Title
Risk Aversion and Demographic Factors Affect Preference Elicitation and Outcomes of a Salary Negotiation

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors
Hale, James
Kim, Peter H
Gratch, Jonathan

Publication Date
2023

Peer reviewed

Risk Aversion and Demographic Factors Affect Preference Elicitation and Outcomes of a Salary Negotiation

James Hale (jahale@usc.edu)[1], Peter Kim (kimpeter@marshall.usc.edu)[1], and Jonathan Gratch (gratch@ict.usc.edu)[2]
[1]University of Southern California, Los Angeles, CA 90089
[2]Institute for Creative Technologies, University of Southern California, Playa Vista, CA 90094

Abstract

Women and minorities obtain lower salaries when negotiating their employment compensation. Some have suggested that automated negotiation and dispute-resolution technology might address such material inequities. These algorithms elicit the multi-criteria preferences of each side of a dispute and arrive at solutions that are efficient and "provably" fair. In a study that explores the potential benefit of these methods, we highlight cognitive factors that may allow inequities to persist despite these methods. Specifically, risk-averse individuals express lower preferences for salary and as risk-aversion is more common in women and minorities, this translates into a "provably" fair lower salary. While this may reflect actual underlying differences in preferences across groups, individuals may be confounding their preferences for salary with their risk preference (i.e., their fear of not reaching an agreement), such that these groups achieve worse outcomes than they should. We further highlight that methodological choices in how negotiation processes are often studied can obscure the magnitude of this effect.

Keywords: negotiation; human-agent; affective computing; non-linear

Introduction

Extensive research has documented that women and minorities obtain lower salaries when negotiating compensation. This has been attributed to a variety of factors. For example, women and minorities are reluctant to ask (Babcock & Laschever, 2009; Lu, 2022), and those that do ask are punished (Amanatullah & Tinsley, 2013). Some have suggested that automated methods might address such disparities by imposing "provably fair" solutions on disputants (Shah, 2017; Thiessen & Soberg, 2003a).

These algorithms elicit the multi-criteria preferences of each side of a dispute (blind to gender and race) and employ decision criteria that maximize social welfare. Assuming individuals honestly report their actual individual preferences, criteria like the Nash Bargaining Solution (Binmore, Rubinstein, & Wolinsky, 1986) arrive at deals that are efficient (in that they are Pareto efficient) and fair (in they maximize joint utility). Tempering this promise, critics have raised concerns as to whether these algorithms accurately elicit individual preferences (Sternlight, 2020), if they are genuinely fair (Sela, 2018), or if disputants will accept the proposed agreements (Lee & Baykal, 2017).

Here we provide evidence that "provably" fair algorithms could perpetuate or even institutionalize salary inequities based on gender, ethnicity, or personality. First, we find that risk aversion systematically influences how people express their preferences for an automated dispute resolution system. This is not reflected in different preferences across issues as suggested by some research – women do not value vacation higher than salary (Croson & Gneezy, 2009). Rather, this influence occurs through differences in the expressed marginal utility for salary (risk-averse individuals claim each additional dollar of salary is worth less than risk-seeking individuals). Second, consistent with prior research (Bontempo, Bottom, & Weber, 1997; Croson & Gneezy, 2009), women and Asians in our sample were more averse to risk. Together, these findings highlight a pathway by which certain groups might be disadvantaged by technology that certifies this disparity as "fair". Interestingly, this pathway is often overlooked by laboratory research on negotiation, which tends to assign preferences (i.e., telling, rather than eliciting what they should achieve in a negation) and assign preferences that are linear in the utility of money (Northcraft, Preston, Neale, Kim, & Thomas-Hunt, 1998).

Cognitive science has an uneven influence over the design of automated dispute-mediation algorithms. Building on psychological findings on multi-attribute decision problems, automated approaches employ sophisticated methods to elicit and represent human preferences over different outcomes. This includes recognizing that within a single issue (e.g., the salary level in a compensation package), people may hold nonlinear preferences in the level for that issue. For example, many show diminishing marginal utility for salary, meaning that each additional salary unit translates into an ever-smaller increase in subjective value. These nonlinearities have important consequences for how people negotiate and the ease with which they can find an agreement (Northcraft et al., 1998). People also have conditional preferences across different issues, meaning the relative importance of two issues may change based on the level obtained on a third (Boutilier, Brafman, Domshlak, Hoos, & Poole, 2004). For example, people might prefer salary to stock for short-term employment but greater preference for stock in a longer-term job. But despite this sophistication in representing nonlinear preferences, automatic approaches use relatively simple, classical criteria Nash Bargaining Solution to decide the fairness of potential agreements (e.g., taking elicited preferences at face value, calculating Pareto efficient solution, and imposing this on disputants).

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

In contrast, psychological research on the fairness of negotiated outcomes takes a more nuanced approach to fairness but has largely ignored cognitive science findings on the structure of human preferences. For example, laboratory studies have successfully replicated real-world findings that women tend to obtain lower salaries, and this has spurred systematic exploration of the pathways through which such effects can occur. For example, job postings often contain language that can put women on the defensive and lead them to lower their aspirations (Tang et al., 2017). Separate from this, women may face more aggressive opening offers and greater use of deception by their counterparts (Amanatullah & Tinsley, 2013). Yet these studies largely use assigned preferences, meaning each participant is given a predefined payoff matrix that determines the value of potential negotiated agreements (these values are typically represented as points translated into lottery tickets to motivate performance). By assigning preferences, it is no longer possible to ask if, for example, women might make different trade-offs across issues than men (Croson & Gneezy, 2009). Further, assigned preference assumes away the nonlinearities found when preferences are elicited. For example, the most common exercise used to study salary negotiations in the laboratory (Neale, 1997) assumes the subjective value of salary increases linearly with the dollar amount, and that the value obtained on one issue is independent of the value obtained on others. Here, we find most of the variance in the salary obtained by automatic algorithms arises from nonlinearities in participants' preferences that these assigned payoff matrices fail to capture.

Hypotheses and Research Questions

Motivated by a convenience sample of business students with many White and Asian participants, we focus on the impact of gender and Asian ethnicity on compensation outcomes in a simulated employment context. Prior research has suggested salary disparity in women might be linked to risk aversion – primarily through its influences on willingness to negotiate (Hernandez-Arenaz & Iriberri, 2019; Marks & Harold, 2011). Though less studied in the context of negotiations, prior work suggests that Asians have a greater propensity towards risk aversion as well (Bontempo et al., 1997). Thus, we focus on how risk might impact elicited preferences in these two groups.

- **RQ1:** How do demographic factors (e.g., gender and race) and individual differences (e.g., risk aversion) impact elicited preferences?

Considering these research questions, we make the following hypotheses:

- **H1a:** Women will exhibit a greater tendency towards risk aversion as a group.

- **H1b:** Asians will exhibit a greater tendency towards risk aversion as a group.

- **H2:** Risk-seeking participants will express tougher within-issue preferences than risk-averse participants.

Experimental Setup

Methods

Participants We use a convenience sample of 170 undergraduates from a west-coast U.S. University with a suitable level of female and Asian participants previously collected by us (Hale, Kim, & Gratch, 2022). Participants were removed if their survey response was incomplete (14 were removed via this criterion). We also removed participants if they completed the task in less than five minutes (27 removed via this), indicating low attention to the task[1]. The demographic breakdown of the remaining 129 participants follows: 64% male, 34% female, 2% other; self-reported race was 8% Hispanic, 46% Asian, 1% Pacific Islander, 4% Black, 33% White, 6% mixed-race, and 2% other; and 61% were born in the United States.

Design Participants were told to imagine they were seeking employment in a high-tech company. They were provided with a description of the company and company culture. They were then told to imagine being offered the job and to specify their preferences over a set of issues to help determine their negotiated compensation package. To analyze the potential of conditional preferences, participants rated their preferences twice, once imagining the job was a one-year contract and once imagining the job was a five-year contract (order counterbalanced). This allows us to examine if the contract length influenced preferences for the other issues.[2] After the experiment was completed, all preference curves were provided to a "fair" automatic-mediation approach as described below.

The experimental setup followed a simple within-participants design (one-year contract versus 5-year contract). Participants entered their preferences in a package already with a one-year contract, then a five-year contract (or vice-versa, as we used counter-balancing). Specifically, participants answered demographic and personality questions (described next) and read an introduction to and expert opinions on a fictional company. They were then told the contract length and told to input their preferences over three elements of the compensation package (*salary*, *stock*, and *vacation*). These issues have the corresponding ten levels:

- **Salary:** \$70k, \$80k, \$90k, \$100k, \$110k, \$120k, \$130k, \$140k, \$150k, & \$160k

- **Vacation:** 5 days, 6 days, 7 days, 8 days, 9 days, 10 days, 11 days, 12 days, 13 days, & 14 days

[1]The demographics of those removed are as follows: 67% male, 28% female, and 5% not reporting; 38% White, 43% Asian, 2% Hispanic, 2% Black, 2% Other, and 13% not responding; lastly, 60% were U.S. born, 36% were not, and 4% did not respond.

[2]We also independently manipulated the textual description of the company, indicating it was achievement versus family-focused, but this had no impact on the analyzed variables, and we ignore this in the remainder of the paper.

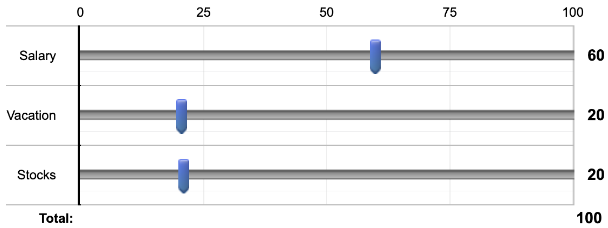

(a) Issue elicitation example

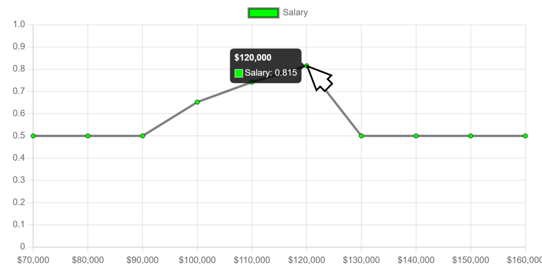

(b) Salary curve example

Figure 1: Depictions of the preference elicitation methods in this study

Preparation: ***Requirements & bottom line***

We value employees that

- Work independently in a fast-paced environment that gives you both great freedom and great responsibility
- Comfortable taking risks and working with ambiguity
- Accomplish amazing amounts of important work
- Inspire others with your thirst for excellence
- Seek what is best for the company, rather than best for yourself or your group

We're like a **pro sports team**, not a family. Company leaders hire and **cut smartly**.

Bottom line: This is not the place for employees seeking job security and a 9 to 5 workday, but for those confident in their skill and willing to put in the effort, the rewards can be substantial. Employees are heavily recruited by other technology companies.

Figure 2: Example job requirements and bottom-line for the fictional company shown to participants before they complete the preference elicitation

- **Stocks:** \$50k, \$60k, \$70k, \$80k, \$90k, \$100k, \$110k, \$120k, \$130k, & \$140k

Participants input their preferences using two stages, inspired by Thiessen & Soberg's SmartSettle, a commercial dispute-mediation platform (Thiessen & Soberg, 2003b); the first intends to capture how much they value each issue (i.e., an issue's *weight*), and the second means to capture a participant's valuation of levels within an issue (i.e., the *curve* over an issue). So, first, a participant will use a slider to allocate 100 points over the three issues (initialized to 0 for each issue), where they must allocate all points and allocate more points to issues more important to them (see Figure 1a). For the second stage, a participant will draw a curve over the levels of each issue by dragging anchors on a 10-point curve, where they can value a level between zero and one, where each level gets initialized to .5 (see Figure 1b). These curves are normalized so the highest point equals one for *toughness* and clustering.

Measures

Individual Differences: We collect self-reported demographic and dispositional information on each participant before the task. We ask each participant to report their race, gender identity, and whether they were born in the United States. We further ask participants to complete several short questionnaires from previously validated scales to gauge several dispositional characteristics.

One's socio-economic status may affect aspirations during a negotiation; as such, we asked participants to complete the MacArthur Scale of Subjective Social Status (Adler, Epel, Castellazzo, & Ickovics, 2000). This questionnaire shows participants an image of a ladder, and are told "[a]t the top of the ladder are the people who are the best off, those who have the most money, most education, and best jobs. At the bottom are the people who are the worst off, those who have the least money, least education, worst jobs, or no job. Please indicate the rung that best represents where you think you stand on the ladder."

Afterward, participants completed another series of questions to evaluate their risk aversion. We use Meertens *et al.*'s 7-item Risk Propensity Scale, which asks respondents to note their agreement to a series of statements (e.g., "I prefer to avoid risks") on a nine-point scale (Meertens & Lion, 2008). To measure risk aversion, we scale this score by negative one (i.e., higher scores, closer to zero, will be more risk averse).

Lastly, people often adopt a *prevention* or *promotion* focus towards goals (Liberman, Idson, Camacho, & Higgins, 1999), of which some research has considered the role in salary outcomes (Tellhed & Björklund, 2011). To measure these two variables, we utilized the Regulatory Focus Questionnaire (Higgins et al., 2001), which asks 11 questions on a five-point scale (e.g., "How often have you accomplished things that got you 'psyched' to work even harder?"). This scale gives two individual scores for *promotion* and *prevention* focus.[3]

[3] Whereas a *prevention* focus emphasizes safety and responsibility, views goals as oughts, and is concerned with non-losses and losses; a *promotion* focus emphasizes hopes and accomplishments, views goals as ideals, and is concerned with gains and non-gains.

Preferences: Participants' preferences for each contract length are completely determined by the weight they assign each issue and the ten values participants provide for each issue (30 parameters total for each contract length). Our primary interest here is how a participant's preferences influence the salary level they obtain. This is determined chiefly by the weight they assign to salary (as opposed to stock or vacation) and the 10-utility values they assign to the salary curve (although the weights and curves of the other issues will have an indirect influence, as we analyze below),

Issue Weight refers to the numeric weight (0 to 100) they provided for a given issue. Unless otherwise stated, issue weight refers to salary.

We operationalize the *toughness* measure (Hale et al., 2022) to capture how the shape of an issue curve tends to influence the level that someone would obtain on the issue if this curve were given to an automated dispute mediation algorithm. This value captures the *toughness* of the curve; i.e., it returns a higher value for curves with more points exclusively at higher levels (e.g., $\langle 0,\ldots,1\rangle$) and a low value for curves with more points only at lower levels (e.g., $\langle 1,\ldots,0\rangle$). In this sense, it collapses the ten curve values into a single number. Tougher salary curves will obtain higher salaries. As we will show below, toughness and issue weight together explain most of the variance in what level someone will get from such an algorithm.

We formally define *toughness* as follows:

$$f_{\text{curve}}(X) = \frac{\sum_{i\in\{1,\ldots,10\}} X_i * i}{\text{sum}(X)}$$

Where X is the vector of elicited value for each level of a curve—e.g., $\langle .5,\ldots,.5\rangle$ would correspond the user valuing each level of an issue equally—and X_i is the value of this issue at level i. This allows the compression of the entire curve to a single value for easy statistical analysis. Of note, we calculate *toughness* using a normalized curve where we divide every element by the maximum element.

Automatically Derived Salary: *Expected Salary* refers to the salary a participant would obtain if their preference 30 parameters (3 issue weights and ten curve points per issue) were given to a fair dispute mediation algorithm. Since the dispute is with a fictitious company, we would have to invent the company's preferences but to avoid any bias in our choice of preference; we simply use the other participant nt's curves to represent a distribution of possibilities to represent the company's interests. Specifically, we flip each participant's curves from left to right (i.e., most participants will assign the most utility to a high salary, so we can assume most companies would assign a high weight to a low salary). To determine the salary a given participant might expect, we repeatedly calculate the Nash Bargaining Solution of the participant's preferences compared with each other participant's flipped preferences. The average salary obtained across all of these mediated solutions is their expected salary.

Results

We initially analyze how people use weights and curves before analyzing how demographic and personality factors shape their use.

Issue Weights

Figure 3 illustrates the average weight given to each issue as a function of the length of the contract. To examine if people exhibit conditional preferences, We performed a two-way ANOVA which uncovers a main effect for *Issue* ($F(1.82, 232.39) = 278.96$): subjects allocated more weight to *salary* ($M = 58.574$, $SD = 16.386$) than *vacation* ($M = 18.558$, $SD = 11.985$) or *stock* ($M = 22.868$, $SD = 14.880$). We find an interaction between *Issue* and *Contract*, $F(1.74, 223.31) = 39.38$ $p < .001$, indicating people hold conditional preferences – i.e., the importance they assign to salary, vacation, and stock depends on the length of the contract. We also find weight distributions vary by certain demographic factors such as gender and race.

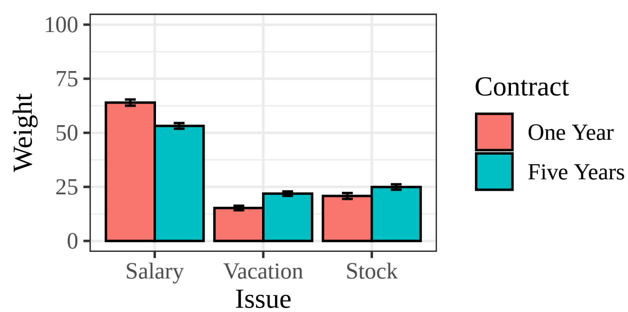

Figure 3: The weight assigned to each issue as a function of contract length. Error bars reflect the standard error of the mean.

A secondary analysis found no evidence that demographic or personality factors influenced the tendency to use conditional preferences; so, for the remainder of the paper, we only examine and report the results for the five-year contracts (the findings for 1-year contracts are almost identical).

Within-issue Curves

We first performed a qualitative analysis to visualize differences in how people express their preferences using the within-issue curves. We performed KML clustering over 10 utilities elicited within an issue (Genolini & Falissard, 2011). KML clustering is designed to find trends in longitudinal data. Figure 5 illustrates the five-curve solution best fitting the data. As with calculating *toughness*, we normalize all curves before running the clustering algorithm, which we run on all curves irrespective of contract or issue.

As seen in Figure 5, 31% of individuals (Curve *A*) used the curves to express diminishing marginal returns for an issue. Closely following this, 26% expressed utility that was linear in the level of the issue (curve *B*). The remainder used

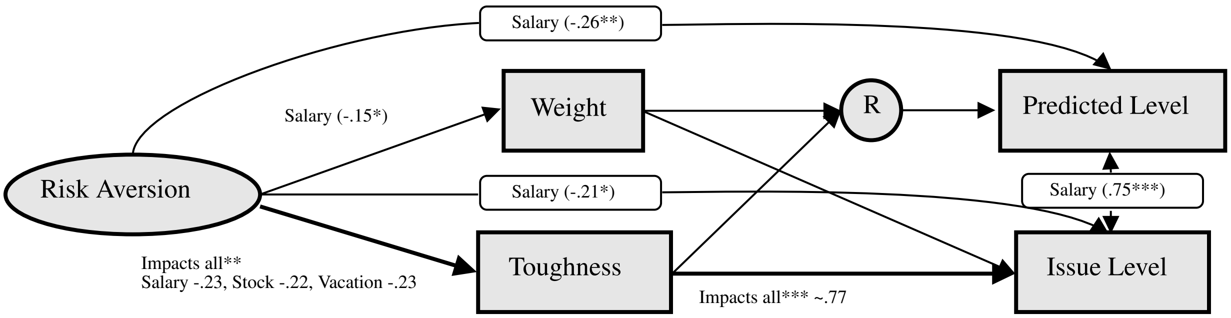

Figure 4: Individual difference effects on preferences (significance levels: * $p \leq .05$, ** $p \leq .01$, & *** $p \leq .001$)

a variety of other curves. Curve clusters differ in their average toughness: curve *A* ($M = 6.32$, $SD = 0.27$), *B* ($M = 7.17$, $SD = 0.40$), *C* ($M = 5.58$, $SD = 0.25$), *D* ($M = 5.70$, $SD = 0.35$), *E* ($M = 5.06$, $SD = 0.74$). Those using curve *B* expressed the toughest preferences.

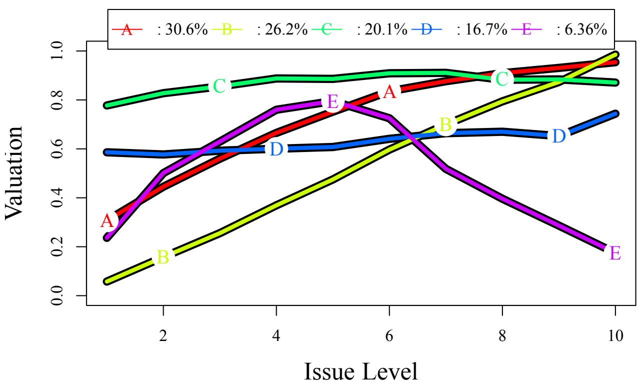

Figure 5: Five elicited curve types found with KML clustering (Genolini & Falissard, 2011)

Do *Weight* and *Toughness* explain outcome?

Before analyzing the impact of demographics and personality on salary, we examine if the Issue Weight and the Toughness of a curve capture most of the variance in how people express their preferences. If so, it suffices to analyze how individual differences impact these parameters.

We perform linear regression analyses of three configurations—considering (1) *toughness*, (2) *weight*, and (3) *weight* + *toughness*—on the three issues (*salary*, *vacation*, and *stock*). These regression models leverage a user's elicited preference profile to predict the corresponding average issue level when using NB, a "fair" mediation strategy, against other users' profiles. Prediction greatly improves when incorporating *toughness*, which suggests the simple *toughness* metric quantifies the impact of entire curves on solutions well. Since issue weight alone predicts poorly the level obtained on a given issue, we need both issue *weight* and curve *toughness*. Figure 6 illustrates the adjusted

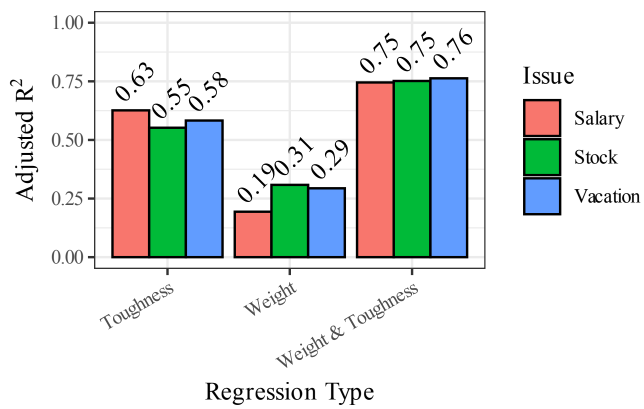

Figure 6: Illustrates the adjusted R^2 values for regressions

R^2 values for these regressions colored by issue type and grouped by the regression configuration.

Demographic Effects on Risk Aversion and Issues

Using a two-tailed Welch's t-test (*M* is *Mean*, and *SD* is *Standard Deviation*) after testing normality with the Shapiro-Wilk test, we investigate whether differences in risk aversion exist between groups in our participant pool. We see significant differences between risk aversion and other factors; e.g., women ($M = -30.02$) have significantly ($p = .041$) higher risk aversion than men ($M = -33.46$); and White people ($M = -34.81$) have significantly ($p = .013$) lower risk aversion than Asian people ($M = -30.00$). **H1a** and **H1b** posit demographic characteristics will affect risk aversion; specifically that women and Asians will show greater risk aversion. By these results, **H1 is supported.**

Risk Aversion's Impact of Salary Outcome

Here, we examine how risk aversion (and other dispositional factors) affects issue outcomes through elicited preferences. While subsequent analyses only focus on the 5-year contract condition for brevity, similar results come from the 1-year condition. We previously established issue weight and toughness suffice to predict the obtained outcome of an issue using NB as a mediation method. First, risk aversion trends to predict Salary issue weight ($r = -.15$, $p = .088$), but does not strongly predict stock ($r = .10$, $p = .271$) or vacation ($r = .07$, $p = .399$). However, we see a greater effect of risk aversion on the elicited curves (the *toughness* metric) where it significantly predicts *toughness* on salary ($r = -.23$, $p = .009$), vacation ($r = -.23$, $p = .008$), and stock ($r = -.22$, $p = .011$). **H2** posits those with greater risk aversion will post less tough preferences relating to salary. By these correlation tests, **H2 is supported.** Additionally, there exists a trend for promotion—another dispositional measure we captured for each participant—to impact salary *toughness* ($r = .16$, $p = .067$).

As we find risk aversion predicts elicited preferences, we also wish to demonstrate this effect carries through the mediation algorithms (NB). Focusing on just the *salary* issue, we find there exists a direct effect of risk propensity on the level of *salary* achieved through NB ($r = -.21$, $p = .019$), but also indirect effects through *weight* and *toughness*, we partially showed this indirect effect previously (risk aversion affecting elicitation). Finishing this, we find *Salary* issue *weight* ($r = .45$, $p < .001$) and curve *toughness* ($r = .79$, $p < .001$) both correlate with salary level.

We further conduct a mediation analysis to test whether *weight* and *toughness* mediate this effect using a bootstrapping procedure with 500 samples. The analysis shows full mediation exists through *toughness* but not *weight*. So, this direct effect vanishes when controlling for these mediators. Figure 7 illustrates this mediation test. Further, Figure 4 illustrates the correlations described in the previous two sections.

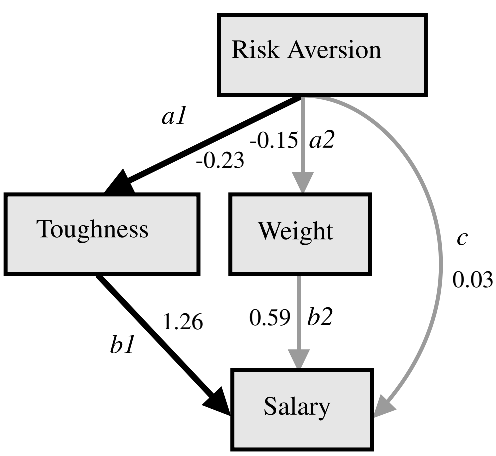

Figure 7: Results of regression analyses for the impact of risk propensity on obtained salary (**5 years**), showing the direct and indirect pathways within the multiple mediation models. The numbers are standardized regression coefficients:

Discussion and Future Work

Overall, our data suggest that risk aversion strongly impacts the salary level an individual will likely obtain from standard dispute resolution algorithms. Women and Asian participants exhibiting greater risk aversion (as found in this study and other literature) suggest these groups may receive lower salaries by such "provably fair" approaches. These findings are primarily due to how individuals express the marginal utility of money. Risk-averse individuals report diminishing marginal utility, whereas risk-seeking individuals tend to express linear utility in the salary level. We find several instances of non-monotonic curves from participants; while unintuitive, scenarios exist where this may be sensible. For example, one may input a non-monotonic curve over the *vacation* issue to a potential employer as they may fear doing otherwise would signal a lack of commitment.

When designing such resolution platforms, one should consider how these impacts affect the proposed solutions of “fair” mediation approaches and negotiation tasks, specifically regarding salary level. Further, should an AI mediator take elicited preferences at face value, accepting that economic inequities reflect systematic differences in the utility of money across different groups, or might these differences in expressed utility arise through structural biases in the populations we studied?

These findings do not necessarily mean these algorithms are unfair. Assuming people honestly report their preferences, the algorithms simply discover that different groups value issues differently. However, the fact that these differences manifest in diminishing marginal utility for money raises some concern as prior research has failed to find such systematic differences (Booij & Van de Kuilen, 2009; Fehr-Duda, De Gennaro, & Schubert, 2006). Instead, this may reflect something unique to an employment context. For example, one interpretation of our results is that people are confounding the utility they assign to different salaries because they fear they might not reach an agreement if they express their true preferences. There is good reason to believe that certain groups, like women, are justified in these fears (Amanatullah & Tinsley, 2013). Indeed, prior work demonstrates women (Babcock & Laschever, 2009) and Asians (Lu, 2022)—show salary disparity, partly due to their aversion to asking for a greater salary.

Further research is needed to confirm our conclusion. First, studies should explicitly measure if salary preference encodes some risk of rejection. For example, we could elicit preferences in an isomorphic context where the fear of being denied an offer is elevated or removed. If confirmed, this suggests preference elicitation methods should be altered to better separate actual preferences from other utility sources (Brown & Curhan, 2012). Another potential study could examine whether risk-averse disputants adjust their stated preferences out of fear of non-agreement by having them select between a raw offer and one adjusted for their risk aversion (i.e., just using a linear issue curve). In this case, their preferring the revised proposal would imply inaccurate self-reported preferences. Lastly, one limitation of our work stems from not having *boss* preferences and instead flipping the within issue preferences from an employee perspective; for future analysis, we should survey actual employers to solicit more realistic preferences from their viewpoint.

References

Adler, N. E., Epel, E. S., Castellazzo, G., & Ickovics, J. R. (2000). Relationship of subjective and objective social status with psychological and physiological functioning: Preliminary data in healthy, white women. *Health psychology*, *19*(6), 586.

Amanatullah, E. T., & Tinsley, C. H. (2013). Punishing female negotiators for asserting too much… or not enough: Exploring why advocacy moderates backlash against assertive female negotiators. *Organizational Behavior and Human Decision Processes*, *120*(1), 110–122.

Babcock, L., & Laschever, S. (2009). Women don't ask. In *Women don't ask.* Princeton University Press.

Binmore, K., Rubinstein, A., & Wolinsky, A. (1986). The nash bargaining solution in economic modelling. *The RAND Journal of Economics*, 176–188.

Bontempo, R. N., Bottom, W. P., & Weber, E. U. (1997). Cross-cultural differences in risk perception: A model-based approach. *Risk analysis*, *17*(4), 479–488.

Booij, A. S., & Van de Kuilen, G. (2009). A parameter-free analysis of the utility of money for the general population under prospect theory. *Journal of Economic psychology*, *30*(4), 651–666.

Boutilier, C., Brafman, R. I., Domshlak, C., Hoos, H. H., & Poole, D. (2004). Cp-nets: A tool for representing and reasoning withconditional ceteris paribus preference statements. *Journal of artificial intelligence research*, *21*, 135–191.

Brown, A. D., & Curhan, J. R. (2012). The utility of relationships in negotiation. In *The oxford handbook of economic conflict resolution* (p. 138). Oxford University Press.

Croson, R., & Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic literature*, *47*(2), 448–474.

Fehr-Duda, H., De Gennaro, M., & Schubert, R. (2006). Gender, financial risk, and probability weights. *Theory and decision*, *60*, 283–313.

Genolini, C., & Falissard, B. (2011). Kml: A package to cluster longitudinal data. *Computer Methods and Programs in Biomedicine*, *104*(3), e112-e121. doi: https://doi.org/10.1016/j.cmpb.2011.05.008

Hale, J., Kim, P., & Gratch, J. (2022). Preference interdependencies in a multi-issue salary negotiation. In *Proceedings of the 22nd acm international conference on intelligent virtual agents.* New York, NY, USA: Association for Computing Machinery.

Hernandez-Arenaz, I., & Iriberri, N. (2019). A review of gender differences in negotiation. *Oxford Research Encyclopedia of Economics and Finance*.

Higgins, E. T., Friedman, R. S., Harlow, R. E., Idson, L. C., Ayduk, O. N., & Taylor, A. (2001). Achievement orientations from subjective histories of success: Promotion pride versus prevention pride. *European Journal of Social Psychology*, *31*(1), 3-23. doi: https://doi.org/10.1002/ejsp.27

Lee, M. K., & Baykal, S. (2017). Algorithmic mediation in group decisions: Fairness perceptions of algorithmically mediated vs. discussion-based social division. In *Proceedings of the 2017 acm conference on computer supported cooperative work and social computing* (pp. 1035–1048).

Liberman, N., Idson, L. C., Camacho, C. J., & Higgins, E. T. (1999). Promotion and prevention choices between stability and change. *Journal of Personality and Social Psychology*, *77*(6), 1135-1145.

Lu, J. G. (2022). Asians don't ask? relational concerns, negotiation propensity, and starting salaries. *Journal of Applied Psychology*.

Marks, M., & Harold, C. (2011). Who asks and who receives in salary negotiation. *Journal of Organizational Behavior*, *32*(3), 371-394.

Meertens, R. M., & Lion, R. (2008). Measuring an individual's tendency to take risks: The risk propensity scale1. *Journal of Applied Social Psychology*, *38*(6), 1506-1520. doi: https://doi.org/10.1111/j.1559-1816.2008.00357.x

Neale, M. A. (1997). New recruit. *Teaching materials for negotiations and decision making. Evanston, IL: Northwestern University, Dispute Resolution Research Center*.

Northcraft, G. B., Preston, J. N., Neale, M. A., Kim, P. H., & Thomas-Hunt, M. C. (1998). Non-linear preference functions and negotiated outcomes. *Organizational Behavior and Human Decision Processes*, *73*(1), 54–75.

Sela, A. (2018). Can computers be fair: how automated and human-powered online dispute resolution affect procedural justice in mediation and arbitration. *Ohio St. J. on Disp. Resol.*, *33*, 91.

Shah, N. (2017). Spliddit: two years of making the world fairer. *XRDS: Crossroads, The ACM Magazine for Students*, *24*(1), 24–28.

Sternlight, J. R. (2020). Pouring a little psychological cold water on online dispute resolution. *J. Disp. Resol.*, 1.

Tang, S., Zhang, X., Cryan, J., Metzger, M. J., Zheng, H., & Zhao, B. Y. (2017). Gender bias in the job market: A longitudinal analysis. *Proceedings of the ACM on Human-Computer Interaction*, *1*(CSCW), 1–19.

Tellhed, U., & Björklund, F. (2011). Stereotype threat in salary negotiations is mediated by reservation salary. *Scandinavian Journal of Psychology*, *52*(2), 185-195.

Thiessen, E. M., & Soberg, A. (2003a). Smartsettle described with the montreal taxonomy. *Group Decision and Negotiation*, *12*(2), 165.

Thiessen, E. M., & Soberg, A. (2003b). Smartsettle described with the montreal taxonomy. *Group Decision and Negotiation*, *12*(2), 165.