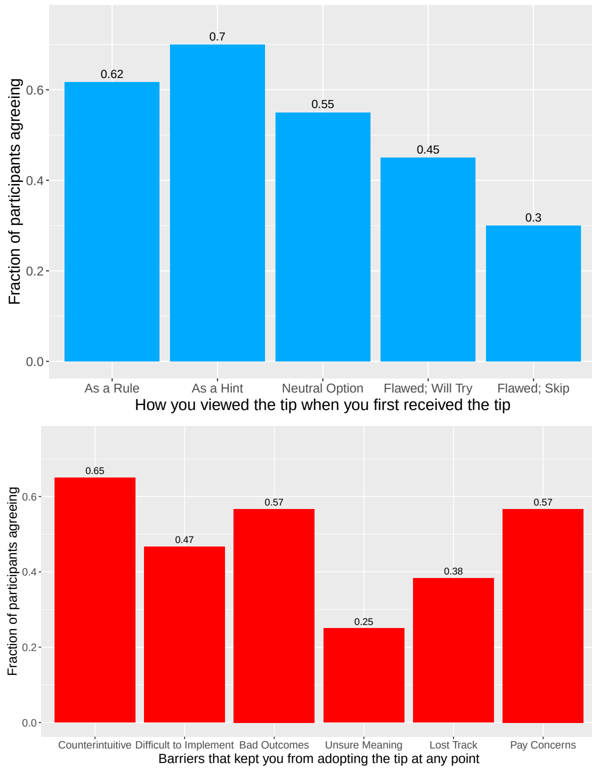

Figure 3: Fractions of participants agreeing to questions about (a) their initial views on the tip, and (b) the barriers that kept them from adopting the tip at any point.

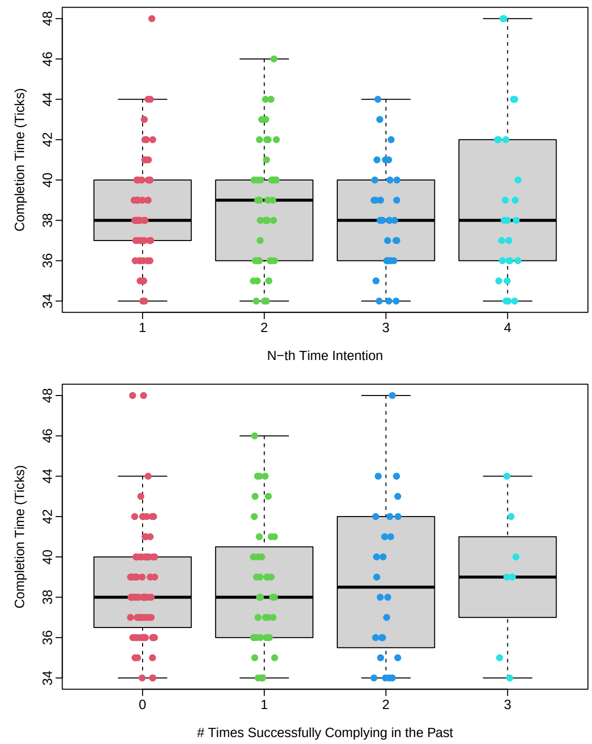

Figure 4: Performance across participants by how many times they (a) expressed an intent to follow the tip, or (b) succeeded in complying with the tip

for Orientation to Change, participants were asked to indicate whether they identified more with "Explorers" or "Developers" (see **Figure 2**). This single-item adaptation allowed us to efficiently capture style variation without significantly increasing task burden. We discuss the limitations of this simplification in **Section 6**, though we believe it still provides meaningful insights for an exploratory analysis.

5.2. Tip perceptions, barriers, and behaviors

Prevalence of tip views and barriers experienced: Figure 3 shows that participants held a mix of positive, neutral, and negative views toward the tip, with most starting out positive or neutral. Interestingly, even among those who initially viewed the tip negatively (responded "Somewhat/Strongly Agree" to *"I viewed the tip negatively as likely flawed, but still planned to try it"* and *"I viewed the tip negatively as likely flawed, so did not intend to follow it"*), a majority (57% and 47%, respectively) still agreed that the tip highlighted new ideas they would not have thought of otherwise, validating the observation that tips can provide benefits to workers even when initially rejected, e.g., by highlighting areas of the solution space that inform worker sense-making as seen in the qualitative results. All six barriers were reported by a non-trivial portion of participants, with the most common being that the tip felt counterintuitive, was difficult to implement, and led to poor outcomes. **Table 2** presents correlations between these barriers and participants' self-reported intent to follow the tip in each disrupted round. The three most impactful barriers were significantly associated with lower intent from Round 4 onward. These three barriers stood out across all analyses (**Figure 5**).

The challenge of operationalizing tips: Despite many participants expressing intent to follow the tip, actual compliance (i.e., having the server cook twice) only correlated significantly with intent in Rounds 5 ($r(58) = 0.29, p = 0.0267$) and 6 ($r(58) = 0.32, p = 0.0126$). Regarding performance, we observe a wide range of completion times even among participants who reported intent to follow the tip across multiple rounds (**Figure 4a**). Similarly, among those who complied with the tip (i.e., had the server cook twice), performance still varied substantially (**Figure 4b**). Although "server cooks twice" is the optimal recommendation for aligning human strategies with the algorithmic

	Intent (Round 3)	Intent (Round 4)	Intent (Round 5)	Intent (Round 6)
Counterintuitive	-0.14 ($p = 0.3357$)	-0.33 ($p = 0.0094$)	-0.31 ($p = 0.0143$)	-0.34 ($p = 0.0071$)
Difficult to Implement	-0.10 ($p = 0.5201$)	-0.3 ($p = 0.0215$)	-0.33 ($p = 0.0102$)	-0.26 ($p = 0.0488$)
Bad Outcomes	-0.12 ($p = 0.4336$)	-0.36 ($p = 0.0042$)	-0.46 ($p = 0.0002$)	-0.34 ($p = 0.0074$)

Table 2: Correlations and p-values among worker intent to follow the tip in Rounds 3 through 6 and the three barriers (*counterintuitive*, *difficult to implement*, and *bad outcomes*).

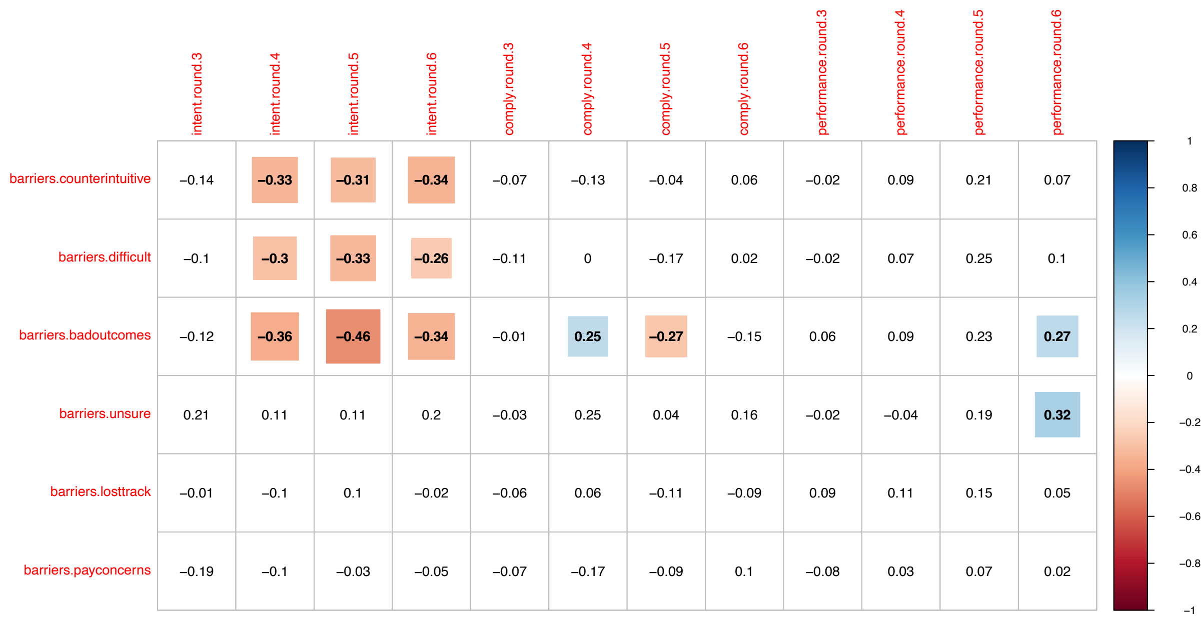

Figure 5: Correlations between barriers and intent to adopt, compliance, and performance. Highlighted cells are those that were statistically significant at $p < 0.05$.

policy, it can still lead to non-performing paths. This variation persists despite post-processing steps in the tip-generation algorithm designed to filter out misleading or high-variance advice.

Which barriers most impacted tip adoption? Three barriers stood out as having the strongest relationship to reduced intent to follow the tip. Participants who found the tip *counterintuitive*, *difficult to implement*, or associated it with *bad outcomes* were significantly less likely to express intent to follow it in Rounds 4–6 ($p < 0.05$). None of the other barriers showed consistent, strong effects (see **Figure 5**).

The influence of "Orientation to Change" on barriers and tip adoption: Participants who identified more strongly as "Developers" on the Orientation to Change dimension (i.e., those who prefer structure and benefit from rules) showed distinct behavioral patterns compared to "Explorers," who tend to view structure as limiting (see **Figure 2**). Developer types were significantly less likely to view the tip as counterintuitive ($r(58) = -0.28$, $p = 0.0320$) or as a threat to their performance-based payment ($r(58) = -0.38$, $p = 0.0031$), and were more likely to comply with the tip right away ($r(58) = 0.30$, $p = 0.0219$). Although not all results were statistically significant, directional trends suggest that Developers generally showed higher intent, compliance, and performance in Rounds 3 and 5. **Figure 6** summarizes correlations between problem-solving style and key outcomes. While these findings are preliminary, given the simplified version of the VIEW instrument, they highlight the potential role of problem-solving style in shaping responses to algorithmic advice. We return to this point in **Section 7**.

6. Limitations

Our analysis provides a richer picture of workers experiences with algorithmic tips, but some of our findings may reflect our specific task context in which workers aim to achieve a known optimal objective with the support of tips taking the form of best practice constraints on the overll strategy. Although

	barriers.counterintuitive	barriers.difficult	barriers.badoutcomes	barriers.unsure	barriers.losttrack	barriers.payconcerns	intent.round.3	intent.round.4	intent.round.5	intent.round.6	comply.round.3	comply.round.4	comply.round.5	comply.round.6	performance.round.3	performance.round.4	performance.round.5	performance.round.6
style.change	**−0.31**	−0.17	−0.22	−0.06	0.04	**−0.4**	0.22	0.01	0.17	0.19	**0.27**	−0.04	0.2	0.11	0.16	0.1	0.16	0.02
style.decide	−0.09	−0.1	−0.14	−0.13	−0.1	−0.22	0.24	0.12	0.01	0.05	0.05	−0.06	−0.01	−0.23	0.03	0.1	0.08	−0.03
style.process	−0.21	−0.15	−0.12	−0.12	−0.14	**−0.41**	0.12	−0.09	−0.14	−0.11	0.05	−0.12	**−0.26**	−0.07	0	0.08	−0.01	−0.06

Figure 6: Correlations between problem-solving style and barriers, intent to adopt, compliance, and performance. The size of each background block represents a statistically significant correlation at $p < 0.05$.

such conditions mirror many real-world contexts where managers optimize complex sets of decisions towards known industry benchmarks, additional research across diverse tasks would deepen generalizability. Further, our exploration of problem solving style only used a highly simplified version of the VIEW framework due to experimental constraints making it impractical to add the full VIEW assessment containing 34 survey items.

7. Concluding Discussion

Addressing diverse reasons for lack of trust: Those who start with more optimistic views of tips seem to be less likely to perceive obstacles encountered as barriers. Given that counterintuitive tips and bad outcomes were two significant trust-related barriers, one promising direction is to explore how one might *enhance tip explainability* or *increase worker confidence* by citing statistics of others or paths others took to go from bad outcomes to successful implementations.

But even if one is not able to boost worker confidence in tips, one may be able to *leverage the fact that tips can provide value as focal points* for later exploration, shifting the goal from convincing people to adopt tips to revealing solution spaces that we want to influence worker sense-making and experimentation.

The fact that tips could lead people down worse paths also raises questions about what it means for a tip to be “optimal”. For example, while the server needed to cook twice to achieve the optimal duration of “34”, it was also possible to get to a duration of “35” with the server cooking once. If the optimal solution for the former case is harder to discover, then the short-term value of “server cooks once” could actually be higher. How might we *design tips that consider the paths people might go down* while searching for the ideal implementation?

Beyond trust to usability and systems: Lack of trust was not the only obstacle preventing people from benefiting from tips. Tip usability, such as a lack of clarity and implementation challenges, is also an important barrier in complex decision-making contexts where tips are more like best practices than step-by-step actions. Broader environmental factors such as misaligned incentives affecting willingness to experiment, can add to implementation barriers. How might one design algorithmic decision support systems that integrate interactive feedback mechanisms, provide actionable guidance for implementation, and incentivize appropriate exploration?

Problem-solving style and collaboration: Our findings reflect different approaches people take to problem solving. Some are logicians reasoning about strategies. Others are experimenters engaging in trial and error. These different approaches affect whether people are willing to try counterintuitive tips, how they are impacted by encountered barriers, and the kind of design interventions that would be effective. People’s “Orientation to Change” may be particularly related to peoples’ perceptions of tips and their experience of barriers. How might teams with different combinations of personalities use tips? How might one form teams or facilitate collaborative interactions that lead to more effective use of tips? We see this as a particularly interesting direction because collaboration might be an effective way to help workers reason about counterintuitive tips or to figure out what a tip means and how to implement it. With the rapidly increasing use of AI across society, it is important to continue developing a richer view of human-AI interaction that can inform more human-centric development of AI algorithms.

References

Bastani, H., Bastani, O., & Sinchaisri, W. P. (2025). Improving human sequential decision making with reinforcement learning. *Management Science*.

Bigman, Y. E., & Gray, K. (2018). People are averse to machines making moral decisions. *Cognition*, *181*, 21–34.

Bogert, E., Schecter, A., & Watson, R. T. (2021). Humans rely more on algorithms than social influence as a task becomes more difficult. *Scientific reports*, *11*(1), 1–9.

Buçinca, Z., Malaya, M. B., & Gajos, K. Z. (2021). To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proc. ACM Hum.-Comput. Interact.*, *5*(CSCW1), 1–21.

Bussone, A., Stumpf, S., & O'Sullivan, D. (2015). The role of explanations on trust and reliance in clinical decision support systems. *2015 International Conference on Healthcare Informatics*, 160–169.

Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research*, *56*(5), 809–825.

Dietvorst, B. J., & Bharti, S. (2020). People reject algorithms in uncertain decision domains because they have diminishing sensitivity to forecasting error. *Psychological science*, *31*(10), 1302–1314.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, *144*(1), 114.

Houtz, J. C. S. (2009). Problem solving style, creative thinking, and problem solving confidence. *Educational Research Quarterly; West Monroe*, *33*(1), 18–30.

Jacobs, M., Pradier, M. F., McCoy, T. H., Perlis, R. H., Doshi-Velez, F., & Gajos, K. Z. (2021). How machine-learning recommendations influence clinician treatment selections: The example of antidepressant selection. *Transl. Psychiatry*, *11*(1), 1–9.

Kawaguchi, K. (2021). When will workers follow an algorithm? a field experiment with a retail business. *Management Science*, *67*(3), 1670–1695.

Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., & Faisal, A. A. (2018). The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature medicine*, *24*(11), 1716–1720.

Kononenko, I. (2001). Machine learning for medical diagnosis: History, state of the art and perspective. *Artif. Intell. Med.*, *23*(1), 89–109.

Lai, V., & Tan, C. (2019). On human predictions with explanations and predictions of machine learning models: A case study on deception detection. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 29–38.

Lee, M. H., Siewiorek, D. P., Smailagic, A., Bernardino, A., & Bermúdez i Badia, S. (2021). A human-ai collaborative approach for clinical decision making on rehabilitation assessment. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–14.

Mahmud, H., Islam, A. N., Ahmed, S. I., & Smolander, K. (2022). What influences algorithmic decision-making? a systematic literature review on algorithm aversion. *Technological Forecasting and Social Change*, *175*, 121390.

Mallari, K., Inkpen, K., Johns, P., Tan, S., Ramesh, D., & Kamar, E. (2020). Do i look like a criminal? examining how race presentation impacts human judgement of recidivism. *Proceedings of the 2020 Chi conference on human factors in computing systems*, 1–13.

McIlroy-Young, R., Sen, S., Kleinberg, J., & Anderson, A. (2020). Aligning superhuman ai with human behavior: Chess as a model system. *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1677–1687.

Mosier, K. L., Skitka, L. J., Heers, S., & Burdick, M. (2017). Automation bias: Decision making and performance in high-tech cockpits. In *Decision making in aviation* (pp. 271–288). Routledge.

Poursabzi-Sangdeh, F., Goldstein, D. G., Hofman, J. M., Wortman Vaughan, J. W., & Wallach, H. (2021). Manipulating and measuring model interpretability. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, (Article 237), 1–52.

Schelling, T. C. (1980). *The strategy of conflict: With a new preface by the author*. Harvard University Press.

Selby, E. C., Treffinger, D. J., Isaksen, S. G., & Lauer, K. J. (2004). Defining and assessing problem-solving style: Design and development of a new tool. *J. Creat. Behav.*, *38*(4), 221–243.

Treffinger, D. J., Selby, E. C., & Isaksen, S. G. (2008). Understanding individual problem-solving style: A key to learning and applying creative problem solving. *Learn. Individ. Differ.*, *18*(4), 390–401.