

Intertemporal Discount Factors as a Measure of Trustworthiness in Electronic Commerce

Christopher J. Hazard and Munindar P. Singh *Fellow, IEEE*

Abstract—In multiagent interactions, such as e-commerce and file sharing, being able to accurately assess the trustworthiness of others is important for agents to protect themselves from losing utility. Focusing on rational agents in e-commerce, we prove that an agent's *discount factor* (time preference of utility) is a direct measure of the agent's trustworthiness for a set of reasonably general assumptions and definitions. We propose a general list of desiderata for trust systems and discuss how discount factors as trustworthiness meet these desiderata. We discuss how discount factors are a robust measure when entering commitments that exhibit moral hazards. Using an online market as a motivating example, we derive some analytical methods both for measuring discount factors and for aggregating the measurements.

Index Terms—Trust, Reputation, Intertemporal Discounting

1 INTRODUCTION

TRUST is an important and broad topic. It incorporates elements of cognition, emotion, social relationships [1], information security [2], and (economic) rationality. We address the narrow, but practically highly valuable, scope of e-commerce and related decision support. Here, the perspective of rationality dominates since the trustworthiness of a business partner is largely governed by its self-interest.

It has been long recognized that trust and reputation are important in achieving desirable and sustainably beneficial outcomes in online auctions and automated business-to-business transactions. Doing business with an agent that has a reputation for being trustworthy generally has the benefit of reducing the risk of a poor outcome. Agents' reputations and perceived trustworthiness can significantly affect the demand and price an agent will receive in domains such as online auctions [3] and supply chains [4]. As business transactions increasingly become automated [5] and autonomous agents become more crucial components of business strategy, successful agents will need to model trustworthiness effectively.

Ideally, from a trust perspective, the mechanisms under which agents interact would be incentive compatible (IC), meaning agents' optimal strategies would be to be honest and truthful. Whereas IC mechanisms can be designed for a variety of interaction models [6], often maximizing profit for the agent running the mechanism is a higher priority. For example, eBay's (<http://ebay.com>) reputation system exhibits a bias toward transaction volume to maximize profit [7], because sellers can game the reputation system [8]. Additionally, implementing an IC mechanism can be infeasible in certain settings in terms of computation or communication [9].

Trustworthiness reflects the worthiness of a *trustee* to aid or protect a *trustor*. For example, a trustworthy trustee will properly fulfill some task for a trustor or refrain from inappropriately revealing a secret. As trustor a learns more about trustee b , the amount of trust that a places in b should ideally approach the amount of trust of which b is truly worthy.

A key intuition is that a trustworthy agent is patient, i.e., interested in long-term relationships: for example, we expect a store for local residents to sell better wares than a tourist trap. In general, anything is worth less in the future than now. With exceptions such as for storage, degradation, and depreciation, having money or a usable item is generally worth more now than later for reasons such as the uncertainty of the future and opportunity to use the item or money in the mean time. For example, most people would prefer \$100 today over \$100.01 next week. But one's premium for immediacy is bounded: typically, most people would prefer \$1,000 tomorrow to \$10 today. An agent's *intertemporal discount factor* reflects its break even point for the present versus the next time unit. For example, if you are neutral between \$90 today and \$100 tomorrow, then your discount factor is 0.90 (per day).

Further, trustworthiness and patience can vary with the *context*: a nearly bankrupt business facing its creditors may sell items without sufficient quality checks. We use context to refer to the risk environment that an agent facing, such as facing a pending bankruptcy or succeeding in a steady market. Outside of the mathematical use with respect to variables, we use *domain* to refer to a type of interaction, such as the role of a provider in a web services market versus the role of a seller in an online auction.

Definition 1: An agent employs *exponential intertemporal discounting* in some context when its utility gain, U , from some event in that context at time t is $U = \gamma^t u$, where u is the utility the agent would have perceived had the event occurred at the present time ($t = 0$) and $\gamma \in [0, 1]$ is the agent's *discount factor*.

An agent's discount factor captures how much it would value something at future points in time relative to the present. A discount factor can be influenced by intrinsic factors, such as patience for an individual, or extrinsic factors, such as cost of capital for a firm driven by market rates. A higher discount factor can yield a greater payoff because the agent is not myopically optimizing, but this rule has exceptions [10], [11].

Even though we intuitively associate trustworthiness with the expectation of future long-term relationships, most current

approaches do not necessarily reflect this intuition. Existing measures of trustworthiness [12] typically use arbitrary ratings or are highly dependent on the domain, distribution, and manner of interactions. A small body of related work has discussed some aspects of the relationship between discount factors and trust [13]–[15]. However, with two exceptions [11], [16], we are unaware of related work directly employing discount factors as a measure of trustworthiness.

Contributions

We develop a model of trustworthiness as discount factor that naturally captures the above intuitions. First, we formalize key technical assumptions typically left implicit: *comparison* (by a trustor of trustees); *strength* (by a trustee of its tasks), and *stability* (of the trustee’s behavior). Second, we demonstrate discount factor as an objective measure of trustworthiness, isolated from subjective effects. We prove that any scalar trust measure is isomorphic to the discount factor. Third, we show that our model is the only approach that meets crucial desiderata for a computational approach to trust.

Fourth, we develop an approach by which an agent may infer the trustworthiness of another based on the latter’s actions. To this end, we consider a series of e-commerce situations where buyers and sellers estimate each other’s trustworthiness based on signals such as the quality of products sold, prices offered and accepted, and eagerness to conclude a transaction. Lastly, we show how information on trustworthiness may be aggregated and estimated, and conclude with a discussion of some practical ramifications.

2 MOTIVATION

The need for trust systems arises in two situations: *adverse selection* and *moral hazard* [17]. Adverse selection occurs with *typed* agents, meaning an agent is predisposed to some course of action due to its one or more (fairly constant) attributes. An agent’s type can range from a strict behavior regimen, such as accepting every offer or always producing high-quality items or being patient, to a parameter the agent uses in evaluating its utility, such as its willingness-to-pay for some item. The presence of typed agents means that agents may be able to improve their utility by determining which agents are of what type, and interacting only with agents of a favorable type. An example of a typed agent would be an agent selling faulty electronics at high prices. The agent may be unable to change the quality or price of the goods it sells, and other agents may do best to avoid purchasing from this agent. The ability of an agent to improve its utility by choosing with whom to interact is strongly affected by the interaction mechanism. An example where an agent may not be able to choose which other agents with which it will interact is an auction setting with perfectly substitutable goods where buyers and sellers are randomly matched by the auctioneer at a set price.

Conversely, in a setting where agents choose trading partners, if agent a manufactures poorer quality items than the other agents, knowing that a manufactures poorer quality items can enable some other agent, b , to increase its own utility by not purchasing from a . Determining trust with adverse

selection can be framed as a multiagent learning problem, as the agents perform signalling behavior to increase the accuracy of their beliefs of other agents’ types.

Moral hazards are created when agents do not bear the full cost of their actions and are thus incentivized to perform actions that may harm the utility of others. For example, a seller who deals with a gullible buyer has the moral hazard of falsely advertising its goods. To address moral hazards, trust systems attach sanctions to unwanted behavior. If agent a performs some unwanted behavior, then a trust system can attach some information to a . This information can be used by a centralized mechanism or individual agents to sanction or avoid interacting with a , with the effect that a would have an incentive to alter its behavior.

Motivating Example: Online Market

An online auction is a practical motivating scenario for a trust system. As our running example, we outline the general mechanics of this scenario to motivate our results and formally analyze it in Section 6. The auction is continually cleared, with buyers choosing which sellers’ offers to accept, if any. Exactly one buyer or seller moves at a time. The order of buyers’ and sellers’ turns are chosen from a stochastic process to simulate realistic market transactions, but each agent gets one turn per unit time. Each agent’s goal is to maximize its expected utility and, to account for time preference, is endowed with a privately known discount factor.

Sellers post or update offers to sell items. A seller’s costs are private but follow publicly known probability distributions. Cost are initialized before the auction begins. An offer states the asking price and the (true or exaggerated) quality of the item. We define quality as the probability density function (PDF) that an item will irreparably fail as a function of time. The expected lifetime of an item is its mean time to failure (MTTF). Section 6.1 considers the cases where a seller can (1) only produce a fixed quality and (2) control the quality of its items.

Buyers see all current offers and choose which and when to accept. When deciding what to purchase, buyers can see the seller’s offer as well as a history of “comments” by other buyers about the seller’s discount factor, price, and quality. After each transaction, a buyer can post a public comment on the seller. In our formal approach, a comment is a numerical observation of one or more of quality, valuation, time, or discount factor. A comment is formulated as a measurement or inequality, such as “I observed the good offered by agent a at price 5.29 to be of a quality that held up for 1 week of use before breaking” or “agent a ’s discount factor is greater than 0.60.” Sellers can see what price and quality other sellers are currently offering and update their offers accordingly.

Each buyer has its own expected utility gain per unit time for having each additional item, a willingness-to-pay per unit time. Buyers have a price sensitivity with respect to quality, based on the expected useful life of an item coupled with the agent’s discount factor and willingness-to-pay. When a buyer makes a purchase, it loses the utility of the amount of the purchase price and gains utility for each unit time that the item is functional.

We note that because the seller controls the price, our model’s descending Dutch auction style resembles Craigslist (<http://craigslist.org>) and the retail presence on eBay, where the seller’s “minimum bid” is effectively the ask price. This is in contrast to the ascending English auction commonly associated with consumer-to-consumer transactions on Amazon and eBay. We choose to examine the seller-price auction because the analysis yields somewhat simpler results and is therefore easier to discuss in the cases of interest.

3 DEFINING TRUSTWORTHINESS

To define trustworthiness, we first must have definitions of how agents interact. We define an *event*, i , as a pair, $\langle u_i, t_i \rangle$ consisting of a change in utility, u_i , by the agent performing the event at a specified time, t_i . We define an event as an isolated change in utility, given all externalities, conditions, and decisions that create the event. An event may have additional side effects, such as altering the utility of another agent, but as these are not essential to our discussion and formalisms, we exclude them in our notation and define an event as a pair for clarity.

We use the following notation. Each agent, a , has a total expected utility function, U , that returns the agent’s total utility given its trustworthiness and a set of events of utility changes. The function may be written more formally as $U : \Gamma \times \{\langle \mathcal{R}, \mathcal{R} \rangle^*\} \mapsto \mathcal{R}$, meaning that the total utility function takes in a real value of trustworthiness, $\gamma_a \in \Gamma$, and a set of events, $I \in \{\langle \mathcal{R}, \mathcal{R} \rangle^*\}$, and returns a real number of total utility of the events. We write it in the form $U(\gamma_a, I)$.

In our running example, an event is a cash flow or an change in ownership or status of a good. At the time when a seller transfers the ownership of the item to the buyer, the buyer receives some utility at that time. The utility gain that the buyer receives may be an expected value if the buyer is planning on reselling the item, perhaps after additional manufacturing or configuration, for a profit. When the seller receives money for the good or service, the event is to add money to the seller’s account at the time when the buyer pays.

As in the running example, we restrict our attention to trust with respect to future actions. This would eliminate some English uses of the word “trust” such as “I trust book reviews on Amazon,” because there is no future action there. It would allow “I trust Amazon to send me the book on time,” which involves a future action.

3.1 Assumptions

We assume trustworthiness is reasonably fixed for the time frame in which the agents act. This is reasonable because if trustworthiness changed quickly, for example, if sellers frequently and unpredictably changed their type, a measure of trustworthiness would not be useful for predicting outcomes.

This does not mean that trustworthiness is fixed for a given agent. Models in which agents’ types change [18] are compatible with our approach. Assumption 1 merely requires that the rate of change for agent types is sufficiently lower than the rate of interactions so that knowing another agent’s type is useful in an agent’s decision model.

Assumption 1: An agent’s trustworthiness are consistent enough to be meaningful across interactions; recent measurements of an agent’s trustworthiness, if accurate, should usually reflect the agent’s current trustworthiness.

Utility theory lies at the core of e-commerce and postulates that agents have valuations for goods or services. A common currency is obviously desirable for commerce [19], and enables agents to compare their valuations.

Assumption 2: A utility loss or gain by one agent can be directly compared to the utility loss or gain of another agent.

Quasilinearity, that total utility can be closely approximated by summing the utilities of individual events, is frequently assumed in consumer theory and e-commerce [20].

Assumption 3: Each agent has quasilinear utility; given two events yielding utilities at the present time of u_1 and u_2 , the agent’s total utility, U , is $U = u_1 + u_2$.

Individual rationality means that an agent will not enter into nor fulfill a commitment unless doing so maximizes the agent’s utility. A buyer will not normally purchase an item that is greater than its willingness to pay for that item. Individual rationality is a core foundation of autonomous agents in much of the e-commerce literature [21].

Assumption 4: Agents are individually rational.

3.2 Intuitions about Trustworthiness

Trustworthiness inherently involves settings where agents directly or indirectly engage in behavior that affects each others’ utilities. The concept of a *commitment* helps capture this relationship. A *debtor* (agent) commits to a *creditor* (agent) to bring about an *event* [22]. In essence, a commitment reflects a dependence of the creditor on the debtor.

Definition 2: $C(b, a, i)$ is a *commitment* from debtor b to creditor a that b will bring about an event i at time t_i yielding a positive real utility to a and a negative utility, u_i , to b .

We restrict attention to commitments that require a negative utility for the debtor simply because commitments that yield positive utility to all parties with no risks does not require trust in our sense. In other words, we seek to capture the intuition about a debtor’s trustworthiness based on the troubles it will go through to fulfill its commitments.

Often, in e-commerce, commitments would occur in complementary pairs so the overall situation would be win-win. For example, when a buyer commits to paying a seller and the seller to providing goods to the buyer both benefit from the transaction. Indeed, given individual rationality (Assumption 4), every commitment that an agent enters must entail the expectation of a complementary commitment, such that the expected sum of the utilities is positive. Agent a may have beliefs as to how it will be repaid, such as having a 50% chance of b deciding on event i and a 50% chance of b deciding on event i' . When evaluating its total utility function, a should evaluate this as the expected value $\frac{U(\gamma_a, i) + U(\gamma_a, i')}{2}$, which holds due to Assumption 3.

The success or failure of a commitment provides a basis for the creditor to measure the trustworthiness of the debtor. For example, b may commit to deliver an item of a specified quality to a . If b fulfills a commitment $C(b, a, i)$, a neutrally or

positively updates its view of the relationship between a and b . If b fails to fulfill this commitment, a negatively updates its view of the relationship between a and b .

We now motivate some key intuitions regarding trustworthiness, which we then combine in our proposed definition of trustworthiness. Figure 1 illustrates the intuitions except *scalar*.




Fig. 1. Illustration of intuitions about trustworthiness.

Scalar. Representing trustworthiness as a single value is a natural convention. One may ask, “How much do you trust b to produce and deliver some item with quality of at least X ?” and receive a reply of “A lot.” Such a value can be quantified; many online services provide ratings as points or percent of customers satisfied. A scalar representation does not preclude an agent from holding additional beliefs of the value or accuracy of trustworthiness, such as a probability distribution, nor from requiring additional information when making a decision of whether to trust, such as how much the trustee values something. Further, we can use different scalars for each context. More formally, we say that the set of trustworthiness values is the set of real numbers, $\Gamma = \mathbb{R}$.

Comparison. A trustor a can compare two trustees b and c . Specifically, a considers b more trustworthy than c if, all else equal, b would suffer a greater hit than c would to fulfill a commitment to a . In essence, a must know something about the valuations and costs incurred by both b and c and be able to compare these values as supported by Assumption 2. This does not mean that a will receive more utility from b ’s commitment than c ’s commitment, only that b is fulfilling a more costly commitment. Formally, agent a would consider agent b more trustworthy than agent c if, all else equal, for some event i with positive utility to a , there exist commitments $C_b = C(b, a, i)$ and $C_c = C(c, a, i)$ such that b would fulfill C_b and c would not fulfill C_c .

If c does not fulfill its commitments to a , by our definitions, this necessarily entails the loss of expected utility by a . If a paid c to deliver an item at a specified quality and c fails to deliver the item or provides an item of low quality, a will have gained less utility than it expected and incurred a negative net utility. This decrease in net utility causes strain on the relationship, causing a to either retaliate against c , such as by posting negative comments about c causing other agents to avoid transactions with c , or to avoid future loss by reducing its involvement with c by not making further purchases from c . In either case, c will initially have greater utility from incurring less cost by providing a lower quality item, but possibly lose

more utility over the long term.

Strength. The behavior of each agent is internally consistent. Given equal impact on a relationship, if an agent is willing to do something difficult to keep a commitment, it should be willing to do something easy. If an agent is willing to deliver 1,000 gallons of kerosene to fulfill a commitment, then the agent should be willing to deliver 600 gallons of kerosene if everything else in the overall commitment stays the same (provided that storing or disposing of the other 400 gallons is not more difficult or costly than delivering it). From the perspective of the debtor, this property does not require actual fulfillment, it only requires that the agent be willing to exert the effort. If an item arrives late due to extenuating circumstances, this does not mean that the seller is necessarily less trustworthy. However, the creditor may only lessen its negative interpretation of an unfulfilled commitment if the creditor has some belief of noise in the signal of whether commitments are fulfilled. Formally, consider events i, j where $u_i \leq u_j$ and agents a, b . If b fulfills $C(b, a, i)$ then b fulfills $C(b, a, j)$.

Stability. The idea of stability is that agents should tend to behave in a manner that reflects a consistent underlying level of trustworthiness, which stems from Assumption 1. This essentially means that an agent, at the present time, considers its trustworthiness to be consistent for modeling future interactions. Using our online market example, an agent should be approximately equally trustworthy if a commitment will be set up now or one month from now, presuming the agent and market remain constant with regard to price, demand, supply, reputations, and reliability of available information. For example, suppose a firm can be trusted now to successfully deliver an order of 20 microphones of a certain quality within two weeks of payment. Then, if all else (e.g., external prices, internal staffing, and such) remains consistent, the firm can be trusted to deliver the same order if it were placed several months later again within two weeks of payment. Suppose the same firm is indifferent between committing to a delivery of 20 microphones and a delivery of 5 speakers today. If again, the environment and agents’ valuations stay the same, the firm should be indifferent between those two commitments if asked again in a month. More formally, if an agent is indifferent between two commitments or sets of events, I_1 and I_2 , then it should also be indifferent if the time is shifted by some arbitrary s . This may be expressed as

$$U(\gamma, I_1) = U(\gamma, I_2) \Rightarrow U(\gamma, \{\langle u_i, t_i + s \rangle : i \in I_1\}) = U(\gamma, \{\langle u_i, t_i + s \rangle : i \in I_2\}). \quad (1)$$

Stability means that an agent should tend to behave in a similar manner across a period of time, but this does not mean that an agent is indifferent between when an event or commitment may happen. An agent may prefer to receive an item sooner rather than later. We are simply stating that, given identical circumstances, an agent would enter the same commitments if they were shifted by some time because the agent is stable. If properties of the environment, agents’ valuations, or agents’ trustworthiness change, the agents may

model such changes and factor them into their decision making however appropriate.

Definition 3: The *trustworthiness* of agent b is a belief by another agent a that takes on a *scalar* value, is relatively stable across time (*stability*), and is used to compare agents (*comparison*) to determine which would be willing to exert more effort (*strength*) to fulfill a commitment.

4 TRUSTWORTHINESS AND DISCOUNT FACTOR ISOMORPHISM

We now derive our main result: an agent's discount factor is a direct measure of its trustworthiness given assumptions.

Because previous changes of utility are accounted for in an agent's current utility, it is only useful to evaluate the impact of future changes to utility. We therefore restrict the domain of t_i to $[0, \infty)$.

Theorem 1: Given commitment as in Definition 2, trustworthiness as in Definition 3, and Assumptions 1, 2, 3, and 4, the representation of trustworthiness satisfying these definitions is isomorphic to an intertemporal discount factor.

Proof: By Definition 2, the utilities of any two events i and j are independent. This definition, coupled with Assumption 3 of quasilinearity, implies that an agent's total utility, U , is a summation of some utility function for each event, f , over all of the events, with $\frac{\partial f}{\partial u_i} > 0$. With trustworthiness γ and the set of events I , this is given by

$$U(\gamma, I) = \sum_{i \in I} f(\gamma, u_i, t_i), \quad (2)$$

Given *comparison* (supported by Assumption 2) and *strength*, an agent, b , is considered more trustworthy than another, c , if b will fulfill a commitment requiring a larger expenditure than c . This implies there is a commitment of some cost that b will fulfill and c will not; below this cost, both agents would fulfill the commitment. We only need to examine an individual event, and can restate this property using the event utility function, f .

Let us evaluate agents b and c with trustworthiness γ_b and γ_c respectively. Let agent a expect a commitment, $\langle u_1, t_1 \rangle$, to be fulfilled by the agent in question where, by Definition 2, $u_1 < 0$. Further, suppose that if the commitment is fulfilled, a will provide some utility back to the respective agent in the continued relationship: as Section 3.2 explains, at least two complementary commitments are required for agents to enter into commitments. We examine the simplest case, where this returned utility is expressed by a single event, $\langle u_2, t_2 \rangle$, such that $u_2 > 0$ and $t_2 > t_1$.

From Assumption 4, $f(\gamma, u_1, t_1) + f(\gamma, u_2, t_2) > 0$ for b and c ; otherwise the relationship is destructive and rational agents would not engage in the commitments. Suppose b chooses to fulfill its commitment and c chooses to not fulfill its commitment. Their decisions show $U(\gamma_b, \{\langle u_1, t_1 \rangle, \langle u_2, t_2 \rangle\}) > U(\gamma_b, \emptyset)$ and $U(\gamma_c, \{\langle u_1, t_1 \rangle, \langle u_2, t_2 \rangle\}) \leq U(\gamma_c, \emptyset)$. If no events occur to change an agent's future utility, the agent's utility does not change, so $U(\gamma_b, \emptyset) = U(\gamma_c, \emptyset) = 0$. This implies, given the above assumptions of the two-event interaction set, that

$$U(\gamma_b, \{\langle u_1, t_1 \rangle, \langle u_2, t_2 \rangle\}) > U(\gamma_c, \{\langle u_1, t_1 \rangle, \langle u_2, t_2 \rangle\}). \quad (3)$$

Because b fulfilled a commitment that was larger than c would fulfill, by *comparison* and *strength*, b is more trustworthy than c . If b is more trustworthy than c , then its trustworthiness value is higher, meaning $\gamma_b > \gamma_c$. We can take the limit as $(\gamma_b - \gamma_c) \rightarrow 0$, to find that

$$\frac{\partial U}{\partial \gamma} \geq 0 \quad (4)$$

holds in this scenario with two events. This means that more trustworthy agents, when their trustworthiness is common knowledge, attain higher expected utility than untrustworthy agents in two-event scenarios, all else being equal.

Stability, supported by Assumption 1, entails that agents are consistent in their trustworthiness. The outer operation of U in (1) is a summation, and the number of terms in each summation (the number of events in each set of events) are not necessarily equal. Therefore, the only two possibilities that allow both equalities to hold are that time has no effect on events' utilities or that a change in time results in a constant multiplicative factor across all terms in a summation independent of the utilities.

First, we consider the case where a change in time results in a constant multiplicative factor. The event utility function f must contain a multiplicand of the form x^t . This is because, given $x \geq 0$, x^t exhibits the appropriate behavior of $x^{t+s} = x^s \cdot x^t$ with x^s being constant for a constant time s . The second case, where time has no effect on f , can be represented by the first case with $x = 1$.

At this point, x remains an undefined attribute that affects the utility evaluation. Supposing x did not affect the trustworthiness of an agent, if b is more trustworthy than c , then (3) must hold. Setting $x = 0$ for agent b would violate this inequality. As this contradicts the assumption that x cannot affect the trustworthiness of the agent, x therefore directly affects the trustworthiness of an agent.

Given *scalar*, only one attribute may affect the trustworthiness of an agent. We now check to make sure that x satisfies the constraints of γ . In the two-event scenario, when $U > 0$ as given by Assumption 4, $\frac{\partial U}{\partial x} = t_1 u_1 x^{t_1-1} + t_2 u_2 x^{t_2-1}$. Because $x \geq 0$, $t_2 \geq t_1$, $u_1 < 0$, $u_2 > 0$, and $U = u_1 x^{t_1} + u_2 x^{t_2} > 0$, we can solve U for $u_2 > -u_1 x^{t_1-t_2}$, and substitute the infimum of u_2 in this expression (and any greater number) into the expression for $\frac{\partial U}{\partial x}$ to find $\frac{\partial U}{\partial x} \geq 0$. This satisfies (4), thus satisfying *strength* and *comparison* (x came out of a derivation of *stability*).

Substituting γ for x and rewriting in the form of (2), we find $U(\gamma, I) = \sum_{i \in I} \gamma^{t_i} u_i$. Revisiting (4), $\frac{\partial U}{\partial \gamma} = \sum_{i \in I} t_i \cdot \gamma^{t_i-1} u_i$. To prevent imaginary terms for events with $t_i < 1$, the constraint of $\gamma \geq 0$ is required. This final utility equation coupled with the domain of γ is, by Definition 1, exponential intertemporal discounting. \square

5 DESIDERATA FOR TRUST SYSTEMS

Devising optimal designs of general-sum multiplayer games is a difficult and domain-dependent problem. However, general desiderata can help guide interaction design. Such desiderata include individual rationality, guarantee of attaining a minimum payoff, guarantee of payoff to be within some ϵ within a

best response strategy, and Pareto optimality when an agent is playing against its own strategy [23]. However, the desiderata for trust and reputation systems are not as straightforward [24] because trust and reputation are supplemental to *primary interaction mechanisms*. A primary interaction mechanism is one, such as a market, that affects agents' utilities directly.

A key motivation for work on trust is that the primary interaction mechanism is not incentive compatible. Were it so, the agents would act honestly out of self interest. Our desiderata not only apply well when the primary mechanism is not IC, but also work when it is IC. IC is highly desirable for mechanism design, but achieving IC may not be computationally feasible [9]. Further, an IC mechanism may not be in the best interest of the agent or firm running the mechanism, because an IC mechanism may not maximize profit.

Many papers on trust propose desiderata [12], [25]–[28]. Dingedine et al.'s [24] desiderata list is the only comprehensive one we have found, but even their desiderata list focuses on aspects that are specific to certain kinds of reputation systems. We now propose desiderata that apply even when no central authority is available to enforce interactions or sanctions, and which focus on top-level goals that directly benefit the agents or system. A desirable system must be:

EVIDENTIAL. An agent should use evidence-based trustworthiness measurements to predict future behavior. This is the essence of a trust system, with an agent rationally assessing others' behavior and acting upon its knowledge. Evidence also includes temporal relevance; new evidence that an agent has successfully changed its type, if credible, should indicate to another agent that old evidence may no longer be relevant. In the online market example, an agent should measure trustworthiness in a quantifiable and repeatable manner based on the quality of goods and timeliness of their offers, to determine how to best engage in future interactions.

AGGREGABLE. Trustworthiness measurements should be accurate, precise, and possible to aggregate. This is key because aggregation enables an agent to communicate about trustworthiness and to put together indirect information obtained from other agents to increase knowledge of other agents' trustworthiness. In the market model, this aggregation involves reading others' comments, albeit with skepticism, to maximize the information considered.

VIABLE. The system should be practical in its computation and communication requirements. An approach that requires an exponentially large number of messages among buyers and sellers or requires each agent to perform an NP-Hard computation on a large dataset would not be tractable.

ROBUST. Measurements should be robust against manipulation; agents may signal or sanction to determine which agents are of what type and to resist strategic manipulation of the measurements. Manipulation can come in many forms, such as building up a reputation and then spending it, opening many pseudonymous accounts to communicate an inflated reputation to legitimate agents (Sybil attack), and opening a new account to expunge a bad reputation. We do not assume an incentive compatible reputation mechanism [6]. IC would be ideal for ROBUSTNESS, but can be impractical in

some problem domains, either because of computational or communicational complexity conflicting with VIABILITY, or because of unenforceability if agents can deviate from the specified mechanism without credible consequences.

FLEXIBLE. Trustworthiness should be applicable across multiple situations within the same context. Trustworthiness measurements should carry over across products, services, and even interaction mechanisms. Suppose a seller is running a web service from which buyers can purchase directly, but also sells some of its items in a simultaneous ascending auction run by a third party. A buyer should be able to carry over knowledge of trustworthiness about the seller from direct sales to infer information about the quality of the items sold on the third party's auction and vice versa, even though the mechanisms are different. If a buyer becomes a seller, its reputation as a buyer should be indicative of its behavior as a seller, provided other agents can infer some knowledge of valuations, capabilities, and beliefs in the new domain.

PRIVACY ENHANCING. The system should maximize agents' privacy by minimizing the collection of information. The implications on a system can be quite broad. We use privacy in this sense to indicate that the public exposition of an agent's attributes is minimized. We differentiate privacy from anonymity. Anonymity is the antithesis of reputation; an agent must be (at least pseudonymously) identifiable in order for others to learn about its trustworthiness. Privacy can prevent an agent's identity outside of the system from being known. Thus maximal privacy would reduce the burden of an agent entering or leaving the system. This is because some cost is incurred by an agent divulging its identity in the system, such as the opportunity cost of preventing the agent from assuming a new identity within the system when its reputation is bad. Less privacy can also imply that the agent has some external account or information that the system could use to sanction it. In this sense, privacy acts as a liability limitation much in the way that a firm partially disassociates liability from its employees. The benefits of privacy are that agents have reduced friction of entering and leaving the system. The drawbacks include 1) a possible influx of unfavorably typed agents and 2) agents with bad reputations reverting to a neutral reputation. Both drawbacks are dependent on how other agents measure and handle trustworthiness.

5.1 Existing Trust Systems

Yu et al. [29] provide a method for discovering peers and communicating reputations that maintains accuracy against noisy ratings and malicious peers. However, Yu et al.'s mechanism is weak against ROBUSTNESS because it measures other agent's quality of service (QoS) and only requires that the aggregate QoS be above a certain threshold. This creates a moral hazard wherein strategic agents will maintain reputations just above the threshold. Their mechanism does not meet FLEXIBILITY well, because it is not clear how to weight and aggregate QoS across domains of interaction.

Teacy et al. [27], Jøsang [30], and Huynh et al. [25] present methods of aggregating trustworthiness from peers that can account for uncertainty. Kamvar et al. [26] propose

a self-policing peer-to-peer reputation system that is highly distributed. Like the work of Yu et al. above, the trust measurements and communications of these three works take into account neither the possibility of different domains nor of different utilities involved, thus violating FLEXIBILITY. For example, their methods do not account for whether an agent is trustworthy enough to deliver a single order of a million items if the agent was known previously to be trustworthy to deliver one item. Similarly, these methods assume agents have a specific type and always perform the same actions, at least on a probabilistic basis, regardless of the other agents and situations involved, thus violating ROBUSTNESS. Such an assumption can be reasonable when one agent is interacting with many anonymous agents, such as a company selling a particular brand of food, but often do not hold under nonanonymity when the agents are rational and fewer, or can precisely control their interactions with others.

Zacharia and Maes' [28] mechanism seeks to achieve low-level behavioral goals, such as enabling agents with higher reputations to have more influence on others' beliefs. However, their subjective trustworthiness measures only weakly achieve AGGREGABILITY. Like the aforementioned trust and reputation systems, their measures are highly specific to the interaction domain, which does not meet FLEXIBILITY. Zacharia and Maes tested their system only against malicious agents that build up reputation and then spend it, and do not examine strategic agents, so we are unable to assess how well their system meets ROBUSTNESS.

Saha et al. [31] support EVIDENTIALITY, because their method uses agents' reputations to directly evaluate the future expectations of utility that would be achieved by each possible interaction. However, their method does not meet AGGREGABILITY because agents cannot aggregate information from sources other than their own interactions. Saha et al.'s method is also potentially weak against ROBUSTNESS if agents can easily change identities and exploit favors offered to unknown agents. Further, their method does not meet VIABILITY because agents cannot communicate their knowledge.

Resnick and Sami [32] focus on preventing various types of reputation manipulations, supporting ROBUSTNESS. Whereas their model appears to meet most of the rest of the desiderata, their model discards potentially useful information, partially conflicting with AGGREGABLE. This is particularly limiting in the case when information on a particular product or agent can change, and the system is slow to adapt because of the sudden increase in information entropy.

5.2 Discount Factor and Desiderata

Using an agent's discount factor as its trustworthiness meets EVIDENTIALITY by definition, because each agent can measure others' discount factors and apply them in a direct manner to evaluate its optimal strategy.

Discount factor measurements meet AGGREGABILITY because they can be combined to increase accuracy and precision. The measurements consist of a range or PDF of another agent's discount factor, and can be combined via probability theory to yield further accuracy [11]. The only difficulty with

discount factor measurements is that the measuring agent must account for its best understanding of what the measured agent is experiencing, and must account for the measured agent's best response. Computing the best responses to find the Bayes-Nash equilibria can be a hard computational problem [9]. In our model, we have found the computational complexity of some discount measurements to be relatively simple or readily approximatable, such as when sellers are slowly dropping their prices in a market with more demand than supply. However, in other situations, such as when an agent is aggregating and deciding the validity of many conflicting reports about one agent from other agents, the computational complexity may be high, yielding a potential conflict with computational efficiency in VIABILITY. Further study is required to find the computational complexity for computing other agents' discount factors in various interactions and to determine whether efficient algorithms exist.

In general, agents would not demonstrate a discount factor lower than their actual unless they are competing with others for limited resources. Agents have difficulty credibly demonstrating discount factors above their own because their impatience prevents them from waiting for the postponed, larger utility. For these reasons, discount factors as trustworthiness measures are aligned with ROBUSTNESS. Further, discount factors are strongly influential in many different domains and situations, such as an agents' desire for quality, the rate at which sellers drop their prices, and how quickly agents come to an agreement in negotiation, discount factors. Whereas the exact method of measuring discount factors changes with the role and situation, discount factors as trustworthiness can maintain their strengths with other desiderata across these domains, regardless of domain-specific valuations and capabilities, thus satisfying FLEXIBILITY.

Discount factors' ability to cope with an open system facilitate PRIVACY ENHANCING in the sense that they offer a low barrier to entry and generally do not require external information to be revealed. If one agent knows nothing about another agent, the maximum entropy distribution of the other agent's discount factor is a uniform distribution on $[0, 1)$, which offers some protection against unknown agents as the expected discount factor is $\frac{1}{2}$. If an agent has a priori knowledge of the distribution of discount factors of agents to be encountered, it may use that distribution instead. If unfavorably typed agents repeatedly assume new identities to expunge poor reputations, or attempt to open a large number of pseudonymous accounts to bolster their own reputation (Sybil attacks), then an a priori distribution can be sufficiently pessimistic in a new agent's discount factor at the expense of how quickly an agent can recognize a new but favorably typed agent. Using discount factors as trustworthiness does not prevent implementations from requiring agents to reveal valuation information, and agents may have some ability to evaluate others' valuations. Therefore, discount factors do not maximize this desideratum, but do not directly violate it.

6 MARKET MODEL EXAMPLES

This section illustrates examples of how discount factor may be measured and utilized with our example online market.

The formalization of the full complex model is beyond the scope of this paper, and we leave extensions involving multiple buyers, sellers, and items simultaneously for future work. We include these basic results to motivate our central thesis of the effectiveness of modeling trustworthiness as discount factors. The first two involve typical trust settings. The third and fourth show how agents can gain knowledge of discount factors outside of something that would normally be measured as trustworthiness while contributing to the agents' knowledge of trustworthiness.

We focus more attention on measuring discount factors than on using them in decision models, as the former has received considerably less attention whereas the latter has been widely used [6], [10], [11], [17], [31]. Ben Zion et al. [33] measure the discount factors of people directly by asking them specific questions. Although it is useful to determine an individual's discount factors from an economic perspective, such measurements may not work in a strategic setting. The literature on measuring private discount factors in strategic interactions is rather sparse. To the best of our knowledge, the following works represent most of what is currently known. The models developed by both Rubenstein [34] and Güth et al. [35] yield equilibrium strategies for bargaining between agents when the agents have private discount factors. However, both models require the agents' discount factors to be one of two discrete values. Smith and desJardins [16] measure the minimal upper bound of an agents' discount factors, although their model requires the assumption that agents only reason with one level of mutual information, rather than assuming agents' actions are common knowledge.

6.1 Discount Factors and Production Quality

To demonstrate both how a seller's discount factor can be measured and how a seller may use its discount factor directly in decision making, we employ the frequently studied grim trigger strategy [36], where an agent permanently stops interacting with another after a bad interaction. This strategy is typical of some trustworthiness settings, particularly when many other agents supply a substitutable alternative. For example, people may not return after a bad experience at a restaurant or may not purchase a replacement printer from the same manufacturer if their previous printer required frequent maintenance. Agents in these settings would have preferred to have avoided these bad transactions in the first place.

Consider buyer b deciding whether to purchase an item advertised at a high quality from seller s for some specific price. The seller will make $\bar{\pi}$ profit on a low-quality item, and $\underline{\pi}$ profit on a high-quality item, where $\bar{\pi} > \underline{\pi}$. This means $\bar{\pi}/\underline{\pi} - 1$ is the percent increase in profit by selling the low-quality item. Suppose s knows b communicates with a set of other agents, B , that also buy from s , where B is common knowledge. If s is found selling items below its advertised quality, buyers in B will avoid purchasing from the seller, causing the seller to indefinitely lose a total of $|B| \cdot \underline{\pi}$ utility worth of potential profit every time interval.

We assume b would prefer to not buy the item than to pay the current price for a low-quality item. Buyer b can use

its knowledge of the seller's discount factor, γ_s , to evaluate whether the seller will produce an item at the advertised high quality. If b believes the seller will produce a high-quality item, then it should proceed with the purchase. The seller will produce the high-quality item if it is more profitable, if

$$\underline{\pi} > \bar{\pi} - \sum_{t=1}^{\infty} \gamma_s^t \cdot |B| \cdot \underline{\pi}. \quad (5)$$

In making its decision whether to purchase the item, b will also evaluate (5) using its current knowledge of all of the values involved. If b makes the purchase and finds the item to be of high quality, then b additionally learns that $\bar{\pi}/\underline{\pi} - 1 < \gamma_s/(1 - \gamma_s)|B|$. If the item were of low quality, then the inequality operator would be reversed. Whereas b may not know the value of $\bar{\pi}/\underline{\pi} - 1$, which is effectively the percent increase in profit, with some reasoning b can still find a discount factor measurement. First, b can use a maximum likelihood estimator based on any other information b has available to find the range in profit, similar to how s may estimate the magnitude of B . Second, b can look at comments and feedback from other agents in B to see what types of products and services were offered to previous agents. If the goods or services offered by s match, then it is likely that the ratio of profitability is the same, and b can substitute any change of B .

If s is a typed agent, always producing high-quality goods, then buyers' expectation of γ_s will approach the highest possible value given the valuations involved. If the difference in profit between a high-quality and low-quality item ($\bar{\pi} - \underline{\pi}$) is large, then γ_s will be observed to be close to 1. And, a typed agent producing low-quality goods will attain a low γ_s .

Example 1: Suppose seller s is offering a high-quality item that cost it \$4 at \$5, but could substitute a low-quality item that costs it only \$1. Further suppose it is common knowledge that if the product turns out to be of low quality, the one-time buyer, b , will tell three other agents that each normally purchase one high-quality item per unit time. If b buys the item, then from (5), the discount factor where s would be indifferent between offering high and low quality is $\frac{5-1}{5-4} - 1 = \gamma_s/(1 - \gamma_s) \cdot 3$, yielding $\gamma_s = 1/2$. Agents observing this transaction would see that b reported $\gamma_s \geq 1/2$ if s provided a high-quality item and $\gamma_s \leq 1/2$ if s provided a low-quality item. Given no other information about s 's discount factor and using the maximum entropy distribution (i.e., uniform) yields an expected value of $E(\gamma_s) = 3/4$ if s provides a high-quality item, and $E(\gamma_s) = 1/4$ otherwise.

To illustrate how trustworthiness can be aggregated, consider another potential buyer, c , reading a comment left by b of obtaining a high-quality item noting $\gamma_s \geq 1/2$, and a comment left by another buyer that $\gamma_s \leq 7/8$. If c believes these comments, c believes $\gamma_s \in [1/2, 7/8]$, with an expected value of $11/16$. Figure 2 illustrates the PDFs for this belief. Now, suppose c is deciding whether to buy a different item from s for \$10, and that s must decide between producing a high-quality item at a cost of \$7 or a low-quality item at a cost of \$4. Buyer c will only influence one other agent not to buy from s if it receives a low-quality item. By evaluating and




Fig. 2. PDFs of discount factors given observations.

simplifying (5) as $\frac{10-4}{10-7} - 1 > \frac{11/16}{1-11/16} \cdot 1$ yields $1 > 11/5$. Because this inequality does not hold, c concludes s will provide a low-quality item and therefore it should not buy from s .

6.2 Discount Factor and Product Choice

Measuring a buyer's trustworthiness can be important in a number of settings. If the buyer does not pay after the seller delivers the item, then the best the seller can do is refuse to sell to the buyer in the future and warn other sellers about the buyer. This sanctioning is the same as discussed in Section 6.1, only with roles reversed. If collusion is policed in the system, but imperfectly so, an untrustworthy buyer would be more likely to collude with other agents because it heavily discounts the utility loss of being caught. Colluding buyers could extort a seller into selling at low price because they could leverage their numbers to produce bad reviews for the seller and thus reduce the seller's future revenue. Whereas other agents may eventually discover the collusion, a large number of bad reviews could still harm some of the seller's future revenue.

We investigate one subtle method of measuring buyers' discount factors. We examine what can be inferred about a buyer's discount factor given its purchasing choice between different items. Because a buyer's valuation is private information, the results here do not give a direct measurement of the buyer's discount factor. However, the results give a constraint between the buyer's valuation and discount factor. These constraints can be used to refine existing information about an agent's valuations and discount factor.

Example 2: Suppose agent a purchases tires for a fleet of delivery vehicles. If a purchases tires with a mean expected life of 5 years rather than tires with a mean expected life of 10 years for an additional 80% higher cost, another agent b simply cannot infer that the agent has a low discount factor. If b has a belief about a 's valuations or current financial situation, b may be able to qualitatively infer that either a has a low discount factor, or a is currently in a difficult financial situation, or some combination of both situations apply. Even though a 's actual state remains ambiguous to b , b still knows more about a after having observed a 's choice.

From our motivating example, we assume that the only reliability information provided is mean time to failure (MTTF), which we represent as q . The maximum entropy distribution,

assuming discrete time intervals, is the geometric distribution with the cumulative distribution function (CDF) $Q(t) = 1 - (1 - 1/q)^{t+1}$, where the probability that the item will fail at each time step is $1/q$. We represent buyer b 's expected utility gain from an item per unit time, that is, its willingness to pay per unit time, as, w_b . The buyer's expected utility of purchasing an item k at price p_k with a failure rate CDF of $Q_k(t)$, $E(U_b(k))$, can be represented as

$$E(U_b(k)) = -p_k + \sum_{t=0}^{\infty} \gamma_b^t (1 - Q_k(t)) w_b. \quad (6)$$

Using a geometric distribution of Q_k with mean q_k thus simplifies to $E(U_b(k)) = -p_k + w_b(1 - 1/q_k)/(1 - \gamma_b(1 - 1/q_k))$.

Whereas (6) determines a buyer's utility for obtaining an item, it may also be used by other agents to infer information about a buyer's discount factor or willingness-to-pay. Consider a buyer, b , deciding between two items: k_1 , at price p_1 with an MTTF of q_1 , and k_2 , at price p_2 with an MTTF of q_2 . Say, b purchases k_1 . If item k_1 is universally superior to item k_2 , that is, it is cheaper ($p_1 \leq p_2$) and longer-lasting ($q_1 \geq q_2$), then the only information gained by other agents is that

$$E(U_b(k_1)) > 0. \quad (7)$$

This information can still be useful because it puts a constraint on b 's possible values for its discount factor and willingness-to-pay. If some other agent, a , believes b 's discount factor to be particularly low, then a can use this assumption to infer that $w_b \lesssim p$. Alternatively, if a has knowledge of b 's willingness-to-pay, a can use this knowledge to gain bounds on b 's discount factor by solving Inequality (7) for the desired variable. In some cases, such as when an agent has a high discount factor or a willingness-to-pay greater than the ask price, no further information is revealed because the bounds are less restrictive than the domain of the variable.




Fig. 3. Expected utility of three purchases.

Now, we consider what information b would have revealed to other agents if k_1 was not universally superior to k_2 . In this case, we know that $E(U_b(k_1)) \geq E(U_b(k_2))$. Solving this inequality for either γ_b or w_b can yield zero or one values, and potentially more with distributions other than geometric. Each value is the end of a boundary within which b 's variable lies. Solving for these boundaries may be generalized for b choosing between multiple items. By finding the values at which b would be indifferent between each pair of items and then finding the range where k_1 yields the highest utility, another agent can obtain bounds on γ_b or w_b . Figure 3 shows

an example of choosing between three items, where if an agent knows that $w_b = 30$, then the agent gains the knowledge that $\gamma_b \in [0.14, 0.80]$.

Because b knows its actions are monitored by other agents and b desires to have a perceived discount factor greater than its own, b has an incentive to buy an item that makes it appear as if it had a larger discount factor. Similarly, b may prefer to reveal a lower value for w_b to sellers in order to bring the price down to a lower value faster. Despite these incentives, purchasing actions must be both credible and utility maximizing for b . Except in certain seemingly rare situations, such as where excessive reliance on communication causes b to have an inflated reputation, we have generally found that an agent's optimal strategy is to play in a manner such that other agents will measure its discount factor to be in a truthful range. In our models and previous work [11], the cost for an agent to over-inflate its reputation typically exceeds the benefit of being able to exploit the reputation in the future, influenced by damage that would be done to its reputation by being inconsistent.

6.3 Measuring Discount Factor By Price

A key benefit of using discount factors as trustworthiness is that further information can be obtained in some settings that normally would not involve trustworthiness directly. Suppose a seller, s , will be selling an item in our market model, but has uncertainty about what price it can obtain. We examine what can be learned about a seller's discount factor in a single seller, single item, single impatient buyer market. This simplification yields a negotiation, and if valuations and discount factors for both agents were all public knowledge, the agents could agree on a price without this delay [37].

Example 3: The website Craigslist (<http://craigslist.org>) is a good example of the scenario we formalize in this section. If an agent is selling a used snowblower in the fair weathered Los Angeles market, information on what price the market will bear would likely be scarce. The seller may believe that a few people might be looking for a snowblower for a distant vacation home in the mountains, for a prop in a movie, or just for spare parts. Using these beliefs, along with the knowledge of what a new snowblower would cost to be shipped to the LA area, the seller might start off at a moderately high price and slowly lower the price if no bids are received. The rate that the seller drops the price can be an indication of the seller's discount factor. Even if the seller undergoes a significant valuation change, such as needing to sell the snowblower because of an unexpected move to a smaller location, examining multiple observations of price drop rates can provide information regarding the valuation changes.

Suppose we have one buyer with a willingness-to-pay of w , drawn from a probability distribution. Further suppose that either the distribution of w accounts for the buyers' discount factors or that the buyers have a low enough discount factor such that w is approximately what they would pay. The seller knows its own discount factor, γ_s . The seller can update its asking price once per unit time, its strategy being to price the item at σ_t at time t , and we denote the complete strategy as

$\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_\infty\}$. The seller's expected utility, $U_s(\sigma)$, can be written as (for notational convenience, we set σ_{-1} to the supremum of the distribution of w)

$$U_s(\sigma) = \sum_{t=0}^{\infty} (\gamma_s^t \cdot P(\sigma_t \leq w \cap w < \sigma_{t-1}) \cdot \sigma_t). \quad (8)$$

The seller's optimal strategy is that which satisfies $\text{argmax}_\sigma U_s(\sigma)$. Figure 4 shows results of numerical solutions for the seller's optimal ask price at each time given discount factors of 0.40, 0.60, 0.80, 0.90, and 0.95. In this example, the buyer's willingness-to-pay distribution is exponential with mean 50.




Fig. 4. Optimal ask price schedules for example sellers.

To find the seller's optimal strategy analytically, assuming myopic buyers, we can view each σ_t as an independent variable and maximize the expected utility in (8). We express the distribution of the buyers' willingness-to-pay by the probability density function (PDF), $v(w)$, and the cumulative distribution (CDF), $V(w) = \int_{-\infty}^w v(x)dx$. Because the variables are mutually independent, we can maximize (8) by setting $\forall \sigma_t \in \sigma : dU_s/d\sigma_t = 0$. The initial case, $t = 0$, is separate from the general case yielding the equations for $t > 0$ as

$$\sigma_1 = \frac{V(\sigma_0) + v(\sigma_0)\sigma_0 - 1}{\gamma_s v(\sigma_0)} \quad \text{and} \quad (9)$$

$$\sigma_{t+1} = \frac{v(\sigma_t)\sigma_t - V(\sigma_{t-1}) + V(\sigma_t)}{\gamma_s v(\sigma_t)}. \quad (10)$$

If we assume a uniform distribution of valuation between 0 and some maximum value, \bar{w} , and assume $\lim_{t \rightarrow \infty} \sigma_t = 0$, we can solve the linear constant coefficient homogeneous recurrence relation created by applying (9) and (10) as

$$\sigma_t = \frac{\bar{w}}{(1 + \sqrt{1 - \gamma_s})} \left(\frac{1 - \sqrt{1 - \gamma_s}}{\gamma_s} \right)^t. \quad (11)$$

We have found this result to match our numeric results (computed in http://www4.ncsu.edu/~cjhazard/research/buyer_delay_discount_factor.xls) as shown in Figure 4. The derivation can be found in the Appendix.

Note that the optimal ask prices decrease exponentially over time based on the discount factor at a constant rate; we have also found this numerically with an exponential distribution of w . From this information, a buyer could predict a seller's discount factor based on a small number of asks. When not in steady-state, the seller will also need to model its payoff based on its belief of the buyers' beliefs of its discount factor in case any buyers erroneously believe the seller's discount factor is significantly higher or lower than it really is. Nevertheless, this result provides a lower bound on a single seller's discount factor. Cramton [38] analyzes a similar situation of delay

in bargaining, except when the discount factors are publicly known and valuations are unknown.

6.4 Measuring Discount Factor By Delay

Like the sellers, each buyer has its own discount factor and is trying to maximize its utility. This section also focuses on just one buyer and one seller. We model what the seller can learn about an interested buyer's discount factor and willingness-to-pay, assuming both are constant over time, if the buyer does not purchase at the current asking price, but waits for the seller to lower the price.

A buyer's utility, $U(t)$ is a function of the time it accepts a seller's offer of price σ_t . The buyer's willingness-to-pay, w_b , and discount factor, γ_b , can be used to write its utility as

$$U_b(t) = \gamma_b^t(w_b - \sigma_t). \quad (12)$$

The buyer will have the opportunity to continually reevaluate its optimal time to accept the seller's offer, but the optimal absolute time does not change. This can be seen for some time offset, x , as $U_b(t+x) = \gamma_b^{t+x}(w_b - \sigma_{t+x}) = \gamma_b^x \gamma_b^t(w_b - \sigma_{t+x})$. Because the comparative difference between utilities at different times is scaled by the constant based on the time difference, γ_b^x , the acceptance time that maximizes utility is the same regardless of when the buyer is reevaluating, making the optimal strategy a subgame perfect solution concept.

When a buyer makes a purchase, the seller observes that at the time of purchase, T , the buyer's utility was the largest. Because the price schedule is strictly decreasing, the decisions at $T-1$ and $T+1$ yield the tightest bounds. The corresponding inequalities are $U_b(T) > U_b(T-1)$ and $U_b(T) \geq U_b(T+1)$. If the seller does not have any information on neither the buyer's discount factor nor the buyer's utility, then the seller only observes a relationship between the two. This observed relationship can be expressed as

$$\frac{w_b - \sigma_{T-1}}{w_b - \sigma_T} < \gamma_b \leq \frac{w_b - \sigma_T}{w_b - \sigma_{T+1}}, \quad (13)$$

or alternatively as

$$\frac{\sigma_T - \gamma_b \sigma_{T+1}}{1 - \gamma_b} < w_b \leq \frac{\sigma_{T-1} - \gamma_b \sigma_T}{1 - \gamma_b}. \quad (14)$$

The seller can use its beliefs of the distributions of w or γ along with (13) and (14) to obtain a PDF of the opposite variable, as we will discuss in Section 7.

Competition brought by multiple buyers decreases the delays that buyers are willing to incur to wait for reduced prices from sellers. For example, if two buyers are waiting for two sellers to decrease their ask prices, and the buyer with the higher willingness-to-pay waits long enough such that the price falls below the other buyer's willingness-to-pay, then the item may be taken by the other buyer. The first buyer must then wait until the seller with the higher discount factor gradually brings its ask price down. Not only does the delay incur lost opportunity to the buyer with the higher willingness-to-pay, but the seller with the higher discount factor will use smaller price decrements and the said buyer's optimal strategy may include paying a higher price than the first item.

As the number of buyers increases in proportion to the number of items sold, the ability of a patient buyer to successfully employ strategic delay decreases. Having more buyers means that the difference between a buyer's willingness-to-pay and the next highest willingness-to-pay decreases, increasing the chance that a drop in price will bring the item within range of more buyers. In the same way that an excess of supply pushes the price of items to 0, the limit as the number of buyers goes to infinity is that the expected profit of buyers goes to 0. In this case, the market is undersupplied, and even buyers with large discount factors rationally behave as myopic buyers.

7 AGGREGATING OBSERVATIONS

Because our discount factor measurements 1) employ Jeffrey-like probability conditioning by admitting overlapping observations that do not necessarily cover the full probability space and 2) encompass the full probability space under the assumption that the measurement is accurate, we can employ Bayesian inference interchangeably with the principal of maximum entropy, obtaining the same results [39]. This means we can use the principle of maximum entropy to find agent's initial uninformed beliefs, then use Bayesian inference to update the probability distributions representing agents' beliefs of others' discount factors and willingness-to-pay. These mathematical tools allow agents to aggregate information about other agents' discount factors and valuations from a variety of different measures, including those we discussed in Section 6. We generalize the aggregation of beliefs depicted Figure 2 across probability distributions and types of observations.

Given no a priori knowledge or beliefs about another agent's discount factor, the maximum entropy distribution is uniform on the range of $[0, 1)$. Suppose agent s observes agent b perform an action that would require b 's discount factor, γ_b , to be between 0 and $3/4$ inclusive. The cumulative distribution function (CDF)¹ of b 's discount factor, as a function of discount factor x , is $F_{\gamma_b}(x) = P(\gamma_b \leq x) = x$, yielding $P(\gamma_b \leq 3/4) = 1$ and $P(\gamma_b > 3/4) = 0$. Using conditional probability, the new CDF in the range of $[0, 3/4]$ becomes $F_{\gamma_b}(x) = P(\gamma_b \leq x | \gamma_b \in [0, 3/4]) = P(\gamma_b \leq x \cap \gamma_b \in [0, 3/4]) / P(\gamma_b \in [0, 3/4]) = 4x/3$.

If agent s observes b perform an action, but s can only observe a relationship between b 's discount factor and its willingness-to-pay rather than a direct observation of either, s can still gain some information about both of b 's attributes. Consider the case in Section 6.4, where the observed relation between the willingness to pay and discount factor follow an inequality. We rewrite the relation $\gamma_b \leq (w_b - \sigma_T) / (w_b - \sigma_{T+1})$ in a more general form to encompass other possible observations, dropping the subscripts for convenience, as $\gamma \leq h(w)$. We use the random variable H to represent a random variable on the range of h that is isomorphic to the random variable of the agent's willingness-to-pay. As long as the function h is monotonic, we can map between the CDF

1. By standard definition, a CDF is a nondecreasing function with domain $(-\infty, \infty)$ and range $[0, 1]$. If a random variable's domain is a subset of $(-\infty, \infty)$, then the CDF is defined as a piecewise function to yield 0 below the random variable's domain and 1 above the domain.

of w , F_w , and the CDF of this transformation, F_H , for some willingness-to-pay of x using h as $F_w(x) = F_H(h(x))$. The probability density function (PDF), f_H , may be found in the usual fashion as $f_H = \frac{dF_H}{dx}$.

Given the relationship $\gamma \leq h(w)$, agent s would like to update its beliefs about the observed agent's γ and w . We use the CDFs F_γ and F_H to denote the current beliefs of γ and w respectively, and the CDFs F'_γ and F'_H to represent the beliefs after the new observation has been taken into account. By the definition of conditional probability,

$$F'_\gamma(x) = P(\gamma \leq x | \gamma \leq H) = \frac{P(\gamma \leq x \cap \gamma \leq H)}{P(\gamma \leq H)} \quad (15)$$

and

$$F'_H(x) = P(H \leq x | \gamma \leq H) = \frac{P(H \leq x \cap \gamma \leq H)}{P(\gamma \leq H)}. \quad (16)$$

Simplifying, we have

$$F'_\gamma(x) = \frac{\int_{-\infty}^x f_\gamma(y) \cdot (1 - F_H(y)) dy}{\int_{-\infty}^{\infty} f_\gamma(y) \cdot F_\gamma(y) dy} \quad \text{and} \quad (17)$$

$$F'_H(x) = \frac{\int_{-\infty}^x f_H(y) \cdot F_\gamma(y) dy}{\int_{-\infty}^{\infty} f_H(y) \cdot F_\gamma(y) dy}. \quad (18)$$

After observing an inequality relation between discount factor and a function of willingness-to-pay, (17) and (18) indicate how an agent's beliefs of another agent should be updated. If the observation yielded an equality relation, such as in Section 6.3, similar results can be derived by simply substituting equalities for the inequalities in the initial formulation, leading to the use of PDF functions in place of the CDF (and 1 minus CDF) functions in (17) and (18).

Example 4: Agent s is selling an item as described in Section 6.4. Buyer b has received extremely accurate information about s from other buyers. However, s has no a priori knowledge about b other than b 's willingness to pay follows an exponential distribution with mean of \$1.00, yielding $F_w(x) = 1 - e^{-1 \cdot x}$. With no a priori knowledge of b 's discount factor, s assumes the maximum entropy distribution, the uniform distribution, yielding $F_\gamma(x) = x$.

The seller's initial asking price is \$1.50. Given b 's knowledge of the seller's discount factor, b predicts that the seller's next utility maximizing price will be \$1.40. Just as s asks \$1.50, b purchases the item, because b would attain more utility by purchasing the item now at \$1.40 than waiting for the price to decrease further due to b 's discount factor. The seller observes the second half of the inequality expressed by (13) as discussed earlier in this section, with $\sigma_T = \$1.60$, and $\sigma_{T+1} = \$1.40$.




Fig. 5. Change in PDF of believed discount factor.

By updating its knowledge via (17) and taking the derivative to convert to the PDF, s 's belief of γ_b is expressed by the PDF $f_{\gamma_b}(x) = 1.38398e^{0.1/(x-1)}$, as shown in Figure 5. If s makes another observation about b , s will also need to compute the updated PDF for b 's willingness to pay, and use both of these functions and combine this with the new observation.

The aggregation methods presented in this section will work in many situations, as long as the prior beliefs and observations follow the principal of maximum entropy. If noise and error in signalling are introduced, the beliefs will need to account for the probability of error. If an agent's willingness-to-pay or discount factor can change via a certain process, then the distributions must be recomputed as time progresses and the entropy must be increased according to the uncertainty from the process of change.

8 DISCUSSION

Agents with low measured discount factors behave in ways that are generally considered untrustworthy. An agent with a low discount factor would produce poor quality items, exert low effort on service tasks, and not offer or return favors. In each case, the agent will prefer smaller utility gain now to a larger gain in the future. If an agent a with a low discount factor were entrusted with a secret by agent b , perhaps for mutual benefit, a would not have a strong incentive to keep the secret. Agent a would not put much value on its future relationship with b , and would reveal the secret to some third agent, c , if agent c offered a little short term gain. Having a low discount factor means an agent is myopic and impatient. Under our definitions and assumptions, trustworthiness is therefore roughly equated to patience.

Agents with high measured discount factors often behave in a trustworthy manner. However, the way discount factors as trustworthiness can depart from intuition is when an agent with a high discount factor faces a moral hazard where it does not expect sanctioning to be effective. The agent with the high discount factor would not necessarily be honest when it is not being observed. It is possible for an agent that steals items from other agents to have a high discount factor if the agent believes that the probability of being caught or the utility loss due to punishment will be sufficiently low. One scenario is the agent's beliefs are wrong and other agents observe the undesirable behavior, attributing the behavior to lower valuations or a lower discount factor. Conversely, if the agent's beliefs are accurate and other agents cannot differentiate an agent that is always altruistic (strongly typed) from an agent that is only altruistic when observed (purely utility maximizing), then no objective trust system could measure this.

The discount factor method requires each agent to model another agent's valuations in addition to its trustworthiness. This model affords the first agent an analytically predictive model of the second. Almost any trust model can be tailored to different domains and contexts, such as automobile repair and cooking. However, discount factors can model a single trustworthiness value across the domains, as long as sufficient information is available about the agent's valuations and capabilities (as defined by the value an agent will receive

from another's action) in the different domains. This means even if an agent repairs automobiles well but cooks poorly, its trustworthiness can be consistent across the domains as long as the contexts are equivalent and the agent's valuations and beliefs can be modeled. Even if information is scarce, agents can have mutual information about the information scarcity and attribute nontrusting behaviors to the scarcity of information.

Using its expectation of another agent's valuations in decision models helps an agent evaluate the trustworthiness of agents in complex situations. Suppose b regularly purchases cheap office supplies from s , and always finds them to be of good quality. In this context, s is trustworthy. Because the profit margins on the items are small, b is only able to know that, for example, $\gamma_s > 0.9$. Now suppose a is looking to buy an expensive office chair. The discount factor that b reports may not indicate that s will sell a high-quality office chair in the different setting, depending on the possible profits. If s focuses on office supplies, it may not have the economies of scale to make larger profits on high-quality office chairs, increasing the incentive to provide one of low quality. Note that discount factors coupled with valuations also can work in the reverse; a supplier of expensive niche items may not be able to efficiently offer cheap bulk goods, and may not experience much sanctioning if it were to provide poor quality goods to an unknown single-transaction customer.

Agents' discount factors may change along with Assumption 1 due to various reasons. External factors include a change in the market or the agent's ownership, and internal factors include an agent deciding to leave a given market at a specified future date.

Practical examples of multiagent interactions involving many individuals, firms, and other organizations can exhibit a range of behaviors, including agents that act strategically, agents that behave in a consistent manner directed by a rigid set of beliefs, and agents that fall between the extremes. Discount factors offer a measurement of trustworthiness that is applicable to the range of agent behavior where both adverse selection and moral hazards exist.

Whereas we focus on discrete events in this paper, the total utility function U can also be extended to evaluate events with continuous utility. Markets typically involve agents receiving utility instantaneously as cash flows, but sometimes agents accrue utility over time, such as when leasing a piece of equipment for a specified duration. This continuous utility can be modeled with discrete events in two ways. The example of leasing equipment demonstrates one way, in that the equipment may simply add value to other events occurring in a duration, such as a piece of equipment that can improve the quality of a manufactured good. The second way to model continuous utility is better explained by the example of a person watching a play or movie. Throughout the entertainment, the person is continuously gaining utility. This can be computed as the limit as the changes in utility become infinitesimally small, but take place over an infinitesimally small time period, by effectively using a Riemann sum to compute the utility as the area under the curve specified by the events' utilities.

The above concepts and techniques open up some important

avenues for future research. One is to expand the discount factor measurements to markets with multiple buyers and sellers and more complex market interactions. For example, in a market with more supply than demand, a seller may optimally adopt a discount factor lower than its own in order to increase its chances of selling an item. From this example, a second avenue is to determine when a rational agent will act according to its discount factor and when it will act as if it had another discount factor in response to its current situation. Combining discount factor measurements with other techniques, such as statistics or other reputation measures, for determining other properties about agents, such as their capabilities and valuations, comprises a third avenue to expand the work into additional applications in and beyond e-commerce. A fourth avenue is to apply the discount factors model and Bayesian inference to comments left by one agent about another, to check for conflicting information in the comments, examine discount factors of those agents posting, and evaluate each comment from the best knowledge of the situation at the time the comment was written. We have begun to study how to evaluate the trustworthiness of such information [11], but much remains to be done.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful comments.

REFERENCES

- [1] C. Castelfranchi and R. Falcone, "Principles of trust for MAS: cognitive anatomy, social importance, and quantification," in *Proc. 3rd Int'l Conf. on Multiagent Systems*. IEEE Computer Society Press, 1998, pp. 72–79.
- [2] S. Weeks, "Understanding trust management systems," in *Proc. of the IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2001, pp. 94–105.
- [3] D. Houser and J. Wooders, "Reputation in auctions: Theory and evidence from eBay," *Journal of Economics & Management Strategy*, vol. 15, no. 2, pp. 353–369, 2006.
- [4] S. Sen and D. Banerjee, "Monopolizing markets by exploiting trust," in *Proc. of the Fifth International Joint Conf. on Autonomous Agents and Multiagent Systems*, 2006, pp. 1249–1256.
- [5] C. Subramaniam and M. J. Shaw, "A study of the value and impact of B2B e-commerce: The case of web-based procurement," *International Journal of Electronic Commerce*, vol. 6, no. 4, pp. 19–40, 2002.
- [6] R. Jurca and B. Faltings, "Obtaining reliable feedback for sanctioning reputation mechanisms," *Journal of Artificial Intelligence Research*, vol. 29, pp. 391–419, August 2007.
- [7] S. Rubin, M. Christodorescu, V. Ganapathy, J. T. Giffin, L. Kruger, and H. Wang, "An auctioning reputation system based on anomaly detection," in *12th ACM Conf. on Computer and Communications Security*, 2005, pp. 270–279.
- [8] T. Khopkar, X. Li, and P. Resnick, "Self-selection, slipping, salvaging, slacking, and stoning: the impacts of negative feedback at eBay," in *Proc. of the 6th ACM Conf. on Electronic Commerce*, pp. 223–231.
- [9] V. Conitzer and T. Sandholm, "Computational criticisms of the revelation principle," in *Proc. of the 5th ACM Conf. on Electronic Commerce*, 2004, pp. 262–263.
- [10] J. C. Ely and J. Välimäki, "Bad reputation," *The Quarterly Journal of Economics*, vol. 118, no. 3, pp. 785–814, March 2003.
- [11] C. J. Hazard, "¿Por favor? Favor reciprocation when agents have private discounting," in *AAAI Workshop on Coordination, Organizations, Institutions and Norms (COIN)*, July 2008, pp. 9–16.
- [12] S. D. Ramchurn, D. Huynh, and N. R. Jennings, "Trust in multi-agent systems," *The Knowledge Engineering Review*, vol. 19, pp. 1–25, 2004.
- [13] T. Addison and S. M. Murshed, "Credibility and reputation in peacemaking," *Journal of Peace Research*, vol. 39, no. 4, pp. 487–501, 2002.
- [14] M. Deutsch, *The Resolution of Conflict*. Yale University Press., 1973.

- [15] J. M. Whitmeyer, "Effects of positive reputation systems," *Social Science Research*, vol. 29, pp. 188–207, 2000.
- [16] M. J. Smith and M. desJardins, "Learning to trust in the competence and commitment of agents," *Autonomous Agents and Multi-Agent Systems*, vol. 18, no. 1, pp. 36–82, February 2009.
- [17] C. Dellarocas, "Reputation mechanism design in online trading environments with pure moral hazard," *Information Systems Research*, vol. 16, no. 2, pp. 209–230, June 2005.
- [18] G. J. Mailath and L. Samuelson, *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, 2006, ch. 18, pp. 568–579.
- [19] S. Willmott, M. Calisti, and E. Rollon, "Challenges in large-scale open agent mediated economics," in *Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2002, vol. 2531, pp. 131–150.
- [20] W. E. Walsh and M. P. Wellman, "Decentralized supply chain formation: A market protocol and competitive equilibrium analysis," *Journal of Artificial Intelligence Research*, vol. 19, pp. 513–567, 2003.
- [21] M. P. Wellman, "Market-oriented programming: Some early lessons," in *Market-Based Control: A Paradigm for Distributed Resource Allocation*, S. Clearwater, Ed. World Scientific, 1996, ch. 4, pp. 74–95.
- [22] M. P. Singh, "An ontology for commitments in multiagent systems: Toward a unification of normative concepts," *Artificial Intelligence and Law*, vol. 7, no. 1, pp. 97–113, March 1999.
- [23] T. Vu, R. Powers, and Y. Shoham, "Learning against multiple opponents," in *Proc. of the Fifth International Joint Conf. on Autonomous Agents and Multi-Agent Systems*, 2006, pp. 752–759.
- [24] R. Dingleline, M. J. Freedman, and D. Molnar, "Accountability measures for peer-to-peer systems," in *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, A. Oram, Ed. O'Reilly, February 2000, ch. 16, pp. 271–340.
- [25] T. D. Huynh, N. R. Jennings, and N. R. Shadbolt, "An integrated trust and reputation model for open multi-agent systems," *J. of Autonomous Agents and MultiAgent Systems*, vol. 13, no. 2, pp. 119–154, 2006.
- [26] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "The eigentrust algorithm for reputation management in P2P networks," in *Proc. of the 12th International Conf. on World Wide Web*, 2003, pp. 640–651.
- [27] W. T. Teacy, J. Patel, N. R. Jennings, and M. Luck, "Travos: Trust and reputation in the context of inaccurate information sources," *Autonomous Agents and Multi-Agent Systems*, vol. 12, no. 2, pp. 183–198, 2006.
- [28] G. Zacharia and P. Maes, "Trust management through reputation mechanisms," *Applied Artificial Intelligence*, vol. 14, no. 9, pp. 881–907, October 2000.
- [29] B. Yu, M. P. Singh, and K. Sycara, "Developing trust in large-scale peer-to-peer systems," in *2004 IEEE First Symposium on Multi-Agent Security and Survivability*, 2004, pp. 1–10.
- [30] A. Jøsang, "A subjective metric of authentication," in *Proceedings of the 5th European Symposium on Research in Computer Security*. London, UK: Springer-Verlag, 1998, pp. 329–344.
- [31] S. Saha, S. Sen, and P. S. Dutta, "Helping based on future expectations," in *Proc. of the Second International Joint Conf. on Autonomous Agents and Multiagent Systems*, July 2003, pp. 289–296.
- [32] P. Resnick and R. Sami, "The influence limiter: Provably manipulation resistant recommender systems," in *Proc. of ACM RecSys 07 Conf. on Recommender Systems*, 2007, pp. 25–32.
- [33] U. Benzion, A. Rapoport, and J. Yagil, "Discount rates inferred from decisions: An experimental study," *Management Science*, vol. 35, no. 3, pp. 270–284, March 1989.
- [34] A. Rubinstein, "A bargaining model with incomplete information about time preferences," *Econometrica*, vol. 53, no. 5, pp. 1151–1172, 1985.
- [35] W. Güth, S. Kröger, and H.-T. Normann, "Durable-goods monopoly with privately known impatience: A theoretical and experimental study," *Economic Inquiry*, vol. 42, no. 3, pp. 413–424, 2004.
- [36] R. Axelrod, "On six advances in cooperation theory," *Analyse & Kritik*, vol. 22, pp. 130–151, July 2000.
- [37] A. Rubinstein, "Perfect equilibrium in a bargaining model," *Econometrica*, vol. 50, no. 1, pp. 97–109, January 1982.
- [38] P. C. Cramton, "Strategic delay in bargaining with two-sided uncertainty," *The Review of Economic Studies*, vol. 59, no. 1, pp. 205–225, January 1992.
- [39] P. D. Grünwald and J. Y. Halpern, "Updating probabilities," *Journal of Artificial Intelligence Research*, vol. 19, pp. 243–278, 2003.



Christopher J. Hazard Christopher J. Hazard is a PhD candidate in the Department of Computer Science, North Carolina State University, Raleigh. His research interests include multiagent systems and e-commerce, particularly on the topics of trust, reputation, game theory, and dynamical systems theory. He is also the founder and president of Hazardous Software, which makes multitemporal simulation software for training and entertainment, being showcased in their upcoming PC game *Achron*. After graduating from Valparaiso University in 2001, he designed CDMA infrastructure software at Motorola. He has also worked on robot coordination algorithms with Kiva Systems and NCSU, won an Outstanding Teaching Assistant award for organizing and teaching e-commerce classes at NCSU, and won a best paper award at ICEC '07.



Munindar P. Singh Munindar P. Singh is a Professor in the Department of Computer Science, North Carolina State University, Raleigh. His research interests include multiagent systems and service-oriented computing with a special interest in the challenges of trust, service selection, and business processes in large-scale open environments. His books include the coauthored *Service-Oriented Computing* (New York: Wiley, 2005). Dr. Singh was the Editor-in-Chief of the *IEEE Internet Computing* from 1999 to 2002.

He is a member of the editorial boards of *IEEE Internet Computing*, *Autonomous Agents and Multiagent Systems*, *Web Semantics*, and *Service-Oriented Computing and Applications*.