

The Market for Evaluations

Christopher Avery, Paul Resnick, and Richard Zeckhauser*

Abstract

Recent developments in computer networks have driven the cost of distributing information virtually to zero, creating extraordinary opportunities for sharing product evaluations. We present pricing and subsidy mechanisms that operate through a computerized market and induce the efficient provision of evaluations. The mechanisms overcome three major challenges: first, evaluations, which are public goods, are likely to be underprovided; second, an inefficient ordering of evaluators may arise; third, the optimal quantity of evaluations depends on what is learned from the initial evaluations. Keywords: evaluations, information sharing, product quality, computer network, market (JEL D70, D83, H41, L15)

Subjective evaluations by others are a valuable tool for consumers who are choosing which products to buy or how to spend their time. For example, we read magazines devoted to product evaluation before purchasing cars and appliances. We ask our friends and read reviews by professional critics when selecting movies and restaurants. Professional colleagues recommend articles. Product evaluations are provided by friends, other consumers, brokers, and frequently even product suppliers.

Yet the use of evaluations is severely limited by today's production, collection, and distribution systems. The production of evaluations is costly, requiring money or time for people to purchase and try a product and then to communicate their reactions to it. The collection of evaluations and their distribution to others is also costly. It is even costly for individuals to process evaluations, especially contradictory evaluations. These transaction costs reduce the use of evaluations, although they are still used frequently when they are entertaining (as with movie reviews) or may influence an expensive purchase (as with assessments of cars).

Computers, which reduce the costs of collecting and distributing information, create new opportunities for evaluation sharing. A reader can enter a numeric evaluation of a product with a single keystroke. That information can be swiftly and cheaply transferred to other computers. Those computers, acting as agents, can process the information for their owners and, if advice is requested, recommend purchase. Computer-based evaluation services have the significant advantage that they can tailor recommendations to each individual's tastes. For example, Internet services keep track of which books, movies, audio CDs, or bulletin board messages each subscriber likes and dislikes (Paul Resnick and Hal Varian 1997). The services perform statistical analysis to match users whose preferences correlate with one another, and then make personalized recommendations, with evaluations by those with similar tastes weighted more heavily.¹ Eventually, we expect to see evaluation services for many products, such as restaurants, and journal articles, and even for service providers, such as doctors, lawyers, and landscapers. As the number of small vendors on the information superhighway grows, it will probably be beneficial to distribute evaluations of vendors, in an expanded form of services already provided by Better Business Bureaus.

Even if computers and networks can make the costs of entering and distributing evaluations trivial, there is still the cost of purchasing a product and evaluating it. These costs may outweigh the consumer's expected benefit from consuming it. Even where product price is a trivial consideration, evaluation costs are still lost when the product is disliked. It is sometimes socially efficient, however, for an individual to try a product despite its negative expected payoff, so that others can benefit from her evaluation.² To achieve social efficiency without coercion, therefore, the expected or actual gains from an efficient sequence of evaluations must be redistributed; that is, those who choose later whether or not to buy the product must compensate those who evaluate it earlier.

This paper proposes a market with cash payments to coordinate production schedules and cost allocation. Our analysis deals with people who are already in a particular goods market; e.g., they wish to hire a lawyer. For some products, say bulletin board messages

or professional journal articles, many individuals will be perpetually in the market. The decision to enter the market is beyond this analysis. Reliance on pricing to coordinate production runs counter to the Internet ethos, which discourages monetary payments for information or services. Yet barter and free provision often lead to woefully inefficient outcomes. We suspect that monetary payment for material provided over the Internet will increase dramatically, in part because methods will be worked out to secure payments and maintain their confidentiality (Jeffrey K. Mackie-Mason and Hal Varian 1994).

Section I lays out the theoretical background for this paper, identifying the special properties of evaluations and the markets that would coordinate their production, distribution, and consumption. Section II sets out a formal game model in which each evaluation provides additional information about the likelihood that future consumers will like the product. Section III considers allocation mechanisms. It begins with two-person examples that illustrate the need for pricing and the difficulties with simplistic pricing schemes. It then presents pricing schemes that secure the socially optimal order and quantity of evaluations. Beyond efficiency, we look for schemes that balance the budget, charge the same price to all individuals taking the same action, and secure voluntary player participation. It proves possible to guarantee any two of these properties, but not all three simultaneously. In our base model, individuals differ in their benefits and costs from products they like or dislike, but are identical in tastes and their ability to make informative evaluations. Our results extend readily to an expanded model with several classes of individuals, differentiated by tastes and evaluation skills.

I. Theoretical Background

Evaluations are unusual commodities; they can not be efficiently produced in a standard market. They possess three distinctive properties:

Evaluations may be treated as public goods. Evaluations are nonrival if the commodity being evaluated has elastic supply, and each person can benefit from an evaluation without reducing its value to anyone else. For example, the benefits of reading a book or buying an appliance are rarely affected by the number of other readers or buyers.³ The voluntary provision of public goods leads to a suboptimal supply, since no individual takes account of the benefits that her provision gives to others. If information is costly to acquire, as it often is, too little will be provided.

Current production and future consumption are antagonistic. There is an opportunity cost to trying a product now (and producing an evaluation) rather than waiting for further evaluations before deciding whether to consume.⁴ Any mechanism that elicits information solely about the direct costs of producing evaluations will not be efficient because it does not take into account opportunity costs, which differ among consumers.

Production plans are contingent. Each individual's production plan is contingent both on the outcome of early evaluations and expectations about others' production. For example, an individual may consume after a favorable initial evaluation, but wait after an unfavorable one, expecting

that someone else will consume and further inform him whether or not he should consume. Thus, the opportunity cost calculation is complex and requires information about contingent future actions, not merely current preferences. Adjusting the amount of a public good procured in response to late-breaking information is analogous to tailoring the size of a posse to the danger of the desperado, an approach that is clearly possible with many public goods, but is extremely rare in practice.

The Importance of Computer Networks. These three properties suggest the need for complex incentive structures to procure efficient production of evaluations. Computer networks facilitate market-based solutions in three ways. First, a computer program acting as a centralized broker can perform complex calculations to resolve bids. A computerized broker can also keep comprehensive records of past purchases and satisfactions, and can undertake data-intensive statistical calculations to estimate the benefits particular individuals will get from different products. Second, the mechanical and verifiable character of automated processes may make electronic brokers seem more trustworthy than human arbiters or market makers.⁵ Third, since computer programs can act as agent on behalf of people, and thereby facilitate the use of bidding mechanisms that otherwise would be too cumbersome.

Take an extreme case, electronic bulletin board messages. The purchase cost is zero, but the evaluation costs can be high since people cannot possibly sift through all the messages that might interest them. A market for evaluations could coordinate decisions about which people should read and evaluate particular messages. People would be unlikely to make explicit cost-benefit analyses when deciding whether to read a message, but software programs acting on their behalf could easily do so.⁶ Human effort would only be required to evaluate the messages; the market would be fully automated.

For items where the purchase price is significant, as with automobiles, we would expect less than full automation of the market. The computer agents would merely recommend purchases to their consumer owners, who would then determine whether to buy. If they did buy, they would then evaluate and inform others.

II. The Evaluation Acquisition Game

Each player in a group faces a single decision, whether to consume a product.⁷ We assume that someone who consumes a product incurs no additional cost to tell everyone else her evaluation; i.e., evaluations become public knowledge.⁸ If one player likes the product, it raises the next player's expected payoff and makes him more likely to consume it as well. This is the evaluation acquisition game.

We make five simplifying assumptions throughout. First, players are risk neutral, so that they are concerned only with expected payoffs. Second, consuming a particular product provides the same benefit (or cost) any time during the game. Thus, waiting is a weakly dominant strategy. Third, if a player is sure to consume the product eventually, we sometimes assume a sliver of discounting or altruism to break ties and have the player consume earlier. Fourth, players report evaluations honestly. Fifth, each person can gain value from a product without preventing the use or diminishing the benefits of others.

Our base model assumes that all players are equally informative as evaluators, and that all recipients get informed equally. There is some probability, ρ , that the next consumer will perceive the product to be Good.⁹ The uncertainty may result from some combination of randomness in the underlying product (occasionally, the chef at a restaurant has a bad day) and randomness in the consumer's perceptions (occasionally, the consumer is in a bad mood and dislikes a superbly prepared meal.) Players may differ in the intensities of their preferences; that is, the payoffs from consuming products they perceive to be “Good” or “Bad”. Following Bayes’ Rule, when someone consumes a product and reports a positive evaluation the assessed value of ρ increases; a negative evaluation reduces its value. If consumers do not reliably report their evaluations, because they sometimes initially misdiagnose their own evaluation of the product or because of occasional data entry errors, the updating rule for ρ is a slightly more complicated application of Bayes' Rule.

None of our results depend on a specific source of uncertainty in ρ or evaluation reporting, so long as ρ and the reliability of reports are common knowledge. Our results do depend on the sources of uncertainty being identically distributed for all consumers, so that an evaluation from any consumer is equally informative. Allowing for systematic, correlated differences in tastes, for variable expertise, or for reliability of reporting requires the more elaborate model discussed in Section III.E.

To facilitate exposition and intuition, we describe a special case of the base model in which uncertainty derives only from the consumer's evaluation process. We use this special case in examples, but rely on the more general formulation for all of our propositions and proofs. In the special case model, there are two underlying quality states for the product, “good” and “bad,” with an initial probability p that the product is good. Each evaluator imperfectly perceives the true state of the product, and hence provides imperfect information to future potential consumers. The level of expertise of evaluators is modeled by two parameters, the probability that a player perceives a good product as Good, and the probability that she perceives a bad product as Bad. We call these parameters g and b , where:

$$g = \text{pr}(\text{perceive product Good} \mid \text{product is good}), \text{ and } b = \text{pr}(\text{perceive product Bad} \mid \text{product is bad})$$

(We assume that $g > 1 - b$.) The players’ evaluations are independent, conditional on the true state of the product. Note that a probability p of a good product does not imply that a person will perceive the product to be Good with that probability. In fact, a person may perceive a product to be Good either by correctly classifying a good product, or by misclassifying a bad product, so that $\rho = pg + (1 - p)(1 - b)$. After an evaluation is received, participants can use Bayes’ Theorem to update p , and hence update ρ . We assume that g , b , and the current value of p are common knowledge, implying that ρ is as well.

There are two critical parameters for each potential consumer -- the payoffs from consuming a product perceived to be Good and one perceived to be Bad. A player who consumes a product incurs a cost; even if the product itself is inexpensive or free, time is a scarce resource, and time spent consuming one product takes away from time to

consume others or to do something else. Call r_i the value of consuming a product that the person perceives to be Good and s_i the value of consuming a product the person perceives to be Bad. Typically, the value gained from consuming a good product will outweigh the cost, so that $r_i > 0$, but for a bad product, $s_i < 0$. We call r_i and s_i a player's intensity values. A person who does not consume a product receives a payoff of 0.

The expected payoff from consuming immediately is:

$$(1) r_i \rho + s_i (1 - \rho).$$

Absent the opportunity to wait for evaluations from others, a player would consume if this expected payoff is non-negative. If other evaluations are forthcoming, however, there is an opportunity cost, since a player would be no worse off, and quite possibly better off waiting until someone else's evaluation yielded a more informed estimate of ρ .

We wish to maximize the sum of expected utilities, scaled to some common metric, for all of the potential consumers. The principal challenge is to arrange for the optimal quantity and sequencing of evaluations. In our model, the broker calculates the efficient allocation given the players' intensity values (the r_i and s_i values), and then offers payments that induce players to choose actions consistent with the efficient allocation.¹⁰ In the pricing schemes we propose, the players need not know each other's intensities and the broker need only know the pool of (r_i, s_i) values, not those of each individual player. (The sole exception arises when the broker must know the players' identities in order to price discriminate; see Proposition 6.)

Our assumption of full and honest evaluations is a limitation that merits further attention. In each of our pricing mechanisms, an evaluator's expected payment is independent of her report. As a result, it is a weakly dominant strategy to report evaluations honestly. We assume that effort is not a choice variable.¹¹

This analysis addresses two contrasting forms of allocation problem. In the first, the batch mode game, there are just two possible rounds of consuming, and multiple people can consume in each round. In the second, the one-at-a-time game, there are multiple rounds, each with one evaluator. For each game, we show how to compute the socially optimal allocation.

For products that are consumed regularly, e.g., movies or bulletin board articles, the batch mode may prove to be of greater practical import than the one-at-a-time game, since it will be easier to coordinate the activities of evaluators if they evaluate in batches. For example, the first round might close 24 hours following the posting of a bulletin board article, or three days after the opening of a movie. First round consumers would have some flexibility on timing their task. The batch mode lowers coordination costs, since prices and consuming schedules are updated periodically -- once in the two-round game -- rather than after each evaluation. In the one-at-a-time game, which might be more appropriate for products such as doctors or automobiles, all players would have to be regularly available, and willing to accept the time delays of sequential evaluation.

A. The Batch Mode Game

Consider first a two-round batch mode game, where first-round consumers provide evaluations that help second-round players make more informed decisions about whether to consume. Figure 1 outlines the steps in the game.

---Figure 1 approximately here---

The Batch Mode Social Optimum. The optimal choice of evaluators in round 1 depends on the difference between the marginal value of an evaluation and the marginal cost of its production. The marginal value of an evaluation is the incremental benefit it provides to the players who wait until the second round to decide whether to consume the product. The marginal cost includes the evaluator's expected gain or loss from consuming immediately, as well as her opportunity cost of not using the information produced by others' evaluations.

More formally:

a_i = player i 's expected value of consuming in the first round. Thus $a_i = r_i \rho + s_i(1 - \rho)$.

$b_i(n)$ = player i 's expected payoff of deciding whether to consume, after n people have provided evaluations.

$x_i(n) = b_i(n) - a_i$ is player i 's expected advantage of deciding after n more evaluations rather than consuming now. It represents the cost for player i of agreeing to take another participant's place as the n th evaluator.

Note that all the values a_i , $b_i(n)$, and $x_i(n)$ are expected values that can be computed before anyone consumes the product. Thus, it is also possible to compute which players should consume immediately and which should wait to decide in the second round. The efficient allocation maximizes the sum of all players' expected payoffs:

$$(2) \quad \sum_{\text{evaluators}} a_i + \sum_{\text{waiters}} b_i(n) = \sum_{\text{evaluators}} a_i + \sum_{\text{waiters}} (x_i(n) + a_i) = \sum_{\text{all}} a_i + \sum_{\text{waiters}} x_i(n)$$

The sum of the a_i 's is a constant, so maximizing equation (2) is equivalent to maximizing the sum of benefits to those who receive the initial evaluations:

$$(3) \quad \max \sum_{\text{waiters}} x_i(n)$$

There are m players in all. For any fixed number of evaluations n , it is optimal to select the people with the $m - n$ largest $x_i(n)$ values to wait. That is, we select the n least-cost initial evaluators, where cost includes the opportunity cost of not waiting for more information. More formally, $x_{(j)}(n)$ is the j^{th} order statistic of $x_i(n)$. That is, for each n ,

array the $x_i(n)$ values in ascending order, and let $x_{(j)}(n)$ be the j^{th} smallest. The optimal number of evaluations from equation (3) reduces to:

$$(4) \max_n \sum_{j=n+1}^m x_{(j)}(n).$$

Moving one more person, i , from the waiting group to the initial consuming group has two effects. First, it increases the expected value for those who remain in the waiting group because they receive an additional evaluation. Second, i gives up the expected value of waiting until the second round. The difference between these values is the difference between the marginal value and marginal cost of i 's evaluation.

Equation (4) maximizes the net benefits of evaluations but it does not necessarily minimize the immediate costs of their production. One implication of Equation (4) is that it may be socially beneficial for some players who expect to lose personal value if they consume in the first round to do so. The initial evaluators may discover that the product is surprisingly valuable, a discovery that benefits the remaining players. Another implication is that while it might seem natural to pick players with the lowest direct costs as first-round evaluators, such a scheme neglects opportunity costs. Even if a player has a high initial expected value from consuming the product, she may benefit greatly from waiting for more information; it could be socially optimal for someone else to consume first -- perhaps even someone with a negative expected value.

B. *The One-at-a-Time Game*

----Figure 2 approximately here-----

In the one-at-a-time game, evaluations are communicated quickly to all players, over many rounds. In each round, one person evaluates, while others wait to consume the new information. Figure 2 summarizes this process. Choosing the number and order of evaluators presents two complexities not found in a conventional sequential sampling problem with varying costs per draw. First, there is no coercive social planner; instead, the pricing scheme must induce the social outcome from the confluence of individual optimizing decisions. Second, the cost of an evaluation is the sum of the direct costs of production (i.e., the product price plus the evaluation cost minus the expected value) and an opportunity cost (i.e., the foregone benefit of receiving evaluations from others which might lead to a better estimate of the expected value and hence a more informed decision about whether to consume the product).

The One-at-a-Time Social Optimum. In the one-at-a-time game, the efficient production schedule specifies a complete ordering of consumers contingent on each possible sequence of evaluations. For example, if the first consumer likes the product, then it may be optimal for player B to consume the product next, while if the first consumer dislikes the product, it may be optimal for player C to consume it next, or perhaps for no one else to consume it.

Formally, an allocation can be represented in a binary evaluation tree, as shown below in Figure 4. Each node represents an evaluation by a player (the root of the tree represents the first evaluation), and each branch emanating from that node represents a possible outcome of that evaluation. A path terminates when no one else is to consume the product. Each player consumes the product at most once, so no player can appear more than once along a single path.¹²

The value to a player of any node in the tree is the expected value of consuming the product, given the updated value of p at that node of tree. The expected social value of an allocation is the sum, over all nodes in the tree, of each node's value to the player it is assigned to, weighted by the probability that it is reached. For a fixed number of players n , the number of possible allocations is finite, implying that there is always an efficient allocation.

C. Sample Calculations

Consider the three-person game defined by the payoffs shown in Figure 3. The initial probability that a product is good is $2/3$, and evaluators correctly perceive the state of a product with probabilities $g = 3/4$ and $b = 3/4$.

---Figure 3 approximately here-----

The efficient batch mode allocation is intuitively clear: since player C benefits enormously from any additional information, both A and B should consume in the first round.

Figure 4 shows the efficient allocation in the one-at-a-time game. Player A consumes the product first. If that evaluation is Good, then player C consumes the product next, while if it is Bad, player B consumes it.

---Figure 4 approximately here-----

At each node, f represents the probability of reaching that node, given the initial values for g , b , and p . Here, p represents the probability that the product is actually good, conditioned on the sequence of evaluations leading to that node. Even though the initial probability that a product is good is $2/3$, the first consumer will perceive it to be Good with probability only $\rho = pg + (1-p)(1-b) = 2/3 * 3/4 + 1/3 * 1/4 = 7/12$. Hence $f = 7/12$ at the node following an initial positive evaluation. Other f values are calculated in like manner; p is updated using Bayes' Theorem.

Intuitively, the prime consideration is to optimize the outcome for player C, since her payoffs, both positive and negative, are so large. While player C would be willing to consume the product immediately, it is better for her to await the results of others' evaluations. If the first evaluation is positive, player C will want to consume the product, no matter what the outcome of the second evaluation. (After a positive and a negative evaluation, the product will still be good with probability $2/3$.) Hence, if the first evaluation is Good, player C should consume the product immediately. If, however, the first evaluation is Bad, player C would be influenced by the outcome of the second

evaluation, so another player should provide it. Overall, player C should consume the product as soon as there is one positive evaluation, but if both players A and B dislike the product, then player C should not consume it. Interestingly, although the sequences (Good, Bad) and (Bad, Good) produce the same amount of information about the product, they lead to a different total number of players consuming of the product. After the sequence (Good, Bad), the remaining player, player B, does not consume the product; after the sequence (Bad, Good), the remaining player, player C, does.

This example suggests that simple rules-of-thumb are not sufficient to determine the efficient ordering of evaluators. Some general rules that do hold at the optimum, however, are:

1. Anyone who is certain to consume the product should evaluate it at the outset of the game;
2. If player A's payoff is always more favorable than player B's ($r_A \geq r_B$, $s_A \geq s_B$ and at least one of these holds with a strict inequality), then A should consume before B.
3. If additional players are added to the game, then there should be (weakly) more evaluations in all circumstances.¹³

D. Discussion of the Games

To sum up, the batch mode game well illustrates the first two distinctive properties of evaluations: their nature as public goods, and the antagonism between current and future production. The one-at-a-time game highlights the third property: future production is contingent on early evaluations. The example above shows the advantage that contingent production offers. In the two-stage batch mode allocation, the number and order of evaluators cannot depend on the content of the first evaluation; rather, the social optimum requires initial evaluations by both players A and B. The one-at-a-time game allows a more nuanced approach; the choice of the second evaluator depends on the outcome of the first evaluation.

III. Inducing the Efficient Allocation

Left to their own devices, individual decision makers usually do not generate the socially efficient allocation. They may produce too few evaluations because they ignore the positive externality their evaluations provide to others. The wrong players may provide early evaluations if those who could benefit most from the evaluations are also most willing to provide them. Players may dissipate the available social surplus that evaluations generate by trying to claim larger individual shares, say by waiting for others to evaluate first. But a central broker that knows the distribution of intensities of participants' preferences -- that is, the pool of r_i and s_i values -- can calculate the efficient allocation of evaluators in either the batch mode or one-at-a-time protocol. The broker can offer side payments to induce players to choose actions generating the social optimum. We assume that the broker pursues efficiency, and seeks no profit from his role.

A. Examples Yielding Inefficiencies

Before proceeding to more general results, we present three instructive two-person batch-mode games that illustrate three reasons to shift away from the laissez-faire equilibrium.¹⁴ To simplify the arithmetic in this section, we assume that a single evaluation is perfectly informative in identifying the state of the product: $g = 1$ and $b = 1$. That is, if one player perceives the product as Good, the other player will too. In this situation, $\rho = p$, and the expected payoff calculation from formula (1) reduces to

$$(5) \quad r_i p + s_i (1 - p).$$

In all three games, it is equally likely that the product is good or bad ($p = .5$). Hence, the utility of consuming the product in the first round is $.5(r_i + s_i)$. After a first-round evaluation by one player, the other player can make a perfectly informed decision: consume the product after a positive evaluation but not after a negative one (assuming $r_i > 0 > s_i$). The expected payoff to player i of waiting while the other provides an evaluation is $.5(r_i + 0)$.

The payoffs to each person vary in the three games (see Figure 5). For example, in game (i), player A benefits by ten units ($r_A = +10$) if she consumes the product and likes it, but loses twelve units ($s_A = -12$) if she consumes it and dislikes it. These varied payoffs lead to different classes of inefficiency.

---Figure 5 approximately here-----

Each player can choose from two initial strategies: consume immediately (C) and wait (W). We assume the payoffs are common knowledge. Figure 6 shows the games in the 2x2 normal form after converting individual outcomes to expected payoffs. The equilibrium in each game is denoted by E. An asterisk (*) denotes an efficient outcome.

---Figure 6 approximately here----

In game (i), it would be optimal for one person to consume immediately and the other to wait. The consumer has an expected loss of 1, since $a_i = r_i p + s_i (1 - p) = .5(10 - 12) = -1$. The value to the waiter of the information generated by that consuming is $b_i(1) = .5(10) = 5$. Unfortunately, waiting is a strictly dominant strategy for both players; no matter what player A does, player B gains by not consuming the product initially, and vice versa. The Nash equilibrium is (W,W), giving each a payoff of 0. Thus, game (i) demonstrates the natural tendency to underprovide information in equilibrium, since neither player cares about the value his or her information provides to the other player.

In game (ii), player B's stakes are higher than player A's: he can gain more if the product is good but his cost of consuming a bad product is also higher. Intuitively, it might seem socially optimal for B to consume first, since his expected value of consuming is positive, while that for A is negative. However, this reasoning neglects B's opportunity cost of not waiting for better information. The social optimum is (C,W). Player A consumes the product first and then, if she finds it Good, recommends it to

player B. In effect, player A acts as the king's taster; if the taster does not get sick, it is safe for the king to eat. Of course, if left to herself, player A will refuse this role (an expected payoff of -1), preferring to wait (at least a 0 payoff). Player B realizes that player A will not provide an evaluation and thus (W,C) is the unique Nash equilibrium, assuming minimal altruism on the part of B induces him to evaluate in the first round. The expected payoff is (5,10). In contrast, the expected payoff of (-1,20) when player A consumes first yields a higher total value. Game (ii) demonstrates that the uncoordinated equilibrium may involve the optimal quantity of evaluations but a nonoptimal order of evaluators.

In game (iii), the players' payoffs are identical and both would be willing to consume immediately, if not for the opportunity to get more information by waiting for the other's evaluation. At the social optimum, one consumes and the other waits. Asymmetric allocations of this sort are difficult to coordinate on a voluntary basis. Without any coordinating mechanism, the players are likely to engage in a costly game of waiting, not unlike a war of attrition, hoping to be the last to decide.

With a common discount factor, δ , for the two players, there is some impetus to consume immediately in game (iii). Then, the unique symmetric equilibrium calls for each player to consume immediately with probability $\frac{1-\delta}{5\delta}$, yielding expected payoffs of 1 each. The discount factor induces the players to sometimes consume in the first round, but the surplus gained is exactly offset by the losses from discounting when both players wait until the second round. For any δ , attempts to claim the surplus dissipate all of the potential value from acquiring information.

B. *Connections to Statistical Herding*

Underprovision of evaluations is similar to but offers an interesting contrast with the phenomenon of statistical herding that arises in models of sequential decision making such as those by Sushil Bikhchandani, David Hirshleifer and Ivo Welch (1989) (BHW) and Abhijit Banerjee (1992). In these herding models, players make decisions in a predetermined order about whether to take some action, such as entering a restaurant. Each person holds distinct information relevant to the decisions of the others. Whenever anyone makes a decision, her choice (and hence imputed private information) becomes public knowledge (the restaurant has a large window), but the actual payoff is not revealed. Strikingly, the group converges on a single decision in finite time because, eventually, the accumulated public information dominates the private information held by any individual. Everyone who decides after this point follows the public consensus, and hence reveals no private information. Since in the end much private information is not revealed, the consensus may be wrong.

Premature consensus can arise in either information dissemination (herding) or information acquisition (the market for evaluations), but such errors in the market for evaluations can occur only in negative decisions. Herding can occur in both positive and negative decisions, because once players reach a premature consensus in either direction no further information is disseminated. By contrast, the inefficiency in evaluation acquisition is asymmetric. Good products can easily get buried beneath negative initial

reviews (as any talented but unsuccessful playwright will tell you). If the product is bad, however, that information will come out in the long run, despite positive early evaluations.

C. Payment Schemes

Transfer payments among the players could achieve full efficiency in the evaluation acquisition game, overcoming the problems of underprovision, wrong ordering, and wasteful surplus claiming. Absent transaction costs, private bargaining will lead to efficient allocations of resources despite private information and externalities, as the Coase Theorem (1960) tells us. In game (i), for instance, player A could pay player B any amount between 1 and 5 to evaluate the product, thereby restoring social efficiency.

Since transactions costs grow with the number of players, however, an agreement on dividing the surplus may be elusive. One solution is for the players to rely on a central broker to coordinate the production of evaluations and distribution of the surplus. The broker offers subsidies to initial evaluators to ensure an efficient outcome. In game (i), for example, the broker could offer 6 units to the initial evaluator, so that each player gains 5 units in expectation. (In the next section we consider how the broker might demand contributions from the players to cover the subsidies, but for now we assume some outside source of funds.)

When the players (or their agents) have full knowledge of each other's preferences, it may suffice for the broker to announce merely an equilibrium price subsidy. For instance, a subsidy of 7 in game (ii) redresses the imbalance in incentives so that A prefers to evaluate the product rather than to wait. Player B waits, secure in the knowledge that A will evaluate the product first. Even with a large number of players who know each other's preferences, the announcement of equilibrium prices can achieve the social optimum if each player employs backwards induction to determine the likely actions of others.

However, if preferences are not known or players cannot be relied upon to solve a complicated backwards induction problem, announcing only equilibrium prices is insufficient. For example, in game (ii), each player's decision depends on the expected action of the other. Player A will refuse a subsidy of 3 if she thinks player B would be willing to accept it. But if she thinks player B would refuse the subsidy of 3, A will take it. Information beyond prices must be announced.

One possibility would be to both announce a subsidy and assign actions to players. For example, a subsidy of 3 together with the assignment of player A to consume first will be sufficient: A knows that B will wait, so A prefers to consume ($-1 + \text{subsidy} = 2$) rather than wait (0). Given the subsidy, the announced assignment is a Nash equilibrium. This scheme suffers two disadvantages. First, it requires the broker to know each individual's intensities, not merely the pool of all players' intensities. Second, it puts the broker in the authoritarian role of assigning actions rather than merely offering subsidies to any player willing to perform the actions.

Fortunately, the broker can merely announce a price and an equilibrium production schedule, but let players choose their own positions. For example, in game (ii) the broker

can announce a subsidy of 7 for the first round and tell each player that the price has been chosen so that if she does not accept it, the other one is sure to do so. Player A now chooses between consuming ($-1 + \text{subsidy} = 6$) or waiting while B consumes (5); hence A prefers consuming. Player B chooses between consuming ($10 + \text{subsidy} = 17$) or waiting while A consumes (20); hence B prefers waiting.

More generally, in the multi-player batch mode game, the broker will announce a price and a quantity n . In order to calculate rational actions for themselves, players need not know each other's intensity values. Each player knows only that in equilibrium exactly n players are sure to accept the subsidy to consume in the first round, and decides whether to be one of them. Similarly, in the one-at-a-time game, the broker will announce subsidies for roles that players can take in an evaluation tree. A role consists of a combination of nodes, no more than one on each path. A player who accepts the subsidy for a role commits to provide an evaluation at each node that is part of that role, should the sequence of evaluations leading to that node actually occur.

D. Desirable Traits for a Pricing Scheme

In addition to ensuring efficient allocation, there are a number of additional properties that might be desirable in a pricing scheme. We consider three. First, players providing (or consuming) the same service should receive (or pay) the same price; we call this *same-action same-price*, or SASP. SASP is required given our assumption that the broker knows the distribution of types, but can't identify a player's type. In the batch mode, SASP requires that every player be offered the same subsidy for making a first-round evaluation, and thus precludes price discrimination. In the one-at-a-time context, SASP requires that for each role, every player is offered the same price.¹⁵ Second, the amount collected from all players should equal the amount paid out in subsidies; this is called *budget balance*. The third property is *voluntary participation*: a player, knowing his preferences and the announced charges and subsidies, will have a positive expected value from participating in the game.

Unfortunately, no efficient payment mechanism can simultaneously satisfy all three desiderata. It is possible, however, to satisfy any pair (see Figure 7).

----Figure 7 approximately here----

Our first general result is a negative one, showing that all three goals can not be achieved in either the batch mode or one-at-a-time game. Intuitively, to satisfy budget balance and voluntary participation, it may be necessary to charge beneficiaries of evaluations the full value that they gain, but this charge may differ from consumer to consumer, and hence SASP cannot be satisfied.

Proposition 1: In the efficient provision of evaluations, SASP, budget balance and voluntary participation can not in general be satisfied simultaneously.

Proof: We proceed by counterexample. Suppose that, in a batch-mode game with three players, the social optimum is for player A to consume in round 1 at a cost of 7 (expected payoff -7) in order to provide benefits of 5 and 3 to players B and C

respectively, who will consume in round 2 should A find the product Good. Player A must be paid at least 7, but C will not participate at a cost greater than 3. If, following SASP, B pays the same as C, the total revenue is at most 6, which violates budget balance.

The counterexample for the one-at-a-time game employs the following parameter values: $p = .5$; $g = b = .9$; $r_D = s_D = -65$; $r_E = +225$, $s_E = -813 \frac{8}{9}$; $r_F = 325$, $s_F = -913 \frac{8}{9}$.¹⁶ The social optimum has D consume, then E if and only if the evaluation is Good, then F if and only if both evaluations are Good. It costs D 65 to evaluate the product. Figure 8 lists the expected payoffs to players E and F at the second and third nodes in the tree; the first payoff is computed at the start of the tree and the second payoff, in parentheses, is conditional on reaching the given node. In the social optimum, the expected payoffs are -65 for D, 60 for E and 41 for F.

F will pay an up-front charge of at most 41 for the right to consume at node 3 (equivalent to an actual charge of 100 conditional on the first two evaluations being positive). For E to prefer node 2 rather than F's role, the up-front charge for node 2 must be at least 22 less. Therefore, it can be at most 19. Total revenue, then, is at most $60 = 41 + 19$. It is impossible to satisfy budget balance, since D requires 65 to consume at the root node.

----Figure 8 approximately here----

Efficient Allocation With an Unbalanced Budget. We first demonstrate the broker's ability to produce an efficient allocation while maintaining voluntary participation and SASP. In the batch mode game, the broker seeks to maximize the expected net benefits from the initial evaluations. If the optimum calls for n^* evaluations, the broker wishes to arrange for the n^* people who would benefit the least from these evaluations to provide them. The broker can achieve that by announcing the price $x_{(n^*+1)}(n^*)$ as the subsidy to first-round evaluators. This price is calculated as the subsidy required to make the n^*+1 st player indifferent between evaluating the product immediately and waiting to receive the n^* evaluations.

Proposition 2: In the two-stage batch mode game, if the broker offers the price $\pi^* = x_{(n^*+1)}(n^*)$ to all players who evaluate in the first round, the efficient allocation is a Nash equilibrium.

Proof: See appendix.

By announcing the price π^* , the broker implements the natural outcome of an n^*+1 st price auction—the generalization of a Vickrey second price auction—without actually conducting the auction. A Vickrey auction is a special case of a Groves-Clarke payment mechanism (William Vickrey 1961; Edward H. Clarke 1971; Theodore Groves 1973). Frequently, such a mechanism is the only way to achieve the first-best outcome in dominant strategies (Jerry Green and Jean-Jacques Laffont 1977).¹⁷ In our analysis, since

the broker knows the pool of (r_i, s_i) values, it can identify the optimal quantity n^* without asking the players to report their intensity values. By announcing the market-clearing price, the broker elicits just enough information to identify the set of players who should evaluate immediately. In contrast, an auction would also reveal the value of the n^* evaluations to each player. Another advantage of announcing the market clearing price is that it can not be inflated by collusive behavior, as the price could be in a full Groves-Clarke mechanism.

Our market -- in effect an auction -- is unusual because it incorporates opportunity costs, which depend on the number of other players who will accept the first-round subsidy. Hence, the broker announces the quantity n^* in addition to the price π^* . The announcement is helpful even through the equilibrium is unique. Without an announced quantity, players may have trouble identifying the equilibrium (and efficient) allocation. Suppose, for instance, that a single evaluation is worth 6 to player A and two evaluations are worth 9 to her. If the broker announces only the price 7, A cannot tell whether the equilibrium has her as the only first-round evaluator or whether the equilibrium has her wait while two other players provide first-round evaluations. With the announcement of a quantity as well as a price, players can easily choose actions that lead to the equilibrium allocation.

In the one-at-a-time game, the optimal quantity is contingent on the outcome of early evaluations. The broker announces an evaluation tree and a vector giving the subsidy for each of m different roles. The broker chooses the set of available roles to correspond to the roles assigned to the players in the social optimum. The broker assures everyone that the subsidies will induce a general equilibrium outcome in which exactly one player will accept the subsidy for each role, and the available roles combine to cover all the nodes of the tree that call for an evaluation. This gives players sufficient information about others' actions to make rational choices. Each player acts in response to the price vector by choosing one role and committing to that role.

Proposition 3: In the one-at-a-time game, there is a vector of subsidies on roles that induces each player to evaluate at exactly the nodes called for in some efficient allocation.

Proof: See appendix.

This vector of subsidies achieves the optimal outcome through general equilibrium pricing, where each role (set of nodes) in the evaluation tree represents a different good available to consumers.¹⁸

Efficient Allocation With Up-Front Subscription Fees and Ex ante Voluntary Participation. A broker can structure its charges and payments so that players pay an up-front subscription charge. If the players do not know their intensities at the time they must subscribe, but the broker knows the pool of intensity values, we demonstrate that the broker can induce an efficient allocation while maintaining SASP and paying out the exact amount of the subscription fees, thus balancing the budget. Since the expected social value of the game is always non-negative, and players do not know their intensity values initially, each expects an equal share of the social value. If the budget will balance,

then each player is willing to pay the up-front subscription fee. We call this ex ante voluntary participation. We do not have full voluntary participation because, if players were allowed to find out their intensities before paying the subscription fee, some might refuse to do so.

In the batch mode, each player is charged an equal share of the subsidies to be doled out. The analogous approach to the one-at-a-time game will not work, because the total subsidy may depend on the evaluations received. It is possible, however, to adjust the subsidies to maintain incentive compatibility and to assure that the sum is the same on every branch of the tree. This allows the broker to charge a fixed subscription fee up front.

Proposition 4: In both the two-stage batch mode game, and the one-at-a-time game, with up-front subscription fees and subsequent subsidies, the broker can induce the efficient allocation while satisfying budget balance and SASP.

Proof: Available from authors.

Efficient Allocation with Price Discrimination. Relaxing the SASP requirement is only helpful if the broker can identify the players' types. Then the broker can balance the budget by demanding higher payments from the players who benefit more; in effect, the broker price discriminates. It also permits the broker to levy exorbitant charges (effectively reductions in price) for actions that deviate from the social optimum. Unlike our previous mechanisms, this requires the broker to know which players have which intensity values, not merely the distribution within the pool of values.¹⁹

Proposition 5: In either the batch mode or the one-at-a-time game, efficient allocation can be achieved while satisfying the budget balance and voluntary participation constraints if the broker knows the players' types.

Proof: Available from authors.

E. Expanded Model—Players Differ in Taste and Evaluation Skills

The base model assumes that the players differ only in intensities of preference. We now expand and generalize the model to allow the players to differ from each other in three more substantive ways. First, the players may differ in their evaluation expertise; some are more informative than others. Second, the informativeness of each evaluator may vary for different recipients; thus, suburbanites and yuppies may each inform their own class about automobiles, just as economists are more helpful in recommending bulletin board messages to economists, and computer scientists to other computer scientists. These two features make the evaluators differentially informative. Third, tastes may be correlated. For example, the evaluations of employees of a firm regarding a product used in the firm's production process may be highly correlated. A second evaluation by an employee of that firm is likely to agree with the first evaluation and is thus less valuable than an independent evaluation of equal expertise.

Formally, we assume that there are T classes of people and that the classification of each individual is known by the broker. The information about a product is now a T -tuple $(\rho_1, \rho_2, \dots, \rho_T)$, where entry ρ_t represents the probability that a person in class t will like the product. An evaluation from a player in one class may provide information to players in other classes, an update to the entire T -tuple $(\rho_1, \rho_2, \dots, \rho_T)$. The informativeness of each evaluator in a class is always the same, meaning that a particular evaluation from any player in a class generates the same update.²⁰ Evaluators in different classes may have different functions. Moreover, individuals in each class may have correlated evaluations, meaning that one evaluation from an evaluator in that class may cause a large update in the assessment of $(\rho_1, \rho_2, \dots, \rho_T)$, but that additional evaluations from players in that class would offer smaller incremental impacts. Individuals in the same class may vary in intensities of preference, as indicated by their r_i and s_i values. The broker knows the actual distribution of (r_i, s_i) values within each class of people, but not necessarily who possesses which values.

Once individuals differ in more than their intensities, however, we can no longer induce players to choose behavior consistent with the efficient allocation through the simple SASP-supporting mechanisms developed earlier.

Proposition 6: It is not always possible to achieve efficient allocation and SASP when individuals fall into more than one class in terms of expertise, informativeness for different recipients, or correlation with each other's evaluations.

Proof: Suppose it were possible to attain efficient allocation and SASP. The following two counterexamples yield contradictions.

Expertise: Suppose that individuals fall into one of two homogeneous classes, either "senior" or "junior": the senior players have more intense preferences (i.e., larger in absolute value whether positive or negative) and are more informative. Assume that it is optimal to acquire a single first-round evaluation by a senior player, which is very informative, and none by junior players. Given higher preference intensities, it will be more expensive to get a senior player to evaluate. Any subsidy that lures a senior will lure every junior. Hence, there is no price-supported Nash equilibrium that induces the efficient production of evaluations.

Informativeness for Recipients or Correlation: Suppose that there are two classes of people, A and B, differentiated either by informativeness (each is perfectly informative for her class, but of no relevance to the other class) or by correlation (the first evaluation from either class is equally informative to both types, but evaluations within a class are perfectly correlated). Assume that members of A have substantially lower intensities of preference and hence lower net costs of evaluating. The optimum for either situation is to induce one A and one B to consume. But for either situation the price of getting a B to consume may be well above the price for the second A; hence the optimum is not attainable without separate prices for As and Bs. Q.E.D.

Fortunately, if we are permitted to pay different prices to individuals in different classes, our earlier results on achieving efficiency alongside other requirements can still be attained. We call this condition "type-SASP." It requires that players within a given class receive (or pay) the same price for providing (or consuming) the same service. All of our earlier results, both negative and positive, generalize to meet the three heterogeneity conditions of our expanded model if we replace SASP with type-SASP, assuming that the broker can identify each player's type.

Proposition 2 generalizes naturally to yield efficient allocations with voluntary participation and type-SASP. Formally, we now represent an allocation of evaluators in the two-stage batch mode as a vector (n_1, n_2, \dots, n_T) , where n_t indicates the number of evaluators from class t in the first stage. We define $x_i(n_1, n_2, \dots, n_T)$ as the net value to player i of waiting until the second round rather than consuming immediately, given that (n_1, n_2, \dots, n_T) will be produced in the first round. Since the number of possible allocations is finite, the broker can clearly calculate an efficient one, $(n_1^*, n_2^*, \dots, n_T^*)$. The broker helps to coordinate the actions of individuals by announcing the efficient allocation as well as the price, thereby facilitating the players' decisions of whether to evaluate.

Proposition 7: In the two-stage batch mode game, if the broker offers the price $\pi_t^* = x_{(n_t^*+1)}(n_1^*, n_2^*, \dots, n_T^*)$ for first-round evaluations by individuals in class $t \in (1, 2, \dots, T)$, then the efficient allocation is a Nash equilibrium.

Proof sketch: The proof follows, on a class-by-class basis, the exact logic used in the proof of Proposition 2. Given the fixed set of evaluations in the first period from other classes of individuals, the individuals in class t will self-select themselves to produce n_t^* evaluations, from those who would benefit least from waiting.

In the one-at-a-time game, it is also necessary to offer a separate set of prices for each class of individuals to achieve type-SASP and voluntary participation. The broker calculates the social optimum and then induces the individuals in a given class to allocate themselves efficiently. As in the base model, the broker announces an evaluation tree to everyone but here the broker offers a restricted set of options and subsidies to each class of individuals: the broker only offers individuals in class t an option to perform a role in the tree if the social optimum assigns an individual from that class to that role.²¹

Proposition 8: For the efficient allocation in the one-at-a-time game with multiple classes of individuals, there is a vector of subsidies that induces each player to evaluate at precisely the nodes called for in that allocation. That vector of subsidies consists of a set of T subsidies for each role, one for each class of players.

Proof sketch: See appendix.

Propositions 7 and 8 show that, in either the batch-mode or one-at-a-time game, the broker can use pricing schemes to induce the social optimum while satisfying type-SASP and voluntary participation. It is then possible to transform those pricing schemes into

alternate pricing schemes that achieve ex ante voluntary participation, type-SASP and budget balance while inducing the social optimum, just as in the case of single player type.²² In the batch mode game, the broker imposes a fixed subscription fee of

$\frac{n_t^* \pi_t^*}{\text{no. of people in class } t}$ to each person in class t . The method of determining optimal fees in

the one-at-a-time game is similar, but requires additional care to assure that the subscription fees collected from each type of player are paid out in full to players of that type on every path through the contingency tree.

The model with multiple classes of players makes computation of efficient allocations more complex, because an evaluation from a player in one class provides different information to different classes. Multiple classes of players, however, pose no fundamental problems in inducing optimal behavior through prices. Since we do not aim to avoid envy between groups, the prices need only induce optimal player choices within each class, thereby reducing the multiple class pricing problem to pricing for a single class.

IV. Conclusion

Personal experience with products is enormously powerful in informing consumers' decisions. Personal recommendation networks and consumer publications help consumers learn from the experience of others. The explosive growth of the information superhighway, combined with the enormous data handling and computational capabilities of computers, has dramatically enhanced the ability of consumers to benefit from the experience of others. We study the potential of a market for evaluations --a mechanism for eliciting, sharing, and paying for information-- to facilitate such learning. A computer-based and network-linked market for evaluations can provide timely guidance to consumers who act on a voluntary basis.

Although computers offer an inexpensive way to distribute evaluations of products, two difficulties remain. The first is the classic problem of underprovision found with many public goods. Products may be expensive to purchase, evaluate, or both. If evaluations are distributed freely, those who can benefit from them may try to free ride on other consumers' willingness to pay. Insufficient money is collected to subsidize the early evaluators.

The second difficulty is more subtle, and more apparent with products whose evaluation costs are lower, such as journal articles. A vast amount of evaluative information is already held by individuals, but it almost always remains privately held. If such evaluations are instead distributed to guide consumers, the benefits of waiting before consuming soar. Those who now consume early may choose to wait once evaluation-sharing technology is widely available.

A market for evaluations may stumble over either of these difficulties. We have shown, however, that carefully designed mechanisms can satisfy any pair of three desirable properties: voluntary participation, budget balance, and an equal treatment

condition labeled Type-SASP. Type-SASP requires that all individuals of a type are permitted to choose among the Same Actions at the Same Price.

Any such mechanisms to efficiently share evaluations would capitalize on the computational power of computers, and the virtually costless communication among them. A software agent, which can meet extreme computational demands, is an important representative of each person in the marketplace. Human effort is required only to purchase and evaluate products, not to coordinate the production and exchange of evaluations. In the not too distant future, people will consult personalized scoreboards that highlight possible items to purchase or peruse. Moreover, a market for evaluations will allow many people to be part-time professional critics, charging for informative early evaluations on some products, while paying for evaluations of others. Such a market creates value by arranging the optimal quantity and order of evaluators.

Appendix

Proposition 2: In the two-stage batch mode game, if the broker offers the price $\pi^* = x_{(n^*+1)}(n^*)$ to all players who evaluate in the first round, the efficient allocation is a Nash equilibrium.

Proof: Consider a player i who is assigned to the waiting group in the efficient allocation, for a payoff $b_i(n^*)$. He has one of the largest $x_i(n^*)$, so $b_i(n^*) = a_i + x_i(n^*) \geq a_i + x_{(n^*+1)}(n^*) = a_i + \pi^*$. Hence, player i would rather not switch to consuming right away because the subsidy is less than the value of consuming n^* evaluations.

Now consider a player j who is assigned to consume. By waiting instead, j can get $b_j(n^* - 1)$, since there will still be $n^* - 1$ other consumers. But $b_j(n^* - 1) \leq b_j(n^*) = a_j + x_j(n^*) \leq a_j + x_{(n^*+1)}(n^*) = a_j + \pi^*$. The first inequality holds because it is preferable to have more information. The second holds because j was assigned to the consuming group, revealing that she must have one of the smallest $x_i(n^*)$ values. The subsidy exceeds player j 's net benefit of waiting, so she would prefer not to wait.

Hence, the efficient allocation is a Nash equilibrium. Note that voluntary participation is satisfied since players always have the option of waiting with non-negative value. Q.E.D.

Proposition 3: In the one-at-a-time game, there is a vector of subsidies on roles that induces each player to evaluate at exactly the nodes called for in some efficient allocation.

Proof: We adapt the technique that Koopmans and Beckmann (1957) employed to find the optimal allocation of m people to m separate positions; in operations research, this is known as "The Assignment Problem". First the broker computes an optimal allocation tree and assignment of players to nodes of the tree. Each player's assignment to a combination of nodes defines one of the m roles for which the broker will offer prices.²³

For each role j and player i , define $x_{ij} = 1$ if player i is assigned to role j , and $x_{ij} = 0$ otherwise. Define u_{ij} to be player i 's expected utility from committing to evaluate at the nodes comprising role j (and otherwise waiting until the end of the evaluation tree to decide whether to consume the product).

The optimality condition means that no other tree yields a higher sum of expected payoffs to the players. It also means that no other assignment of players to roles that cover the chosen allocation tree yields a higher sum of expected payoffs. Thus, the chosen

assignment of players solves the maximization problem

$\max \sum_{i=1}^m \sum_{j=1}^m u_{ij} x_{ij}$, subject to the constraints:

1. $\forall i \sum_j x_{ij} = 1$. (Each player is required to select exactly one role.)
2. $\forall j, \sum_i x_{ij} = 1$. (Each role is allocated to exactly one player.)

To match the standard linear programming format, however, we consider a related linear programming problem. Allow $x'_{ij} \in [0,1]$ rather than requiring discrete values 0 and 1.

Define $K = -\min_{i,j} u_{ij} + 1$ and $u'_{ij} = u_{ij} + K$ so that all of the utility values are positive. In

addition, relax all of the equality constraints to be inequalities. The related problem, then, is

$\max \sum_{i=1}^m \sum_{j=1}^m u'_{ij} x'_{ij}$ subject to the constraints:

3. $\forall i \sum_j x'_{ij} \leq 1$. (Each player can select no more than one role.)
4. $\forall j, \sum_i x'_{ij} \leq 1$. (Each role is allocated at most once.)

Given that every player's payoff for every role is now positive, constraint 3 must be binding. Given that constraint 3 is binding, the total x'_{ij} weight is m , and hence constraint 4 must be binding for every role. The solution to the dual problem gives shadow prices for each constraint in the primal problem. The set of shadow prices for constraint 4 is a set of prices p'_j for nodes j that induces players to choose optimal allocation values x'_{ij} as a decentralized pricing equilibrium. Note that there may be a range of solutions to the dual problem. Within this range, each set of prices achieves the desired property of inducing an optimal allocation as a pricing equilibrium.²⁴

Birkhoff and von Neumann showed that every mixed allocation assignment is the convex combination of pure allocation assignments (this result is presented as Theorem 8.6 in Schrijver, 1986). Thus, if a mixed allocation assignment is an optimal solution to the primal problem, then each of its component pure allocations is also an optimal solution to that problem.²⁵ Any set of prices that induces the mixed allocation as an equilibrium would also induce any of the component pure allocations as an equilibrium (since the players must be indifferent between all options chosen with positive probability in a mixed allocation equilibrium). Without loss of generality, then, we assume the players choose an optimal allocation with x'_{ij} values restricted to 0 and 1.

We convert the shadow prices p'_j into subsidies $s_j = -p'_j + K$ that are large enough to sustain voluntary participation with the unadjusted utilities u_{ij} . These subsidies still induce the same allocation as an equilibrium, because the relative values of the roles are unchanged by adding the constant subsidy K . The allocation induced must also be optimal for the original maximization problem, since in any allocation each player will receive exactly K more in the second maximization problem than in the first, leaving the relative values of complete allocations unchanged. Q.E.D.

Proposition 8: For each optimal allocation in the one-at-a-time game with multiple classes of individuals, there is a vector of subsidies that induces each player to evaluate at precisely the nodes called for in that allocation. That vector of subsidies consists of a set of T subsidies for each node, one for each class of players.

Proof sketch: This proof follows the same logic as that for Proposition 3. The broker solves a separate linear programming problem for each class of players (only nodes designated for that player's type are available). As in the proof of Proposition 3, the broker can produce a set of subsidies for each class that induces the assignment of individuals in that class to their appropriate nodes in the full evaluation tree. The combination of choices of all classes of individuals produces the social optimum. Q.E.D.

References

- Banerjee, Abhijit. "A Simple Model of Herd Behavior." Quarterly Journal of Economics, August 1992, 107(3), pp. 797-818.
- Banerjee, Abhijit. "The Economics of Rumours." Review of Economic Studies, April 1993, 60(2), pp. 309-327.
- Bikhchandani, Suchil, Hirshleifer, David and Welch, Ivo. "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades." Journal of Political Economy, October 1989, 100(5), pp. 992-1026.
- Clarke, Edward H. "Multipart Pricing of Public Goods." Public Choice, 1971, 11(1), pp. 17-33.
- Coase, Ronald. "The Problem of Social Cost." Journal of Law and Economics, 1960, 3, pp. 1-44.
- Debreu, Gerard. "New Concepts and Techniques for Equilibrium Analysis." International Economic Review, 1962, 3, pp. 257-273.
- Fudenberg, Drew and Tirole, Jean. Game Theory. Cambridge, MA: MIT Press, 1991.
- Glazer, Jacob and Ma, Ching-To Albert. "Efficient Allocation of a Prize -- King Solomon's Dilemma." Games and Economic Behavior, September 1989, 1(3), pp. 222-233.
- Green, Jerry and Laffont, Jean-Jacques. "Characterizations of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods." Econometrica, March 1977, 45(2), pp. 427-438.
- Green, Jerry and Laffont, Jean-Jacques. Incentives in Public Decision-Making. New York: North-Holland Publishing Co., 1979.
- Groves, Theodore. "Incentives in Teams." Econometrica, July 1973, 41(4), pp. 617-631.
- Hill, Will, Stead, Larry and Rosenstein, Mark. "Recommending and Evaluating Choices in a Virtual Community of Use." Proceedings of CHI 95 Conference on Human Factors in Computing Systems. Denver, ACM, 1995, pp. 194-201.
- Koopmans, Tjalling C., and Beckmann, Martin. "Assignment Problems and the Location of Economic Activities." Econometrica, January 1957, 25(1), pp. 53-76.
- Leonard, Herman B. "Elicitation of Honest Preferences for the Assignment of Individuals to Positions." Journal of Political Economy, June 1983, 91(3), pp. 461-479.

Mackie-Mason, Jeffrey K. and Varian, Hal. "Economic FAQs About the Internet." Journal of Economic Perspectives, Summer 1994, 8(3), pp. 75-96.

Prendergast, Canice. "A Theory of Yes Men." American Economic Review, September 1993, 83(4), pp. 757-770.

Resnick, Paul and Varian, Hal. "Recommender Systems." Communications of the ACM, 1997, 40(3), pp. 56-58.

Samuelson, Paul. "Pure Theory of Public Expenditure." Review of Economics and Statistics, 1954, 36(4), pp. 387-389.

Samuelson, Paul. "Diagrammatic Exposition of a Theory of Public Expenditures." Review of Economics and Statistics, 1955, 37(4), pp. 350-356.

Schrijver, Alexander. Theory of Linear and Integer Programming. New York: Wiley, 1986.

Takayama, Akira. Mathematical Economics. Cambridge: Cambridge University Press, 1985.


Van Zandt, Timothy. "Information Overload in a Network of Targeted Communication." TM 11210-900627-05, AT&T Bell Labs, June 27, 1990.

Varian, Hal. "Equity, Envy, and Efficiency." Journal of Economic Theory, 1974, 9(1), pp. 63-91.


Vickrey, William. "Counterspeculation, Auctions and Competitive Sealed Tenders." Journal of Finance, 1961, 16(1), pp. 8-37.

1. Process intensities
2. Determine prices
3. Begin Round 1

3a. Some players consume product and send evaluations to broker



3b. Broker distributes evaluations to remaining players




4. Begin Round 2: remaining players decide whether to consume product
5. Make payments


Figure 1. The Batch Mode Game

1. Process intensities
2. Determine prices
3. Begin round

3a. One player consumes product and sends evaluation to broker



3b. Broker distributes evaluation to remaining players



4. Begin another round if any player is willing to evaluate
5. Make payments

Figure 2. The One-at-a-Time Game

<u>Player A payoffs</u>		<u>Player B payoffs</u>		<u>Player C payoffs</u>	
Good	Bad	Good	Bad	Good	Bad
12	-24	12	-24	1000	-1000

Figure 3. Payoffs in a Three-Person Game




Figure 4. Efficient Allocation: One-at-a-Time Example

Game	<u>Player A payoffs</u>		<u>Player B payoffs</u>		<u>Inefficiency</u>
	Good	Bad	Good	Bad	
i)	10	-12	10	-12	underprovision
ii)	10	-12	40	-20	wrong ordering
iii)	12	-10	12	-10	wasteful surplus claiming

Figure 5. Payoffs from Consuming in Three Games

		<u>Player B Action</u>										
		Consume		Wait		Consume		Wait		Consume		Wait
<u>Player A Action</u>	Consume	-1,-1	-1,5 *	-1,10	-1,20 *	1,1	1,6 *					
	Wait	5,-1 *	0,0 E	5,10 E	0,10	6,1 *	1,1 E					
		(i)			(ii)			(iii)				
		underprovision		wrong ordering		wasteful surplus claiming						

Figure 6. Expected Payoffs for the Games in Normal Form

Desiderata	Is efficiency possible?
Budget balance + SASP + voluntary participation	no (Proposition 1)
SASP + voluntary participation	yes (Propositions 2 & 3)
Budget balance + SASP	yes (Proposition 4)
Budget balance + voluntary participation	yes (Proposition 5)

Figure 7. Possible Traits for an Efficient Pricing Scheme

Node	Pr(reaching node)	Player E payoff	Player F payoff
2 nd : after Good	0.5	60 [120]*	10 [20]*
3 rd : after Good, Good	0.41	82 [200]*	41 [100]*

* The first value in each cell is the expected value taken at the root of the tree, before anyone evaluates the product. The second value [in brackets] is the expectation conditional on reaching and having that player consume at the node.

Figure 8. Payoffs for One-at-a-Time Counterexample in Proposition 1

* Avery: Harvard Kennedy School of Government, Cambridge, MA 02138; 617-496-4063; chrisa@ksg1.harvard.edu. Resnick: University of Michigan School of Information, 314 West Hall, Ann Arbor, MI 48109; 734-647-9458; presnick@umich.edu. Zeckhauser: Harvard Kennedy School of Government, Cambridge, MA 02138; 617-495-1174; richard_zeckhauser@harvard.edu. Financial support from the corporate sponsors of the MIT Center for Coordination Science (Resnick) and the Decision, Risk, and Management Science Division of the National Science Foundation (Zeckhauser and Avery) is gratefully acknowledged, as is hospitality of the MIT Economics Department (Zeckhauser). Miriam Avins, Sushil Bikhchandani, Jerry Green, John Pratt, Herb Scarf, Hal Varian, and anonymous referees provided helpful comments. Participants in seminars at the Harvard Economics department, the Kennedy School, the MIT Sloan School, the Harvard Law School, the Yale Economics department, the University of Michigan Economics Department, the Harvard Conference on the Internet and Society, and the Berkeley workshop on collaborative filtering helped to refine the ideas and presentation.

¹ Users of these services have suggested that matching people whose evaluations correlate strongly could be an effective way to run a dating service (Will Hill, Larry Stead and Mark Rosenstein 1995).

² We refer to a generic player as female. When two players are involved, we refer to the second as male. Computerized agents and brokers are neuter.

³ Some evaluations, such as those of stocks, whose prices may respond to demand, or restaurants, which could become crowded, are not public goods. Hal Varian (personal communication) has pointed out that with information goods, unlike true public goods, there can be a credible threat of exclusion. We rely on this threat in some of our pricing schemes: players may be forced to pay if they wish to access others' evaluations.

⁴ However, if one can consume a product several times without reducing its value, and one's own evaluation does not provide full information about one's interest in the product, there will be some value to subsequent evaluations by others. For example, if the product to evaluate is a vacation in Paris, an evaluator who did not enjoy the city on her first visit may be influenced to try again if subsequent evaluators provide rave reviews.

⁵ In this analysis all decisions are personal. In subsequent work, we expect to study situations where individuals provide evaluations for a collective decision, say whether an organization should pursue an initiative. In eliciting evaluations for such decisions, computers may have a valuable advantage in being able to hide information, revealing it only when doing so benefits the principal who provides it.

⁶ As one referee observed, “The problem of bounded rationality is diminished by the use of cheap silicon to do most of the processing.”

⁷ For simplicity, we often refer generically to *players* without explicitly acknowledging the computerized agents that make recommendations to or decisions on behalf of the human players.

⁸ Universal transmission of evaluations is critical to efficiency; Banerjee (1993) shows that when information is exchanged haphazardly, individuals can come to disbelieve honest evaluations.

⁹ We capitalize “Good” and “Bad” to indicate perceptions of players. Of course, players may not perceive products as purely Good or purely Bad, but additional evaluation outcomes would add complexity to our model without yielding new insights.

¹⁰ For expositional ease, we refer to a single efficient allocation as “the efficient allocation,” although it is possible that different allocations could be tied. In this case, the broker selects one arbitrarily.

¹¹ If evaluation effort were variable, pricing schemes could induce effort by rewarding players for evaluations that matched those of others. To do so, however, would encourage both collusion and a reluctance to state idiosyncratic opinions, producing a breed of “Yes Men” (Canice Prendergast 1993). We expect to return to the effort inducement problem in future work. One intriguing possibility for deterring collusion is that individuals could be rewarded for matching others, but all would be punished for a degree of agreement far beyond the statistical norm.

¹² An evaluation tree simply specifies the nodes at which evaluations will be made, while an allocation specifies both a tree and which player will evaluate at each node.

¹³ Interestingly, the addition of a single player can increase the number of evaluators by more than one on average in the social optimum, even though all players are identical. For example, say there is a single player who barely prefers not to consume the product. With the addition of a second player, it could be that player 2 should now evaluate immediately and that player 1 should also buy the product if player 2 likes it.

¹⁴ For two-person games, the batch mode and one-at-a-time processes are nearly identical. The batch mode permits both players to consume in the first round, but that is never an optimal outcome.

¹⁵ SASP is related to a requirement of anonymity, which is a common objective in studies of mechanisms to optimally provide public goods (Jerry Green and Jean-Jacques Laffont 1979). An anonymous provision mechanism permits price discrimination based on preferences, but not based on identity; if two consumers interchange their preferences, their payments are interchanged as well. SASP is a stronger requirement, because it

requires that consumers with different preferences who consider the same action be offered the same price. In essence, SASP is a version of a no-envy condition far removed from its usual fair division context (Hal Varian 1974; Jacob Glazer and Ching-To Albert Ma 1989).

¹⁶ The fractional s values were chosen to produce integers for expected values.

¹⁷ Much recent literature in mechanism design forgoes dominant strategies and implements desirable outcomes using Bayesian assumptions (Drew Fudenberg and Jean Tirole 1991, p. 270-284).

¹⁸ Paul Samuelson (1954, 1955) developed the connection between optimal public goods provision and general equilibrium pricing. Gerard Debreu (1962) provides more general efficiency results for general equilibrium subject to restricted consumption sets for consumers.

¹⁹ Given knowledge of the pool of values, the fanciful world of mechanism design offers a “shoot-them-all” procedure to elicit individuals’ honest intensity values (r_i, s_i) . The broker requests individual reports and threatens to call off the entire evaluation game if the set of reports does not match the known pool. Since everyone prefers that the game go on, honest reporting, a focal point, will be the natural equilibrium.

²⁰ In Section II, we illustrated the updating process for a model with only one class of evaluators and two product states (good and bad). There, the updating of ρ required a straightforward application of Bayes' Rule. Here, to account for taste and expertise differences, we would need more than two product states and a more complicated application of Bayes' Rule.

²¹ This yields an equivalent outcome to a pricing scheme that includes a large penalty for choosing a role not in the restricted set of options.

²² Note that the broker can achieve voluntary participation and budget balance, just as in Proposition 4, regardless of whether all individuals fall into one class or multiple classes.

²³ It is an open question whether the linear programming techniques used in this proof can be extended to the case where players are offered a choice of all possible combinations of nodes in the tree, restricted only to choosing no more than one position on each path through the tree. Unfortunately, the linear programming formulation for arbitrary choices of nodes in a binary allocation tree, rather than pre-defined roles, admits the possibility of mixed strategy solutions that have greater value than any pure assignment of players to roles. In such cases, the optimal mixed allocations can not be realized through any actual evaluation process, and no set of prices on nodes can cause the market to clear with a pure assignment of each node to exactly one player.

²⁴ Herman Leonard (1983) extends the analysis of Koopmans and Beckmann to consider a one-dimensional assignment problem where individual preferences are privately known. Among the solutions to the dual problem in that case, the solution that minimizes the sum of the shadow prices corresponds to Groves-Clarke charges and also induces honest revelation of the individual values.

²⁵ The difficulty in extending Proposition 3 from prices on roles to prices on individual nodes arises with this step. A feasible mixed allocation of nodes can be represented as the convex combination of pure allocations of nodes, but it is not guaranteed that each of these pure allocations is itself feasible (e.g., a single player may be required to evaluate at more than one node on a particular path).