Can Markets Help?: Applying Market Mechanisms to Improve Synchronous Communication

Gary Hsieh[1], Robert Kraut[1], Scott E. Hudson[1], Roberto Weber[2]

[1]Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15213
{garyh, robert.kraut, scott.hudson}@cs.cmu.edu

[2]Department of Social & Decision Sciences
Carnegie Mellon University
Pittsburgh, PA 15213
{rweber}@andrew.cmu.edu

ABSTRACT

There is a growing interest in applying market mechanisms to tackle everyday communication problems such as communication interruptions and communication overload. Prior analytic proofs have shown that a signaling and screening mechanism can make senders and recipients of messages better off. However, these proofs make certain assumptions that do not hold in real world environments. For example, these prior works assume that there are no transaction costs in a communication market and that monetary incentives are the only motivators in communication between strangers.

This research builds upon prior analytic work and empirically tests the validity of the claim that signaling and screening mechanisms will improve communication welfare. Our results show that while these types of markets can indeed improve communication welfare, a simpler, less expressive fixed-price market can lead to higher welfare than a more expressive, variable pricing and screening mechanism. Findings from this study also provide valuable insights for technology designs. For example, these results suggest the need to reduce cognitive overhead in using communication markets.

Author Keywords

Computer mediated communication, market mechanisms, empirical studies, economics

ACM Classification Keywords

H5.3. Information interfaces and presentation (e.g., HCI): Group and Organization Interfaces – *computer supported cooperative work*

INTRODUCTION

Modern synchronous communication technologies, such as instant messengers (IMs) and mobile phones, allow for immediate, real-time communication. They support both formal and informal communications that are vital to the workplace and everyday communication [8,12,17,19].

CSCW'08, November 8–12, 2008, San Diego, California, USA.

However, current synchronous communication technologies are far from perfect. A major problem is information asymmetry when starting a communication. Senders (initiators of communication) know what the communication is about but do not know the availability and the responsiveness of the receivers (recipients of communication requests). On the other hand, receivers know their own availability but do not know the purpose and content of communication. This information asymmetry leads to inefficient allocation of time and attention, creating both undesired and costly interruptions and important and urgent messages that are ignored.

One solution to the problem of information asymmetry is to use economic markets. Markets, through prices and incentives, can be a valuable mechanism for allocating resources efficiently, even in situations where information is highly asymmetric [7]. Much of the existing work on applying markets to communication has focused on the problem of spam. It has been shown that charging a tax or postage to communicate causes senders to become more selective in their communication requests, therefore reducing the amount of undesired communication [4,14,15]. More recent research has focused on market mechanisms that allow for senders to purchase the right to communicate with receivers [6,10]. These types of markets promote valuable communication that would have otherwise been screened out by taxes, postage or other simple screening mechanisms.

Specifically, Loder, van Alstyne and Wash present the Attention Bond Mechanism [17]. With this mechanism, receivers post a take-it-or-leave-it “price” for their attention, and senders can attach payments along with their communication requests. Messages with payment amounts less than the posted price are rejected. This type of market leads to efficient outcomes by ensuring that both parties benefit. Senders do not commit to paying more than they believe the communication is worth, while receivers only attend to communication requests if the communications are more valuable than their current task at hand. Using analytic modeling, Loder et al. demonstrated that such a design frequently led to higher welfare (measured by financial gains) than a perfectly applied tax and even over a perfect spam filter in certain scenarios.

However, research on communication markets has primarily been analytic and makes unrealistic assumptions about the humans using these markets. They assume that humans are rational, trying to maximize their utility, and that utility can be represented by monetary gain. They also assume that senders can accurately estimate the value of a communication to them and that receivers can accurate estimate interruption costs. Finally, they also assume that both the estimations and the calculations of utility are themselves costless, requiring no mental effort. Without a study of humans' interactions with this communication market in real time, we are faced with many questions. Can people actually use markets well enough under time constraints to improve their utility? How elaborate must a market be for it to benefit communication? If there are costs associated with decisions, how does that affect market mechanism designs for communication?

To explore these questions, we compared three communication systems in a laboratory experiment, modeled on the common communication practice of question and answer exchanges. Participants are provided with one of three synchronous communication systems to ask for help and to provide help to other players. One is a no-market, baseline condition; the sender sends a request with no financial incentives attached and the receiver has to decide whether to respond. Another is a variable-price bidding and screening system, similar to the Attention Bond Mechanism proposed by Loder et al. Here senders offer to pay an individually set price for an answer, and receivers accept communication only if this price exceeds their individually set reservation price. The third is a fixed-price system that is a less expressive version of the variable-price system, in which senders pay receivers a fixed price if communication occurs. Our results show that participants in both the fixed and variable-price market mechanisms performed better than participants in the no-market baseline condition. However, the less expressive fixed-price market surprisingly led to higher earnings than the more expressive variable-price market.

This work provides three major contributions. First, it empirically demonstrates that applying markets in synchronous communication can improve communication efficiency. Second, while markets may facilitate communication, markets are not used perfectly. Our findings provide a more accurate depiction, relative to the assumptions of economic theories, of how communication decisions are made. Third, by showing that fixed-price markets work better than variable-pricing ones, the research offers insight into the simplifications needed to design appropriate market-based communication systems

COMMUNICATION MARKETS AND MECHANISMS

Information asymmetry is an underlying problem when starting a communication. Each party involved knows how the communication may affect him or herself, but not how it may affect the other person. This imbalance of information prevents senders from interrupting receivers at optimal times and prevents receivers from handling communication requests more effectively.

While technology and legal solutions exist to combat unwanted communication [*e.g.* 3,20], our work focuses on using economic markets. Technology and regulatory solutions, such as filters and whitelists/blacklists, work by blocking wasteful communication, whereas markets may encourage valuable communication. For example, if a sender could earn $10 from a communication that provides the receiver no benefit, filters and whitelist/blacklists would attempt to block this communication because it is an "unwanted" communication from the receiver's viewpoint. Under an economic approach, however, the sender may be able to allocate some of his monetary gain to the receiver, making it worthwhile for a receiver to respond to the communication. Such outcomes would benefit both communication parties and increase overall welfare.

Seriosity (www.seriosity.com) is a startup company that attempts to apply market mechanisms to email. Using *serios* – a point or currency-like system – Seriosity's productivity application supports email use by allowing the users to indicate the relative value and importance of a message. While success is yet to be determined, there is undoubtedly a growing interest in leveraging markets to facilitate communication. For these services to be successful, we need a systematic understanding of the effects of market mechanisms in everyday communication. With this knowledge, designers and implementers can build better technology to benefit users.

Using Market Mechanisms to Improve Communication

Using market mechanisms to facilitate decentralized allocation of resources is one of the fundamental principles of economics [7]. In a market, senders use prices to express their estimate of the value of the communication. Receivers use the price information to better decide how to handle the communication request. Senders and receivers do not need to have full information disclosure prior to communication nor is any one party responsible for deciding if the communication should happen.

There are many different approaches in deploying market mechanisms for communication. Research on spam has explored the use of stamps, surcharges and auctions to reduce unwanted communications [4,14,15]. These types of mechanisms shift the burden of identifying unwanted communication to senders, who are knowledgeable about the content of communication [23]. Incremental costs for sending a message force the senders to be selective, sending messages only if they believe the value of the message reaching a receiver is more than the sending price.

However, like filters and regulations, these mechanisms are one-sided and do not simultaneously take into account both parties' communication contexts. What is valuable to the sender may not be valuable to the receiver; the sender's willingness to pay more for a communication does not mean that communication is more desirable to the receivers.

Furthermore, communication that is potentially valuable to the receivers may be deterred because senders do not want to pay the surcharges. This is perhaps why Kraut et al.'s empirical study on using variable rate postage to reduce spam showed that while such mechanism reduced communication, receivers did not see postage as a signal of communication value [14].

Therefore, in our work, we explore the use of two-sided markets where senders pay receivers for their valuable resources. Having monetary transfers allows the senders to make binding bids for communication (bonds in Loder et al's terminology), which will be paid to receivers if communication occurs. This property allows the senders to provide incentives to receivers to communicate; a communication that originally is valuable only to the sender can become valuable to both parties.

Limitations in Existing Work

Analytic proofs demonstrating the effectiveness of communication markets rely on debatable assumptions on the human communicators [10,17,23]. They assume that communicators are rational money-maximizers – senders will not offer to pay more than their financial value for communicating and receivers will not offer to provide help if they are not given sufficient financial reward. They also assume that at the time of communication, each party can estimate the value of the communication, in financial terms, in an accurate and unbiased manner (for example, instantaneously and without mental effort or error).

However, these assumptions do not necessarily hold in real world communication. People often lend assistance to others for reasons other than monetary reward [1]. Given that people may respond to communication requests without any financial incentives, it is unclear whether the additional explicit reward that is central to the effectiveness of market mechanisms is likely to achieve a noticeable difference from the existing intrinsic motivation. Furthermore, human cognitive limitations may mitigate the effectiveness of markets as allocation mechanisms [13, 21], by introduction "transactions costs" to the use of markets [2]. These potential issues motivate the need to explore what happens when payment-based market mechanisms are applied to real-time human communication. If the assumptions underlying analytical models regarding human behavior and motivations do not hold in communication markets, will the application of these mechanisms still improve communication? Answering this question requires carefully exploring the actual behavior of communicating parties under a market mechanism.

Synchronous Communication Markets

In this study, we focus on a specific, but common use of synchronous communication—the general help-seeking scenario [19]. As with all markets, there are buyers and sellers: the *sender*, or the person who initiates the communication requesting for help, acts as the buyer, while the *receiver*, or the person who is the recipient of the help request, acts as the seller. Instead of apples and oranges, in our scenario, the goods in transaction are the solutions to one's task. We assume the solution transacted between them is indivisible. This implies that communication is between two parties (dyadic) and that when it is finished, the sender's task is completed. To keep our market realistic, we allow each sender to request communication from many receivers (a one-buyer and multiple-seller market). Harper et al. have shown the value of market mechanisms in real-world question-and-answer scenarios similar to the ones we model here [9].

In many real-world communication markets, the actors would take into account both prior and future relationships with specific individuals in the market. For example, in the Loder et al. implementation described previously, sellers may decide not to charge a fee to friends or others with whom they have had successful communication in the past. In the experiment described below, communication requests are anonymous to control for these relationship factors.

To empirically study communication markets, we compare three different communication systems: a no market baseline, a variable-price market, and a fixed-price market:

No market condition (*no market*). To study communication markets, we must start with a no intervention baseline. As with most current synchronous technologies, senders in the *no market* condition can request help without financial cost and receivers have no financial incentives to answer help requests.

SV_0: $Gains_{Help}$

RV_0: $-Costs_{Opportunity}$

Equation 1: Sender's (SV_0) and Receiver's (RV_0) Financial Valuation of a Help Communication

To maximize welfare, in our help-seeking scenario, senders should ask for help if their financial valuation of help (benefits minus costs) is greater than zero. Similarly, receivers should provide help only if their financial valuation of giving help is greater than zero. Therefore, in our help-seeking scenario, where communications guarantee aid to the senders with their tasks, senders in the *no market* condition who are extrinsically motivated should always ask for help (gain from help is greater than 0, Eq. 1). On the other hand, receivers in the *no market* condition should never respond to incoming help requests since giving help prevents them from working on their own tasks and gives neither reputation benefits nor direct reward (there is always an opportunity cost, Eq. 1).

However, significant prior research in economics and psychology shows that even without extrinsic incentives, altruism and other intrinsic motivations cause people to voluntarily help even anonymous others [1,11]. These intrinsic motivators may add additional costs to senders asking for help but may also motivate receivers to voluntarily provide help. With the addition of these costs and gains, we would expect that in the *no market* condition, senders will refrain from bombarding receivers with excessive help requests (i.e. help requests will be less than 100%) and that receivers will offer some help (Equation 2).

SV_{NM}: $Gains_{Help} - Costs_{Intrinsic}$

RV_{NM}: $Gains_{Intrinsic} - Costs_{Opportunity}$

Equation 2: Simple Model of Sender's (SV_{NM}) and Receiver's (RV_{NM}) Valuation of Help Considering Intrinsic Motivators

In terms of individual sender and receiver welfare, we would expect financially inefficient outcomes to result from the use of this communication mechanism. On one hand, help exchanges that are extremely beneficial to the senders may not occur because receivers' opportunity costs may limit their help offers. On the other hand, when receivers do provide help due to intrinsic motivations, they are usually worse off financially because of it. This is precisely the problem with existing synchronous help.

Variable-price payment market (*variable market*). A variable-price payment market is comparable to the Attention Bond Mechanism in that the communication may occur at many different prices, allowing senders and receivers to use a continuous scale to express their communication value. In an Attention Bond Mechanism, receivers post a take-it-or-leave-it price and the senders decide whether or not to pay that price for communication. In our variable-price mechanism, senders place a bid on how much they would pay for the communication, and receivers set a reservation price on how much they need to be paid to respond to a communication request. Communication occurs when the bid is higher than or equal to the reservation price. The payment amount is the reservation price as set by the receiver. We chose this market design as it seems more appropriate for synchronous communication. Receivers should only be required to update the take-it-or-leave-it price at the moment of a request for help. Furthermore, there may be privacy concerns with posting take-it-or-leave-it prices, because receivers' prices signal their business or value of their time. With our bidding and screening design, neither party needs to disclose this information to the other.

SV_{M}: $Gains_{Help} - Costs_{Intrinsic} - Payment_{Help}$

RV_{M}: $Gains_{Intrinsic} - Costs_{Opportunity} + Payment_{Help}$

Equation 3: Simple Model of Sender's (SV_{M}) and Receiver's (RV_{M}) Valuation of Help in a Payment Transfer Market

Rational senders in the *variable market* should place bids on help requests based on the value the communication has for them. For example, if they gain $0.12 from receiving help, senders should offer to pay up to $0.12, minus the costs discussed previously. In our setup, we allow senders to place $0 help request bids. Since help requests in the *no market* condition are essentially $0 bids, the *variable market* condition should result in about the same number of help requests as is in the *no market* condition. On the other hand, rational receivers should dynamically adjust their reservation price to match their net cost from providing help. Because receivers in the *variable market* condition receive greater financial compensation when they offer help than in the *no market* condition, help should occur more frequently. Senders and receivers should be able to use the market rationally; both getting help and giving help will improve their welfare (Equation 3).

Fixed-price payment market (*fixed market*). In a *fixed market* condition, senders must offer to pay receivers a fixed-price for each completed communication. In our study, the fixed price is set to $0.20. The fixed-price condition represents a less expressive and less flexible version of the variable-price market. The payment value is restricted to just one value, instead of being opened to the continuous range of values available under the variable-price mechanism. Therefore, senders and receivers can less precisely express the value and cost of the communication.

In the *fixed market* condition, rational senders should ask for help as long as the value of the communication exceeds the fixed payment threshold and receivers should offer help only if their opportunity cost to communicate is lower than the fixed payment. Similar to the variable-price condition, we expect the market to be used rationally and both getting help and giving help will lead to improvements in task performance. But because it is less fine-grained, the frequency of help produced by the mechanism will be lower. For example, in our fixed-price market, senders who would gain $0.19 from receiving help would not ask for help because their gain is less than the $0.20 fixed price. In contrast, in the *variable market*, senders can offer less than $0.19 and enable the welfare-increasing help to occur.

However, the *variable market* may also be less effective than the *fixed market*, producing lower overall welfare. Cognitive limitation may interact with the additional complexity of the variable market to reverse the potential gains from increased expressiveness. The fine-grained decision that must be made in the *variable market* condition is much more complicated and requires more time and attention for the decision than the coarse-grained decision in the *fixed market* condition. Instead of the binary decision of "should I pay $0.20 for help?" senders in the *variable market* condition are instead faced with two decisions—should I pay for help?" and "how much should I pay?" Similarly, rather than simply deciding whether or not to accept a $0.20 payment for responding to a communication request, as in the *fixed market* condition, receivers in the *variable market* condition must determine the precise cost to interruption. These additional costs for using the market mechanism are a type of transaction cost incur during the decision process, so they are hard to model as an independent factor. They are better considered as meta-level costs associated with using more complex systems such as the *variable market*.

Predictions for Communication Markets

Based on our reasoning presented above, we have the following hypotheses for our markets.

Market Use:

H1. The percentage of help requests will be higher in both *no market* and *variable market* conditions compare to the *fixed market* condition; the percentage of help requests will not differ between *no market* and *variable market* conditions.

H2. The percentage of help exchanges will be highest in the *variable market* condition, second highest in the *fixed market* condition, and lowest in the *no market* condition.

H3. Senders in all conditions will benefit financially from using the communication system; receiving help will correlate with improvements in welfare.

H4. Receivers will benefit financially from using the communication system in the market conditions, but not in the *no market* condition. In the *no market* condition, providing help will negatively correlate with financial welfare.

Market Efficiency:

H5. Market mechanisms (*fixed* and *variable market* conditions) will lead to higher welfare than *no market* condition.

H6. A *fixed*-price (less expressive, less complicated) system will lead to higher welfare than a *variable*-price (more expressive, more complicated) system.

DESIGN OF EXPERIMENT

To test these hypotheses, we developed a study in which participants work on a task independently in four-person sessions. Participants can broadcast help request and provide help to other participants in their session.

Each session was assigned to one of three different communication mechanisms (*variable market*, *fixed market*, *no market*) and each participant played both the roles of senders and receivers concurrently. Tasks were solving memory/concentration puzzles (see Figure 1).

The goal of the puzzle was to find the locations of the matching celebrity faces from a set of cards containing pairs of celebrity faces. At the start of each puzzle, the cards were placed face down. Participants were allowed to flip over two cards at once. If the two cards matched, they would remain face up; otherwise, they would be turned face down automatically. We selected this game because: 1) most participants are familiar with it; 2) it provides the participants a good sense of task progress; 3) it is quick, allowing us to collect data from repeated plays, and 4) most importantly, interruptions lead to costs similar to real life communication (players cannot continue with their work when communicating).

Participants earned one cent for each pair of faces they matched. Participants also earned a puzzle bonus for solving the whole puzzle (*i.e.* matching all the faces in the set). The bonus was \$0.25, \$0.50 or \$0.75 and was randomly selected and made visible at the start of each puzzle. Variations in puzzle bonus allow us to examine the effects of task value on participant's assessment of the value of communication. Participants had 90 seconds to work on each puzzle. After participants solved a puzzle or if the time expired before solving one, the puzzle refreshed itself (all cards turned face down with location randomized). The size of individual puzzles varied randomly from 26 to 40 cards. Puzzle value and puzzle size were not correlated, and participants were informed of this independence.

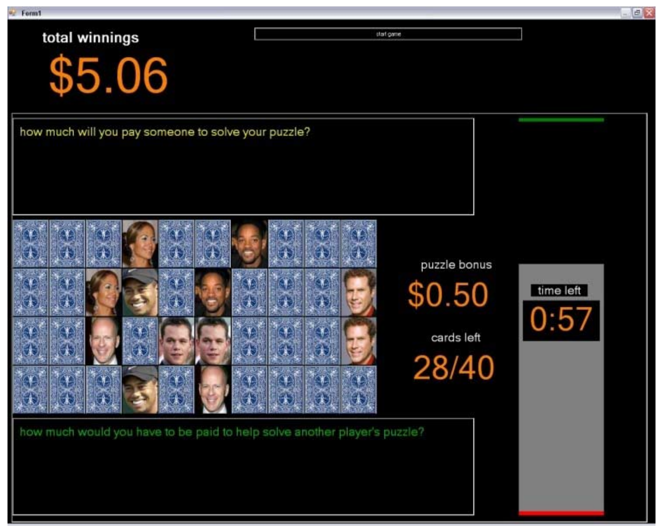

Figure 1: Game interface

Communication Mechanism Manipulations

During each puzzle, participants had one opportunity to ask for help from other participants in their session. The time for this opportunity was randomly selected, from the start of the puzzle up until 20 seconds remained on that particular puzzle. When this "bidding opportunity" occurred, that individual participant's puzzle paused. Prior to continuing work on their puzzle, participants had to decide if they wanted to ask for help. If they chose to ask for help, their help request was broadcast to all other players anonymously.

Participants were given one randomly-timed opportunity to ask for help during each puzzle. This design is different from what we might expect from real-world communication, where people can request for and decide to provide help at any point in their task. We chose this design because it allows us to measure the time spent on making each decision and because it allows us to randomly sample choices made at several points during the puzzle-completion process.

If a request was made, each receiver's puzzle paused. Receivers then had to decide how to handle the request. When receivers decided to help the sender, the software controlled the help interaction: all help exchanges were computerized and took 15 seconds. This controlled help guaranteed help exchanges when requests were accepted and enabled receivers to know beforehand the exact time cost for providing help. As soon as any one receiver agreed to help, all help requests to the other players were canceled. This design made the setup more representative of the real time nature of the communication—if a sender is helped by someone, any subsequent help from others will add no value to the sender.

(a)
(b)
(c)

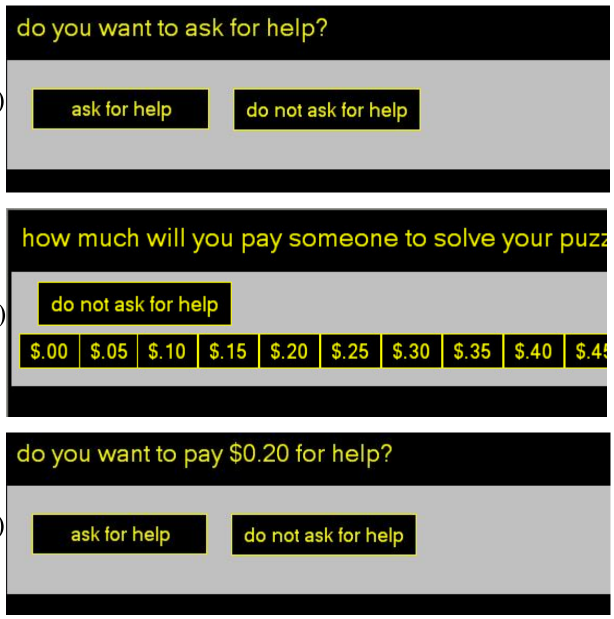

Figure 2: Help requesting options for (a) *no market*, (b) *variable market*, and (c) *fixed market*

In *no market* condition, the interface manipulation was simple. At the bidding opportunity, participants were asked to select either "ask for help" or "do not ask for help" (Figure 2a). On the receiver's side, receivers were asked to choose between provide "help" or "do not provide help."

In *fixed market* condition, the interface manipulation was similar to that of *no market*. At the bidding opportunity, participants were asked to choose to pay $0.20 to "ask for help" or "do not ask for help" (Figure 2b). On the receiver's side, receivers were asked to choose between providing help for $0.20, or not providing help. For the study, selecting a different fixed-price value should not influence hypotheses; any fixed price is a less expressive version of the *variable market* condition.

In the *variable market* condition, senders had the option to bid an amount from $0.00 up to the current puzzle bonus at $0.05 increments, and a separate option to select not to bid (Figure 2c). Unless a reject help option was selected, a help request was broadcast to the other players. Upon reception of a help request, receivers did not see the bid. Instead, they selected a reservation price between $0.00 and $0.75 at $0.05 increments. Receivers could also directly choose not to provide help.

We chose to restrict the number of options to $0.05 increments to keep the interface consistent. Even with this simplification, the *variable market* condition is still a much more expressive system than the *fixed market* condition.

Participants & Procedures

108 students from a university participated in this study for monetary reward (36 per condition). Participants were randomly assigned into one of the three conditions (between subjects).

Each session was one hour long. Each session had four participants working independently on their task—the memory/concentration puzzle. After a brief introduction and the completion of consent forms, experimental procedures and the task were explained to the participants in detail. Participants then played the puzzle game without communication or help for a couple of minutes to get used to the interface. They then worked on the task with the help mechanism for twenty-seven minutes. At the end of the experiment they were given a final questionnaire.

Measures

Primary dependent measures were overall earnings and time spent on decision during the study. Other task related information, such as number of puzzles solved, number of help requests and help exchanges were also logged.

RESULTS

H1-Help Requests

We expected senders in both *no market* and *variable market* conditions to ask for help with roughly the same help request percentage, since senders in both conditions can broadcast the minimum value of $0. In contrast, players in the *fixed market* condition would ask for help only for a subset of request opportunities (when their communication value was greater than $0.20).

	no market	*variable market*	*fixed market*
Request Opportunities	19.6	19.9	21.1
Actual Help Requests % help requests/opportunities	11.4 (58%)	11.5 (58%)	9.1 (43%)
Instances of Help % help exchanges/requests	6.2 (54%)	6.1 (53%)	8.4 (92%)

Table 1: Help requests and exchanges breakdown

As we hypothesized, participants made more help requests in the *no market* and *variable market* conditions than the *fixed market* ($F(2, 105)=2.39$, $p=0.096$). It is interesting to note that only about 60% of all request opportunities in the no market conditions led to help requests (Table 1). This may indicate either the strength of intrinsic motivations or the belief that help requests may go unanswered.

H2-Help Exchanges
We predicted that *variable market* condition would result in the highest percentage of help exchanges followed by the *fixed market* condition, and then the *no market* condition. Surprisingly, as shown in Table 1 line 3, our results disconfirm this hypothesis. *Fixed market* led to the highest percentage of help exchanges while *no market* and *variable market* had about the same percentage of puzzles helped.

The results are surprising since the *fixed market* condition, the less expressive market, led to more help exchanges than the *variable market* condition ($F(2,71)=25.6$, $p<0.0001$). A possible explanation as we will discuss later, deals with the idea that the more fine-grained decisions in the *variable market* leave less room for error in decision making, and may have resulted in more missed help opportunities.

H3 & H4-Financially Rational Use of the System
Given that help exchanges occurred in these systems, did participants use these exchanges in a financially rational manner? We predicted that all senders would use the exchange mechanisms to improve their welfare (H3), but that only receivers in the market conditions would use the exchange mechanism to improve their welfare (H4).

To analyze this, we used repeated measures analysis of variance. Individuals' earning (welfare) is the dependent variable and the number of puzzles solved by self, number of puzzles in which help was received, number of puzzles in which help was given and the conditions were repeated. We included 2-way interactions for our analysis. For this analysis, session was modeled as a random effect.

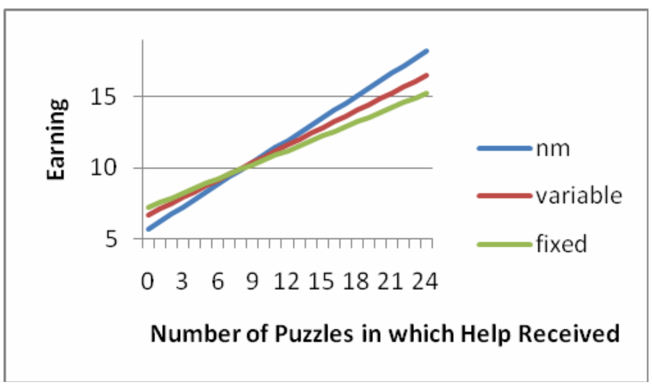

Figure 3: Graph of Interaction Effects between Receiving Help and Individual Earning in our Model

Regarding receiving help, our results show that there is a positive correlation between receiving help and welfare; the more times an individual receives help on a puzzle, the higher the individual's welfare ($t(95)=21.02$, $p<0.0001$). There are also significant interaction effects—receiving help in *no market* led to significantly higher gains in individual's welfare ($t(95)=4$, $p=0.0001$), while receiving help in *fixed market* led to significantly lower gains in individual welfare ($t(95)=21.02$, $p<0.0001$) (Figure 3).

As for giving help in the market conditions, giving help improved welfare in both market conditions, although it led to significantly higher welfare gains in *fixed market* ($t(92)=3.01$, $p=0.003$). On the other hand, giving help in the *no market* condition decreased welfare ($t(88)=-4.62$, $p<0.0001$) (Figure 4).

These results confirm what we had hypothesized. Participants in market systems utilize the exchange to improve their welfare, while receivers in *no market* provided help even though they incurred financial costs.

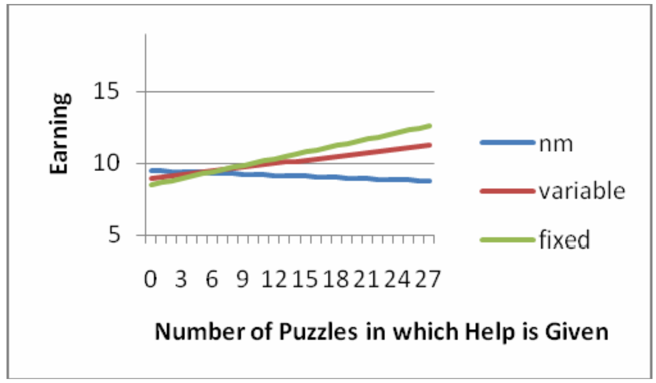

Figure 4: Graph of Interaction Effects between Giving Help and Individual Earning in our Model

H5-Efficiency of Market versus No Market
Given that players in market conditions were able to use the market to improve their welfare, did the market conditions lead to higher welfare (overall earnings) when compared to the *no market* condition?

When grouping the market conditions together, markets did perform better than no market by a dollar ($9.80 to $8.80, $F(1,106) = 3.30$, $p<0.07$). This suggests some modest efficiency gains from applying market mechanisms in real world communication.

H6-*Variable Market* versus *Fixed Market*

	no market	*variable market*	*fixed market*
Earning from puzzle bonus	$3.52	$3.81	$3.75
Earning from matching pairs	$2.20	$2.11	$2.17
Earnings from getting help	$3.10	$2.31	$2.78
Earnings from giving help	$0.00	$1.05	$1.67
Total Earnings	$8.81	$9.28	$10.37

Table 2: Breakdown of earnings by components

When analyzing the markets individually, it becomes apparent that the difference in earning between the market conditions and *no market* condition is mainly due to the significant difference between *fixed market* and *no market*, and not *variable market* and *no market* (*no market*=$8.81, *fixed market*=$10.37, *variable market*=$9.28). Pair-wise analysis shows that earning in *fixed market* is significantly higher than *no market* ($F(1,105)=5.94$, $p=0.02$) and marginally higher than *variable market* ($F(1,105)=2.90$,

p=0.09), whereas there is no significant difference in earnings between *variable market* and *no market* ($F(1,105)=0.5$, $p=0.5$).

When separately examining four components of earnings, we see that the overall difference between the conditions is a result of money players earn from using the communication mechanism (Table 2).

One possible explanation for why *fixed market* led to higher welfare than *variable market* is that because it is a much simpler system, it requires less use and familiarity. Perhaps participants need to learn to better use *variable market* with experience. We therefore compared average earnings from each condition between the first and second half of the study (13.5 minutes each). We see that there is a general increasing trend in earnings in all conditions. At each stage *fixed market* performed the best, then *variable market*, then *no market* (Figure 5), indicating that the advantage of *fixed market* does not disappear with experience.

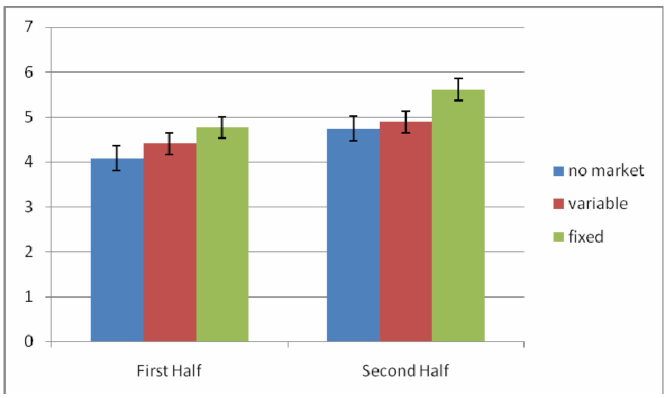

Figure 5: Breakdown of earnings by components

Another explanation for this performance difference is transaction costs due to cognitive limitations. Transaction costs can have two different types of effects on the communication market. First, as mentioned in our analysis of help exchanges, the more complicated and finer-grained decision leaves less room for error by both senders and receivers. Finer-grained decisions give senders more chances to overbid for help while receivers are also more likely to overprice their help reservation value. The greater complexity of the *variable market* may simply add error to the behavior of senders and receivers, thus leading to fewer mutually beneficial help exchanges in the *variable market* condition than we had hypothesized. Unfortunately, given the subjective component of gains and costs for both senders and receivers, it is impossible to precisely identify such error in their behavior.

Another possible transaction cost effect is that more time and attention may be required to make a decision. Participants in *fixed market* only have to decide whether or not to request/offer help for $0.20, instead of determining the exact communication value in *variable market*. The simple decision could significantly reduce the time required for decisions. We, therefore, compared the time players spent on each decision and found a significant difference in decision time (Table 3). *Fixed market* is lowest at 1.774 seconds, then *no market* at 2.165 seconds, and then *variable market* at 2.881 seconds ($F(2,105)=18.45$, $p<0.0001$). If we compare the total time used on decisions between the *variable market* condition and the *fixed market* condition, on average participants spend 70 seconds in more on their decisions in the *variable market* condition. Based on *fixed market* participant performance, having 70 extra seconds can be an additional $0.40 in earning. While this alone may not explain all of the earning difference between *fixed market* and *variable market*, it does contribute to some of it.

	no market	*variable market*	*fixed market*
Time Spent on Decisions	2.165s	2.881s	1.774s

Table 3: Time spent on decision

It is interesting to note that decision time in *fixed market* was also significantly faster than *no market*. One possible explanation is that the decision in market mechanisms may be dominated by weighing extrinsic costs and benefits, while the decision in *no market*, based on intrinsic motivation, may be more convoluted and require more time.

DISCUSSION

This study provides empirical evidence that communication markets can improve communication efficiency and productivity. However, our work also highlights some effects of real humans interacting in a communication market. Prior work in this domain has focused primarily on extrinsic gains and costs, but our work shows the importance of factoring both intrinsic motivations and transaction cost produced by cognitive limitations.

While this work has focused on a synchronous communication market, some of our findings may also apply more broadly to depict human behaviors in market environments more broadly. This understanding can provide some design implications for communication and human-market interactions, which is vital for future technologies leveraging economic markets.

Intrinsic Motivations in Help Communications

If humans are rational self-interested agents, then communication would never occur when receivers have nothing to gain. In our scenario of questions and answers, no help would be offered as the sender is the only party benefiting financially from such interaction.

But as our study shows, help in the *no market* condition occurred as frequently as in the expressive market condition (*variable market*). This means that even between anonymous strangers, people do not act to simply maximize their own personal, immediate monetary gain. Prior work on public goods and altruistic behavior provide evidence that people voluntarily help others, especially when they expect others to do so. In this kind of problem, players face a decision to help others, but at a personal cost that exceeds one's personal benefit [11,16]. Our results, like results from these prior studies, demonstrate that the strong predictions of self-interest and rationality (where no voluntary help will occur) are wrong.

Thus, individuals' willingness to voluntarily help others should be an important factor to remember as we evaluate market-based communication mechanisms. Often, we use the *no market* condition as a baseline comparison point, but communication behavior in the *no market* condition is not as straightforward and inefficient as traditional economic analyses suggest. Even without explicit incentives, implicit incentives can significantly impact communication decisions. Moreover, while our results assumed a simple additive effect between extrinsic and intrinsic motivations, explicit incentives may sometimes "crowd" out implicit incentives, thereby reducing the willingness of individuals to voluntarily help [5]. Further research is required to explore such effect in this communication market domain.

Transaction Costs in Communication Markets

Compared to prior work, one of the most surprising findings in this study is that the fixed-price mechanism led to the most exchanges of help and the highest overall earnings. We expected *fixed market* to perform better than *no market*, as *fixed market* allowed senders a mechanism to share their communication gains, thus providing an incentive for communication by receivers. But we did not expect *fixed market* to perform better than *variable market*. The richer communication available through *variable market* should have led to better allocation of resources.

We put forth two possible explanations for *fixed market* outperforming *variable market*. One is that the additional complexity and extra options available in *variable market* may have led to more errors in requesting or granting help. The other is that the *variable market* may have introduced additional decision costs, which we indirectly measure by considering the time it took to make decisions. As we demonstrated in the results, the more complex decision took longer, resulting in a $0.40 earnings difference, which explains some of the overall earnings difference.

Implications for Technology Design

One of the underlying motivations of this work was to understand whether or not users can use market mechanisms in real-time to improve communication efficiency. Our work demonstrates that people can, to a certain degree, use prices as signals to reflect their communication value and reservation prices to filter out unwanted communication. However, the fact that a fixed-price mechanism led to better performance than a variable pricing mechanism poses many implications for human-market communication interactions. First, for real time synchronous communication, a fixed-contract or a limited option market design may be more suitable to reduce the transaction cost. Otherwise having a more expressive market might not be able to offset the loss in time spent on decision making. More generally, if highly expressive markets are to be applied to facilitate resource management in our everyday technologies, designers of human-market communication interfaces must consider human decision costs such as time.

Another set of technology implications deals with non-optimal use of markets. To maximize the gain of efficiency from markets, humans need to act like agents from economic models. When they do not, technology interventions may be introduced to provide guidance for how markets should be used. Given possible human errors in calculating gains and costs in the communication market, can we use technology to reduce these valuation errors?

It is important to note that our goal is not to claim that economic markets, if designed well, are the ultimate solution for allocation of resources. Nor do we believe that more expressive markets are generally worse than less expressive ones. Rather, we hope our findings will highlight the benefit and potential flaws of different communication systems and that the right design depends on user characteristics and features of the environment and task. From an individual user's perspective, market mechanisms are valuable because they can ensure that communication is beneficial to all individuals involved. This is especially useful when receivers are likely to receive communication requests from unrecognizable senders. However, for a communication media to be used within an organization or work team, not having a market mechanism may have benefits. Without the monetary focus, communicators will be more aware of the social implications for their actions. This might increase the level of social capital that is vital in establishing long-term relationships.

LIMITATIONS AND GENERALIZABILITY

The stylized task used in our study exhibits many features representative of everyday tasks (deadline, task value, and noticeable incremental progress). While it lacks realism, the simplification does allow us to model the communication costs and benefits and clearly analyze the influences of the mechanisms on task performance.

The study focuses on a specific communication scenario. In reality, senders may have multiple opportunities to ask for help, from many people, for any given task. Also, senders do not need to be the one desiring the communication (a sender might be simply returning a prior communication request) nor are communicators limited to either sender or receiver roles in any given communication. It is important to point out here that the *variable market* mechanism used in this study can be adapted to allow for bi-directional transfer of wealth as it is supported by Attention Bond Mechanism. The basic idea is that after the communication occurs, a receiver could choose to release the payment depending on how valuable the communication is. We did not complicate the *variable market* mechanism in our study to include such a feature, because it was not necessary in our communication scenario. Lastly, this work does not factor in the relationships between communicators, which may significantly impact how a market mechanism is used and the efficiencies that may result.

Even though this study uses a messaging mechanism, we believe that the findings can extend to other types of synchronous communication such as cell phone use. The effectiveness of pricing should be independent of what modality of synchronous communication they use. Most of the findings may also extend to more general forms of communication, although the importance of time would be greatly reduced in asynchronous setting.

CONCLUSION AND FUTURE WORK

In this work, we explored the use of market mechanisms in synchronous communication. We show that, compared to a baseline mechanism that does note utilize market pricing, markets lead to higher productivity. However, our results also indicate an important tradeoff between market expressiveness and decision costs; using a market with one pricing option may be more beneficial than using a market with many pricing options.

Our analysis of market use provides better insights into what communication may be like if we apply markets to our everyday communication. Our discussion points out some strengths and weaknesses of a communication market, and suggests that market mechanisms may not always be ideal for managing communication under all contexts.

One logical next step is to extend from theory to practice. For example, just how can such markets be incorporated into current communication media? If the market is to be expressive, what would an actual design look like in which prices can be used easily and efficiently? We leave these questions for future research.

ACKNOWLEDGMENTS

We would like to thank Daniel Avrahami, Wendy Chang, Laura Dabbish, Susan Fussell, Ian Li, Bilge Mutlu, Carol Pai, Tuomas Sandholm, Karen Tang and Joe Tullio for their suggestions and help throughout this work. This work is supported by NSF grants ITR-032535 and CNS-0551554.

REFERENCES

1. Batson, C. D., Harris, A.C., McCaul, K.D., Davis, M., & Schmidt, T. (1979). Compassion or Compliance: Alternative Dispositional Attributions for One's Helping Behavior. *Social Psychology Quarterly*, Vol. 42, No. 4, 405-409.
2. Coase, R. H. (1937). The Nature of the Firm. *Economica*, Vol. 4, No. 16, 386-405
3. Cranor, L.F. & LaMacchia, B.A. (1998). Spam! *Communications of the ACM*, 41(8), 74–83.
4. Dwork, C., & Naor, M. (1993). Pricing via processing for combating junk mail. In *Advances in Cryptology-Crypto* '92: Lecture notes in computer science 740, 139–147. New York: Springer-Verlag.
5. Fehr, E. & Gaechter, S., Do Incentive Contracts Crowd out Voluntary Cooperation? *IEW - Working Papers iewwp034*, Institute for Empirical Research in Economics - IEW.
6. Fahlman, S.E. (2002). Selling Interrupt Rights: A Way to Control Unwanted E-Mail and Telephone Calls. *IBM Systems Journal*, 41 (4), 759-766.
7. Hayek, F A. (1945). The uses of knowledge in society. *American Economic Review*, 35, 519–530.
8. Handel, M. & Herbsleb, J.D. (2002). What is Chat doing in the workplace? In *Proc. CSCW*, 1-10.
9. Harper, M., Raban, D., Rafaeli, S., Konstan, J. (2008). Predictors of Answer Quality in Online Q&A Sites. In *Proc. CHI*, 865-874.
10. Hermalin, B.E. and Katz, M.L. (2004). Sender or Receiver: Who Should Pay to Exchange an Electronic Message? *RAND Journal of Economics*, 35(3), 423–447.
11. Hoffman, E., McCabe, K., Shachat, K., & Smith, V. (1994). Preferences, property rights and anonymity in bargaining games. *Games and Economic Behavior*, 7: 346-380.
12. Isaacs, E., Walendowski, A., Whittaker, S., Schiano, D.J., & Kamm, C. (2002). The Character, Functions, and Styles of Instant Messaging in the Workplace. In *Proc. CSCW*, 11-22.
13. Kahneman, Daniel. (2003). Maps of Bounded Rationality: Psychology for Behavioral Economics. *The American Economic Review*, Vol. 93, No. 5: 1449-1475.
14. Kraut, R. E., Sunder, S., Telang, R., & Morris, J. (2005). Pricing electronic mail to solve the problem of spam. *Human Computer Interaction*, 20, 195-223.
15. Krishnamurthy, B., SHRED: Spam Harassment Reduction via Economic Disincentives. Working Paper, AT&T Research.
16. Ledyard, J. (1995). Public Goods: A Survey of Experimental Research. In *The Handbook of Experimental Economics*. Roth and Kagel, eds. Princeton, NJ: Princeton University Press.
17. Loder, T.; Van Alstyne, M.; and Wash, R. (2006). An Economic Response to Unsolicited Communication. *Advances in Economic Analysis & Policy*: Vol. 6: Iss. 1, Article 2.
18. Nardi, B. A., Whittaker, S., & Bradner, E. (2000). Interaction and outeraction: Instant messaging in action. In *Proc. CSCW*. 79-88.
19. Quan-Haase, A., Cothrel, J., and Wellman, B. (2005). Instant messaging for collaboration: A case study of a high-tech firm. *Journal of Computer-Mediated Communication, 10*(4), article 13.
20. Rainie, L., & Fallows, D. (2004). The impact of CAN-SPAM legislation (Pew Internet and American Life Project).
21. Thaler, Richard H. (1994). *The Winner's Curse: Paradoxes and Anomalies of Economic Life.* Princeton NJ: Princeton University Press.
22. Trivers, R.L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology*. 46, 35-57.
23. van Zandt, T. (2004). Information Overload in a Network of Targeted Communication. *RAND Journal of Economics*, 35 (3), 542-560.