

Algorithmic Wage Negotiations: Applications to Paid Crowdsourcing

John J. Horton
Harvard University
383 Pforzheimer Mail Center
56 Linnaean Street
Cambridge, MA 02138

john.joseph.horton@gmail.com

Richard J. Zeckhauser
Harvard Kennedy School
79 JFK Street
Mailbox 41
Cambridge, MA 02138

richard_zeckhauser@harvard.edu

ABSTRACT

This paper highlights the problem that buyers and sellers encounter in arriving at prices in a distributed labor market. We argue that automated negotiation might partly remedy this problem and introduce a program, “hagglebot,” that we used to negotiate payment rates for an image-labeling task with workers at Amazon’s Mechanical Turk. We report the results of an initial pilot experiment that used a simple bargaining game and randomly assigned subjects to receive either a low (1 cent) or high (5 cents) offer to perform a follow-on task after they completed an initial fixed-price task. Subjects were generally reluctant to make counter-offers or end negotiations. As a result, subjects who received the 1-cent offer ended up with far lower average wages than subjects in the 5-cents group.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics; J.m [Computer Applications]: Miscellaneous

General Terms

Human Factors, Economics, Experimentation

Keywords

Crowdsourcing, Amazon’s Mechanical Turk, Human Computation

1. INTRODUCTION

Would-be buyers and sellers of labor have to agree about wages and the quantity of work to be performed. This general problem is common to all labor markets and is often resolved through negotiation. However, negotiation is not practicable in all situations. In particular, when stakes are low, negotiation is unattractive because it is time-consuming; a prolonged negotiation can quickly dissipate any surplus that might be gained through trade. In a thick market (with many buyers and sellers for a good), posted prices and auction processes are often preferable to negotiation, and vice versa in a thin market (with few buyers and sellers). The relationship between stakes and market thickness and parties’ willingness to bargain appears in many contexts: people rou-

tinely bargain over houses, cars and salaries—they (at least in the U.S.) rarely bargain over groceries or haircuts.

In distributed labor markets, the work performed by any single worker is often vanishingly small. For this reason, buyers’ take-it-or-leave-it prices tend to predominate, and negotiations are rare. However, as we discuss below, posted prices have many drawbacks when used in labor markets. Negotiation might yield better net results, if it could be done more cheaply. One potentially technical way to lower the costs of negotiation would be to have it performed by a machine. An automated bot that can bargain on behalf of buyers and/or sellers has essentially no time-based opportunity cost and could help parties reach mutually beneficial agreements more efficiently.

In this note, we describe a new software tool called hagglebot that we designed to serve as an automated negotiating agent for buyers of labor. Our research using this tool is at a very early stage. This note is limited to describing hagglebot and our initial pilot study conducted in Amazon’s Mechanical Turk marketplace.

1.1 Prior Work

Negotiation has received a great deal of research attention, both positive and normative [7]. Scholars and practitioners alike have noted the gap between how people actually bargain and how they should bargain. Like most scenarios in which humans have to deal with uncertainty, negotiators make a variety of systematic mistakes: they get anchored by irrelevant numbers, overweight small probabilities, are risk-seeking in the domain of losses [5], place unwarranted emphasis on the status quo [8] and so on. Even setting aside behavioral biases, negotiators make mistakes for the simple reason that negotiations are complex and cognitively demanding.¹

One way to improve both prescriptive and descriptive knowledge is to collect more data on how people actually negotiate in real settings and, critically, to experimentally manipulate factors to determine causality. Unfortunately, creating realistic negotiation scenarios in the laboratory is challenging. However, researchers in a number of disciplines—with computer science leading the way—have begun running experiments online using online labor markets, where far greater realism is possible. Some examples in economics

¹For this reason, there have been attempts to create decision-support tools for negotiation. Tests of some of these tools have shown that they can significantly improve outcomes [2].

include [6], [1] and [3]. Horton, Rand and Zeckhauser ([4]) argue that online experiments can offer a high degree of both internal and external validity.

2. REACHING PRICES FOR LABOR

In any potential economic exchange, both the buyer and the seller have reservation prices: for the buyer, it is the maximum amount they are willing to pay; for the seller, it is the minimum amount they are willing to accept. If the buyer’s reservation price is greater than the seller’s, a mutually beneficial transaction is possible.

Posit that a beneficial exchange exists. Within the range of feasible prices, buyers and sellers are locked in a zero-sum game, with any increase in price benefiting sellers and hurting buyers. As the two players jockey for a larger slice of the pie, they may make strategic demands that prevent a deal, in which case both lose.

For commodities sold in markets with many buyers and sellers, strategic considerations are moot: prices are determined by market-level supply and demand. Would-be buyers and sellers can look to the market to learn prices. In contrast, in thin markets, market prices may be unobservable and uninformative.

Labor markets can be very thin in this sense, as transactions are usually for a specific job and a specific worker. For example, there is no active market in “Joseph Jones will mow my lawn on July 15th for 1 hour,” even though we can get a general sense of the price of semi-skilled manual labor. The absence of a market to reveal prices forces parties to bargain or simply accept the limitations inherent in using posted prices in thin markets.

Prices offered on take-it-or-leave-it terms are the overwhelming norm in distributed labor markets. From a buyer’s perspective, setting those prices is difficult. Without knowing the distribution of workers’ reservation wages for a particular task, the buyer does not know how many workers will accept a given offer; if he wishes to have n people complete the task within m days, he will be at somewhat of a loss. The problem becomes even more complex if the buyer allows workers to complete multiple tasks, as both fatigue and experience become relevant. The power to negotiate—and learn from negotiation—can help buyers avoid the pitfalls of offering too much (getting a bad deal) or too little (not getting the desired amount of work done). Through price discrimination, negotiation also has the potential to get any fixed amount of work done at an overall cheaper price.

3. HAGGLEBOT

The pioneer hagglebot is designed to negotiate over price with a worker at Amazon’s Mechanical Turk (MTurk). Hagglebot can make offers, solicit counter-offers and make decisions about whether to accept counter-offers or end negotiations. It does all of these things through a natural language chat interface, with worker inputs structured in such a way that only counter-offers are entered in free-form text. While any potential bargaining strategy could be pursued by hagglebot, for experimental purposes, we have selected a simplified form of bargaining.

In our experiment, workers first completed an image-labeling task for a fixed price. After completing this initial task, workers negotiated with hagglebot over the price of labeling an additional image. (This avoided the challenge of get-

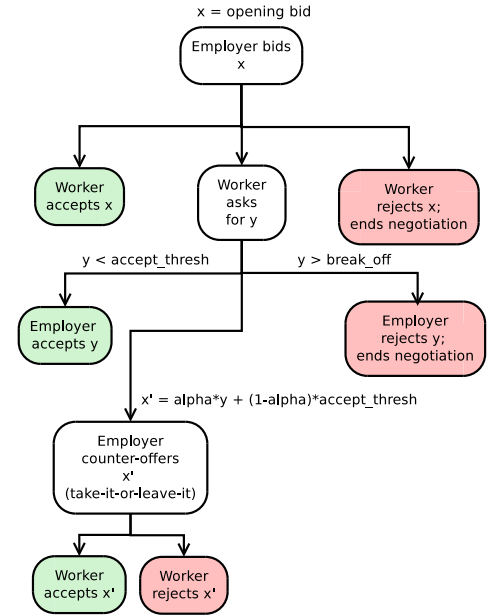


Figure 1: Restricted bargaining space

ting workers to bid on tasks that they had not previously encountered, and hence could not estimate the difficulty in performing.)

3.1 Restricted bargaining space

Unlike auctions or posted prices, negotiations can proceed in any number of ways. There are no formal rules; buyers and sellers are free to make offers, counter-offers and end negotiations. These steps can proceed in any order, and can be conditioned upon what has already transpired. The enormous “bargaining space” implied by this flexibility requires us to narrow the scope of allowed actions in order to make empirical progress, yet the narrowing must preserve realism. For the initial pilot studies, we programmed hagglebot to negotiate in a particular structure, as illustrated in Figure 1. The bargaining had at most three phases: an initial offer from hagglebot, a counter-offer from a worker and a take-it-or-leave-it counter to the worker’s counter-offer from hagglebot.

3.1.1 Parameters

The buyer makes an initial offer of `opening_bid`. A worker can accept the offer, end negotiations or make a counter-offer. If `counter < accept_thresh`, work commences at the price counter. If `counter > break_off`, negotiations end. For values between `accept_thresh` and `break_off`, the hagglebot counter-offers

$$\text{alpha} * \text{accept_thresh} + (1 - \text{alpha}) * \text{counter}$$

If the worker accepts, work commences at this new price. Otherwise, negotiations end and the new work is not undertaken.

3.2 Interface

Figure 2 illustrates the interface where workers negotiate with hagglebot. The paradigm of a chat window was

used and offers were expressed in natural language. Responses from workers were necessarily constrained to their three choices.




Figure 2: Chat interface

3.3 Technical Implementation

Hagglebot is written in Python and runs on the Google App Engine (GAE). Task postings to MTurk and worker payments are done using the Boto toolkit.² Experiment parameters are loaded into hagglebot via a structured text document containing details on payment rates, titles, group parameters, etc. Hagglebot automatically launches an experiment and collects results. The actual haggling algorithm, “hagglorithm,” was written as a stand-alone script that abstracts from the details like the interface or the storage of data. This enables future researchers to easily test other bargaining strategies.

4. EXPERIMENT

For the pilot experiment, our goal was to test the software and determine whether workers are as willing to negotiate with hagglebot as they are with a human. As a paid crowdsourcing task, we used image labeling (workers were asked to give descriptive labels for an image).

1. All workers perform an initial image-labeling task for a fixed payment.
2. Workers are offered the chance to perform another image-labeling task. Negotiation is performed by hagglebot, in accordance with experimentally-assigned parameters.
3. If negotiation succeeds, work commences at the agreed price. Otherwise, negotiation ends and the second task remains unperformed.

²<http://code.google.com/p/boto/>

4.1 Experimental Design

The hagglebot parameters used for the experiment are given in Table 1. There were two experimental groups, HIGH and LOW. The two groups were identical except for the opening bid: in HIGH, the bid was 5 cents; in LOW it was 1 cent. The experiment was run on September 6th, 2010. One hundred subjects were recruited.

Table 1: Hagglebot parameters

Parameters	HIGH	LOW
opening_bid	5	1
accept_thresh	10	10
reject_thresh	20	20
alpha	.5	.5

4.2 Results

Table 2 presents a cross-tabulation of subjects’ responses to the initial offer by treatment group. A greater number of workers in LOW ended negotiations immediately after receiving offers and a greater number made counter-offers compared to what was observed in HIGH. However, a chi-square test gives a p-value of .29, suggesting that we are quite likely to get this pattern of results by chance even if there was no cross-group differences in probabilities. The main conclusion that can be drawn is that a larger sample is needed, as well as a smaller initial offer to obtain more variation in outcomes. We may also conclude that, in this example, we were well-served by starting with a low wage offer.

Table 2: Cross-tabulation of initial responses to the opening bid, by experimental group

	HIGH	LOW
Accepted offer	36	34
Ended negotiations	6	8
Made Counter offer	6	13

Notes: The chi-square test result for the cross-tabulation, with 2 degrees of freedom, is $p = 0.29$.

Because most workers accepted the opening bid regardless of group assignment, there are large differences in the mean wage earned by the different groups. Figure 3 plots the negotiated, second-task price versus experimental group. Points are jittered horizontally to prevent over-plotting. It shows that the mean price in HIGH was considerably higher than the mean price in LOW, with no overlap in the 95% confidence intervals. Confirming what is graphically clear, the regression line is:

$$price = \underbrace{-2.991}_{[0.606]} LOW + \underbrace{5.791}_{[0.311]} \quad (1)$$

with $N = 93$ and $R^2 = 0.2$.

5. CONCLUSION

This note introduces a planned program of research aimed at understanding how workers negotiate and how this knowledge can be used in applications. The simple pilot described




Figure 3: Agreed wage (per image) for successful negotiations (points horizontally jittered) .

already illustrates, for example, the advantages of more flexible pricing: wages were considerably lower in LOW without inducing significant differences in uptake. Potential future extensions include testing different parameter configurations, examining actual label output and conditioning worker behavior on observable characteristics. Haggbot might then be modified to allow for bargaining strategies that adapt in real time to reflect knowledge gained from the outcome of bargaining.

6. REFERENCES

- [1] D. Chandler and A. Kapelner. Breaking monotony with meaning: Motivation in crowdsourcing markets. *University of Chicago mimeo*, 2010.
- [2] E. Chen, R. Vahidov, and G. Kersten. Agent-supported negotiations in the e-marketplace. *International Journal of Electronic Business*, 3(1):28–49, 2005.
- [3] J. Horton and L. Chilton. The labor economics of paid crowdsourcing. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, 2010.
- [4] J. Horton, D. Rand, and R. J. Zeckhauser. The online laboratory: Conducting experiments in a real labor market. *NBER Working Paper w15961*, 2010.
- [5] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, 47(2):263–291, 1979.
- [6] W. Mason and D. Watts. Financial incentives and the ‘performance of crowds’. In *Proceedings of the ACM SIGKDD Workshop on Human Computation (HCOMP)*, 2009.
- [7] H. Raiffa. *The art and science of negotiation*. Belknap Press, 1982.
- [8] W. Samuelson and R. Zeckhauser. Status quo bias in decision making. *Journal of risk and uncertainty*, 1(1):7–59, 1988.